The Inevitable Ascendance of Artificial Intelligence:

A Modern Conceptualization of Social and Scientific Progress

Amanda Ortscheid

Advisor: Dr. Valerie Olson

Feb. 21, 2018

Abstract

Many artificial intelligence (AI) researchers believe that it will soon be possible to create AIs which can reason, make decisions, and act intelligently without human supervision, at which time AIs may very well be beyond the control of their human creators. Within the AI community, this belief—and the certainty with which it is often held—has inspired a renewed commitment to preparing for the ethical challenges of the supposedly inevitable advancement and consequent ubiquity of artificial intelligence. This project asks, What sort of a temporality are these experiences and beliefs embedded in? And what is the relationship between researchers' everyday practices and a belief in the inevitable progress of AI? Drawing on in-depth interviews conducted with academic researchers in AI, I argue that these claims are embedded in a distinctly modern temporality because their logic is underpinned by a classically evolutionist understanding of social progress, in which social progress is driven by advancements in technology. However, AI researchers add their own twist to these beliefs: instead of segmenting society from the natural world, they segment society from technology. This is because the reality of their everyday practice as scientists, in which a particular project's success is uncertain and dependent on the cooperation of other scientists, conflicts with these long-term predictions of AI-wide success. The conceptual segmentation of society from technology is an attempt to reconcile this conflict, somewhat paradoxically, by associating society with certain futures and associating technology with uncertain pasts and presents—while still viewing technological progress as the key to social progress.

Introduction

In March 2016, the artificial intelligence AlphaGo won a 5-game series against the renowned professional Go player Lee Sedol. A brisk wind swept across the artificial intelligence community and its onlookers, simultaneously inducing dread, excitement, and—surely at least for AlphaGo's creators at DeepMind—a powerful sense of self-satisfaction. In *Scientific American*, neuroscientist Christof Koch declared, "An era is over and a new one is beginning. The methods underlying AlphaGo, and its recent victory, have startling implications for the future of machine intelligence" (2016). *Wired* reporter Cade Metz agreed, claiming that "the victory [was] notable because the technologies at the heart of AlphaGo are the future . . . already changing Google and Facebook and Microsoft and Twitter" (2016). Ken Jennings, one of the *Jeopardy!* contestants who lost to IBM's Watson in 2011, sympathized with Lee by noting that the real terror of going head-to-processor with a machine is knowing that one's "opponent isn't just a room full of servers or a few thousand lines of code . . . [but] the Future," because every day we draw closer to the time when humans are finally obsolete (2016). In fact, Watson's superior *Jeopardy!* performance has already convinced Jennings of *his* "own inevitable obsolescence" (2016).

Despite the urgency with which commentators made these pronouncements, however, this story was far from headline-news for anyone outside of the tech community. But since I was a computer science student at the time it broke, predictions that AlphaGo would change the world forever were everywhere *I* turned. They were in grinning professors' pre-lecture announcements, idle chatter in the computer lab, and the resolutions of lunchtime debates. At first, I found these predictions to be merely amusing: the same sort of self-aggrandizing hype that all AI's successful projects, no matter how miniscule their success, have engendered over the last century. Then I learned how seriously some of the brightest—and perhaps more importantly, most well-funded—minds in the tech community were treating predictions of this kind. Over the last few decades, concerns about the impact of soon-to-be-arriving super-intelligent AIs on human society have increasingly opened research centers, built industry coalitions, and shaped the mission statements of non-profits. And in almost every case, as in the case of AlphaGo, their concerns center not just on AI and the forms it might take, but on how these forms of AI might create a given future.

But what exactly do people mean when they invoke notions of "the future" in this way, I wondered? These invocations of the future did not remind me of my favorite science fiction novels; instead, they recalled several influential but now discredited theories anthropologists first made all the way back in the 1800s. Early anthropologists like Lewis Henry Morgan and Sir James Frazer contributed to a body of theories known together as sociocultural evolutionism (Lesser 1952[1939], 134). A common theme of these arguments was that societies progress through predefined, universal stages of development (Lesser 1952[1939], 134). In Morgan's work, each stage was defined by the technology available, and societies progressed by developing new technologies (Morgan, 1877). In Sir James Frazer's work, each stage was defined by a society's way of gathering knowledge about the world, and societies progressed by turning from magic and religion to science (Frazer, 1890). Many people today do not know the details of these theories and yet they stand in the shadows of them, for these theories helped shape many current academic fields, not just anthropology. They left profound marks on sociology, psychology, and biology, for example. And even though these classical theories of unilineal social evolutionism have been shown to be based on faulty, prejudiced data sets (Lesser 1952[1939], 135–136), it seems that every few decades, some scholars determine to find new evidence and resurrect them anyway (Kahn 2001, 657). Most notably, after being condemned in the 1930s and '40s (Lesser 1952[1939]), they were revived in the '50s, by way of modernization theory, which to this day—while highly controversial—still often sets the parameters for debates about what it means to be "modern" (Kahn 2001, 657). Could predictions about the inevitable progress of AI have been shaped by this legacy? Certainly, to me, the truth of these predictions was not at all obvious. They seemed to rise rather crookedly from the historical evidence. In fact, these claims did not accord at all with my understanding of both the history of efforts to solve AlphaGo and the history of AI in general. So, I suspected there might be some distinctly *cultural* force making these claims' truth values seem self-evident.

Thus, this project was born, its goal to use anthropological methods to answer two questions: first, what does "the future" mean for people who have a stake in AI, and, secondly, what role do predictions about the future play in efforts to formulate an ethics of AI? In support of this goal, I conducted eight interviews with researchers active in the subfields of machine learning, general intelligence, and gaming. In this paper, I draw on material from these interviews, academic and popular texts, the ethnographic record, and anthropological theory, to give tentative answers to the questions above. After reviewing the history of AI and relevant anthropological theory, I first present the logic of a popular understanding of the future as a temporality for people concerned about the ethical implications of AI rapidly, and perhaps inevitably, advancing. I put these beliefs in their cultural context, showing how they are embedded in a modern, scientific understanding of time in which social progress is driven by advancements in technology. I then argue that the logic of this temporality, in particular, rests on the conceptual segmentation of "society" from "technology." In this conceptualization of the relationship between society and technology, the two are linked but have their own contrasting temporal qualities. Long-term social progress, driven by technological advancement, is certain and a version of our society with more advanced technology than we currently have, in a way, already exists in the future. Technology, however, as seen separately from the society in which it is used, is uncertain. It is a category whose items are always liable to fail or be supplanted, whose present and past successes are and always were doubtful, because of social and political constraints—but also the very nature of the scientific process. Assigning different temporal qualities to society and technology, within the same overarching modern temporality, thus allows researchers to argue (1) that social progress and scientific progress do not proceed in the same manner, though they are linked to each other, and (2) that our shift to a society with advanced AIs is inevitable.

Historical Background

Let us continue a short while longer with the case-study presented above, so that I may support the claim that predictions about the inevitable rise of artificial intelligence do not agree well with the history of AI. Of course, historical trends do not determine the future, but given that it is scientists who claim to be able to make predictions about AI's future, we might as well think scientifically—that is to say, using inductive logic. The history of efforts to solve AlphaGo is one that shows how, as one researcher told me, AI and "machine learning, in particular, has a lot of boom-and-bust kind of excitement." It is a history which does not, at least to me, seem to show linear or exponential progress, as in the predictions of some figures in the AI community. And it is a history which shows how politics and fragile egos can steer science, sometimes in the direction of discovery, sometimes burying lines of inquiry.

In his 1988 Presidential Address to the Association for the Advancement of Artificial Intelligence, Raj Reddy noted that chess was the game which "ha[d] been studied continuously since the birth of artificial intelligence in 1956" and "play[ed] the same role in artificial intelligence that the studies of E. Coli play in biology" (10). But for at least 60 years before DeepMind built AlphaGo, a perhaps surprisingly numerous and dedicated group of both professional researchers and hobbyists had also been trying to build computer programs that could play Go at a human-level (Hafner, 2002). These researchers quickly realized that solving Go would be even more difficult than solving chess—if they were to follow the popular approach to programming computers to play chess, anyway. A program taking this approach decides which moves to make in much the same way as humans do: by evaluating each possible move by the moves it might allow one's opponent to make (Campbell et al. 2002). But even in chess, this is infeasible to do with certainty, i.e. to compute all the possibilities until the game's end, with our current hardware, because it requires so much memory and computing power. The 1950 estimations by pioneering information theorist Claude Shannon, which indicate that such an approach to solving chess would need thousands of years to decide its first move, are still widely accepted (Shannon 1950, Steinerberger 2015). Nevertheless, chess-playing programs have successfully used a modified version of the approach, in which a program only looks a certain number of moves ahead (Levinovitz, 2014). However, in Go, this is still not effective, because there are even more possible moves on a 19-by-19 Go board than there are on an 8-by-8 chess board (Levinovitz, 2014). Still, these years of work did not take too much of a toll on computer scientists' characteristic optimism. Though many believed that it was highly unlikely that a winning program would be created any time soon, few would have said it was impossible, and the work continued, each program building on the successes of its antecedents (Hafner 2002).

Despite popular media's characterization of AlphaGo's technology as a sharp departure from previous efforts to solve Go, the techniques AlphaGo uses in and of themselves were not newly discovered by the AlphaGo team. Researchers at DeepMind optimized existing techniques and combined them, drawing on decades of previous work. But they also drew from areas of computer science not typically associated with gaming, like computer vision. Key to AlphaGo's success was that it combined decision search techniques already popular in games in general and in Go programs in particular, like those described above (Hafner 2002), with deep neural networks, a technique most famously used in computer vision applications. A deep neural network is a mathematical model of networks of neurons in the brain, which learns by adjusting the parameters to functions—an analogy to the "firing" of neurons (Liu et al. 2017). Research into neural networks actually began in the 1950s with the work of Warren McCulloch and Walter Pitts, but they did not have much success because their networks could not "learn" (Olazaran 1996). The research was taken up most enthusiastically by Frank Rosenblatt, whose mathematical model, the Perceptron Machine, was probably the most influential of its time, because it could learn (Olazaran 1996). Or, at least, it was the most overblown (Olazaran 1996). Annoyed by Rosenblatt's extravagant claims that the machine would soon allow computers to talk and think for themselves, critics pointed out the limitations of the perceptron, and an anti-perceptron sentiment remained strong in the AI community for most of the 1970s (Olazaran 1996). In the 1980s, pioneering work by Geoffrey Hinton and colleagues Yann LeCun and Yoshua Bengio revived some interest in neural networks and laid the groundwork for today's achievements, but it was not clear that their work would result in these great successes. Their work was slow and quiet all the way until the 2000s (Hernandez 2014).

Although it would seem that, serendipitously, all worked out for Go, this history of artificial intelligence and its boom-and-bust pattern made me doubt claims that current achievements are necessarily an indication of greater achievements to come. It is not as though the researchers who make these claims are not aware of the history of artificial intelligence; any introductory course in artificial intelligence will typically also introduce one to the trials and tribulations of AI. And among many technologists and scientists, one of the most famous and well-loved philosophical theories of science is Thomas Kuhn's, which agrees with a view of scientific progress as cyclical and hard-won (Bird 2013, Godfrey-Smith 2003). Kuhn, in *The Structure of Scientific Revolutions*, argues that science develops in phases. First, someone makes an important theoretical contribution that solves a problem no one could previously solve. Then the approach that allowed for that solution is enthusiastically taken up by others. The approach amasses support, both in terms of numbers of researchers and numbers of grants, and a full-fledged paradigm develops and takes over the entire field—forcing its body of theories and evidence onto anyone who wishes to succeed in that field (Bird 2013). People who do work in this paradigm continue along, until they find that there are just too many problems that cannot be solved using a paradigmatic approach. Then there is a crisis, in which all the fundamental assumptions of the paradigm are questioned. If someone manages to make a new contribution that can explain these problems, the cycle will begin again (Bird 2013).

But if most people in the AI community are familiar with this history, and with theories of science such as Kuhn's that do not cast scientific progress as linear or exponential, then how is it that they can have confidence in claims that AI will inevitably, exponentially, progress? A major influence on popular discussions of artificial intelligence has been the Singularity movement, which began in earnest in the 1990s. In 1993, computer scientist and science fiction author Vernor Vinge gave a talk at the VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute, in which he proposed that humans were "on the edge of change comparable to the rise of human life on Earth . . . [due to] the imminent creation by technology of entities with greater than human intelligence," and that the time when this happened could be called the "singularity" (1993). By the 2000s, the Singularity was not just an idea, but a movement, with research into it being carried out at the Singularity Institute For Artificial Intelligence and inventor Ray Kurzweil's books spreading its gospel to the masses (Farman 2010). No one thus far has really been able to satisfactorily describe the precise nature of the relationship between scientists' certainty in progress and the Singularity movement; after all, even in 1965, the cryptologist I.J. Good argued that intelligent computers could lead to an "intelligence explosion" such as that posited by Singularitarians—although, to be fair, he called his predictions "speculations," and it is not clear just how confident he was in them—other than confident enough to publish them, of course (Chalmers 2010, 9). What is important is that these ideas are now mainstream. Even if they are not explicitly accepted by many researchers as scientific fact, they are most certainly influencing discussions about the future of AI and the ethics of AI.

In parallel with predictions about the inevitable progress of AI, and often in direct response to it, debates about the possibility of human-level artificial intelligences and the ethics of creating such artificial intelligences have been gaining urgency for many people within and without the artificial intelligence community. In 2000, Sun Microsystems co-founder Bill Joy became infamous for his controversial article "Why the Future Doesn't Need Us," a direct (and quite pessimistic) response to Singulatarianism. In this article, he admitted that, though his projects' "ethical dimensions' ha[d] concerned [him]" from the start, "it was only in the autumn of 1998 that [he] became anxiously aware of how great are the dangers facing us in the 21st century," because it was then that Kurzweil convinced him of the high probability of a rapid advancement of AI and other associated technologies. More recently, IBM's Chief Science Officer for Cognitive Computing Guru Banavar wrote, "It is no exaggeration to say that in the years ahead, most aspects of work and life as we know it will be influenced by these [artificially intelligent] technologies. And that makes us more than computer scientists. It makes us architects of social change" (2017). Even here, though Banavar does not make explicit reference to Singulatarianism, he still shares Joy's concerns and justifies these concerns with similar reasoning, i.e. he is certain that AI will continue to advance and spread.

Additionally, the 2000s have seen the birth of numerous research centers focusing on identifying and preparing for the risks such advanced AIs might pose to humanity. In 2012, astronomer Professor Lord Martin Rees and Skype co-founder Jaan Tallinn founded the Centre for the Study of Existential Risk at the University of Cambridge, with one of its primary goals being to investigate how to mitigate the existential risk AI will pose to humanity. In 2005, philosopher Nick Bostrom founded the Future of Humanity Institute at the University of Oxford, which investigates both methods for the governance of AI development and methods for creating safe AIs. Outside of academia, in 2016, the Partnership on AI launched, with the mission of addressing ethical issues in AI and bringing AI to bear on social problems, and with founding partners Amazon, Apple, Facebook, Google, Microsoft, and IBM. In 2015, the nonprofit OpenAI began research in artificial general intelligence (AGI) with the goal of making AGIs which would be guaranteed to benefit humanity.

Theoretical Background

This project focuses on the temporal aspect of the future concept in discussions of the ethics of AI. But what does the study of human biocultural diversity have to do with time, and how is it that an anthropologist could study time, anyway? In fact, many anthropologists have claimed that anthropology has always been concerned with time and that the anthropological approach is well-positioned to make unique insights into the relationship between time and human experience. For example, in Nancy Munn's influential review of the anthropology of time, she argues that time is ever-present in our studies because "we and our productions are in some sense always 'in' time (the socioculturally/historically informed time of our activity and our wider world)" (94, 1992). Of course, time being present in our studies does not necessarily mean that it figures in our theories. However, Durkheim's classic *The Elementary Forms of Religious Life* did provide a conceptualization of time within a theoretical framework that places time firmly in the domain of the social sciences. So, one might still consider time to have been a subject of anthropological inquiry since anthropology's very beginnings, even if one disagrees with Munn, because *The Elementary Forms of Religious Life* is one of anthropology's foundational texts (Gell 2002, Munn 1992). In it, Durkheim argued that "we cannot conceive of time except by distinguishing its different moments" (in Gell 2002) and, importantly, that how we distinguish these moments from each other reflects the "rhythms" of social life; that is to say, time is a reflection of social life (Gell 2002). Thus, one can also study time merely by studying social life. Additionally, according to Munn, the Durkheimian paradigm assumes that "social diversity, qualitative heterogeneity, and a conceptual segmentation… connected to activity or motion" characterize time (94–95, 1992). In simple terms, the experience of time differs throughout a person's life according to the social rhythms which segment it.

These core aspects of the Durkheimian paradigm still appear in most studies in the anthropology of time. For example, Nielsen explicitly situates his work within this tradition in order to use his study of the social aspects of house-building in Maputo, Mozambique to argue against a so-called linear view of time and to push the reader to consider that, at least for his informants, time can be non-linear, with presents and futures coexisting (2014). However, while Durkheim continues to inspire, most contemporary anthropological works do not directly follow Durkheim in the manner of Nielsen's. Walford notes that "Durkheim's original insight now sits alongside a myriad of theories and descriptions of social time—collective, subjective, material, or otherwise" (21, 2013). Walford's own work advances Durkheimian theory by arguing that the categories of "natural time" and "social time" are mutually constituted and must be studied together (2013). Others have kept the idea of social time, but argued that "temporalities" should be seen as epistemic rather than ontological qualities (Ringel 2016), or suggested that attention to time should be seen as a theoretical "tool" which should never be decontextualized from specific cultures (Hodges 2002).

In this study, the Durkheimian assumption that there is such a thing as social time is carried forward, but "the future" is taken to be a temporality in the sense that Ringel employs it, as an epistemic quality. My primary concern is how people's subjective experiences of time influence their attitudes toward science as a knowledge-making endeavor, and vice versa. The way that people describe the future is taken to be an indication of their attitudes toward it, providing insights into the ways in which time plays a role in their lives as scientists and researchers. This agrees somewhat with the emic view of many futurists, who are themselves active in making predictions of this kind, because it subsumes the futurists' view that the future concept is a metaphor guiding efforts to create futures (McDermott 1990, Judge 2016).

However, it goes further than the futurists' view, because to anthropologists, a temporality is created not just through rhetoric, but also through practices (and, in turn, a temporality can influence practices).

Furthermore, in this particular case, I consider this future to fall into a certain category of temporality that has already been studied by anthropologists: I consider it to be an example of a sort of "modern time." The first studies of modern time focused on social time within modern institutions—in particular, the political, bureaucratic, and economic institutions that drive capitalism (Bear 2014). Recently, one of the most influential contemporary theorists of modern time, Laura Bear, has argued against this approach and instead suggested that "institutions *mediate* divergent representations, techniques, and rhythms of human and non-human time" (emphasis Bear's); they do not produce modern time (6, Bear 2014). This is offered as explanation as to why recent studies of modern time identify numerous, diverse "characteristics" of modern time (though I wonder whether something can really be called characteristic unless it holds up under comparative analyses). Another explanation might be that the category of "modern" itself is often taken for granted in the anthropology of modern time: not a single work on modern time referenced in this background section explicitly justifies the categorization of its subject matter as "modern." Even Bear assumes that her ethnographic example is modern but does not define modernity, because the events described in the example occurred in the present or, actually, near past (2014). In this way, rather than fit their fieldwork to the modern, these anthropologists fit the modern to their fieldwork, so that they are hard at work constructing the very category of the modern.

This is understandable. After all, as Merrifield and co-authors note, "'modernity' appears at once conceptually straightforward and theoretically elusive. Scholars seem to know it when they see it, yet specific definitions and approaches to the study of modernity rarely seem complementary" (2013). (For a detailed history of the meanings of "modern," see Kahn 2001.) Particularly relevant, as noted by Kahn, is that "while the positive valuation of modernization that derived from its liberal-evolutionist origins fell into widespread disrepute . . . critics shared with modernization theorists the vision of modernity as a process of emancipation and continuous technological change" (2001, 657). The "evolutionist" part of the "liberal-evolutionist origins" Kahn refers to are those theories discussed earlier in this paper: sociocultural evolutionist theories based on limited, corrupted data that argued that societies progress through a set of universal stages—the most advanced stage, of course, being whichever stage the theorists thought their own society belonged to. Additionally, one reason why these modernists saw technological change as a force of progress was that they saw humans as apart from the natural world and modernization was, to them, a process of coming to know and control the natural world (Latour 1993). Most recently, scholars like Michel Foucault and Arturo Escobar have defined modernity as a set of global processes producing or pushing back against the division of peoples into haves and have-nots; they base their theories on the observable consequences of actions taken by those who believed that being in a more "advanced" stage meant they should control those in a "primitive" stage (Merrifield et al. 2013).

Although this project is informed by all of the definitions of modernity described above, the classical definition is most useful for this study. What makes the AI researchers I interviewed modern is, on the one hand, a matter of personal attitude: the belief that their work is technoscientific and that technoscience is a continuous and powerful driving force behind change in society. On the other hand, it is also a matter of being caught up in a modern society: since I exclusively interviewed researchers working in academic settings in the United States, everyone

I interviewed is affiliated with bureaucratic academic institutions operating in a capitalist nation-state. They also practice a strain of AI which is descended from 1950s Euro-American academic research into the simulation of human reasoning, a strain which was born in academia and largely financed in its early days (and even now) by the United States government (McCorduck 1979, Bengio 2016, Forbus 2010).

This sort of modernity, with its roots in modernization theory and sociocultural evolutionism—and, especially, a Morganian evolutionism that sees technological change as the primary method of advancement—is one which implies that scientists' work is that of driving societal change, and thus, future-making. But do many contemporary scientists themselves (still) believe this? And do they often carry out their work with an orientation toward the future? The ethnographic record spells out "Yes" to both of these questions. Adams, Murphy, and Clarke argue that modern time, especially for those who work in the sciences, centers on a belief that the future is knowable, if uncertain, so that a desire to know the future often motivates scientists' present actions, resulting in their experience of the present being marked by "anticipation" (2009). More generally, Holston, drawing on his fieldwork studying modern religion in Brazil, argues that the "presupposition that alternative futures are possible through … active intervention," together with a commitment to future-making that guides present actions, is an essential characteristic of modernity (1999). Bryant argues that when futures cannot be fully anticipated, a feeling of overwhelming present-ness is produced, leading to a feeling of crisis (2015). Guyer argues that in macroeconomics and evangelical Christianity, near futures have been evacuated and instead people orient themselves toward long-term futures (2007).

One notices that in these studies, future-making also often reveals attitudes toward the relationship between agency and morality. For example, the presupposition that "active intervention" can bring about "alternative futures," perhaps better futures, figures a sort of agency which is both temporal and moral. This is highly relevant to the study at hand, because it provides a way into thinking more critically about the role the future concept plays in discussions of the ethics of AI. These studies suggest that feelings of anticipation and crisis become entangled with questions of ethics when they are induced by tension between (1) one's apparent helplessness in the face of the inevitable and (2) one's moral duty to at least *try* to avert impending catastrophes.

But in order to fully understand how these concepts are theoretically useful, we first need to decide what agency *is*. In 2001, Laura Ahearn argued that anthropologists have for too long defined agency ethnocentrically, because they saw agency as the sum of "will" and "power." This encouraged anthropologists studying agency to focus on questions of quantity which are ultimately impossible to answer: for example, Does so-and-so have agency? How much? She offered this definition of agency which has since become influential: "the socioculturally mediated capacity to act" (112, 2001). While I find Ahearn's suggestion to focus on the qualitative, not quantitative study of agency to be helpful, in my case, I think an appropriate measure of attention to will and power is still helpful. After all, it may be ethnocentric to Euro-American cultures, but I am mainly studying discourses which have taken place in Euro-American cultures. Coming back to (1), we might place this feeling in an anthropological framework by thinking of it as a "temporal agency." Moroşanu and Ringel's recent work on temporal agency loosely defines it as "how human beings relate to the temporal dimensions of their lives, and whether they are able to influence them" (17, 2016). They focus on a dimension of this, which they call "time-tricking:" "the many different ways in which people individually and collectively attempt to modify, manage, bend, distort, speed up, slow down or structure the times they are living in" (17, 2016). When AI researchers try to bring to life a vision of the future, they might be said to be engaging in time-tricking. As for (2), if one does accept willing as an aspect of agency, then Mattingly's work on moral willing as the continual re-orientation of one to one's self-narrative becomes useful (2010). Bear connects ethics and agency by following Ramo and, in turn, Aristotle, to define the phronesis of time as "a sense of the world that leads to action or praxis at the right time in the correct manner" (2016).

All this is to situate the project at hand within the anthropological literature. As shown above, the history of AI and the ethnographic record taken together indicate that time, modernity, science, agency, and ethics are, in the case of AI, entangled. Thus, this project uses the concepts of modern time and future as temporality in order to investigate claims of inevitable scientific progress in AI that perpetuate this entanglement.

Popular Claims About the Future of Artificial Intelligence

The logic of many claims that AI will soon advance exponentially is held together by a keystone premise that precisely opposes my own previous argument. These claims depend upon humanity currently being in an era of history unlike any other. This is why proponents of these claims can dismiss the boom-and-bust history of AI as irrelevant to the question of whether or not AI will inevitably advance: they believe that we are but caught in a recent and revolutionary trend of rapid technological advancement.

Let us discuss a prototypical example of this. In 2015, the director of the Future of Humanity Institute and co-founder of the Institute for Ethics and Emerging Technologies, Nick Bostrom, gave a Ted Talk called "What happens when our computers get smarter than we are?" In this talk, before he could answer the question posed by its title, Bostrom first had to justify the use of the word *when*. He did so through several parallel arguments, each of which depends upon a modern understanding of time. Here is the first:

> We are actually recently arrived guests on this planet . . . If Earth was created one year ago, the human species, then, would be ten minutes old. The Industrial Era started two seconds ago.

In this argument, Bostrom contrasts the deepness of natural time against the shallowness of human time in order to make two logical moves, both of which depend upon a modern understanding of time. First, by measuring human time in terms of natural time, he argues that humans are *in* nature and yet, by emphasizing the newness of human time, he argues that humans are *separate from* the rest of the natural world. This also qualifies the supposed rapidness of our most recent technological revolution as a *relative* rapidness. The conceptualization of time underlying these arguments is one that is modern in a way described best by Bear. According to Bear, modern "science and technology tightly link social, human time to external non-human rhythms; frame time as a radically other secular force; and *project a deep history of natural time*" (emphasis mine, 2016). Bostrom can only measure human time in terms of natural time if the two are seen as being tightly linked, with time being an external force which stretches back for billions of years. This is modern because it again has its roots in the same theories as evolutionism and other classically modern ideas. In particular, a "deep" history of time, as based on theories proposed by the geologist Sir Charles Lyell in his treatise *Principles of Geology* (1830–1833), was necessary for at least the theories of biological, if not also social, evolution. Charles Darwin's theories are thought to have been especially influenced by Darwin's reading of

Principles and friendship with Lyell (Anderson 2007, 450). Lyell's theory of uniformitarianism, so called because it states that all geological features were slowly created through the same processes that can be observed in the present, provided a way for the empirical study of geology to make inferences about the past and future (Anderson 2007, 450–451). Deep time was not Lyell's concern—and he was not the first to propose it—but uniformitarianism implies deep time, because, if one accepts that geological formations develop through small, uniform changes, then the development of many features would have taken thousands of years—and uniformitarianism's popularity led to increased acceptance of deep time (Hancock 2006, 194).

Bostrom then explains what makes this era unique, other than its being most recent: we are at the beginning of a trend of advancement which follows an exponential curve. He bases this argument on the observation that over the last 10,000 years, world GDP has grown exponentially:

> Another way to look at this is to think about world GDP over the last 10,000 years . . . [referring to the graph] . . . It's a curious shape for a normal condition. I sure wouldn't want to sit on it . . . Now it's true, technology has accumulated through human history, and right now, technology advances extremely rapidly—that is the proximate cause, that's why we are currently so very productive.

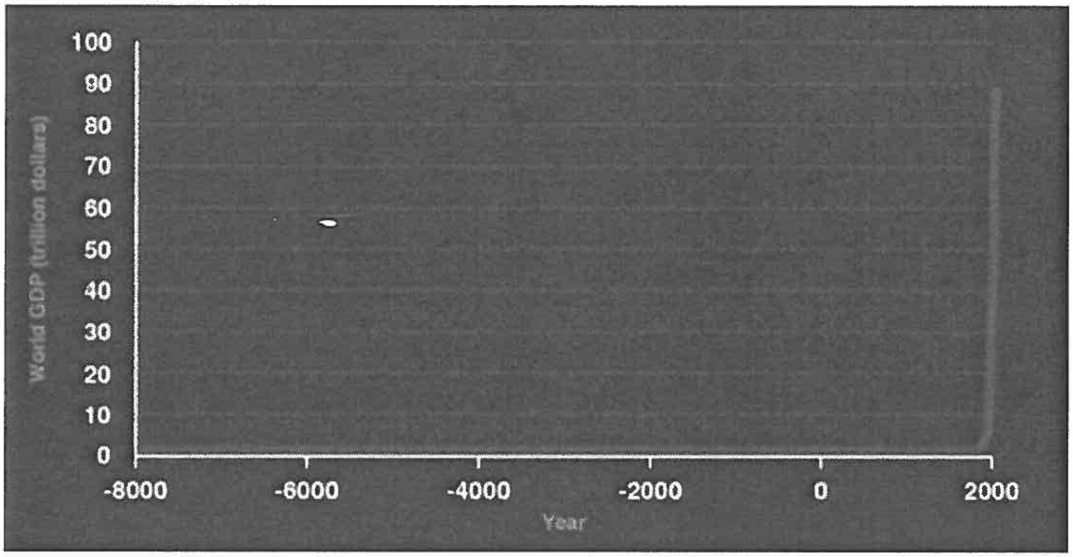

Of course, this is nothing without the assumptions that (1) GDP can accurately measure pre-capitalist economies and (2) the primary cause of GDP growth is technological advancement. Bostrom gives no indication as to how he obtained these figures, but the most famous predictions of historical GDP are Angus Maddison's, which rely on the assumption that per capita output was the same all over the world before the Industrial Era; according to economist Gregory Clark, Maddison assigned "any society without a sophisticated production technology, without significant urbanization, and without a substantial rich class, or just where nothing is known" the same paltry basic subsistence GDP per capita (2009, 1157). Thus, these estimates rely on circular logic: they begin with the assumption that only a certain type of advanced technology can sustain economic growth, calibrate records accordingly, and then use those calibrated records to show that economic growth could only be sustained in the post-Industrial Era. But even putting aside these logical issues with the *projected* historical measures of GDP, a more accurately measured GDP based on current data still was never meant to measure progress of any type that was not related to the production of capital, and economic growth is not always a good indicator of the well-being of all the citizens of a nation, since wealth may be unequally distributed (Giannetti, et. al., 2015). Anthropologists studying development, following Escobar, have long recognized that GDP and other such measures have played more of a role in justifying the postcolonial order than in promoting growth in so-called developed countries (Escobar 1995). Most recently, economic anthropologist Hannah Appel has synthesized these studies to argue that the GDP's primary use has been in making sense of deep inequalities caused by colonialism by shifting the blame for former colonies' poverty from their colonizers to the "natural laws" of economics (2017, 298–299). The underlying cause for all of this is that the GDP is, again, a modern phenomenon: a modern tool really only meant for modern economies, i.e. the economies of a late imperialist world.

But just so that it is clear that Bostrom's ideas are deeply entangled with modern, evolutionist theories of social-scientific progress, Bostrom offers one last argument. He argues that the ultimate cause of intelligence has been the evolution of the brain:

> But I like to think back further, to the ultimate cause. Look at these two highly accomplished gentlemen [photographs of Kanzi and Ed Witten] . . . When we look under the hood, this is what we find. [Renderings of bonobo and human brains] Basically, the same thing. One is a little larger, it maybe also has a few tricks in the exact way it's wired. These invisible differences cannot be too complicated, however, because there have only been 250,000 generations since our last common ancestor. We know that complicated mechanisms take a long time to evolve. So, a bunch of relatively minor changes take us from Kanzi to Witten, from broken-off tree branches to intercontinental ballistic missiles.

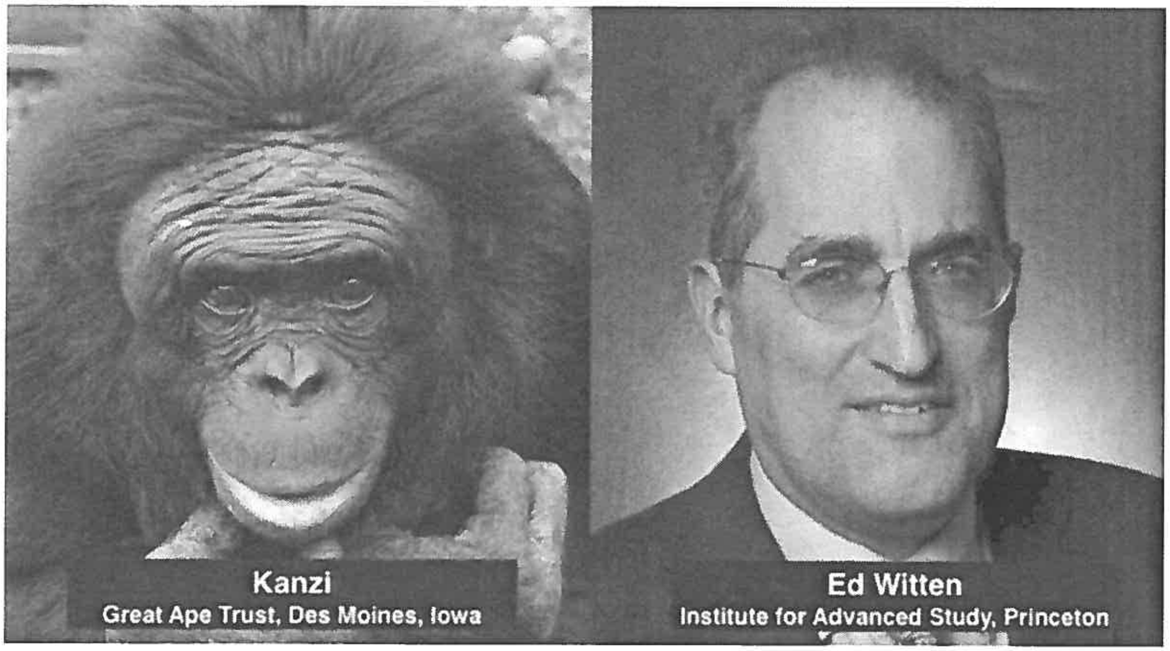

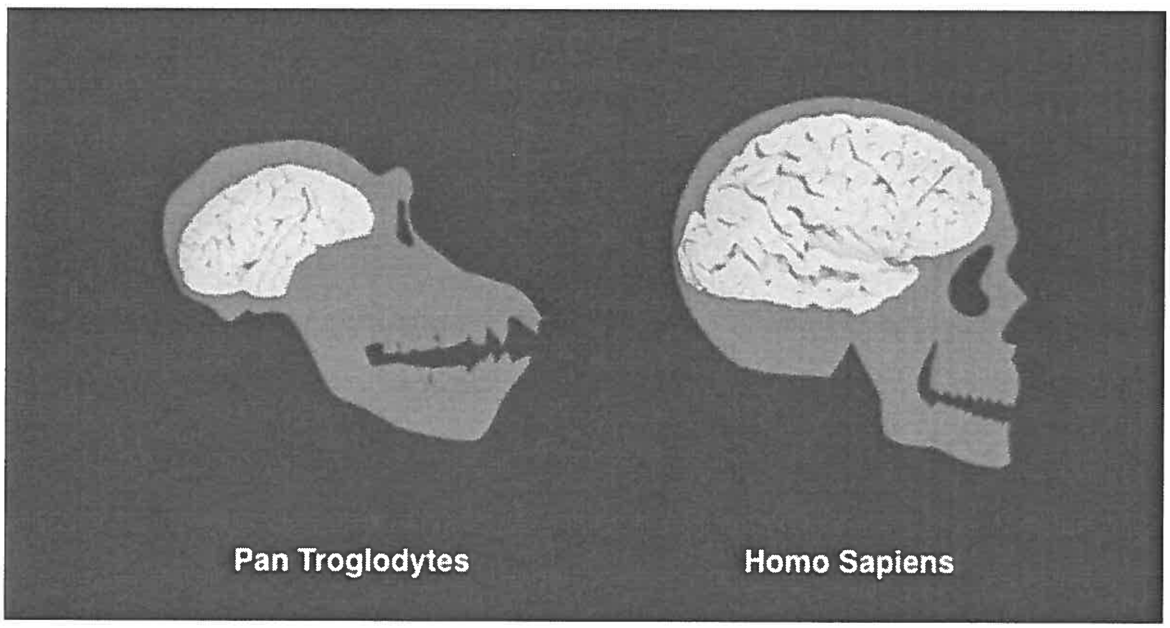

This is the summary of all previous arguments and thus depends on all the assumptions they also depend on. First, it strongly relies on the interpretative principle of uniformitarianism and its implication of deep time; that is what allows Bostrom to say that "we know that complicated mechanisms take a long time to evolve" and that it is only "relatively minor changes" that are necessary for great change. It also relies on the assumption underlying his arguments about the GDP: that technological advancement, born out of the recesses of some great mind, leads directly to the development of and widespread use of tools such as "intercontinental ballistic missiles"—regardless of the social circumstances that push us toward such things as intercontinental ballistic missiles. Of course, none of this is as strong an argument as, say, Morgan's, regarding the relationship between evolution and technological progress. Bostrom does not contend that these stages are pre-defined: only that complexity is superior to simplicity, and that systems should evolve to be more complex over time.

However, the title of his talk implies that he does believe that systems *will* necessarily evolve over time—which already has a famous counterargument in the case of the peppered moth. This classic example shows how quickly natural selection can switch between selecting for certain genes, and how a population may adapt quickly to new pressures in the environment without necessarily "evolving," in the sense that no new or universally "better" alleles of a gene appear in the population. Before the 1800s, most peppered moths in England were light brown, but a few were dark brown. The light brown moths blended in better with tree bark, and thus were less of a target for birds. But when soot from factories turned the trees black, suddenly the dark brown moths were selected for in response to the resulting increase in bird predation; very quickly, the dark brown moths flourished and proliferated in industrialized areas of England (Hesman Saey 2011, 1). No mutations were necessary; it was the proportion of moths which had the gene for melanism— which already existed in the population—for one color or the other which changed in accordance with patterns of bird predation, which in turn was linked to pollution of the environment by humans (Cook et. al 2012, 609–610).

It should not be surprising, given how I have shown how crucial the concept of time is to these claims, that, throughout all this, "the future" is a reference point for Bostrom. First of all, he describes his work as "think[ing] about the future of machine intelligence." His commitment to uniformitarianism allows him to talk about the future as a certainty—however, he still claims that the way we develop technology now could shape the future. This is because, to Bostrom, the same process could result in a set of possibilities, though a finite, somewhat foreseeable set of possibilities. This is exemplified by his statement, "We need to think of intelligence as an optimization process, a process that steers the future into a particular set of configurations." Here, all the configurations of the future already exist, but our actions might determine which configuration is made reality.

Now, the reader may be wondering if such confidence in the efforts of scientists to advance AI is really prototypical. At least in this new industry which produces research into the ethics of AI, it is. Almost all the major players in this new industry of the ethics of AI, many of whom I have already mentioned, have stated in no uncertain terms that they believe strongly that AI will advance, no interpretation necessary. A curious example in which many of the major players all agreed on this at once came in 2017, when the Future of Humanity Institute hosted a conference to bring together scholars and industry leaders in economics, law, ethics, and philosophy to discuss how AI might benefit humanity. There, during the panel, "Superintelligence: Science or Fiction?," which included entrepreneurs Elon Musk (Tesla, SpaceX) and Demis Hassabis (DeepMind), academic researcher Stuart Russell (Berkeley), inventor and futurist Ray Kurzweil (Google), philosopher Nick Bostrom (FHI), neuroscientist Sam Harris, computer scientist Bart Selman (Cornell), philosopher David Chalmers (New York

University) and programmer Jaan Tallinn (CSER, FLI), the cosmologist Max Tegmark asked, "Will it [superintelligence] actually happen? 'Yes,' 'no,' or 'it's complicated'?"

"A little bit complicated, but yes," said Tallinn.

"Yes, and if it doesn't, something terrible has happened to prevent it," said Harris.

"Yes," said Hassabis.

"Probably," said Bostrom.

"Yes," said Chalmers. "Yes," said Kurzweil. "Yes," said Selman. "Yes," said Russell.

"No," said Musk. He wasn't joking. But he laughed in a resigned sort of way.

Admittedly, this panel was not a representative sample of artificial intelligence researchers. The mission of the Future of Humanity Institute is "to shed light on crucial considerations that might shape our future," and it was founded by Nick Bostrom. But it was certainly a sample of the most famous and influential voices in this debate—and leading voices in the ethics debates.

Technology versus Society

But how does this compare to how my interviewees saw the future? What further insights into can be gained from a more holistic analysis of AI researchers' beliefs about the advancement of AI? After interviewing AI researchers about their everyday practice in addition to their beliefs about social-scientific progress, comparisons between the two revealed an interesting difference in their attitudes toward the two. If they did see the future as certain, it was because the technology was (1) already in the world and (2) the obstacle to its mainstream acceptance was not a technological problem, but rather a "societal" problem. In order for the technology to advance, society had to advance, to rearrange itself in such a way that people could use the new technology. From this, it became clear that the "future" was more a certain configuration of society than a time and that what was inevitable was not technological progress, but this societal progress thought to be driven by technological progress.

Let us take an example of a certain future, which several people mentioned to me: the self-driving car. One researcher told me that he thought that self-driving cars are "so inevitable that many AI researchers would consider that almost a solved problem." I asked if it was inevitable because there were already prototypes. The researcher agreed, but added that it was also—if not more so—because there is already "a sort of deployment [of self-driving cars] happening at a small scale . . . [and] the obstacles to deployment [are] not the technical AI problems [but] the legal problems . . . the societal views of what it means [for cars] to self-drive and . . . [the] moral question also, like, if a car gets into an accident, who do you blame? . . . Which is sort of outside the realm of AI." Several researchers pointed out that the real threat to people is economic. To one researcher, what people "seem to be most afraid of is superintelligence, and the existential threat to humanity." However, he thought that was "ridiculously overblown." The real concern, for him, was that "technology is relentlessly replacing certain kinds of jobs and there's strong economic incentives to, say, replace . . . the three million [currently employed] truck drivers with self-driving trucks."

When the concepts of technology and society are kept separate from one another, and the problem of progress is seen to be one of societal change, then the ethical problems technologies may pose can be made society's problems rather than AI's. One researcher said, "Rather than AI sort of becoming evil . . . the other reason people are scared about AI is that it's going to take over their jobs [which is] maybe a little bit more legitimate of a concern [but] that's not a problem with AI, that's just a problem with progress. Whenever there is progress, whether it's AI or not, many jobs get obsolete and new jobs come up." Some explicitly pointed to social elements of progress when discussing debates on the ethics of AI. One researcher suggested that an important question to the debates was "To what extent is that reluctance or fear or discomfort with artificial intelligence going to get in the way and for how long?" This researcher was not so sure that: "Maybe it is inevitable that someday we'll get over it, just like someday we saw the benefit of blood transfusions and got over the fear that, you know, we're going to get somebody else's soul into our body." Another researcher agreed that whether or not something was an ethical problem had to do with the extent of the impact of it on society: "So far, even if you were ignoring the ethical repercussions, the impact was so limited that it didn't really matter. But in the coming years, machine-learning and AI is going to have such a huge impact, that the ethical problems are going to be front-facing and obvious to us."

This assumption is held out in the opposite way by opponents. Those who did not think that progress in the future was inevitable flipped the script and focused on the technological problems. When the focus is on the technological, researchers tend to express more doubt, showing that understandings of scientific progress such as Kuhn's are indeed influential. In fact, one researcher told me explicitly that he found Kuhn's model of science as problem-solving, in which new paradigms may solve some new problems while undoing solutions to old problems, to be the most convincing. Another researcher told me, "I'm not a singularity optimist in the sense [that] I don't believe it's inevitable . . . it's a scientific question as to whether [the assumptions underlying predictions of inevitable progress] are going to prove to be true or not . . . I don't believe there's any evidence that it's possible to have systems that are exponentially smarter than people. To some extent, people are optimal responses to their environment." By "optimal response," he means that we have adapted as best as we can to the pressures currently on us, and, in some ways, our intelligence is thus not entirely "general." He mentioned that he agreed to some extent with the criticisms levelled by *Wired* founding editor Kevin Kelly in his article "The AI Cargo Cult: The Myth of a Superhuman AI," in which Kelly argues against several assumptions underlying claims about the inevitable, exponential advancement of AIs. In particular, he agreed with Kelly that we do not yet know whether there are bounds on intelligence; that is a scientific question that has not yet been answered.

When asked about the present (and the process of making new technology they were currently engaged in), all the AI researchers I interviewed tended to focus on the technological and express these doubts. "Most of the smaller, if they're more independent projects . . . they have a kind of life of their own . . . You think you're doing one thing, and you'll start working on it, and it just, that doesn't work or something else becomes more interesting and you end up chasing that." On the commercial success of one of his projects, a researcher told me that "I think it's fair to say it [the success] was a hope. I don't think it was a, you could say it was an expectation. We could see that if we were very, very lucky and worked very, very hard, it might wind up [happening]." Other ethnographic studies furnish examples of people using narrative to make meaning out of the past, believing that events turned out the way they were meant to (Mattingly 2010). However, in this case, we can see that, actually, when forced to reflect on the past, researchers easily picked out the uncertainties. Perhaps they were simply trying to be humble. Nevertheless, this shows that they are able to see many ways that a project could have gone.

Everyone I interviewed acknowledged that there are factors that determine success in artificial intelligence research beyond sheer brilliance. Many told me that a strong influence on their work were other people around them. A guiding force on their careers had been having to compromise because of the interests of collaborators and the structure of academia, wherein the grunt work of most projects actually get carried out by graduate students. For example, one researcher told me, "If I am starting a new project, it's not just out of the blue, it's from some conversation with a PhD student or someone at that level and they have to be interested in it, they have to think it's important." Everyone said that they worked closely with other people.

While almost everyone I interviewed did not pose enormous ethical risk, time was still an important part of their arguments regarding this. Almost all the researchers I interviewed agreed that, whatever ethical problems AI posed, they would not be urgent until the future, specifically calling on a meaning of the future more linked to time than that of people who see the future as, say, a "set of configurations," like Nick Bostrom. However, while they did not see the future as being one in which AI was necessarily widespread in society, they still shared concerns with those interested in the ethics of AI about the spread of AI. For example, one researcher said, "As far as the technology that I'm involved in creating, right now I don't have a lot of worries about that because it's not highly-transitioned. If it does become transitioned into, you know, common use, the hard thing for me ethically is knowing that I've helped to develop something that is meant for good. All of my things are meant for good. All of my things are meant to help people. However, they could be used for bad." It is important to note the emphasis on technology being "transitioned," which is also important to people who are Singularity optimists. In fact, fear of the "transition" undergirds a lot of Singularitarians' concerns as well, and the difference is not that they think that we do not have time until tech is transitioned, but that it might be too late to stop it by the time it is transitioned and clearly poses a threat, but by then, it might be too late. For example, one of Nick Bostrom's most convincing arguments for the need for ethical development of AI with an eye to the future is his analogy to how we are all kind of dependent on the internet now because it is transitioned (Bostrom 2014).

This, again comes back to my central claim that most of these arguments depend on seeing technology and society as in competition with each other and putting the onus for ethics and the consequences of technology on "society." Technology can only pose a problem if it becomes part of society—society has the ethical problems, not technology. Another researcher told me, "I think the ethical concerns are really far off in the future . . . for the field as a whole." Of course, when one takes the concept of the inevitable in the denotative sense, whether or not something is far off in the future does not matter to its inevitability. Inevitability is about causality. And yet, inevitability as used by people to describe tech seems to embed it in this understanding of time wherein the inevitable is far off in the future because it is the society and future that is certain, and the past, present, scientific that is uncertain. This has actually been a common theme in discussion of technological progress since at least the 1980s. For example, in the 1980s, psychologist Sherry Turkle studied changes in attitudes toward the use of simulation among faculty and students at MIT who study architecture, civil engineering, chemistry, and physics respectively. During her time there, there were many controversies surrounding the Athena Project, an experiment which introduced computing to the general MIT population in the 1980s. Although many students and professors were unhappy at having to use CAD tools in class, a student told Turkle that those in charge of the Athena project "just assumed that the computer would win" (12). Turkle argues that the way students coped with tensions created by predictions that computers would revolutionize the design process while still letting design be firmly under the control of human designers was "to declare problematic effects inevitable but to relegate them to a far-distant future" (18).

Only one interviewee said that he dealt with ethical problems in his work at the current moment. But the ethical issues were not specific to AI. He created AI agents that acted as lawyers, so the ethical problems were not new problems, but the same ones human lawyers grapple with. To solve these problems, he "talked to legal ethicists and . . . read a lot of what the ethics around what lawyers can and can't do." He did not explicitly draw on any knowledge from computer science. This is in line with the reason most people gave for not focusing on ethics: they did not feel prepared to deal with ethical problems. "I feel I am not very well equipped to deal with ethical repercussions, I haven't been trained in it. I am not even sure I can evaluate what they would be or how to resolve them." This concern is shared by the researchers behind new research centers in the ethics of AI; the feeling that AI researchers (or, for that matter, anyone) are not equipped to deal with questions of ethics in AI is what drove them to open these centers.

Conclusion

This data indicates that, for AI researchers, success in the future is seen as more certain than in the present. Future success is, to many, inevitable. This is a consequence of researchers orienting themselves toward long-term futures, an orientation which is justified and encouraged by a classically evolutionist understanding of modern time. Within this temporality, "society" and "technology" are held apart, although technological advancements are thought to drive social change. In line with classically evolutionist social theories that underpin modernity, technology is thought to have a sort of power over people: it can force them into the necessary configurations for the technology to succeed.

However, though technological and social progress are thought to be related, different qualities are attributed to them. Technological progress is uncertain because of the social and political constraints on the scientific process in academic contexts. Research relies on the proclivities of graduate students, grant funding, and the abilities of all the researchers involved. Societal progress, paradoxically, is not subject to the same social constraints because it is constituted by the throwing off of those constraints. Thus, researchers can make claims, such as in the superintelligence argument, that social progress happens by snowballing: once a technology is deployed, people will adopt it and soon it will be widespread.

It is the certainty of advancement in the future that makes "technology" a category in which items are liable to fail or be supplanted, because there could be no future progress if we have already reached the pinnacle of technological advancement. Because advanced technology only exists in the future, its effects must also be in the future. That is why "the future" is central to discussions of the ramifications of advanced AI.

Two caveats. All this is not to say that the AI researchers making these claims must be wrong because they are biased by their own cultural backgrounds. The presence of bias does not necessarily invalidate claims. I have pointed out these biases merely to expand the scope of the discussion. Secondly, the diversity of attitudes toward progress in AI among AI researchers has to be acknowledged. Many more views exist than just those mentioned in this paper. In fact, the diversity of these attitudes is such that in the process of conducting this pilot study, it became clear that for a study like mine to have any use, its focus must be much narrower and its data gained from long-term ethnographic research with one small group of researchers.

References

Adams, Vincanne, Michelle Murphy and Adele E. Clarke. 2009. "Anticipation: Technoscience, life, affect, temporality." *Subjectivity: International Journal of Critical Psychology* 28(1): 246–265.

Ahearn, Laura M. 2001. "Language and Agency." *Annual Review of Anthropology* 30: 109–137.

Anderson, Owen. 2007. "Charles Lyell, Uniformitarianism, and Interpretive Principles." *Zygon: Journal of Religion & Science* 42(2): 449–462.

Appel, Hannah. 2017. "Toward an Ethnography of the National Economy." *Cultural Anthropology* 32(2): 294–322.

Banavar, Guru. "The Science of AI and the Art of Social Responsibility." *Huffington Post* (New York City, NY), Feb. 15, 2017 (updated Feb. 16, 2017). http://www.huffingtonpost.com/entry/the-science-of-ai-and-the-art-of-social-responsibility_us_58a4b920e4b0b0e1e0e205b9

Bear, Laura. 2014. "Doubt, conflict, mediation: the anthropology of modern time." *Journal of the Royal Anthropological Institute* 20: 3–30.

——. 2016. "Time as Technique." *Annual Review of Anthropology* 45: 487–502.

Bengio, Yoshua. 2016. "Machines Who Learn." *Scientific American* 314(6): 46-51.

Bird, Alexander. 2013. "Thomas Kuhn." *The Stanford Encyclopedia of* Philosophy (Fall 2013 Edition), edited by Edward N. Zalta, https://plato.stanford.edu/archives/fall2013/entries/thomas-kuhn/.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies.* Oxford: Oxford University Press.

——. 2015. "What happens when our computers get smarter than we are?" https://www.ted.com/talks/nick_bostrom_what_happens_when_our_computers_get_smarter_than_we_are. Accessed Dec. 21, 2017.

Bryant, Rebecca. 2016. "On Critical Times: Return, Repetition, and the Uncanny Present." *History and Anthropology* 27, no. 1: 19–31.

Campbell, Murray, A. Joseph Hoane Jr., and Feng-hsiung Hsu. 2002. "Deep Blue." *Artificial Intelligence* 134: 57–83.

Chalmers, David. 2010. "The singularity: A philosophical analysis." *Journal of Consciousness Studies* 17(9–10): 9–10.

Clark, Gregory. 2009. Review of *Contours of the World Economy, 1-2030 AD: Essays in Macro-Economic History. Journal of Economic History* 69(4): 1156–1161.

Cook, L.M., B.S. Grant, I.J. Saccheri, and J. Mallet. 2012. "Evolutionary biology: Selective bird predation on the peppered moth: the last experiment of Michael Majerus." *Biology Letters* 8(4): 609–612. doi:10.1098/rsbl.2011.1136

Escobar, Arturo. 1995. *Encountering Development: The Making and Unmaking of the Third World.* Princeton: Princeton University Press.

Farman, Abou. 2010. "The Intelligent Universe." *Maisonneuve* 36. Available online: https://maisonneuve.org/article/2010/08/2/intelligent-universe/.

Forbus, Kenneth D. 2010. "AI and Cognitive Science: The Past and Next 30 Years." *Topics In Cognitive Science* 2(3): 345–356.

Frazer, Sir James George. 1890. *The Golden Bough: A Study in Comparative Religion.* London: MacMillan Publishers.

Future of Life Institute, "Superintelligence: Science or Fiction? | Elon Musk & Other Great Minds," YouTube Video, 1:00:14, Jan. 30, 2017, https://www.youtube.com/watch?v=h0962biiZa4.

Gell, Alfred. 1992. *The Anthropology of Time: Cultural Constructions of Temporal maps and Images*. Oxford: Berg.

Giannetti, B.F., F. Agostinho, C.M.V.B. Almeida, and D. Huisingh. 2015. "A review of limitations of GDP and alternative indices to monitor human wellbeing and to manage eco-system functionality." *Journal of Cleaner Production* 87: 11–25.

Godfrey-Smith, Peter. 2003. *Theory and reality: an introduction to the philosophy of science.* The University of Chicago Press: Chicago.

Guyer, Jane I. 2007. "Prophecy and the near future: Thoughts on macroeconomic, evangelical, and punctuated time." *American Ethnologist* 34(3): 409–421. doi: 10.1525/ae.2007.34.3.409.

Hafner, Katie. "In An Ancient Game, Computing's Future: As Programmers Battle to Teach a PC to Play Go, They Unlock More Secrets of Artificial Intelligence." *New York Times*, (New York, NY), Aug. 1, 2002.

Hancock, P.A. 2006. "On Time and the Origin of the Theory of Evolution." *KronoScope* 6(2): 192–203.

Hernandez, Daniela. "Meet the Man Google Hired to Make AI a Reality." *Wired* (San Francisco, CA), Jan. 16, 2014.

Hesman Saey, Tina. "Moth Mutation Explains Classic Example of Evolution." *Wired* (San Francisco, CA), Apr. 15, 2011.

Hodges, Matt. 2002. "Time and Modernity in the Mediterranean: A Case from Languedoc." *Journal of Mediterranean Studies* 12, no. 1: 191–221.

Holston, James. 1999. "Alternative Modernities: Statecraft and Religious Imagination in the Valley of the Dawn." *American Ethnologist* 26, no. 3: 605–631.

Jennings, Ken. "The Go Champion, the Grand Master, and Me." Slate (New York, NY), Mar. 15, 2016.

Judge, Anthony J.N. 1993. "Metaphor and the language of futures." *Futures* 25(3): 275–288.

Kahn, Joel S. 2001. "Anthropology and Modernity." *Current Anthropology* 42(5): 651–680.

Kelly, Kevin. "The AI Cargo Cult: The Myth of a Superhuman AI." *Backchannel,* (San Francisco, CA), Apr. 25, 2017.

Koch, Christof. "How the Computer Beat the Go Master." *Scientific American*, (New York, NY), Mar. 19, 2016.

Latour, Bruno. Translated by Catherine Porter. 1993. *We Have Never Been Modern.* Cambridge: Harvard University Press.

Lesser, Alexander. 1952. "Evolution in Social Anthropology." *Southwestern Journal of Anthropology* 8 (2): 134–146.

Levinovitz, Alan. "The Mystery of Go, the Ancient Game that Computers Still Can't Win." *Wired* (San Francisco, CA), Apr. 12, 2014.

Liu, Weibo, and Zidong Wang, Xiaohui Liu, Nianyin Zeng, Yurong Liu and Fuad E. Alsaadi. 2017. "A survey of deep neural network architectures and their applications." *Neurocomputing* 234:11–26.

Mattingly, Cheryl. 2010. "Moral Willing As Narrative Re-Envisioning." In *Toward an Anthropology of the Will*, edited by Keith M. Murphy and C. Jason Throop, 50–68. Stanford: Stanford University Press.

McCorduck, Pamela. 1979. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. San Francisco: W. H. Freeman and Company.

McDermott, W.B. 1990. "The future as metaphor." *Futures* 22(2): 191–202.

Merrifield, Caroline, Erik Harms, Ryan Cecil Jobson, and Sayd Randle. 2013. "Modernity of This Century: Four Ethnographic Perspectives." *Anthropology of this Century* 8: http://aotcpress.com/articles/modernity-century-ethnographic-perspectives/

Metz, Cade. "In Two Moves, AlphaGo and Lee Sedol Redefined the Future." *Wired* (San Francisco, CA), Mar. 16, 2016.

Morgan, Lewis H. 1877. *Ancient Society*. New York: Henry Holt and Company.

Moroşanu R. and F. Ringel. 2016. "Time-tricking: a general introduction." *Cambridge Journal of Anthropology* 34: 17–21.

Munn, Nancy. 1992. "The Cultural Anthropology of Time: A Critical Essay." *Annual Review of Anthropology* 21(1): 91–123.

Nielsen, Morten. 2014. "A wedge of time: futures in the present and presents without futures in Maputo, Mozambique." *Journal of the Royal Anthropological Institute* 20: 166–182.

Olazaran, Mikel. 1996. "A Sociological Study of the Official History of the Perceptrons Controversy." *Social Studies of Science* 26, 611–659.

Reddy, Raj. 1988. "Foundations and Grand Challenges of Artificial Intelligence." *AI Magazine*: 9–21.

Ringel, Felix. 2016. “Beyond Temporality: Notes on the anthropology of time from a shrinking fieldsite.” *Anthropological Theory* 16(4): 390–412.

Shannon, Claude E. 1950. “Programming a Computer for Playing Chess.” *Philosophical Magazine* 7 41(314): 256–275.

Steinerberger, Stefan. 2015. “On the number of positions in chess without promotion.” *International Journal of Game Theory* 44: 761–767.

Turkle, Sherry. 2009. With William J. Clancy, Stefan Helmreich, Yanni Loukissas, and Natasha Myers. “Simulation and Its Discontents.” In *Simulation and Its Discontents*. Cambridge: MIT.

Vernor, Vinge. 1993. “The Coming Technological Singularity: How to Survive in the Post-Human Era.” https://edoras.sdsu.edu/~vinge/misc/singularity.html. Accessed Dec. 19, 2017.

Walford, Antonia. 2013. “Limits and Limitlessness: Exploring Time in Scientific Practice.” *Social Analysis* 57(1): 20–33.