JOHN GERRARD
NEURAL EXCHANGE

LACMA

John Gerrard

Neural Exchange (Leaf Covered Figure) 2017

Simulation; dimensions variable

Producer: Werner Poetzelberger
Programmer: Helmut Bressler
3-D modeler: Max Loegler
Dancer: Esther Balfe
Motion-research actor: Christoph Gasgeb
Motion-capture suit: Xsens/7 Reasons
Neural network: TensorFlow
Game engine: Unigine

The Art + Technology Lab is presented by Hyundai.

The Art + Technology Lab is made possible by Accenture
with additional support from Google and SpaceX.

The Lab is part of The Hyundai Project: Art + Technology at LACMA,
a joint initiative exploring the convergence of art and technology.

JOHN GERRARD
NEURAL EXCHANGE

Introduction

Joel Ferree and Amy Heibel

Since 2014, the Art + Technology Lab at the Los Angeles County Museum of Art (LACMA) has supported artist experiments that engage emerging technology. Through its technology sponsors, the Lab provides grants, in-kind support, and facilities at the museum to artists to develop new projects.

An Art + Technology Lab grant is not a commission. It is not an exhibition. It is a form of support that allows artists to pursue new directions in their practices, to take purposeful risks, and to develop projects that wouldn't be done otherwise. In a highly quantified world that places a premium on deliverables, it sets out to create a space that emphasizes process over product and teases out the conversations that result from such experimentation.

John Gerrard came into the Lab program with a proposal to use gaming technology to render a large-scale image in real time. After meeting with advisors John Suh of Hyundai, Brian Mulford of Google, and Bryan Catanzaro of NVIDIA, Gerrard became interested in neural networks. Shortly thereafter, he submitted a new project proposal, this one for what would eventually become *Neural Exchange (Leaf Covered Figure) 2017.*

Artificial neural networks mimic the network of neurons in the brain, but on a smaller scale. In short, they allow computers to learn from data. First, a training set of data is introduced, then the computer uses the training set to identify or generate new data based on what it has already "learned." For example, a training set of 100 cat images can be introduced and "learned" by a machine to help it identify new cat images. Gerrard originally sought to put this kind of machine learning into practice in his work, "hand-building" a digital model similar to a character found in a battle-simulation game and employing the character's vernacular martial gestures as a training set to produce new gestures. The artist considered this to be a machine-generated militaristic choreography.

Following discussions with Lab advisors—specifically at Google—the decision was made to move beyond "found" animations and to generate unique captures using a special suit embedded with sensors. This converts action into data at multiple points in space and time. In addition, the focus moved from military-simulation games to martial arts performances—specifically kata in karate, in which a precise series of gestures is performed by a group simultaneously. Later still over the year-long research period, the results of these captures were deemed too aggressive, and the studio began to work with a ballet dancer, Esther Balfe.

Working with Balfe and accounting for the computing needs and processing time necessary to train a neural network, a group of four core movements was developed. It includes two that travel backward and two forward—each easily identifiable. This training set was performed by Balfe hundreds of times, as repetitions, captured as data and used to train the network over multiple iterations. Once learned, the neural network uses these actions to develop a choreography—a sort of perpetual choreographic generator existing outside of traditional animation techniques. It is this idea of a choreographic generator that has become the core area of research for Gerrard's studio with respect to the Art + Technology Lab grant.

Toward the end of the research period, developing a basic kind of interaction between two dancing figures became important. As a neural network cannot be programmed—but must learn—a new training set had to be developed, which two characters are captured responding to each other's actions. Usually this meant a backward movement was matched or responded to with a forward one. The research has resulted in these two outcomes: one performance of a perpetual dance generated by the neural network and another in which two characters also dance with no duration, but in which a basic interaction can be witnessed.

The performers in the final work are figures dressed in leaf-covered "Gillie" suits, a type of camouflage clothing that the artist also saw as referring to the old European tradition of the "green" or "wild" man. Gerrard's team went out into the woods of Vienna to shoot a model dressed in this clothing. The images were the basis for a virtual 3-D model—the "leaf-covered figure" used in the final performance.

Within the parameters of the neural network there is an unpredictability in the final movements that are generated. Natural human variation in the core training set of four motions means that their interpretation and representation are producing some novel results and unexpected outcomes. The artist has stated:

> I have this strange feeling that I am equipping my engine—via the [neural network]—with a kind of an imagination. I frame that imagination through the kind of actions that are captured; however, looking at the leaf-covered figure perform, as I am doing now, something else is emerging...

Joel Ferree, Program Director, Art + Technology Lab
Amy Heibel, Adjunct Curator, Art + Technology Lab

Neural Exchange

Bryan Catanzaro

1) In both theory and practice, what is a neural network?

A neural network is a type of mathematical function. A function is just a mapping between an input and an output—for every input point, the function defines the output. Neural networks can be effectively constructed to map very complicated input spaces to very complicated output spaces. For example, imagine a function where the input is a recording of a person speaking and the output is a transcription of what the person said. This is an extraordinarily complex function: it has thousands or millions of inputs representing the recording, and tens or hundreds of outputs representing the words that the person said. This function has to be able to ignore the unimportant details of the recording, such as background noise that might be present, whether the person has a high voice or a low voice, how quickly or slowly the person is speaking, or the particular accent they use to pronounce their words. We know that this function exists because we can implement it, albeit tediously, by asking a person to listen to the recording and write down what the recording said. However, finding the mathematical function that implements this mapping is a very difficult problem.

Neural networks are simple ways of building these complex functions. They are built out of layers stacked on top of one another, where each layer has a set of parameters that manipulates the input to each layer in a simple way in order to produce a refined output. Neural networks have been so successful because it's possible to "train" them—to find a function that manages to solve the problem we care about while ignoring all the distracting details that make these problems so difficult for computers to solve. As you train a neural network, each layer learns to progressively construct a representation of the input that captures its most important characteristics and ignores the things that don't matter. During training, we use large amounts of data—the inputs to a problem with their associated outputs—to compare what the neural network produced with what the true output should have been. We then update the parameters of the neural network to make it do better on the data samples. With large amounts of data, we are able effectively to search for a function that actually solves a problem important to humans—like speech recognition, which is changing the way humans interact with computers. This search is incredibly computationally intensive. Training a state-of-the-art speech-recognition model can take tens of billions of billions of simple math operations like additions and multiplications. Most importantly and perhaps surprisingly, this training process allows us to find functions that generalize: that can ignore the distracting details and can solve problems on data the neural network has never seen before.

2) What are the potentials and the aims of this technology?

Artificial intelligence (AI) augments human intelligence by creating tools that assist humans in intellectual work. The Industrial Revolution gave us machines that can perform physical work, and the rise of computing gave us tools that were good at solving logical problems step-by-step. But it's always been difficult for computers to deal with data that humans understand almost intuitively, like the contents of images or speech. Artificial intelligence is now opening up new ways to automate many tedious tasks, thereby giving us new tools to do our work. For example, AI is giving us higher-level tools for image editing that understand the contents of the image and how to manipulate images in a coherent way, rather than just changing the colors of indicated pixels. Ultimately, AI aims to increase human productivity, which is how our economy has always grown. It aims to automate tedious tasks that people would rather not do.

Autonomous vehicles are a dream that many of us have had for a long time, and AI is making them real. Driving requires paying attention for long periods of time, which many humans find challenging, and others simply cannot do. Giving people mobility without requiring them to drive would open up many people's worlds, and give us back the time that we currently spend driving or parking. It would free up expensive real estate currently occupied by parking garages—no need to park close to your destination if the car can park itself and return to pick you up later. Autonomous vehicles will open up new kinds of businesses because the cost of delivering things will decline dramatically. The barrier to making autonomous vehicles is making vehicles that understand the world around them and can react appropriately. This is what AI aims to do. Automating tedious tasks that currently can only be done by humans opens up new possibilities in all activities we humans do every day.

3) Conversely, what risks might this technology hold—socially, economically, or otherwise?

Thoughts of superintelligent, malevolent AI may capture the imagination, but I'm worried about two more concrete risks that are closer at hand. First, AI may increase inequality in our society. Because AI is easy to replicate, a few people applying AI can make big changes in how we do many things. This will change the nature of work for many people, and may displace jobs, especially in the short term. For example, truck driver is one of the most common occupations in the United States in 2017, but autonomous long-haul trucks may replace many truck drivers. What will they do for their next job? How will we educate them to do something else? I'm concerned about how people will adjust, even though if I think long-term, this transition will be good for all of us. As a society, we're going to need to figure out how people will find meaning and purpose without the jobs that give people structure today. I believe this is possible, although challenging. It will require new policies, with a lot more redistribution, so that the people benefiting financially from AI are helping those being displaced. I hope it also gives people much more time for art, community building, and the humanities: those things that we know enrich our lives, but that are currently not compensated well monetarily. I spent two years volunteering full-time, trying to make the world a better place, working with individuals personally, and those two years were some of the most purposeful, intentional, and focused years of my life, despite the fact that I wasn't paid. I hope that AI frees up our time so that as a society we can give more people the opportunity to do similar things. I hope AI gives us the freedom to lift our sights and translate more of our idealism into reality, because it will liberate the human capital currently being spent on tedious labor.

Second, AI will create more compelling tools with which to manipulate society. Propaganda will be far more compelling, honed by AI to garner more attention. "Fake news" will include convincing-looking videos of people saying things they didn't say. We will have news scandals driven by fake videos of political candidates saying things they never said that will look extremely convincing. We'll need new ways to authenticate images, video, and audio, because it will become too easy to conjure up fake materials. People will attempt to use AI to manipulate democracy, and differentiating between truth and fiction will likely be even harder. I'm more concerned with the impact of AI on our societal institutions and economy than the possibility of malevolent AI, like the Terminator or even HAL 9000. Truly sentient, superintelligent agents are still far away in a technological sense. Just as there were connections between the Industrial Revolution and the Communist Revolution, the rise of AI will lead to political strains on our systems that will challenge the way we think about the world and our place in it. I'm more worried about the political instability that may arise during this transition than the technology itself.

Bryan Catanzaro is Vice President of Applied Deep Learning Research at NVIDIA, where his mission is to find new ways of improving NVIDIA's work by applying AI. He received his PhD in Electrical Engineering and Computer Sciences from the University of California, Berkeley, after earning a BA in Russian.

Neural Exchange

Brian Mulford

1) In both theory and practice, what is a neural network?

Neural networks (NNs) are everywhere. You can see them in how your Facebook and LinkedIn connections cluster in groups around family, work, friends, and acquaintances. You can see them in the relationship between states and cities on a map, and in the way flowers grow across fields. NNs are anywhere where objects or entities cluster and disburse with explicit relationships to one another. The fact that we see these interconnected relationships occurring in common contexts, it helps stage how mathematical NNs are formed as a series of simple ties between things.

Applied in computer science, NNs extend the concept to indicate a method of analyzing nonlinear numerical tensors (data that forms arrays of arrays) so as to apply concrete classifications to objects—typically in order to measure and proscribe unitary value, which is then used for prediction. When performed electronically on a large scale, we generally call it deep-learning machine learning (DLML). The "deep-learning" portion relates to the many layers of relationships being summed up and passed to subsequent layers.

Let's explore what linear and nonlinear mean in this context. Linear data or lists of numerical factual representations compared to other sets of linear data will produce easily recognized models. For instance, a set of data representing age and another representing height can be plotted along two axes and show a steady left-to-right pattern where height increases with age and then tapers off. This is called linear data, in that it produces a relatively smooth line with values appearing discretely above and below the line division.

Nonlinear data is more complicated. Imagine, instead of a smooth line separating values into two distinct areas, the data points when separated form an S curve with values appearing all over the chart. This is much harder to understand and predict. For example, if we add dietary factors, genetic factors, health, geography, current elevation, and favorite genre of film to our hypothetical model of age/height, it would be nearly impossible to plot in a chart. This is an example of nonlinear data: abstract problems that have so many widely changing variables that standard methods of analysis aren't able to make meaningful inferences from the data.

NNs provide a set of methods (some optimized for data relating to time or sequences, or self-organizing data) for making sense of highly varied data. These methods result in being able to classify (e.g., this picture is a cat) or predict values (e.g., you will probably like this movie) in such a way that isn't readily visible when looking at data using traditional methods (e.g., databases or charts). Machine learning (ML) enables this to occur at a high scale and to store the classification and logical values so that they may be compared with new data or take on new data, creating a cycle where existing data helps organize future data.

Examples of where you see ML being applied range from how Gmail automatically detects spam and junk email, to what movie recommendations you see, to predicting crop yields or weather.

Practically speaking, NNs work by observing structures present within information encoded into representational numerical data sets. In fact, in today's environment, 70 to 80 per cent of ML is tied up in creating numerical representations of data (features) and applying meaningful descriptions (labels) to them. The rest is applying various NN methods for analyzing and testing the data (models).

The core of NNs is in what's called hidden layers. These layers are abstractions of numerical ranges between data sets that over many iterations form groups or clusters of distances between numbers. These abstractions represent logical groups of information and pass it to deeper layers for further clarification. This process is called training—the core of which becomes the recognition model and subsequent output used for rapid identification, object classification, and prediction.

When we think about the graph made by our connections on Facebook (people that we connect to, organized by strongest to weakest), it'll resemble a series of clusters, with densely connected lines and elongated, dispersed lines between clusters. To an NN, the inputs would be all the people in your friends list and all the communications between them and the output would be the graph showing your most important relationships. The model would be the established patterns it forms that would allow any new friends joining your network to be quickly classified into the graph. It even could predict other people not in your friend list who would be strong connections.

With that description in mind, consider how much of the world can be categorized as series of weighted connections. Think about how we as humans interpret conditions based upon prior experience, known interactions, and physical properties, and make predictions as to what it all means. ML makes it interpretable via a computer, which is often many times more accurate than what people can calculate independently. For instance, ML models for image recognition can "memorize" billions of pixel gradients and recall them with perfect accuracy, whereas the human mind will start to lose recall after fewer than a hundred. Trying to recognize a thousand hieroglyphic symbols would be taxing and time consuming for most people, but for ML it would be easy and quick to generalize the shapes and recognize them with accuracy in moments.

2) What are the potentials and the aims of this technology?

With generalized ML, many hard problems (millions of variables) may be solved relatively quickly (albeit expensively) and, more importantly, preserved for future growth and reuse. Computer vision is one area where ML stands out as a particularly well-suited solution. Using NNs with a combinatorial approach of techniques for abstracting large amounts of tensors (recurrent where data is sequentially sensitive, convolutional where object groups can be generalized for rapid recognition or generative and predictive features), models have been created that enable computers to recognize faces, objects, motion (in video), moods, and ethnographic features in near-real time. These models are applied to video to identify scenes, movement, spatial relationships, and a host of other elements that until a few years ago were solely the domain of the human mind.

While NNs and the variants of how to abstract data have been around for 30 or more years, until recently the complexity and necessary computational power required made them exotic tools used sparingly. In November 2016, Google released an open-source platform called TensorFlow, which, when combined with scaled cloud computing (being able to rapidly scale processing power for a period of time necessary to run massive calculations on demand), hides most of the complexities of ML while focusing on making it simple to develop models from uploaded data. It's similar to how database platforms (e.g., Oracle, Microsoft SQL, Teradata, etc.) made it possible for anyone to store and retrieve structured data in the 1980s without having to understand the underlying theory and technology. TensorFlow is among several platforms that enable ML to be a core part of any complex decision making without having to understand or optimize the algorithmic structures or underlying code to develop training sets and models.

In some ways, it's convenient to compare ML and NNs to how computers evolved from highly specific scientific tools to devices stored in our pockets that contain a collective record of human intelligence. What was once an avant-garde computer-aided mathematical exercise is now embedded in your life in unexpected ways—from recommendations you get on Yelp or Netflix, to how you play games, to how your bank detects fraud, to how apples are grown and supplied to your local grocer.

Highly transformative technologies tend to be measured exponentially, doubling upon themselves in iteratively shorter times—for example, the use of electricity, computing power relative to transistor size, the use of the internet, or even biological development, which necessitated billions of years to transform microbes into animals but only a few hundred thousand years to make modern humanoids. ML is a transformative technology that may be the most impactful exponential example we have seen to date.

ML models today are still relatively independent of one another and represent discrete efforts to categorize segments of information: objects in pictures, word meaning, elements moving in video, etc. The models are laboriously created from numerical and representational data (e.g., numerical pixel values contained in a picture, stock prices, biological facts) and generate models that make continual use possible. For instance, Google created an open-source model that recognizes more than 200,000 objects from within images. Private models that are used for determining what people are searching for and quality of video and audio continue to evolve and be applied to complex problems (e.g., to predict what will happen next based upon a series of images, to determine if a statement is sarcastic, or to predict stock market prices).

With TensorFlow and other platforms shifting emphasis on ML to creating numerical representations of data and subsequent models that serve as a basis of collective knowledge, ML has the future potential to model human comprehension that is distributable in hours, versus the 20-plus years of learning the human brain requires. This will possibly occur with greater accuracy and the ability to classify and predict in nanoseconds, thousands of times faster than organically possible.

This leads to the controversy around ML (often called AI in this context) and what is known as the singularity.

3) Conversely, what risks might this technology hold—socially, economically, or otherwise?

The singularity is a concept forwarded by Ray Kurzweil that suggests that at an intersection starting in the mid-2020s, computing power will equal the processing power of a single human brain (2 * 10^16 calculations per second, or approximately the number of neurons that can fire simultaneously in the average brain) for about $1,000. In a nutshell, this means that we'll have reached the point where the cost to produce a processor that can compute on par with a brain will be trivial. While this does not mean that computing is the same as a human brain, it represents a milestone in calculated complexity and what by extension it can achieve (e.g., pilotless air travel, common use of driverless cars, etc.).

By 2050 the calculation capacity shifts to the equivalent of the entire human population for $1,000.

Let that sink in for a moment: all the brain capacity of the human population in something at about the cost of the phone in your pocket. Imagine what happens when we have billions of devices with that level of simultaneous computing complexity at our fingertips! Presently, ML's applicability is gated by computational power, the relatively small amount of models created, and the fact that these tools are expensive to utilize. But when computational power grows to billions of times greater than we have today, without a cost constraint, what *couldn't* be modeled? And by then, what will remain as yet un-modeled?

At some point, the use of quantum computing will come into play, potentially making it possible for ML to recognize and match data and models without human intervention (this is also called non-supervised ML, but it isn't at a point where it is recognizing data-modeling approaches in an efficient way—yet). Combine that with the certainty that computing power will grow and vast amounts of data will be categorized and represented in stored cloud-based models, and we very well may see the science fiction of AI develop into a reality.

At worst, some think that this is the basis for a dystopian, AI-controlled future. Others see it as the end of economic progress, as AIs will be able to make decisions about and control all but the most manual tasks. While these are interesting thoughts to contemplate, there are practical conditions and risks that will present themselves far before we reach that sci-fi vision.

The key problem to consider is that ML is still human-directed and given to bias, mistakes, mislabeling, and a host of other problems that result from unintentional mistakes in creating models. Take, for instance, the example of ML models trained to recognize cancer in x-ray images. In tests of the model's accuracy, it beat experts in visual diagnosis by statistically significant levels, but when it was used with new data, its accuracy plummeted. Why? As it turns out, the names of the hospitals—which often contained the word "cancer"—were used as labels in the training data sets. Logically, people being treated at cancer centers were likely to have cancer; the model factored that data in and biased itself unintentionally.

The point is, ML model creation is subject to and amplifies every human error that exists. Often, it's not possible to find the flaws in the data without substantial trials to determine its accuracy. Because the outcome of ML model predictions will scale to extend to nearly every aspect of human life, the risk is amplified to world-changing levels.

Considering how few real-world things and their related abstract manifestations exist as publically accessible models today, we have a long time before the catalyst of the singularity will result in the conjoined and independent fabrication of an all-mind. However, when that happens, the quality and accuracy of the models we build today will have become the building blocks for future composite models, recursively and incrementally adding more knowledge to itself. If undetected flaws are built into the models, it could be catastrophic, depending on the extent to which humanity becomes reliant upon it for sustenance.

More urgently, our present pursuit of ML representations over the next decade will lead to crime, manipulation, and state-sponsored acts of war. Even now we see shades of this happening today in terms of manipulating inputs in social media to drive ML-derived feed-preference models, affecting what shows up in your newsfeed (stories that are suggested for you, or that are put at the top of your list when friends post them). Imagine a weapon suddenly suspending or corrupting the network of ML models operating infrastructure; entire countries could be cast into silence without a shot being fired or a bomb being dropped. Perhaps control of ML/AI could become the next Cold War, should controls not be put in place to limit their civic impact.

Closer to today's reality, the threat of hackers grows exponentially when, instead of ransoming data, they subtly alter ML models, in turn affecting every subsequent classification and prediction the models drive. When scaled to financial systems, public utility controls, commerce, and food production, it's imaginable entire countries are at risk, to the extent they rely upon models.

Of course, not everything must be a dystopian future. Rather, it's a call to action to insure we are thoughtful in what safeguards we build in now. I'll quote Ray Kurzweil on the risks and potential of exponentially transformative technologies:

> People often go through three stages in examining the impact of future technology: awe and wonderment at its potential to overcome age old problems, then a sense of dread at a new set of grave dangers that accompany these new technologies, followed, finally and hopefully, by the realization that the only viable and responsible path is to set a careful course that can realize the promise while managing the peril.[1]

Brian Mulford works in Global Product Development at Google, where he blends technology, machine learning, analytics, digital anthropology, and usability to develop products that matter. Mulford received a MBA from the University of Southern California's Marshall School of Business after earning a BS in Electrical Engineering and Computer Sciences from Massachusetts Institute of Technology.

1 Ray Kurzweil, "The Intertwined Promise and Peril of Twenty-First Century Technology," *Living with the Genie*, spring 2013, http://www.columbia.edu/cu/genie/t04_02.html.

Visions of a Driverless World

Tea Uglow

Define: a neural network (NN), or machine learning (ML), or artificial intelligence (AI). These are loose terms for a topic that is a bit like quantum physics. It is a transformative form of computing that allows machines to effectively "learn" from huge databases of information called "libraries" until the software itself can "create" new content. "Content" may mean anything from literally driving a truck across America to composing a Bach-like fugue. The outputs of artificial intelligence are limited by the design of the network and the quality of the training library, but capable of a regression analysis, where the neural networks can identify patterns of data that are far beyond the scope of human facilities, leading to outcomes that can seem either eerie or extraordinary—like a truck that drives itself. For many of us, we will understand that they exist, but feel life is too short to care. For us, neural networks might as well literally be quantum physics: undeniably important, definitely real, and mind-numbingly hard to comprehend.

Already we are beginning to see AI's role in driverless transport. This will revolutionize human infrastructure over the next decade, and will certainly be a very obvious benefit of "machine learning." Just one part of that "revolution" will be reducing the 1.3 million deaths (and 20 million injuries) caused each year (mainly) by human drivers. For those still having accidents, the arrival of driverless diagnosis will be transformative to the medical industry, freeing the art of diagnosis from human bias, exhaustion, or simple prejudice—or, at the very least, providing a pretty impressive second opinion.

This is leading to some unusual academic programs. For example, in 2017, a peer-reviewed Stanford article reported 91 per cent accuracy in distinguishing the sexual preference of men using a deep neural network based on facial recognition alone. In other words, it had a 91 per cent "gaydar" hit rate. Humans score around 50:50 in the test used—as one might expect in a test where you choose between two faces. The humans can't tell; the machine can. It has learned some skill that we cannot divine. We cannot ask what it has learned; we can only conjecture while the neural network improves on its statistic.

The study also notes that many countries in the world have criminal statutes regarding homosexuality and are actively pursuing this model of law enforcement:

> The laws in many countries criminalize same-gender sexual behavior, and in eight countries—including Iran, Mauritania, Saudi Arabia, and Yemen—it is punishable by death (UN Human Rights Council, 2015).[1]

The authors of this study, Michal Kosinski and Yilun Wang, argue that if utilized universally, such technology could result in the legal imprisonment or death of LGBTQ people; therefore the accuracy of such technology is of crucial importance to policy makers, lawyers, human rights advocates, and, naturally, the homosexual community. It is perhaps an example of demonstrating the dangers of opening Pandora's box by opening the actual box. The point is not that it is a probable apocalyptic scenario, especially given the number of apocalyptic scenarios that appear more likely at the time of writing. However, for the queer community, it is simply the silent foghorn of heteronormative biases when it comes to machine learning. We are now in a strange place where the potential models of control for future generations are being developed by a tiny subset of a demographic with a singular mindset, low empathic or social skills, and fixed cultural norms: groups of similar minds building artificial minds to learn from data gathered from a digital global hive mind, with all its many prejudices.

At an industry level, this academic naivety is echoed with enthusiasm, devolving decision making to models of pattern recognition that defy analysis or synthesis into human-readable "knowledge." Patterns are based on data that is weighed, considered, or analyzed by a "training process," looking for the optimal number of variables and avoiding omission bias, as if the existing models of behavior online were the epitome of human behavior and intellect. It creates the specter of a world of knowledge that is algorithmically derived and unreadable—a world based on the injustice, idiocy, and entrenched biases of majority-think, with none of the vision, idealism, unilateral capacity, or romance of the solitary human imagination. Sadly, neural networks cannot fix for bigotry, a problem not helped by the sense that there is no place for the humanities in this new world order. Code. Math. Science. Engineering. These hold the keys not only to global economic power but to global culture as well. After all, what are the options? Poets?

There are a number of artists across the centuries whose work has examined both the cyclical and cynical qualities of "power," whether they are celestial, economic, or military, but rarely do they stir us to action; they tend to the reflective, not the imperative. In a present of algorithmic bias, cyber warfare, and drone surveillance, our artists are often more elegiac than prophetic. For progress, we look to industry.

The World Economic Forum lists among its top ten problems facing the world: agriculture food security, economic growth and social inclusion, the future of the global financial system, gender parity, and even the future of the internet. Google CEO Sundar Pichai wrote in the company's 2016 founders' letter that Google had taken "another important step toward creating artificial intelligence that can help us in everything from accomplishing our daily tasks and travels, to eventually tackling even bigger challenges like climate change and cancer diagnosis."[2] So the potential is not insignificant.

Artificial intelligence will, accordingly, be solving cancer, fixing social inequality, or preventing global warming. Currently it is writing screenplays and making music. "Humans making culture"

---

1 Michal Kosinski and Yilun Wang, "Deep Neural Networks Are More Accurate than Humans at Detecting Sexual Orientation from Facial Image," *PsyArXiv Preprints*, September 2017, https://psyarxiv.com/hv28a/. Note that the study used only Caucasian faces and did not attempt to filter for queer, transgender, or bisexual identities.

2 Sundar Pichai, "This Year's Founders' Letter," April 28, 2016, https://www.blog.google/topics/inside-google/this-years-founders-letter/.

is not in the UN's list of the world's top ten problems; neither is distinguishing gay men from straight men. There is no known problem with the creation of "art." We make a lot of art. Some—most—is awful. Some is commercially successful, some critically successful, and some is transcendent genius that allows us to see the world in new ways—art that changes the way society exists and understands itself. In some ways, art is in robust health (while chronically underfunded). But underfunding is the least of its worries.

Let's move for a moment to the secondary and tertiary consequences of our artificial or "driverless" culture. What will MoMA show? And how do our children evolve into artists if there is no economy supporting the early, grunt work–like ages of an artist, when everything we do is kind of bad? Certainly worse than an AI would produce. Who will pay for the bad art? The implications of artificial intelligence disrupting the structure of the creative economy at the entry level is interesting to consider. The endless tide of words, pictures, music, and film currently generated, edited, and curated by humans hides the fact that the humans involved learn simultaneously. That process is likely to get automated. We are drifting past headlines like: "Google's Art Machine Just Wrote Its First Song"[3] and "Robo Tunes: This Is What Music Written by AI Sounds Like."[4] Or, in literature: "Google AI Is Really Good at Writing Emo Poetry."[5] And from the arthouse: "This Short Film Was Written by a Sci-Fi-Hungry Neural Network,"[6] to the multiplex: "IBM's Watson Sorted through Over 100 Film Clips to Create an Algorithmically Perfect Movie Trailer."[7] Considering this is an industry that is not really broken, why "fix" it?

The language used can be unsettlingly anthropomorphic: for example, a reliance on soft phrases like "training" AIs, or the idea that art, music, or scripts were "found" on the internet, then "fed" or "shown" to the computer. A human artist certainly goes through this process, and we cannot tell whether the human algorithm invests anything new in the process—perhaps we only ever derive our outputs too? "No such thing as a new idea," said Mark Twain—and every other writer ever, according to Google. Yet we intuit that humans can synthesize, and are also pretty sure that a mechanical algorithm can only derive. We've finally arrived at a real question that has been a hypothetic sci-fi staple: can an artifact create?

A new generation of artists will emerge that has always worked with machine intelligence, and doubtless to this generation these entities will simply be "tools," analogous to a camera or a light bulb. For a generation or two, these artists will also, necessarily, be computer engineers who have the skills or who can afford to employ a team, through either patronage or funding. It is perhaps not the most democratic of prospects for culture, but this is progress, or maybe a regression to the Renaissance studio model. It is back at the industrial shop floor that the implications are perhaps more complicated. In many futures, the part being automated is the human input. The neural network simply mimics and reverse-engineers historic human creative processes in order to generate cultural content that is equal to or better than human outputs. We might ask: what benefit does this bring to either artists or society? What are the secondary consequences of "libraries" of culture in which the works of Shakespeare need not be attributed or musicians remunerated because the output is a novel "creation"? Even as the educators of our automators, we cannot imagine that there is a glowing future for humanity in that industry. Just as for truck drivers. Or for the queer community.

The effort to literally automate the creation of culture is considerable, yet presumably not the ultimate goal. It is more likely that cultural output is a convenient playground for the "driverless" future. We are not expecting the tech giants to linger long in the playground; it is a second generation of perhaps less well-intentioned corporations that will likely be watching, and quite often consumers who are presented with the bill when it is too late to decline. So we see echoes of the future of algorithmic culture through the spectrum of dancing soldiers, or machine-learned Mozart, or cut-up sci-fi. They are painting an algorithmic future, not because it is needed, but because you can't hurt anyone with a film trailer, whereas automobiles, drones, emergency rooms, and financial services leave a little more space for liability.

And since we're inside this circus of trust, we should admit that most of us won't be able to discern human culture over algorithmic. It will be like picking wine at the restaurant. And if we can't tell, how can we care? Will we build complicated "organic" guidelines, a kitemark, a taste test to ensure our culture isn't just being churned out by a poor, basement-dwelling Chinese supercomputer? It would be nice if humans were cultured enough to say "this is intolerable," but I think we all know we aren't, and we won't—at least not until we can no longer recall what we have lost.

Tea Uglow leads part of Google's Creative Lab, specializing in work with cultural organizations, artists, writers, and producers on experiments using digital technology at the boundaries of traditional cultural practice. Uglow studied fine art at the University of Oxford's Ruskin School of Art before completing two further degrees in book arts and design management at University of the Arts London.

3 Russell Brandom, "Google's Art Machine Just Wrote Its First Song," *The Verge*, June 1, 2016, https://www.theverge.com/2016/6/1/11829678/google-magenta-melody-art-generative-artificial-intelligence.
4 Joon Ian Wong, "Robo Tunes: This Is What Music Written by AI Sounds Like," *Quartz*, May 27, 2016, https://qz.com/694081/this-is-what-music-written-by-ai-sounds-like/.
5 Anna Cafolla, "Google AI Is Really Good at Writing Emo Poetry," *Dazed*, May 16, 2016, http://www.dazeddigital.com/artsandculture/article/31140/1/google-ai-is-really-good-at-writing-emo-poetry/.
6 Brittany Vincent, "his Short Film Was Written by a Sci-Fi-Hungry Neural Network," *Engadget*, June 9, 2016, https://www.engadget.com/2016/06/09/this-short-film-was-written-by-a-sci-fi-hungry-neural-network/.
7 Clinton Nguyen, "IBM's Watson Sorted through Over 100 Film Clips to Create an Algorithmically Perfect Movie Trailer," *Business Insider*, August 31, 2016, http://www.businessinsider.com/ibm-watson-morgan-movie-trailer-2016-8.

John Gerrard in Conversation with Adam Kleinman

Adam Kleinman: Why are you interested in neural networks (NNs)?

John Gerrard: In the mid-1990s, I first began to work with 3-D scanning, which I thought of as a sort of sculptural photograph. A record of the real, but retaining a dimensionality—one could turn it, look around it, even if just in computer space. This struck me as interesting in art historical terms. At that time, I was hearing lots of talk about game engines, but was not able to connect with one properly until 2001, at the Futurelab in Linz, Austria, which had a computer science lab–type structure and where I was an artist in residence. Game engines grew out of military simulations—from flight cockpits and battlefields, for instance—and bled into popular culture as video games. By the early 2000s, they were becoming increasingly visually sophisticated, as seen in games from, for instance, Rockstar Games. On arrival at the Futurelab I was given a working budget and matched with a modeler, programmer, and producer as a matter of course, a structure I still work within. The art school idea of doing everything yourself—which I had struggled with for years—was given short shrift. Within a year I had produced a first and extremely basic artwork, in the game engine, and over the next two years in that lab—where I stayed from 2001 to 2004—I was able to formally develop the outcomes toward a particular cold and polished aesthetic that I was pursuing. I was also able to integrate for the first time the kind of temporal component that I had imagined, in that the worlds I was making could unfold over a year in solar terms.

It came together in 2007 when I produced *Dust Storm (Dalhart, Texas) 2007*, which was displayed at an exhibition curated by Linda Norden entitled *Equal, That Is, To The Real Itself* at Marian Goodman Gallery in New York. This constituted a sort of coming-of-age in terms of discussions around the work and also in terms of the medium for me. Over the last ten years, 2007 through 2017, I have worked exclusively in game-engine space—initially in an early engine called Quest3D and more recently in a Russian one called Unigine—, which is extremely powerful. Works designed in that engine such as *Solar Reserve (Tonopah, Nevada) 2014* and *Western Flag (Spindletop, Texas) 2017* are at times indiscernible from cinema, yet they retain a slippery relation to the real—one hovers between trusting one's eyes and questioning them. Nonetheless, while these works are extremely "good-looking," they are not particularly "smart" in terms of emergent behavior or autonomy of any sort. It was in this context that I began to hear and read more of neural networks, both in the popular press, in regard to self-driving cars, but also through kind of gimmicks such as Google's DeepDream outcomes that went all over the web. Computers dreaming dogs into everything, as there were too many dogs in their training sets, etc. In truth, the level of manipulation that gives rise to DeepDream-type outcomes is very significant—they are almost fictions as such—but, the level of public response to the myth of the dreaming computer was remarkable. Thus, for me, the DeepDream images were one of the first times I was led to inquire: so what is this neural network everyone is talking about?

Cut to early 2016 and my receiving an Art + Technology grant from the Los Angeles County Museum of Art (LACMA). I subsequently found myself in a room in L.A. with working technologists, having brief meetings with all of them. Three of the meetings stood out, with representatives from Google, NVIDIA, and Hyundai, as each entity was working with neural networks across multiple applications—broadly, image recognition and manipulation for Google, speech recognition and production streamlining for NVIDIA, and self-driving cars for Hyundai. Brian Mulford from Google let me know that his company had just released a powerful NN called TensorFlow—and that he would support my engagement with it. Within a month or so, my long-term programming collaborator Helmut Bressler was tentatively entering this very new and unfamiliar space. I have to be clear, however, that fundamentally I had absolutely no idea what a neural network really, actually was, bar the kind of understanding one can derive from Wikipedia. Thus the LACMA grant at the outset really allowed a total novice the means to dedicate studio time to a sort of NN 101. Even with that, I have only really come to understand the method from actually working within the network. One cannot program a neural network as such, no less than you can program a pet—one can *only* train it. I had a strong brain block against understanding that—following more than fifteen years of programming outcomes in game engines—but eventually I got it, which has caused a very profound shift in my thinking about many things. I had been seeking a way to move beyond the rigid, constrained, almost "cinematic" components of my use of game engines (in that all animations were fundamentally instruction-based and timeline-based), and suddenly a new space emerged, which—as I will outline—I had been unwittingly searching for since around 2011.

AK: How and why did you come to the leaf-covered figures, and why are they dancing—and not, say, fighting?

JG: I had developed an idea to produce a choreographic generator using the LACMA grant, and programmer Helmut Bressler spent around six months experimenting with and looking at existing solutions. We found one that had potential (while still looking absolutely awful), but it was a good basis to start. Rewinding a little: if one is looking to move characters in an engine, one needs to *animate* them. I had absorbed 3-D motion scanning into my practice very early on, first using it in *Oil Stick Work (Angelo Martinez/Richfield, Kansas)* in 2008, for instance—where a worker paints a barn black very slowly over a 30-year period, from 2008 to 2038. Motion capture is powerful; you get captured movements from life as data. It suits the wider origin of the work in 3-D scanning and a sort of smooth pipeline from "life" into its representation within the virtual. We scanned all movements, but had to develop very laborious techniques to ensure that the animations would not "jump" as the character transitioned from one action to another. In short, here was a wide-open game-engine medium with enormous scope for what I would describe as temporal play, but locked to the historic timeline if you wished to produce movement. It all got much worse with what I call the Exercise series—a body of work that created semi-fictitious military exercises in real sites.

This began with *Live Fire Exercise 2011* in collaboration with Wayne McGregor for the Royal Ballet, and developed through *Infinite Freedom Exercise (near Abadan, Iran) 2011* and *Exercise (Djibouti) 2012*, ending with *Exercise (Dunhuang) 2014*. Many of these projects involved massive motion-capture sessions and literally years of work processing the animations and streaming them into the engines. It was a kind of desperate and nightmarish work: expensive, frustrating, and kind of disappointing, as the sort of flattening of the lifelike qualities of the motion captures (themselves eye-wateringly expensive to acquire in massive camera-based capture labs in Prague and elsewhere) was very profound. In retrospect, each of the Exercise series was an artwork in search of a neural network, but I did not have one, nor were NNs in the mix in 2011 for me, and there it is.

Jumping from 2016 to many years earlier, I had read of leaf-covered figures operating within folklore in *The White Goddess: A Historical Grammar of Poetic Myth* by Robert Graves, and spontaneously decided to remake the form. I grew up in rural Southern Ireland, so had plenty of green spring leaves to work with, alongside my sister, Esther, who was happy to collaborate. So, around 2003, I produced these strange throwback images of a "lost" character, cloaked in leaves, standing around my old childhood home.

In essence, the leaf-covered figure relates to what is known as the "Green Man," often carved into medieval cathedrals: a leaf-covered face peering down from pillars and roofs. However, my interest in it was to do with the historic acknowledgment of the change of the seasons, the return of the sun, and the solar panels of leaves—of many sorts—providing in time a great bounty for the local human population. I have always interpreted the leaf-covered figure as a symbol of the exchange between the natural world and the power of the sun to the human population—something that electricity, refrigeration, and, in particular, petroleum has made extraordinarily distant by the late 20th and early 21st centuries. I wanted to see this character—referred to only in text by Graves—again, so I remade her. In so doing I was struck by a sort of curious latent power in the presence of the character. The character made sense, somehow, and was very familiar. I thought it was important, but in the end left it there until 2016—13 years later—when I revived the character to perform in a research work titled *Neural Exchange (Leaf Covered Figure) 2017*, for LACMA.

To do so, my studio and I traveled to the Wienerwald in early summer 2017 and dressed an actor in leaves, closely photographing the results and the materials. The modeler, Max Loegler, then spent the summer making virtual leaves and twigs and—using the original pictures as a guide—re-dressing a virtual character to look the same. We were also able to develop some basic physical properties for the leaves, so they could bounce and move using local physical data.

Outside of that basic presence, the idea of exchange is absolutely fundamental to the decision to use the leaf-covered figure as a character. In essence, the figure is a mythic/historic one whose presence recalls the fundamental exchange between the sun, vegetation, and humanity. The neural network evokes an exchange again between the organic world—in this case the mammalian brain—and the human population, in that the neural network copies and mimics the functioning of the brain, and in so doing has unlocked an extraordinarily powerful new methodology. These techniques sit behind much progress in the tech world: image recognition allowing for self-driving cars and mass surveillance of social networks for commercial or even political ends is only possible with neural networks. The list goes on.

In my work and research, I wanted to align the leaf-covered figure as an acknowledgment of a locale as such, and an exchange. The NN probably needs to do a little of the same, or at least we have to be aware of the power of this technology and from whence it emerged. We derive huge benefit from copying the structure of the brain—and I wanted this strange figure to be animated by this benefit in movement terms—while also speaking of the wider benefits of our sustaining a relation to the "natural" world. To square a circle of development from when super-powerful entities were once personified by trees to now, when trees are sort of quaint things on the edge of the techno-campus, but where that campus is nonetheless deeply invested in replicating organic forms for power, progress, and profit.

In relation to fighting or dancing: somehow, in the search for a vocabulary of movement for my leaf-covered figure, the idea of martial arts came in. In particular, I was interested in karate kata, in which a very formal vocabulary of movement exists, which can be motion-captured and identified as such if it emerges from the neural network's choreographic generator. In the end, working within the martial arts arena brought an excessively aggressive feel to the research, so we backed away from those practitioners and worked instead with a ballet dancer, Esther Balfe. As a dancer she can both repeat movements many times, on account of her ballet training, and also develop a movement vocabulary and reproduce it, which is invaluable. So we have now moved only to work with Esther for the research, and it is her four basic movements that have been used to train our NN. These actions are processed by the NN and go on to animate the leaf figure—which, surprisingly, is working. What is most powerful within that is that this figure, which exists as a piece of software, no longer has a performance duration; she, animated by the NN, can perform into perpetuity—smoothly and, while it is a small data set, innovatively. This is kind of amazing for me.

AK: What did you learn from training an NN?

JG: Fundamentally we developed a generator written in Python using Tensorflow. It's designed to create human character animations based on sets of data acquired through motion capturing. This outputs a perpetual dance performance of a limited sort. As outlined, doing this has allowed me to come—slowly—to an understanding of what this thing actually is, which was an epiphany. And to understand that the more data and the more computational power one has access to (we have very little here in the production space), the results—which are already very encouraging—can develop in leaps and bounds. There is no real limit to the NN potential—only the size of the training data and the size of the cloud that can process that data. I went into the research seeking a generator that could animate a character seamlessly, with no duration and without animation stringing; however, I emerged with much, much more. Coming to the second question, what I sense I am seeing is this strange feeling that I am equipping my engine—via the NN—with a kind of imagination. I frame that imagination through the kind of actions that are captured; however, looking at the *Leaf Covered Figure* perform—as I am doing now—something else is emerging.

AK: What do you think will be the future relationship between NNs and what normatively is called human creativity?

JG: To a degree, I suspect that the NN may allow a sort of computational imagination to emerge. That is the feeling that I am getting from this very early work in the form using motion capture as training set. From this limited set of four movements, the more we train the network, the better the motions get and the more human the character becomes, in that the small quirks, movements, and pauses are "learned" more fully and emerge in the character from time to time. In addition, I am seeing new actions that were not captured—these are, in a sense, either approximations of movement that happen to work or a new movement dreamed up by the machine from the training set. I do not feel threatened by this at all—more the opposite: I am artistically exhilarated by it. The machine could not make this character or imagine this research, but it can free up the space(s) for performance within it, and, within a limited frame of parameters, enrich it, much as a close collaborator might.

John Gerrard (b. 1974) is best known for his commitment to large-scale works that take the form of real-time computer simulations, created in painstaking detail over the course of many months or years. Often exploring geographically isolated locations, the works frequently refer to structures of power and networks of energy that have made possible the expansion of human endeavor in the past century. Recent solo presentations of Gerrard's work include *Western Flag*, commissioned by Channel 4, London (2017); *X. laevis (Spacelab)*, commissioned by Wellcome Collection, London (2017); *Power.Play,* Ullens Center for Contemporary Art, Beijing (2016); *Solar Reserve*, Lincoln Center in association with Public Art Fund, New York (2014); *Sow Farm*, Rat Hole Gallery, Tokyo (2014); *Exercise*, Borusan Contemporary, Istanbul (2014); *Pulp Press (Kistefos) 2013*, a permanent installation for Kistefos Museet, Norway (2013). *Exercise (Djibouti) 2012*, Modern Art Oxford, Oxford (2012); and *Infinite Freedom Exercise*, Manchester International Festival, Manchester (2011). His works are in the collections of the Fondation Louis Vuitton, Paris; Irish Museum of Modern Art, Dublin; Irish Arts Council Collection, Dublin; Los Angeles County Museum of Art, Los Angeles; Museum of Modern Art, New York; Tate, London; Pinakothek der Moderne, Munich; Hirshhorn Museum and Sculpture Garden, Washington D.C.; and Inhotim Collection, Brumadinho, Brazil. John Gerrard lives and works in Dublin and Vienna.
He is represented by Thomas Dane Gallery, London, and Simon Preston Gallery, New York. For further information see johngerrard.net.

Adam Kleinman (b. 1978) is a writer, curator, educator, and sometime performer. Former dOCUMENTA (13) agent for public programming, he is currently editor in chief and adjunct curator at Witte de With Center for Contemporary Art. At Witte de With, Kleinman co-runs the arts and politics journal *WdW Review*, and has led the curation of numerous exhibitions and public programs. He was curator at Lower Manhattan Cultural Council, where he created the interpretative humanities program Access Restricted and developed LentSpace, a cultural venue and garden design by Interboro Partners, which repurposed an entire vacant Manhattan block. He is a frequent contributor to multiple books and magazines including, *Art Agenda*, *Artforum*, *e-flux journal*, *frieze*, *Mousse*, and *Texte zur Kunst*.

This catalogue is published on the occassion of the performance *John Gerrard: Neural Exchange* at the Bing Theater, Los Angeles County Museum of Art, November 1, 2017.

Thanks to LACMA, Michael Govan, Amy Heibel, Joel Ferree, Jennie King, Rita Gonzalez, Miranda Carroll, Kenric McDowell (Google), Andrea Held (Google), Brian Mulford (Google), John Suh (Hyundai), Bryan Catanzaro (NVIDIA), Chris Malachowsky (NVIDIA), Unigine and Günther Weinlinger/7reasons Medien GmbH.

Authors: Joel Ferree, Amy Heibel, Brian Catanzaro, Brian Mulford, Tea Uglow, Adam Kleinman and John Gerrard
Editor: Adam Kleinman
Copy Editor: Claire Crighton
Design/Layout: Caroline Ecker

Images: Stills from *Neural Exchange (Leaf Covered Figure) 2017*

Printing: Druckerei Jentzsch, Vienna