arts

MDPI

Foreword

In 1843, the renowned mathematician Lady Ada Lovelace – often known as the world's first computer programmer – published her famous "Notes" (Lovelace 1843), and which, remarkably, included some insightful visions on computational art.

The "Analytical Engine" was a mechanical computer designed by the inventor Charles Babbage in 1837. It was only partially built in their time, but this did not stop Lovelace from designing programs for it, and theorizing on its potential. Furthermore, she had the insight to see that it could go beyond just acting on numbers to solve equations, but could as well do symbolic manipulation to perform true general purpose computing. Not only that – she had the vision to foresee that the Analytical Engine could be used to compose generative music:

> Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.

Or even that it could create generative imagery: "We may say most aptly, that the Analytical Engine *weaves algebraic patterns*, just as the Jacquard loom weaves flowers and leaves."

But it's another of her rather controversial statements that I'd like to recall:

> The Analytical Engine has no pretensions whatever to *originate* anything. It can do whatever we *know how to order it* to perform. It can *follow* analysis; but it has no power of *anticipating* any analytical relations or truths. Its province is to assist us in making *available* what we are already acquainted with.

And almost two centuries later, we are still grappling with this statement, and still trying to understand our relationship with the machine. Is it simply assisting us to make available what we are already acquainted with? Or can it originate anything? Can it anticipate any analytical relationships or truths?

And of course, we are not the first to revisit these questions. In 1950, over a century after Lovelace published her "Notes", the famed computer scientist Alan Turing also addressed these topics in his seminal paper "Computing Machinery and Intelligence" (Turing 1950). He reframed the question in the context of surprise, asking whether a machine could ever surprise us. And, he added, "Machines take me by surprise with great frequency." But his main proposition – echoing his collaborator Douglas Hartree – was that in order for a machine to really create something original, it should have a property which would not have been available to Lovelace or Babbage. And that property, he concluded, was the ability to learn: "Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain."

And this idea – now known as *machine learning* – is the very concept that underlies the recent surge in the fields known as *artificial intelligence* (AI) and *AI art* (though this latter label has not been adopted by everybody – including myself – for reasons which will soon become apparent).

The academic field of artificial intelligence is rooted in computer science but spans many other disciplines such as psychology, neuroscience, statistics and philosophy. Since the 1950s, AI researchers have been thinking about the properties required to be able to create 'intelligent' machines, and they've been designing and building computational models of such systems.

For a similar amount of time, artists have been independently investigating the role of computers in art, giving birth to many overlapping subgenres such as computational art, generative art, (new) media art etc. As early as the 1960s, artists such as Harold Cohen were already engaging with AI (McCorduck 1990), and as early as the 1980s artists were using machine learning (such as evolutionary algorithms) for their artworks.

However, just as the mainstream emergence of the internet in the 1990s gave birth to a new era in computational art – known as net art – the last few years also marks a new era in computational art. As I previously mentioned, this new wave is colloquially referred to by some as 'AI art'. But since it only refers to very specific, recent technologies, it would probably be more accurate to describe it as (the much less catchy) 'deep learning art'.

Deep learning is a form of machine learning based on large artificial neural networks and massive amounts of data. The algorithms date back to the 1980s (or even earlier depending on who you ask), but it's only recently that the technology has been broadly accessible and practically useful at scale. This is due to the vast increases in computing power that we have recently developed, and the massive amounts of data now available required to 'train' these huge neural networks. But of course the underlying reason for this recent explosion in deep learning is the political and social climate that requires, supports and funds this research; and needless to say, the primary purveyors of deep learning research are also the primary purveyors of the data economy and mass surveillance – both state-sponsored and commercial – such as Google, Facebook, Amazon etc.

This situates deep learning, whether as a medium, tool, or subject matter, in a very unique position, and with very particular challenges and opportunities, for artists.

Any topic surrounding big-data immediately raises questions regarding privacy and ownership. On one end of the spectrum we have organizations infringing on the rights of their customers by stealthily harvesting, selling and otherwise exploiting their data. On the other end of the spectrum we have individuals (such as artists) making work using data belonging to others. And in between, we have combinations that present more complicated ethical and legal challenges, such as the use of a neural network designed by one person or group, implemented by a second person, modified by a third person, trained by yet a fourth on data owned by a variety of individuals and collected by a fifth person, and using scripts written by a yet a sixth (true story!). This makes it very clear that our old concepts of 'ownership' or even 'authorship', and our legal and economic systems built on these concepts, are becoming obsolete in these new digital ecosystems.

Given, furthermore, that deep learning can be thought of as a technology which attempts to extract meaningful information from vast amounts of big-data, any progress in deep learning thus has the potential to impact any enterprise which is big-data driven. And in this day and age, practically all of our enterprises are big-data driven, from physics, chemistry and biology to finance, psychology and even politics.

And these deep learning algorithms are notoriously inscrutable. Often referred to as black boxes, they are incredibly difficult to decipher, and meaningfully control or correct if they give undesired outcomes. As Alan Turing himself said in his 1950 paper, "An important feature of a learning machine is that its teacher will often be very largely ignorant of quite what is going on inside."

But this is also why they are so powerful. This is precisely what makes these big-data driven learning algorithms both terribly exciting, but also desperately terrifying. We are already seeing the unexpected negative consequences of for-profit companies using closed-source, closed-data, proprietary software to make critical decisions affecting thousands, if not millions or billions of lives. We already have algorithms in use deciding which job adverts to show people – and not only that, but actually learning from current salary schedules to be sexist, by showing higher paid jobs to men only. We already have platforms, such as YouTube or Facebook, algorithmically spreading targeted misinformation and increasingly extreme propaganda. And extrapolating into the future, the potential damage these algorithms can inflict doesn't require a vivid dystopian imagination.

But these are also the technologies that will allow us to see beyond what we would otherwise be capable of seeing, just as the telescope allowed Galileo – quite literally – to look at the sky and see the solar system in a new light. By helping us see and understand patterns and meaningful information in vast amounts of data, these big-data driven learning algorithms can contribute to breakthroughs in many different fields, from helping us cure our most awful diseases, to early warnings for earthquakes or other natural disasters, to suggesting solutions for the ecological crisis we are facing. And some of these breakthroughs are quite likely to be unimaginable by our current norms, like houses that can grow organically, or cars that can photosynthesize!

The future is not written yet, and it is up to us to write it. Ultimately, technologies are not separate entities which are external to us. They are part of us. They are extensions of our bodies, extensions of our minds, and extensions of our values. They are embedded within us, and we are embedded within them – the users who use them, the researchers that develop them, the organizations that fund them – nor can they be totally separated from the motivations behind their development, and the values of those who support and promote them.

The future is not written yet, and it is up to us to write it. And as artists, we are in a unique position to help shape – or at least envision – potential futures. As artists, we try to trod untrodden paths. We try to imagine alternative realities and futures. We try to see – and share – different perspectives. We even try to *feel* those different perspectives, and enable others to feel them as well. To quote Golan Levin, Professor of Electronic Art at Carnegie Mellon University, we try to create art that "comforts the afflicted or afflicts the comfortable".

I myself, as a computational artist practicing for almost two decades, have been thinking about such topics for many years, and I have decided to focus specifically on these big-data driven learning systems – with meaningful human control! – for my PhD. And the authors of the essays in this book, collectively, have centuries (!) of experience between them in computational art and creativity. We realize that we are together entering a new chapter in this area of inquiry, and – in the context of machines that can learn from vast amounts of data – that revisiting the relationship between artists and machines, and between humans and machines, is more urgent than ever.

Memo Akten

References:

Lovelace, Ada A. 1843. Notes by the Translator [of Luigi F. Menabrea's "Sketch of the Analytical Engine"]. In *Scientific Memoirs*. London: Richard and John E. Taylor, 3. Available online: https://www.fourmilab.ch/babbage/sketch.html (accessed on 4 April 2020).

McCorduck, Pamela. 1990. *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. New York: W.H. Freeman & Company.

Turing, Alan. 1950. Computing Machinery and Intelligence. *Mind*, LIX (236): 433–460. Available online: https://academic.oup.com/mind/article/LIX/236/433/986238 (accessed on 4 April 2020).