

Artificial Intelligence and the Decoupling of Creation and Comprehension

Patrick Hebron

Introduction

Machine learning has catalyzed a radical transformation of how ideas are generated and solutions are found. It is a transformation that in many respects shakes the very foundation of both scientific and cultural discourse. In other ways, it may provide a path to the reunification of scientific and cultural pursuits - a new Renaissance.

In this talk, I would like to consider some of the ways in which artificial intelligence is beginning to extend our effective capabilities beyond what is formally understood in the sciences and then raise the question of what implications this may have in the cultural realm and, most generally, what it means for the future of human creativity.

Solution without History

When a human sits down to design the solution to a problem - be it a problem in the sciences or the design of a new car, chair or evening gown, he or she cannot help but draw upon history. Even if the designer wishes to reject past approaches and put forward a paradigmatically novel one, this must inevitably be done in relation to what has come before. We have lived in the world and have been surrounded by those earlier solutions. In this way, it is very difficult for any field of mathematics, science, engineering, design or art to break free from past assumptions.

In many ways, this is a good thing. We learn from the successes and failures of the past. We build upon the foundations of distilled knowledge that were established by prior generations, enabling us to build ever higher instead of always starting from scratch. The ability to distill knowledge and to transmit it to others through various forms of language is perhaps the defining and distinguishing quality of humankind - the thing that has enabled us to come so far in our understanding of the universe and of ourselves.

We need only to look at Thomas Thwaites' *Toaster Project* - an artistic exercise aimed at discovering whether an individual person could construct a relatively simple household object entirely from scratch, starting with mining of the raw materials - to see just how much we depend upon each other and what has come before. (Thwaites) This drooping pile of materials, crudely approximating the form and scarcely the function of a

toaster, should be proof enough that no one person possesses the time, expertise or access to materials necessary to design and engineer even a basic element of the modern world on their own.

In some respects, though, our dependence upon history as well as the knowledge and design patterns distilled by others can prevent us from seeing beyond vestigial constraints. It has been famously noted that we cannot approach the problem of building a car from the perspective of developing a faster horse. R. Buckminster Fuller illustrated this idea beautifully in his book, *Operating Manual for Spaceship Earth*:

If you are in a shipwreck and all the boats are gone, a piano top buoyant enough to keep you afloat that comes along makes a fortuitous life preserver. But this is not to say that the best way to design a life preserver is in the form of a piano top. I think that we are clinging to a great many piano tops in accepting yesterday's fortuitous contrivings as constituting the only means for solving a given problem. (Fuller 21)

Questioning even one component assumption could lead to new possibilities for the system as a whole. Doing so, however, may require the reworking of numerous embedded assumptions. To warrant this effort, the value of the newly-opened possibility should be substantial. Yet, that nascent value may not be immediately recognizable. The alternate strategy may utilize very different premises or facts about the environment from those we have encountered in our own explorations. If our own experiences have not led us to those same facts or if the strategy's successful outcome cannot be immediately observed, then we are unlikely to discern the value of the strategy a priori. A lengthy process of experimentation, refinement and contextualization may be required to develop the approach and advance it to a level that exceeds the utility and efficacy of the preceding approach.

It may be comforting to know that this problem is not unique to our human efforts and can be found in the work of nature itself. For example, suppose there is a bird with a short beak that is able to utilize a particular food source within its environment. There may be another more abundant food source that can only be accessed by a long beak and as such, the population may naturally evolve towards that expression. If, however, a medium-length beak is not able to utilize either food source, then it is less likely that the bird's beak will evolve through this local minimum in order to arrive at the optimal length associated with the more abundant food source. Instead, the species is likely to stay where it is, utilizing the less abundant food source.

There is always an opportunity cost associated with entertaining alternate strategies. With machine learning systems, though, we can begin to circumvent some of these limiting factors. With respect to the problem of local minima within a problem space and the opportunity cost associated with its exploration, the strength of artificial intelligence is not only a matter of raw intellectual horsepower, but also of how we commit machine intelligence to a particular domain or task. While humans and machines both have limited lifespans and experiential bandwidth, machines are able to transmit learnings in a somewhat different manner from us. When humans communicate ideas to one another, they start from a mental representation and distill it into symbolic language, which is then interpreted by the listener back into a new mental representation. A great deal of the nuance of that communication is lost to the low-fidelity nature of the symbolic representation. With machine learning systems, we can set thousands of learners to interact with a given problem and discover a wide variety of unique approaches. We can then select the most effective of those strategies or in some cases even fuse the mental representations of numerous machines together to produce ever more optimal solutions without having to pass through low-fidelity linguistic distillation.

Though the utilization of additional computational resources does not truly circumvent the concern of opportunity cost, it does at least simplify the challenge. Instead of needing to find the scarce resource of additional human practitioners with the appropriate expertise and contextual knowledge to help out, the matter becomes a purely financial one - we can simply pay for more servers.

Once acquired, these computational resources can be applied in such a way that requires no historical knowledge of the problem space. Reinforcement learning and related systems such as genetic algorithms are able to approach problem-solving through sophisticated forms of trial-and-error. Their design lineage is internal—starting from something random and moving to something better through a succession of small, objective improvements.

These powerful mechanisms for automatically finding the needles in extraordinarily large haystacks, however, cannot function entirely on their own. They require some form of external input to indicate whether the outputs they produce are in line with the ultimate goals of the design processes with which they are engaged. That is, they can figure out for themselves how to design a good solution to a given problem, but they need to be told what constitutes “goodness” in that problem space. In some contexts, articulating the design’s goals to the machine is straightforward. For example, a game-playing system might receive a reward for winning a game or scoring a point and an antenna-creating system might receive a reward for increasing the signal-to-noise ratio. In the case of the former, the rule for when a reward is given is defined by the

rules of the game itself. For the latter, the reward is determined by an equation that models the behaviors of physics in the real world.

For many design problems, though, particularly ones that exist within the cultural sphere, it is difficult to derive an objective metric that can be used to guide the machine's process. If the criteria for assessing the aesthetic quality of a painting were as straightforward as quantifying how much blue paint was used, then we could easily apply this kind of approach to the goal of creating beautiful paintings. But since beauty is a matter of subjective taste, caught up in cultural traditions as well as individual preferences, it does not lend itself to this sort of assessment metric and, as a result, it is less clear what role a machine can play in the task. We will return to this consideration again later.

Solution without Theory

In 2016, DeepMind's AlphaGo system beat eighteen-time world champion Lee Sedol in four games of a five-game Go match. AlphaGo's novel strategies, best exemplified by the now famous "Move 37," defied conventional wisdom developed over several thousand years of human history, sending shockwaves through the Go and machine learning communities.

What came next, though, was perhaps even more significant. The following year, DeepMind published a paper in the journal *Nature*, introducing AlphaGo Zero, a new version of its system that, unlike its predecessors, was devised to learn from "tabula rasa." That is, it was trained entirely through self-play without any hardcoded policies or heuristics and without observing any human-played games.

After three hours of training, AlphaGo Zero played at the level of a human beginner, taking a sub-optimal strategy focused on maximizing the capture of stones. After nineteen hours, it had learned more advanced strategies including life-and-death, influence and territory. After three days, it surpassed AlphaGo Lee, winning one hundred games to zero. After forty days, it beat all previous versions of the software to become the world's best Go player. (Silver and Hassabis)

David Silver, who led the AlphaGo project at DeepMind, noted, "what we started to see was that AlphaGo Zero not only rediscovered the common patterns and openings that humans tend to play [...] it discovered them and ultimately discarded them in preference for its own variants, which humans don't even know about or play at the moment." (DeepMind)

If we look at the rate of knowledge acquisition in AlphaGo Zero in contrast with the progression of human play, we find similarly shaped curves. Each follows a roughly logarithmic trajectory. This seems logically intuitive. When you know almost nothing, every experience is an opportunity to learn something new. Once you have reached an expert level, you become more set in your strategy and an increasingly small number of experiences expose new conditions that might lead to novel insights or the questioning of an established strategy.

It is worth underscoring that while it took only three days for AlphaGo Zero to reach Lee Sedol's level, it took an additional thirty-seven days to reach its current state-of-the-art. If we project forward on the human trajectory with guidance from AlphaGo Zero's trajectory, it is conceivable that without intervention, humans would reach AlphaGo Zero's current level sometime around the year 56,480 AD.

This astounding projection should give us pause. Though there is no reason to suspect that the conventional human approach to the distillation and scaffolding of knowledge is running out of steam, it is quickly becoming tangibly evident that the human rate of advancement simply cannot compete with this new approach to knowledge acquisition by machines.

Silver goes on to say:

If you can achieve tabula rasa learning, you really have an agent that can be transplanted from the game of Go to any other domain. You untie yourself from the specifics of the domain you're in and you come up with an algorithm which is so general that it can be applied anywhere.

For us, the idea of AlphaGo is not to go out and defeat humans, but actually to discover what it means to do science and for a program to be able to learn for itself what knowledge is. (DeepMind)

Indeed, a new methodology for scientific advancement is emerging, one that is forged in the absence of conventional means of knowledge distillation and scaffolding. After releasing AlphaGo Zero, DeepMind went on to develop an even more generalized system, called AlphaZero, that, in addition to Go, could also learn to play chess and shogi at a superhuman levels - all with the same neural network architecture and without human intervention to adapt the system to each of the games.

In many ways, we should celebrate these recent advancements. They mean that we may begin to lift ourselves out of any suboptimal strategies arrived at within the

historical trajectory and advance our effective capacities in a wide range of applied domains by following the intuitive explorations of machines. The ability to find timely solutions to difficult scientific problems should certainly resonate with us given the difficult reality we have all faced over the last two years.

This new methodology also suggests a reformation of how human intellectual capital should be spent. It is a reformation that is perhaps best expressed through DeepMind's incredibly concise yet outlandishly grandiose mission statement to "solve intelligence," and then use it to solve everything else. That is, rather than distributing our collective intellectual focus across many disparate domains of inquiry, it would be more economical and seemingly more effective to concentrate those intellectual resources towards the development of increasingly robust learning systems, which can then in turn be rapidly deployed into each constituent domain of the sciences and to new challenges as they arise.

If an artificial intelligence system were to develop a successful vaccine or treatment for Covid-19, we would rejoice in it and benefit from the machine's achievement, regardless of whether we understood its strategy or methodology. Of course, it might be beneficial for future virology research to be able to incorporate the machine's ideas into academic literature. But, if this were not possible, we could at least set ourselves on the idea that this technical discipline could advance even in the absence of human comprehension. So long as its technical goals continue to be met, perhaps we should have no real cause to require greater human comprehension.

Yet, it is somewhat concerning to imagine the sciences upon which we critically depend advancing far beyond the knowledge base of humans. This is the stuff of science fiction - alien technology descending upon our intellectually inferior society to deliver indecipherable technologies built upon operating principles that are entirely foreign to our current understanding of the physical world.

Does this superior intelligence really have our best interest at heart? Going further, if we live in a world driven by systems we do not even begin to understand, can we even know what constitutes our own best interest?

In his book, *Superintelligence: Paths, Dangers, Strategies*, the philosopher Nick Bostrom points out that a machine-driven world could become disastrous for humans, even without supposing a directly malicious intent on the part of the machine. He refers to the idea of *perverse instantiation* as, "a superintelligence discovering some way of satisfying the criteria of its final goal that violates the intentions of the programmers who defined the goal." (Bostrom 146) He offers the example that if the explicitly stated goal

for a system were to “make us smile,” then a perverse instantiation of a solution might be to “paralyze human facial musculatures into constant beaming smiles.” He goes on to say, “the AI may indeed understand that this is not what we meant. However, its final goal is to make us happy, not to do what the programmers mean when they wrote the code that represents this goal.” Taking his thoughts to even more alarming and explicitly malicious heights, Bostrom worries that “a machine superintelligence might itself be an extremely powerful agent, one that could successfully assert itself against the project that brought it into existence as well as against the rest of the world.” (Bostrom 115)

Even if the dystopian suggestion that artificial intelligence’s intentions for us would somehow drift from our own seems meritless or far-fetched, we may remain concerned for our agency, the future role of our own being. There is, at the very least, the strong implication that the role of the human and of the machine would become reversed with the machines adopting the role of agent - the keepers of ideas and the controllers of destiny - while the humans become helpless, fragile things that require periodic tune-ups and whose best interests are paternalistically set by their wiser, more experienced machine counterparts.

There are many scenarios in which blind faith in the machine’s capacity for knowledge acquisition and problem-solving cannot be sufficient. If the success of the system cannot be measured in purely objective terms, then we need to design systems that will make the machine’s strategy more comprehensible to us - helping us to see for ourselves the machine’s experiences, the premises or facts about the environment, that led the machine to those strategies. Even in cases where the criteria for success is indeed objective in nature, there are still numerous reasons why it remains important for humans to be in the loop - able to comprehend and control the behaviors and direction of inquiries carried out by the machine. Bostrom offers one such rationale for this, saying:

An agent’s ability to shape humanity’s future depends not only on the absolute magnitude of the agent’s own faculties and resources - how smart and energetic it is, how much capital it has, and so forth - but also on the relative magnitude of its capabilities compared with those of other agents with conflicting goals.

In a situation where there are no competing agents, the absolute capability level of a superintelligence, so long as it exceeds a certain minimal threshold, does not matter much, because a system starting out with some sufficient set of capabilities could plot a course of development that will let it acquire any capabilities it initially lacks. (Bostrom 120)

Indeed, the ahistorical and indecipherable functioning of superintelligent agents operating on complex problems that are of great consequence to human existence carries some hazards and mixed blessings of which we should remain mindful.

Convergence of Intuition and Theory

Fortunately, it is not lost on the creators of these powerful technologies that finding ways of bringing human knowledge along with the advancement of machine capabilities is of great importance.

In several recently published papers, DeepMind and similar institutions have begun to expand their focus from the discovery of effective architectures for problem-solving systems to the challenge of making the machines' choices and strategies more comprehensible to human practitioners.

The authors of the paper *Acquisition of Chess Knowledge in AlphaZero* state this intention in their introduction and reason that an important first step in this effort is to ascertain whether the machine's mental representations correspond with the concepts and strategies devised by humans in a given domain of inquiry:

Machine learning systems are commonly believed to learn opaque, uninterpretable representations that have little in common with human understanding in the domain they are trained on. Recently, however, empirical evidence has suggested that at least in some cases neural networks learn human-understandable representations.

[...]

The AlphaZero network provides a rare opportunity for us to study a system that performs at a superhuman level in a complex domain. The ability to interpret AI systems is particularly valuable when these systems exceed human performance on the tasks they were optimized for, as the systems could uncover previously unknown patterns that could prove to be useful beyond the systems themselves. Understanding such complex systems is likely to involve a substantial amount of analytical effort, so it is desirable to first understand whether the system's representations have any correspondence with human concepts. If no such correspondence exists, then further investigations are unlikely to pay off, whereas if they can be found then a deeper investigation may well uncover more. (McGrath et al. 1)

The authors go on to note that, "although the results we show in this paper are far from a complete understanding of the AlphaZero system, we believe that they show strong

evidence for the existence of human-understandable concepts of surprising complexity within AlphaZero's neural network." (McGrath et al. 1)

What follows from there is an analysis of how human-known chess concepts can be located within a neural network and a lengthy discussion of the validity of the techniques used to make these associations. Their findings offer intriguing insights, which seem simultaneously consequential to our understanding of machine learning, the game of chess and even to the mechanics of our own learning processes. For example, the authors describe the sequence of events in AlphaZero's acquisition of knowledge about chess openings, stating that, "Figures 4, 7 and 6 suggest a sequence: that piece value is learned before basic opening knowledge; that once discovered, there is an explosion of basic opening knowledge in a short temporal window; that the network's opening theory is slowly refined over thousands of training steps." (McGrath et al. 19) More generally, they conclude that, "there is a marked difference between AlphaZero's progression of move preferences through its history of training steps, and what is known of the progression of human understanding of chess since the fifteenth century." (McGrath et al. 16)

In another recently published paper, *Advancing Mathematics by Guiding Human Intuition with AI*, researchers at DeepMind and their academic collaborators apply their thinking to a "model for collaboration between the fields of mathematics and artificial intelligence that can achieve surprising results by leveraging the respective strengths of mathematicians and machine learning." (Davies et al. 1)

One of the key premises in this paper, conveyed through a quote from Terence Tao, is the observation that, "It is only with a combination of both rigorous formalism and good intuition that one can tackle complex mathematical problems." (Tao) It is here that we begin to come full circle. It seems intuition - that ephemeral, unquantifiable and seemingly unscientific force of the human intellect - was relevant to scientific and mathematical discovery all along.

Indeed, powerful ideas do not start from their formalization. Rather, they are generally driven to formalization only after intuition sparks our suspicion that there is indeed something to formalize and subsequently guides the process of formalization at many intermittent points along the way.

Perhaps this should be especially evident in the field of mathematics, where some of the most intuitively obvious notions - such as the assertion of the Poincaré conjecture that, to put it simply, a sphere in three dimensions cannot be transformed into a torus

without tearing a hole in its surface - have also turned out to be among the most difficult to prove formally.

The sciences and arts may not be so different after all.

The Future of Creativity

As we look to understand artificial intelligence's likely impact on the arts, it may be instructive to consider painting's trajectory after the advent of photography. Though this comparison will likely prove insufficient in terms of the scale of impact, it may at least tell us something about the nature of it.

In the work of artists such as Chuck Close, Vija Celmins and Gerhard Richter, painting embraced the aesthetic of photography and used it as a referential material that imbued specific cultural cues upon the painted imagery - the sentimental tones and impromptu compositions of Kodachrome snapshots; the scientific and unromanticized portrayal of natural phenomena; the heavy imprint of journalism and historical documentation. Eadweard Muybridge and Doc Edgerton enabled us to visualize the world in a frozen moment - a horse galloping or an apple exploding. Though those realities were always available to us in a sense, their timescales seem not to have become palpable for us until the arrival of the camera. Once seen, they would not be forgotten. They were now a part of the visual vocabulary.

In another camp, artists such as Pablo Picasso and Jackson Pollock were freed from painting's long historical pursuit of optical realism, moving their focus to new goals such as the development of novel perspectival techniques that cannot be achieved by a camera but which are conceivable within and even familiar to the human mind. In truth, nothing ever really stood in the way of painting's divergence from optical realism and indeed numerous artistic traditions had developed around figurative abstraction long before photography disrupted the Western trajectory in painting. Nevertheless, it seems there was some intellectual or emotional obstruction to this turn of focus. Perhaps it can be explained through the McLuhanism that it certainly was not a fish that discovered water. In other words, photography made optical realism an explicit choice - a visual style, if you like - as opposed to an implicit and preordained fact of depiction. Whatever the reason, we can safely say that photography changed how we look at the world. So, it is only natural that it also changed how we look at the world through other artistic mediums.

In a similar vein, there should be little doubt that artificial intelligence will deeply permeate the aesthetic vocabulary of the arts. At the same time, these technologies will

likely help to usher in an artistic movement that sees itself as a return to that which is authentically human - a rejection of automation as well as a celebration of craft and tradition.

With respect to artificial intelligence's permeation of the visual vocabulary, let us briefly turn our attention to generative models, another kind of machine learning system, which differ from the reinforcement learning systems we have so far discussed in numerous ways, but are at least comparable in their significance to human discourse. One key difference to note is that while reinforcement learning systems are arguably capable of extrapolation or the achievement of strategies outside the domain of what has been previously explored, generative models such as GANs can be said to be generally limited to the domain of interpolation or the ability to synthesize any combination of properties that exist within the domain of what has already been explored.

It is perhaps a fortuitous coincidence that generative machine learning models have arrived on the scene just a short time after the postmodern era in art melded all artistic traditions into a plasmatic state of cultural matter. In this state, generative models seem somehow inevitable - the logical conclusion to the historical trajectory; the mathematical embodiment of remixing and reappropriation. We may lament the generative model's limitation to the domain of interpolation and note that none of its outputs actually mean anything, despite having all of the trappings of cultural substance. Yet, these limits perfectly coincide with Arthur Danto's assertion that the historic progression of art has come to an end. "There will go on being art-making," Danto says, "But art-makers, living in what I like to call the post-historical period of art, will bring into existence works which lack the historical importance or meaning we have for a long time come to expect." (Danto 111)

An alternate conclusion we might draw in the case of both the postmodern painting and generative output is that we have not run out of historical meaning, it's just been moved somewhere else. The meaning now lives in the act of the reappropriation. Its subject is interpolation.

This is just the first stage of transformation, though. These things are already in the works. Over the next decade, we will witness a radical transformation through which media will become completely malleable. Image, audio and video editing and synthesis will enjoy every possible degree of freedom. Tools of the trade will offer random access into the possibility space.

Today we call these kinds of outputs "deepfakes," a pejorative term associated with the genuinely nefarious use-cases that have led mainstream headlines. Gradually, though,

it is inevitable that the lack of photographic truth will become culturally assimilated. The political sphere will have no choice but to come to terms with the fact that regulatory action will not be able to prevent all nefarious uses of these already-disseminated technologies and intellectual property law will have to establish precedent to guide law-abiding users in matters such as usage right for a public figure's likeness. It will be a tumultuous transformation. Eventually, though, these technologies will be assimilated.

One important outcome of this transformation will be its impact upon the production requirements associated with creative output. While it once took hundreds or even thousands of people to shape a creative work such as a blockbuster film, the synthetic capacities of generative models will increasingly enable an individual person to perform each and every function themselves - actor, director, cinematographer and so forth. In this sense, rather than nullifying the human creative role, these technologies have the ability to amplify the human creative voice by enabling one person to exact their vision in ways that today are only available to the wealthiest and most established artists.

We must ask, then, if production value becomes cheap, where does the premium go instead? Generative models will progressively commoditize interpolative creation. That is, no expertise will be required to turn to your television and say, "I want to watch a brand new episode of Seinfeld that's been updated for 2028 sensibilities."

We have already seen some of the ways in which reinforcement learning systems can demonstrate extrapolative creativity in the sciences and simulatable or optimizable domains. Can it also be applied as such in the cultural sphere? As earlier discussed, this would require the formulation of reward function or fitness criteria - a vexing thing to produce in the subjective realms of cultural meaning.

In the end, though, unless artificial intelligence goes completely rogue in the most extreme manner anticipated by science fiction, there is an important point of which we should not lose sight:

It will still be us humans who point the machines to specific scientific or cultural domains. Those choices are a reflection of our values and interests, not those of the machine. One prominent example of this phenomenon within the sciences is the tendency for philanthropic spending in medicine to prioritize research that is likely to yield outcomes that are relevant to people living in developed countries as opposed to developing ones. As with the directing of human intellectual resources, the question remains the same with artificial intelligence: where did we choose to point the resource?

Art simply cannot be devoid of cultural context. Its creation may be automated through various means and with various levels of intervention. But if by 'culture' we mean our own, then by definition art can hold no cultural meaning outside of the human world. As a result, we remain the cultural arbiters. The history-flattening capacity of the machine learning model and all of the new aesthetics that come with it must be placed back into history - themselves a trend, a look, a style that will ultimately be superseded by another in the cultural lineage.

Though the pace of human cultural evolution will be accelerated by the ahistorical creative capabilities of artificial intelligence, there is some natural impedance factor in the human capacity for change. That impedance factor in the cultural sphere relates not only to the arts but also to the sciences. The way we speak and the matter of which qualities in a work of art move us can only evolve so quickly. Similarly, the full cultural acceptance of scientific capabilities such as human genetic modification will almost certainly trail significantly behind the technical capacity for it.

We may be able to arrive in new places more quickly, but it is still up to us to determine where we want to go. Making these sorts of decisions almost definitely requires comprehension. At the very least, it takes intuition.

Bibliography

- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.
- Danto, Arthur C. *The Philosophical Disenfranchisement of Art*. Columbia University Press, 2005.
- Davies, Alex, et al. "Advancing mathematics by guiding human intuition with AI." *Nature*, vol. 600, 2021, pp. 70-74. *Nature*, <https://www.nature.com/articles/s41586-021-04086-x>.
- DeepMind, director. *AlphaGo Zero: Discovering New Knowledge*. 2017. *YouTube*, <https://youtu.be/WXHFqTvffSw>.
- Fuller, R. Buckminster. *Operating Manual For Spaceship Earth*. Zürich, Lars Müller Publishers, 2011.
- McGrath, Thomas, et al. "Acquisition of Chess Knowledge in AlphaZero." 2021. *Arxiv*, <https://arxiv.org/abs/2111.09259>.
- Silver, David, and Demis Hassabis. "AlphaGo Zero: Starting from scratch." *DeepMind*, 18 October 2017, <https://deepmind.com/blog/article/alphago-zero-starting-scratch>. Accessed 8 December 2021.
- Tao, Terence. "There's more to mathematics than rigour and proofs." *Blog*, 2016, <https://terrytao.wordpress.com/career-advice/theres-more-to-mathematics-than-rigour-and-proofs>. Accessed 8 December 2021.
- Thwaites, Thomas. "The Toaster Project." *The Toaster Project*, 2009, <http://www.thetoasterproject.org>. Accessed 8 December 2021.