Issue coedited by: Universitat Oberta de Catalunya and Leonardo/The International Society for the Arts, Sciences and Technology

artnodes
E-JOURNAL ON ART, SCIENCE AND TECHNOLOGY

https://artnodes.uoc.edu

ARTICLE

NODE «AI, ARTS & DESIGN: QUESTIONING LEARNING MACHINES»

«Creative AI: From Expressive Mimicry to Critical Inquiry»

Angus G. Forbes
University of California, Santa Cruz

Date of submission: April 2020
Accepted in: June 2020
Published in: July 2020

Recommended citation

Forbes, Angus G. 2020. «Creative AI: From Expressive Mimicry to Critical Inquiry». In: Andrés Burbano; Ruth West (coord.) «AI, Arts & Design: Questioning Learning Machines». *Artnodes*, Nº.26: 1-10. UOC, Leonardo/ISAST. [Accessed: dd/mm/yy]. https://doi.org/10.7238/a.v0i26.3370

SOME RIGHTS RESERVED

The texts published in this journal are – unless otherwise indicated – covered by the Creative Commons International Attribution 4.0 International license. The full text of the license can be consulted here: http://creativecommons.org/licenses/by/4.0/

Abstract

The nascent field of what has come to be known as "creative AI" consists of a range of activities at the intersections of new media arts, human-computer interaction, and artificial intelligence. This article provides an overview of recent projects that emphasise the use of machine learning algorithms as a means to identify, replicate, and modify features in existing media, to facilitate new multimodal mappings between user inputs and media outputs, to push the boundaries of generative art experiences, and to critically investigate the role of feature detection and pattern identification technologies in contemporary life. Despite the proliferation of such projects, recent advances in applied machine learning have not yet been incorporated into or interrogated by creative AI projects, and this article also highlights opportunities for computational artists working in this area. The article concludes by envisioning how creative AI practice could include delineating the boundaries of what can and cannot be learned by extracting features from artefacts and experiences, exploring how new forms of interpretation can be encoded into neural networks, and articulating how the interaction of multiple machine learning algorithms can be used to generate new insight into the intertwining sociotechnical systems that encompass our lives.

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

Keywords
creative AI, machine learning, generative art, new media art

IA creativa: De la mímica expresiva a la investigación crítica

Resumen
El incipiente campo de lo que se conoce como "IA creativa" consiste en una serie de actividades en las intersecciones de las nuevas artes mediáticas, la interacción persona-computadora y la inteligencia artificial. Este artículo proporciona una descripción general de proyectos recientes que enfatizan el uso de algoritmos de aprendizaje automático como un medio para identificar, replicar y modificar características en los medios existentes, para facilitar nuevas asignaciones multimodales entre las entradas del usuario y las salidas de los medios, para ampliar los límites en las experiencias del arte generativo e investigar críticamente el papel de las tecnologías de detección de características e identificación de patrones en la vida contemporánea. A pesar de la proliferación de proyectos de este tipo, los avances recientes en el aprendizaje automático aplicado aún no han sido incorporados o cuestionados por proyectos creativos de IA, y este artículo también destaca las oportunidades para los artistas computacionales que trabajan en esta área. El artículo concluye imaginando cómo la práctica creativa de IA podría incluir e delinear los límites de lo que se puede y no se puede aprender extrayendo características de artefactos y experiencias, explorando sobre cómo las nuevas maneras de interpretación pueden codificarse en redes neuronales y definiendo cómo la interacción de múltiples máquinas con algoritmos de aprendizaje se pueden utilizar para generar una nueva visión de los sistemas sociotécnicos entrelazados presentes en nuestras vidas.

Palabras clave
IA creativa, aprendizaje automático, arte generativo, arte de nuevos medios

The Paper:

1. Introduction

One advantage in using machine learning to extract meaning from data is that it lets the researcher sidestep the need to articulate the low-level details contained in the data, which can be difficult to tease out and hard to define. How do you describe what films you like? It is easier to provide a training set of films that you've rated and let the algorithm discover what features highly related films have in common (Hallinan & Striphas 2016). How do you capture the nuances in meaning when translating a phrase from one language to another? It is more accurate to provide the machine learning system with a vast amount of data in order to infer these subtleties without requiring formal semantics (McCann et al. 2017). How do you best describe the special characteristics of a person so that they can be distinguished from others in an image, no matter where the image was taken, what pose they are in, or what they are wearing? State-of-the-art recognition systems do not require any description whatsoever, only a sufficient number of examples that the deep learning network extrapolates from and encodes as weights within its hidden layers (Taigman et al. 2014, Sun et al. 2014). What strategy do you use to articulate the rules that define an artist's expressivity? Style transfer algorithms effortlessly let you transform any image or video into an impressionist painting, using even a single image of a painting to automatically find the characteristic elements of the artist's style (Gatys et al. 2016).

For many applications, deep learning neural networks are the most effective method to identify useful features in datasets and to use them to interpret new data with similar content. In addition to choosing the most computationally efficient architecture or parameters, a main focus of the data analyst using them becomes to define the space of interpretation by choosing the dataset that represents that space, by selecting an appropriate loss function for training the network, and by deciding what outputs can be returned when querying the network. Learning to interpret the data occurs through a process of encoding hierarchies of features that indicate whether a particular input (or part of that input) belongs to a particular category. Although there has been much work on trying to make sense of what these features "mean" (Olah et al. 2017, Carter et al. 2019), either individually or in aggregate, understanding is enabled through a process of curation rather than by explicit explanation. In this way, machine learning introduces a new approach to making sense of the world in which choosing examples and defining mappings judiciously enables new applications and new forms of creative expression.

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

The Creative Coding Lab at University of California, Santa Cruz[1] investigates the use of machine learning algorithms for scientific research and creative explorations across a range of contexts. One effort, Deep Illumination, explores how deep learning can be used effectively in the graphics pipeline, investigating, for example, how to infer complex lighting models from a large dataset of examples, rather than through expensive rendering calculations, and evaluating how such an approach can provide useful trade-offs between time and memory (Thomas & Forbes, 2017, Elek et al. 2019, Alsaiari et al. 2019). Our lab has also investigated the use of machine learning technologies for a range of practical applications. One project, *CompostNet*, trains a neural network to classify food waste appropriate for available trash and recycling receptacles (Frost et al. 2019a). Another project uses machine learning to predict biker density at dangerous road intersections so that drivers and bikers can experience improved shared road safety (Dubey et al. 2019a). Researchers in the Creative Coding Lab have also investigated creative applications using machine learning. For example, the *Art I Don't Like* project used a novel recommender system that introduces users to artists and art genres that they may be unfamiliar with (Frost et al., 2019b), and the *Data Brushes* art application enables users to interactively paint using specialised brushes that generate output using neural style transfer networks (Dubey et al. 2019b). Much of the architecture for deep learning neural networks was first theorised and implemented in previous decades (Bishop 1995, LeCun et al. 1998, Rumelhart et al. 1996), but the recent explosion of deep learning techniques and applications introduced in the last few years was in part enabled by innovations in GPU technology (LeCun et al. 2015, Krizhevsky et al. 2012). Neural networks are loosely modelled on the behaviour of neurons, and the Creative Coding Lab has been exploring models of computational intelligence inspired by other biological processes. One recent project, developed in collaboration with astrophysicists at University of California, Santa Cruz, emulates properties of the Physarum polycephalum (the "many-headed slime mold") in order to infer a simulation of the dark matter filament structure of the Cosmic Web using only a sparse sampling of astrophysical observations (Burchett et al. 2020). [2]

The term "creative AI" is increasingly used by artists and designers who utilise machine learning to generate creative outputs, or who treat machine learning algorithms as a medium in and of itself in various ways (McCormick et al. 2020). In recent years, creative AI projects have been featured at the NeurIPS Workshop for Creativity and Design, as well as at other arts and computation venues, such as the ACM SIGGRAPH Art Gallery and Art Papers tracks, the IEEE VIS Arts Program, and the International Symposium on Electronic Art. Broadly speaking, creative AI projects involve one or more of the following: mimicking existing data, mapping features found in one dataset onto another, or mapping inputs to outputs in unusual ways, visualising or otherwise probing the inner workings of the algorithm, and analysing or speculating about the societal impact of machine learning systems. These activities can enable new kinds of generative artworks that can either replicate or incorporate existing artworks or create entirely new artistic outputs. They also can be used to design new techniques of more expressively interacting with existing art forms. In doing so, they introduce new ways to analyse and experience cultural artefacts and cultural data. Finally, the machine learning algorithm, its computational architecture, the input it requires, the resulting output, and the analysis framework it is part of can be thought of as a cultural artefact in and of itself, enabling new forms of critical inquiry. In the sections below, I provide an overview of these trends, along with descriptions of related projects, and highlight opportunities for computational artists working in this area.

2. Creative AI as expressive mimicry

Creating software that automatically generates artworks —either in the style of a particular artist, or in an original voice that does not directly reference existing work— is a perennial pursuit in new media practice and generative art. Well-known early examples include Harold Cohen's robot paintings (Cohen 1995) and David Cope's experiments in musical intelligence (Cope 1996). Often in these projects, the visual or audio outputs, while interesting on their own, are a byproduct of the actual artwork, which is the system itself: in Cohen's case, AARON is the artwork; for Cope, his EMI software is the main creative contribution. A more recent example is introduced by Sougwen Chung, who, as part of her *Drawing Operations* series, co-improvises drawings in collaboration with a robotic arm that is controlled via a recurrent neural net that has been previously trained on her own drawings (Chung 2019). Research into techniques that can be used to simulate human expressions, voices, and faces meant to fool users or for other nefarious purposes, also called "deep fakes", shows great creative potential for designing realistic human behavior, perhaps in combination with text generation and speech generation techniques. For example, work by Suwajanakorn et al. (Suwajanakorn et al. 2017) demonstrates how a voice impressionist can create a convincing video of another person speaking words that they never uttered. Thies et al. (Thies et al. 2016, Thies et al. 2019) introduce projects that enable a user to become a kind of virtual puppeteer using their own facial expressions to modify the expressions of another person

1. https://creativecoding.soe.ucsc.edu/
2. An overview of projects from the UCSC Creative Coding Lab was presented in late July 2019, as part of the "AI in the Arts and Design" panel discussion with Erkki Huhtamo, Memo Akten, and Max Sims at ACM SIGGRAPH, organised by Ruth West, Victoria Szabo, and Danielle Siembieda.

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

in a video. Work by Fried et al. (Fried et al. 2019) demonstrates a method to surreptitiously modify a video of a person talking simply by editing the textual transcript of the video. Chan et al. (Chan et al. 2019) introduce a method to transfer the recorded movements of an expert performer onto a new video featuring an amateur performer, appearing to transform novices into professional dancers. This technology exacerbates difficulties in separating facts from opinions, in thinking critically, and in identifying bias and propaganda (Gebru 2019, Jo & Gebru 2020), but it also potentially presents new avenues for exploring these issues and for new forms of creative work.

3. Creative AI as interactive mapping

Machine learning enables the creation of tools that map a range of inputs to new outputs, often in a different modality. By definition, all algorithms require an input that is then processed in some way to produce an output. Neural networks, including deep learning networks, are "tuned" through a training process that encodes an effective mapping of inputs to outputs for a particular dataset (the training set). If successful, and if the training set is representative of the kinds of inputs that will be encountered in the future, then the network can be queried nearly instantly to provide a meaningful output given some new, previously unseen input data. Fiebrink's *Wekinator* tool enables users to quickly train a neural network (or another machine learning algorithm) to recognise, for example, different gestures from a web camera and associate them with sounds or musical instructions (Fiebrink et al. 2016). *KIMA: The Wheel* is a multimedia performance by the art collective Analema Group that uses machine learning to correlate sound and visual parameters, generating a multimodal mapping between voices and visual outputs (Gingrich et al., 2018). Style transfer networks that encode stylistic features of a source image learn to map any image into a transformed version of that image that incorporates those features. Gatys et al. (Gatys et al. 2016) introduced *neural style transfer*, which makes use of a convolutional neural network to identify image patterns that represent a particular painter's "style", and can then transfer it onto any other image, making it possible, for example, to transform a photograph into an image that looks like it was painted by Van Gogh or Kandinsky, to use popular examples.

4. Creative AI as generative art

A range of techniques investigate the neural network as space of possibility. The "deep dream" algorithm, which transforms images into psychedelic quilts was originally created as a tool to highlight which features were being activated when processing an image with a neural network. If a neural network is trained to classify, for example, different species of birds, then a particular patch of a bird image (or an image that contains bird-like objects) will trigger the neurons within the network that have been tuned to respond to that particular bird feature. Often these features, when viewed in isolation, resist easy interpretation, and represent a particular curve or gradient or texture that proved to be useful in detecting a bird within an image (Olah et al. 2017). The *Inceptionism* project takes these features and iteratively integrates them onto the image, allowing us to see which features are observed in a given input image. To continue the example, even if an input image contains no birds at all, and if the network is trained only to recognise bird features, the technique ends up generating a kind of Boschian hellscape of bird parts (Mordvintsev et al. 2015).

Initial breakthroughs in deep learning led to state-of-the-art methods in data classification, identifying items in a photo, automatically tagging people in social media posts, or recommending products or content based on previous interactions or purchases on a website (LeCun et al. 2015, Goodfellow et al. 2016). If a network has been trained to identify particular features in order to, say, decide what category an image belongs to, then that network could also be used to generate new images made up of those features and belonging to that category (Goodfellow et al. 2014). The generative adversarial network (GAN) architecture consists of both a *generator* network and a discriminator network. During the training process, the generator network gets better at producing output, and the *discriminator* network gets better at distinguishing a real image from the training data from a generated image. Once the generator is sufficiently trained, any input to the generator network will produce a realistic output, that is, an output that contains features recognised by the discriminator network as a real image. The input to the generator network is a vector of numbers within a particular range of values describing a "latent space", and slightly changing the values of one of more of the numbers in the input vectors produces images that are similar to each other (Bojanowski et al. 2019). Animations of images created by "drifting" through the latent space (i.e. updating the input vector) produces a morphing between images that resemble the training images, sometimes creating a surreal effect. Artists have been inspired by GAN techniques that make it possible to direct the data generation process (Mirza & Osindero 2014, Radford et al. 2015, Karras et al. 2019). For example, a recent iteration of Refik Anadol's *Machine Hallucination* project uses a GAN trained on 100 million photographic memories of New York City found publicly in social networks to create synthetic representations that envision a possible "near future" (Anadol 2019). Mario Klingemann has created a series of animations using a technique he calls "neural glitch", in which he alters the weights in a trained generator to create intriguing "misinterpretations" that nonetheless retain a coherent style (Klingemann 2018). Casey Reas' *Earthly Delight* series generates what he terms "compressed cinema", using a GAN architecture trained via processed stills from Stan Brakhage's experimental films in which plants are directly placed on

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

top of clear film strips (Menezes 2019). Memo Akten's *Learning to See* processes live camera input, composing images that resemble the shape and structure of this input, but replacing the content with data learned through training a network on particular types of images, transforming, for example, keys and wires into flowers and waves, or faces into galaxies (Akten et al. 2019).

5. Creative AI as critical inquiry

Some recent creative AI projects can be considered as critical inquiries that investigate sociotechnical systems that utilise machine learning. Tom White's influential project *Perception Engines* creates idiosyncratic images made out of a few simple shapes with solid colors and curved dark lines. While at first glance they seem to be vaguely evocative of a particular object or action, upon reading the title of each print (such as "cello", "cabbage", "hammerhead shark", "iron", and "tick"), it becomes hard to see anything else. While the prints create a kind of visual puzzle, they also function as images that return the highest confidence score on different image classification algorithms (often higher even than photographs of those objects), providing insight into what shape features form a "Platonic ideal" of a category encoded in the image recognition network, and representing the "character" of a class more effectively than any one instance (White 2018). Avital Meshi's *Classification Cube* features an interactive surveilled space in which multiple machine learning algorithms are used to classify a participant's behaviours, expressions, age, and gender. In addition to making it clear that some expressions and poses are incorrectly categorised, and that a person's age or gender can be misclassified depending on seemingly minor changes, the project provides a space for reflecting on the ubiquitous automated interpretations that permeate our daily lives (Meshi & Forbes, 2020). While machine learning systems are implicated in algorithmic bias (Diakopoulos 2015, Eubanks 2018), bias of course exists prior to being encoded into datasets and deep learning networks trained on those datasets. Creative interrogations of machine learning systems can help to pinpoint aspects of a data analysis pipeline that introduce bias and spark discussions about the ramifications of weaving machine learning into the fabric of public life.

6. Creative AI opportunities

Novel sophisticated machine learning techniques are presented each year at the International Conference on Computer Vision (ICCV), the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Neural Information Processing Systems (NeurIPS), the ACM Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH), and various other computer science venues. Researchers often put versions of the articles online at the arXiv.org open-access archive and make the code for these projects available in online repositories, enabling anyone to test out their techniques using popular software frameworks, such as TensorFlow and PyTorch. Given their accessibility, there are many opportunities for incorporating contemporary machine learning techniques into creative projects. For instance, Isola et al. and Park et al. introduce architectures that have been used to make interactive demos that infer a reasonable image from only outlines or coloured rectangles (Isola et al. 2017, Park et al. 2019). A more recent project called *GauGAN* lets a user easily modify generated images by "painting" particular features on the image (Bau et al. 2019), and an interactive demo by Liu et al. lets a user edit an existing photo by erasing people or objects, automatically "inpainting", replacing them with relevant elements from the surrounding landscape (Liu 2018). Other generative machine learning projects have appeared over the last few years, many of which are geared toward graphics techniques for visual effects in films and games, but have not yet, to the best of my knowledge, been incorporated into media arts projects or to augment interactive performance. Xie et al. (Xie et al. 2018) showed that realistic motion dynamics could be created and shaped interactively by training a neural network on a database of fluids. Their system learns to generate fine details in explosions, water, or smoke from low-resolution inputs, which speeds up computation and enables visual effects artists to quickly create high-quality animations of different fluids. A number of projects have focused on generating realistic human and animal motion and motion planning strategies for navigating specialised environments, including for rock climbing simulations (Naderi et al. 2017), walking through diverse terrain (Zhang et al. 2018), or in crowds (Amirian et al. 2019). For example, work by Holden et al. (Holden et al. 2017) trains a neural network using a database of human movement captured in a motion capture lab, including walking, jumping, climbing stairs, and crouching. This network is then able to determine the most reasonable motions for a virtual character moving through any scene, finding correlations between the motions stored in the networks and the elements within the scene. Even for scenes with arrangements of terrain and objects that are quite different from the data it was trained on, the network produces synthetic motion outputs that are convincingly realistic.

Many creative AI projects differentiate themselves by curating the data and labels they choose for the training set or as inputs into the network. To take just two examples, Chris Rodley uses a style transfer network to create compelling images of dinosaurs composed out of fruit (Rodley 2017) and Pinar Yanardag and Emily Salvador use generative adversarial networks trained on a database of fashion designs to create new dresses and jewellery (Yanardag and Salvador, 2019). Some intriguing machine learning techniques enable cross-modal mapping, in which data from one domain informs or creates the output in another (Baltrušaitis et al. 2018). Recent techniques

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

automatically provide captions from an image or accurately label subregions in an image (Karpathy & Fei-Fei 2015, Gan et al. 2017), or the reverse, generate images from text (Qiao et al. 2019). For example, Zhang et al. generate accurate images (at least at first glance) of birds from simple descriptions, such as "This bird is red and brown in color, with a stubby beak", enabling users to "paint" with words (Zhang et al. 2017). In addition to encapsulating a form of cognitive blending, in which emergent meanings are constructed from mixing together partial matches in two different domains (Fauconnier & Turner 2003), they illustrate that existing cultural artefacts (such as online field guides for bird watchers, photo collections of flowers, or various forms of social media) contain conceptual analogies that define "unseen" relationships that enable new forms of automated reasoning (Peyre et al. 2019, Yan et al. 2019). Designing machine learning systems that leverage or investigate cultural artefacts presents opportunities for new creative work and cultural insight.

Techniques such as style transfer and inpainting show that there is unexpected information that can be mined from even a small number of input data samples, and which can then be used for creative reinterpretations. Other recent examples include learning 3D information from 2D data, such as a technique to synthesise animations that contain novel views of complex scenes from a set of input images (Mildenhall et al. 2020) and a technique that estimates the depth of elements in an image in order to automatically create a "Ken Burns" animation effect consisting of zooming, panning, and motion parallax (Niklaus et al. 2019). Another technique learns to synthesise frames of future frames from a single image, predicting plausible ways that a scene might change over time (Xue et al. 2018). Techniques such as these can infer unexpected features and relationships between those features, and present many creative possibilities that have yet to be fully explored. Additionally, different types of sensors can expose new features in data. For instance, by using a slow-motion camera with a high temporal resolution, Davis et al. (Davis et al. 2014) were able to recreate sounds in a room by observing the subtle motions of particular objects in that room, such as plants or packaging. In one experiment, they demonstrate that they can retrieve someone singing a nursery rhyme simply by recording and accentuating the vibrational movements on a package of chips. While this has implications for surveillance, it also illustrates how inventive uses of sensors can provide unexpected streams of information in other sensory domains. Incorporating datasets from higher resolution instruments into machine learning systems can lead to new creative applications.

7. The future of Creative AI

Given the continuing breakthroughs, it is worth thinking about what machine learning is not yet able to achieve, and about what components of an artwork cannot be effectively modelled or mimicked. For example, so far, machine learning approaches have not successfully generated convincing dramatic experiences or engaging multimedia performances. These kinds of experiences require contextual information which we do not yet understand how to encode effectively and thoroughly. Narrative, dance, performance, and cinema are inherently more complex than static images or sound recordings, and require integrating many elements simultaneously, such as lighting, editing, acting, narrative, and sound design. Machine learning makes the assumption that all relevant features can be found within the training data, and even if there were a way to gather and label relevant data from, say, a film or a live performance, we bring our knowledge of the world and our expectations about how to interpret particular genres when experiencing art. Moreover, these experiences are ultimately interior and perhaps ineffable, resonating with a rich personal database of our own experiences and our own thoughts and feelings. That is, machine learning algorithms can effectively identify and utilise features in artworks in increasingly sophisticated ways, but do not model how an artwork is perceived or why it is interpreted in a particular way. Media artists, in addition to using new media forms to create new representations and new experiences, also investigate the nature of media itself, and often foreground concept over or alongside aesthetics and technical craftsmanship (Shanken 2002a, Agüera y Arcas 2017, Ackerman et al. 2018). Creative AI practitioners will continue to identify which concepts resist machine learning approaches and to investigate how machine learning tools can make particular interpretations either inescapable or impossible.

Machine learning technologies can be thought of as a type of measuring instrument. Many sensors include a computational component in which data is filtered or otherwise processed to separate out the noise from the signal. Neural networks measure distinguishing features in data, and can provide insight into the system the data is drawn from, as well as about other systems with which it is entwined. For example, observing transportation patterns or analysing pollution levels can be used to provide insight into the economic health of a city (Washington 2020), and interactions on social media can be used to identify personality traits, and then exploited for targeted advertising or disinformation campaigns (Kaiser 2019). Insight into these auxiliary systems then could allow us to infer patterns from yet other interacting systems. The promise of "big data" is not simply that we can collect higher and higher resolution spatiotemporal data, and not only that we can retrieve and analyse data more and more quickly, but that we can make use of all this data to make sense of how systems interact and integrate with each other (Shanken 2002b, Hassad 2020). How should we design the next iteration of machine learning tools that reason about the world holistically by integrating multiple interpretations encoded in text, mined from image and video databases, perceived by sensors, provided by human-computer interactions, and communicated by yet other machine learning tools? Creative AI will continue to be a space in which artists and researchers create

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

personal yet empirical research projects that explore and challenge the logic of how different systems and interpretations of those systems promote or impede each other.

References

Ackerman, Margareta, James Morgan, and Christopher Cassion. "Co-Creative Conceptual Art." *Proceedings of the Ninth International Conference on Computational Creativity* (ICCC), pages 1-8. 2018.

Agüera y Arcas, Blaise. 2017. "Art in the age of machine intelligence." *Arts* 6, no. 4: 18. https://doi.org/10.3390/arts6040018

Akten, Memo, Rebecca Fiebrink, and Mick Grierson. "Learning to See: You Are What You See." *Proceedings of ACM SIGGRAPH Short Art Papers*, pages 1-6. 2019. https://doi.org/10.1145/3306211.3320143

Alsaiari, Abeer, Ridhi Rustagi, Manu Mathew Thomas, and Angus G. Forbes. "Image Denoising Using A Generative Adversarial Network." *Proceedings of the IEEE 2nd International Conference on Information and Computer Technologies* (ICICT), pages 126-132. 2019. https://doi.org/10.1109/INFOCT.2019.8710893

Anadol, Refik. "Machine Hallucination". 2019. Available online at: http://refikanadol.com/works/machine-hallucination/.

Baltrušaitis, Tadas, Chaitanya Ahuja, and Louis-Philippe Morency. "Multimodal machine learning: A survey and taxonomy." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, no. 2 (2018): 423-443. https://doi.org/10.1109/TPAMI.2018.2798607

Bau, David, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. "Semantic photo manipulation with a generative image prior." *ACM Transactions on Graphics* 38, no. 4 (2019): 1-11. https://doi.org/10.1145/3306346.3323023

Bishop, Christopher M. *Neural Networks for Pattern Recognition.* Oxford University Press, 1995. https://doi.org/10.1201/9781420050646.ptb6

Bojanowski, Piotr, Armand Joulin, David Lopez-Paz, and Arthur Szlam. "Optimizing the latent space of generative networks." *arXiv* preprint arXiv:1707.05776, 2019.

Burchett, Joseph N., Oskar Elek, Nicolas Tejos, J. Xavier Prochaska, Todd M. Tripp, Rongmon Bordoloi, and Angus G. Forbes. "Revealing the Dark Threads of the Cosmic Web." *The Astrophysical Journal Letters* 891, no. 2 (2020): L35. https://doi.org/10.3847/2041-8213/ab700c

Carter, Shan, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. "Activation atlas." *Distill* 4, no. 3 (2019): e15. https://doi.org/10.23915/distill.00015

Chan, Caroline, Shiry Ginosar, Tinghui Zhou, and Alexei A. Efros. "Everybody dance now." *Proceedings of the IEEE International Conference on Computer Vision*, pages 5933-5942. 2019. https://doi.org/10.1109/ICCV.2019.00603

Chung, Sougwen. "Artefact 1." Sougwen Chung Portfolio Website. 2019. https://sougwen.com/project/artefact1.

Cohen, Harold. "The further exploits of AARON, painter." *Stanford Humanities Review* 4, no. 2 (1995): 141-158.

Cope, David. *Experiments in musical intelligence*. Middleton WI: A-R Editions, 1996.

Davis, Abe, Michael Rubinstein, Neal Wadhwa, Gautham J. Mysore, Frédo Durand, and William T. Freeman. "The visual microphone: Passive recovery of sound from video." *ACM Transactions on Graphics* 33, no. 4 (2014):79. https://doi.org/10.1145/2601097.2601119

Diakopoulos, Nicholas. "Algorithmic accountability: Journalistic investigation of computational power structures." *Digital Journalism* 3, no. 3 (2015): 398-415. https://doi.org/10.1080/21670811.2014.976411

Dubey, Mahika, Alan Peral Ortiz, Rakshit Agrawal, and Angus G. Forbes. "Predicting Biker Density at Bikeshare Station Intersections in San Francisco." *2019 IEEE Global Humanitarian Technology Conference* (GHTC), pages 1-7. 2019a. https://doi.org/10.1109/GHTC46095.2019.9033019

Dubey, Mahika, Jasmine Otto, and Angus G. Forbes. "Data Brushes: Interactive style transfer for data art." *Proceedings of IEEE VIS Arts Program.* 2019b. https://doi.org/10.1109/VISAP.2019.8900858

Elek, Oskar, Manu M. Thomas, and Angus Forbes. "Learning Patterns in Sample Distributions for Monte Carlo Variance Reduction." *arXiv* preprint arXiv:1906.00124, 2019.

Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* New York: St. Martin's Press, 2018.

Fauconnier, Gilles, and Mark Turner. "Conceptual blending, form and meaning." *Recherches en Communication* 19 (2003): 57-86. https://doi.org/10.14428/rec.v19i19.48413

Fiebrink, Rebecca, and Baptiste Caramiaux. "The machine learning algorithm as creative musical tool." In *The Oxford Handbook of Algorithmic Music*, edited by Roger T. Dean and Alex McLean, chapter 12. Oxford University Press, 2016.

Forbes, Angus G. "Articulating media arts activities in art-science contexts." *Leonardo* 48, no. 4 (2015): 330-337. https://doi.org/10.1145/2810177.2810179

Forbes, Angus G. and Kiyomitsu Odai. "Iterative synaesthetic composing with multimedia signals." *Proceedings of the International Computer Music Conference* (ICMC), pages 573-578. 2012.

Fried, Ohad, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B. Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. "Text-based editing of talking-head video." *ACM Transactions on Graphics* 38 no. 4 (2019): 1-14. https://doi.org/10.1145/3306346.3323028

Frost, Sarah, Bryan Tor, Rakshit Agrawal, and Angus G. Forbes. "CompostNet: An Image Classifier for Meal Waste." *2019 IEEE Global Humanitarian Technology Conference* (GHTC), pages 1-4. 2019a. https://doi.org/10.1109/GHTC46095.2019.9033130

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

Frost, Sara, Manu Mathew Thomas, and Angus G. Forbes. "Art I don't like: An anti-recommender system for visual art." *Proceedings of Museums and the We*b. 2019b.

Gan, Chuang, Zhe Gan, Xiaodong He, Jianfeng Gao, and Li Deng. "StyleNet: Generating attractive visual captions with styles." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pages 3137-3146. 2017. https://doi.org/10.1109/CVPR.2017.108

Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "Image style transfer using convolutional neural networks." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pages 2414-2423. 2016. https://doi.org/10.1109/CVPR.2016.265

Gebru, Timnit. "Race and Gender." *arXiv* preprint arXiv:1908.06165, 2019.

Gingrich, Oliver, Sean Soraghan, Alain Renaud, Evgenia Emets, and Dario Villanueva-Ablanedo. "KIMA: The Wheel—Voice Turned into Vision: A participatory, immersive visual soundscape installation." *Leonardo* (Online - Accepted for publication) (2018): 1-13. https://doi.org/10.1162/leon_a_01698

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." *Advances in Neural Information Processing Systems*, pages 2672-2680. 2014.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning.* MIT press, 2016.

Hallinan, Blake, and Ted Striphas. "Recommended for you: The Netflix Prize and the production of algorithmic culture." *New Media & Society* 18, no. 1 (2016): 117-137. https://doi.org/10.1177/1461444814538646

Hassad, Rossi A. "A foundation for inductive reasoning in harnessing the potential of big data." *Statistics Education Research Journal* 19, no. 1 (2020): 238-258.

Holden, Daniel, Taku Komura, and Jun Saito. "Phase-functioned neural networks for character control." *ACM Transactions on Graphics* 36, no. 4 (2017): 1-13. https://doi.org/10.1145/3072959.3073663

Isola, Phillip, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. "Image-to-image translation with conditional adversarial networks." *Proceedings of the IEEE conference on Conference on Computer Vision and Pattern Recognition*, pages 1125-1134. 2017. https://doi.org/10.1109/CVPR.2017.632

Jo, Eun Seo, and Timnit Gebru. "Lessons from archives: strategies for collecting sociocultural data in machine learning." *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 306-316. 2020.

Kaiser, Brittany. *Targeted: The Cambridge Analytica Whistleblower's Inside Story of How Big Data, Trump, and Facebook Broke Democracy and How It Can Happen Again.* HarperCollins, 2019.

Karpathy, Andrej, and Li Fei-Fei. "Deep visual-semantic alignments for generating image descriptions." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128-3137. 2015. https://doi.org/10.1109/CVPR.2015.7298932

Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401-4410. 2019. https://doi.org/10.1109/CVPR.2019.00453

Klingemann, Mario. "Neural Glitch." *Quasimondo.* October 28, 2018. http://underdestruction.com/2018/10/28/neural-glitch/.

LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner. "Gradient-based learning applied to document recognition." *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324. https://doi.org/10.1109/5.726791

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep learning." *Nature* 521, no. 7553 (2015): 436-444. https://doi.org/10.1038/nature14539

Liu, Guilin, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. "Image inpainting for irregular holes using partial convolutions." *Proceedings of the European Conference on Computer Vision* (ECCV), pages 85-100. 2018. https://doi.org/10.1007/978-3-030-01252-6_6

McCann, Bryan, James Bradbury, Caiming Xiong, and Richard Socher. "Learned in translation: Contextualized word vectors." *Advances in Neural Information Processing Systems*, pages 6294-6305. 2017.

McCormack, Jon, Oliver Bown, Alan Dorin, Jonathan McCabe, Gordon Monro, and Mitchell Whitelaw. "Ten questions concerning generative computer art." *Leonardo* 47, no. 2 (2014): 135-141. https://doi.org/10.1162/LEON_a_00533

McCormack, Jon, Patrick Hutchings, Toby Gifford, Matthew Yee-King, Maria Teresa Llano, and Mark d'Inverno. "Design Considerations for Real-Time Collaboration with Creative Artificial Intelligence." *Organised Sound* 25, no. 1 (2020): 41-52. https://doi.org/10.1017/S1355771819000451

Menezes, Caroline. "Interview with Casey Reas". *Studio International.* May 21, 2019. https://www.studiointernational.com/index.php/casey-reas-interview-computer-art-coding.

Meshi, Avital and Angus G. Forbes. "Stepping inside the Classification Cube: An intimate interaction with an AI system." *Leonardo* 53, no. 4 (2020): 387-393.

Mildenhall, Ben, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis." *arXiv* preprint arXiv:2003.08934, 2020.

Mirza, Mehdi, and Simon Osindero. "Conditional generative adversarial nets." *arXiv* preprint arXiv:1411.1784, 2014.

Mordvintsev, Alexander, Christopher Olah, and Mike Tyka. "Inceptionism: Going Deeper into Neural Networks". Google AI Blog.

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

June 17, 2015. https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html.

Naderi, Kourosh, Joose Rajamäki, and Perttu Hämäläinen. "Discovering and synthesizing humanoid climbing movements." *ACM Transactions on Graphics* 36, no. 4 (2017): 1-11. https://doi.org/10.1145/3072959.3073707

Niklaus, Simon, Long Mai, Jimei Yang, and Feng Liu. "3D Ken Burns effect from a single image." *ACM Transactions on Graphics* 38, no. 6 (2019): 1-15. https://doi.org/10.1145/3355089.3356528

Olah, Chris, Alexander Mordvintsev, and Ludwig Schubert. "Feature visualization." *Distill* 2, no. 11 (2017): e7. https://doi.org/10.23915/distill.00007

Park, Taesung, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. "Semantic image synthesis with spatially-adaptive normalization." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337-2346. 2019. https://doi.org/10.1109/CVPR.2019.00244

Peyre, Julia, Ivan Laptev, Cordelia Schmid, and Josef Sivic. "Detecting Unseen Visual Relations Using Analogies." *Proceedings of the IEEE International Conference on Computer Vision*, pages 1981-1990. 2019. https://doi.org/10.1109/ICCV.2019.00207

Qiao, Tingting, Jing Zhang, Duanqing Xu, and Dacheng Tao. "MirrorGAN: Learning text-to-image generation by redescription." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1505-1514. 2019. https://doi.org/10.1109/CVPR.2019.00160

Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." *arXiv* preprint arXiv:1511.06434, 2015.

Rodley, Chris. "Deep Dinosaur". June 19, 2017. https://chrisrodley.com/2017/06/19/dinosaur-flowers/.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. "Learning representations by back-propagating errors." *Nature* 323, no. 6088 (1986): 533-536. https://doi.org/10.1038/323533a0

Shanken, Edward A. "Art in the information age: Technology and conceptual art." *Leonardo* 35, no. 4 (2002a): 433-438. https://doi.org/10.1162/002409402760181259

Shanken, Edward A. "Cybernetics and art: Cultural convergence in the 1960s." In *From Energy to Information: Representation in Science and Technology, Art, and Literature*, editors Bruce Clarke and Linda D. Henderson, pages 155-177. Stanford University Press, 2002b.

Sun, Yi, Xiaogang Wang, and Xiaoou Tang. "Deep learning face representation from predicting 10,000 classes." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1891-1898. 2014. https://doi.org/10.1109/CVPR.2014.244

Suwajanakorn, Supasorn, Steven M. Seitz, and Ira Kemelmacher-Shlizerman. "Synthesizing Obama: Learning lip sync from audio." *ACM Transactions on Graphics* 36, no. 4 (2017): 1-13. https://doi.org/10.1145/3072959.3073640

Taigman, Yaniv, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf. "Deepface: Closing the gap to human-level performance in face verification." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701-1708. 2014. https://doi.org/10.1109/CVPR.2014.220

Thies, Justus, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. "Neural Voice Puppetry: Audio-driven Facial Reenactment." *arXiv* preprint arXiv:1912.05566, 2019.

Thies, Justus, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. "Face2face: Real-time face capture and reenactment of RGB videos." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2387-2395. 2016. https://doi.org/10.1109/CVPR.2016.262

Thomas, Manu Mathew, and Angus G. Forbes. "Deep Illumination: Approximating Dynamic Global Illumination with Generative Adversarial Network." *arXiv* preprint arXiv:1710.09834, 2017.

Washington, Simon, Matthew G. Karlaftis, Fred Mannering, and Panagiotis Anastasopoulos. *Statistical and Econometric Methods for Transportation Data Analysis*. CRC press, 2020. https://doi.org/10.1201/9780429244018

White, Tom. "Synthetic Abstractions." Aug 23, 2018. https://medium.com/@tom_25234/synthetic-abstractions-8f0e8f69f390.

Xue, Tianfan, Jiajun Wu, Katherine L. Bouman, and William T. Freeman. "Visual dynamics: stochastic future generation via layered cross convolutional networks." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, no. 9 (2018): 2236-2250. https://doi.org/10.1109/TPAMI.2018.2854726

Yan, Chenggang, Liang Li, Chunjie Zhang, Bingtao Liu, Yongdong Zhang, and Qionghai Dai. "Cross-modality bridging and knowledge transferring for image understanding." *IEEE Transactions on Multimedia* 21, no. 10 (2019): 2675-2685. https://doi.org/10.1109/TMM.2019.2903448

Yanardag, Pinar and Emily Salvador. "The Little Black Dress Reimagined by an A.I." 2019. https://lbd-ai.com/.

Zhang, Han, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N. Metaxas. "Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks." *Proceedings of the IEEE International Conference on Computer Vision*, pages 5907-5915. 2017. https://doi.org/10.1109/ICCV.2017.629

Zhang, He, Sebastian Starke, Taku Komura, and Jun Saito. "Mode-adaptive neural networks for quadruped motion control." *ACM Transactions on Graphics* 37, no. 4 (2018): 1-11. https://doi.org/10.1145/3197517.3201366

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST

CV

Angus G. Forbes
University of California, Santa Cruz
angus@ucsc.edu

Angus Forbes is an associate professor in the Department of Computational Media at University of California, Santa Cruz, where he directs the Creative Coding Lab. His research investigates novel techniques for visualising and interacting with complex scientific information, and his interactive artwork has been featured at museums, galleries, and festivals throughout the world. He was the SIGGRAPH 2018 Art Papers Chair and will be the Art Gallery Chair for SIGGRAPH 2021. More information about current projects can be found at https://creativecoding.soe.ucsc.edu/.

Universitat Oberta de Catalunya

LEONARDO
THE INTERNATIONAL SOCIETY FOR THE ARTS, SCIENCES AND TECHNOLOGY

© 2020, Angus G. Forbes
© 2020, of this edition by FUOC and Leonardo/ISAST