IS AI A GENERAL PURPOSE TECHNOLOGY?

A Dissertation
Submitted to
the Temple University Graduate Board

In Partial Fulfillment
of the Requirements for the Degree
DOCTOR OF BUSINESS ADMINSTRATION

---

by
Mark F DeSantis
May 2024

Examining Committee Members:

Shreeram Mudambi, PhD, Advisory Chair, Management
Solon Moreira, PhD, Management
Min-Seok Pang, PhD, Management Information Systems
Mitchell Weiss, External Reader, Harvard Business School

ABSTRACT

A general purpose technology (GPT) is an exceptional link in a complex chain of innovations connected through space and time. Common examples include the steam engine, electrification, and the automobile (Bresnahan & Trajtenberg, 1995; Bresnahan & Trajtenberg, 1995; & Bekar, et. al., 2018). A GPT is set apart from other innovations because of the massive breadth, depth, and duration of their impact on our world and, therefore, worthy of special consideration. (Trajtenberg, 2019; Bekar, et. al., 2018; Strohmaier & Rainer, 2016; Ott, et. al., 2009; & Bresnahan & Trajtenberg, 1995). Yet, the changes GPTs bring can take decades and even centuries to manifest fully and much of GPT research to date is post hoc (Bresnahan & Trajtenberg, 1995; Lipsey, Bekar, & Carlaw, 2005; & Jovanovic & Rousseau, 2005). Yet executives, entrepreneurs, researchers, policymakers, and investors would all benefit from foresight anticipating the next GPT. This research explores this prospect empirically by tracing the 'pervasiveness dimensions' of a candidate GPT before its ultimate economic and societal impact is observed. Our candidate GPT is Artificial Intelligence (AI) and a key enabling technology, Deep Learning (DL). We trace the evolution of AI and DL from their birth in the latter half of the 20th century to the present, with a focus on the diffusion of large-scale applications of AI/DL from 2000 to 2020 (Goldfarb, et. al, 2021). We are aided by a rich base of prior academic research across various disciplines as well as industry analyses and reports. We hope to add a forward-looking tool to the decision-makers toolkit, while also increasing our understanding of AI not only as an innovation but as an extraordinary technological, economic, social, and even political phenomenon.

To my family and friends who pushed, guided, and encouraged me to pursue my curiosity, wherever it takes me.

ACKNOWLEDGMENTS

I want to thank my dissertation committee including my Advisory Chair Shreeram Mudambi, Ph.D., Solon Moreira, Ph.D., Min-Seok Pang, Ph.D., and Mitchell Weiss, Harvard Business School, along with all the excellent instructors and staff at the Fox School of Business. Their advice and support made this work possible. Last, I want to thank my colleagues at Bloomfield who put up with my absences and encouraged me along the way.

TABLE OF CONTENTS

Page

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

> We are victims of one common superstition, the superstition that we understand the changes that are daily taking place in the world because we read about them and know what they are.
>
> Mark Twain[1]

Every decisionmaker would benefit from anticipating the next technological innovation, especially if that innovation possesses the game-changing potential of, say, electrification, the steam engine, the Internet, or the automobile. These and other major innovations through time are the general purpose technologies (GPTs) that have transformed our world. This research explores a possible way the choices of executives, technologists, entrepreneurs, investors and policy-makers can be better informed when they confront a possible GPT. Fortunately, a rich research base exists, though much of it is focused on innovations of the past and their historical impact, and we combine it with empirical data to propose a forward-looking decision tool.

The main challenge with 'seeing over the horizon' is best characterized by Peter Drucker (1969) when he highlighted the near impossibility of predicting the impact of technological innovations:

> [innovations] …tend to be synergetic and the result of several developments, each independent in its origins and the outgrowth of a separate discipline with its own, …are often quite indirect, and by-products rather than main products & …tend to be threshold phenomena; until they reach a certain level, there is no impact at all.

---

[1]Twain, M. (1892) "About All Kinds of Ships."

Our research objective is more realistic than precise predictions of the future. Rather, our proposed tool would generate a meaningful and discerning range of possibilities when evaluating whether an innovation as a possible GPT. Such a tool is not without precedent, as we show in Chapter 4 when we discuss technology forecasting. However, our research is novel in that (1) we focus on the characteristics unique to GPTs, (2) we assess AI as a candidate GPT, and (3) we put AI in a historical and contemporary context.

We utilize GPT theory, along with innovation, economic, and diffusion theory, to assess a contemporary innovation with the potential impact of past GPTs. In choosing an innovation for our research we focused on one defining characteristic of a GPT identified in the literature and discussed later in this chapter.

To that end, several contemporary innovations were considered and, in the end, I chose artificial intelligence (AI), along with the key AI-enabling technology of deep learning (DL). [2] As the literature review reveals, there is a great deal of anecdotal and empirical data supporting the proposition AI and DL ("AI/DL") are presently creating and altering products, services, industries, markets, and economies around the world in a way and to a degree typical of GPTs from past eras. This makes AI/DL a strong candidate for GPT status and a proper choice for our analysis.

[2] In addition to AI, robotics, bio-medical advances, quantum computing, and edge computing we considered and explored as candidate GPTs and an innovation focus for this research.

The Macro-impact of Innovation

Early GPT research was informed by economic theory, with contributions over the decades from sociology, innovation theory, political science, and technology forecasting (Drezner, 2019 & Horowitz 2018). For economists, GPTs are 'engines of economic growth' and 'drivers of seismic (i.e., macroeconomic) shifts' (Strohmaier & Rainer, 2016 & Ott, et. al., 2009) explaining the underlying forces in recent history's sharpest Total Factor Productivity (TFP) shifts (David, 1991).[3] These macro-consequences have previously taken decades to manifest, which is why GPT research is post hoc and explanatory (David, 1991 & Bresnahan & Trajtenberg, 1995).

For other scholars of innovation, GPTs explain the 'micro-dynamics' of the innovation process, shedding light on how and why some innovations evolve to have out-sized impacts and others not so much (Strohmaier & Rainer, 2016; Ott, et. al., 2009; Bresnahan & Trajtenberg, 1995; Lipsey, Bekar & Carlaw, 1998; & Jovanovic & Rousseau, 2005). Despite a growing body of research supporting GPTs, scholars like Robert Gordon remain skeptical of single innovations generating macro-impacts, at least with the relative magnitude of prior innovations (Gordon, 2016; Cowen, 2011; Bloom et al., 2017; & Nordhaus, 2015). We address their critiques in Chapter 2.

In the meantime, it is important to understand some of the differing views about the relationship between economic change and innovation. For instance, not everyone

[3] TFP is the measure of the ratio of aggregate output (e.g., GDP) to aggregate inputs (Sickles & Zelenyuk, 2019).

believes contemporary GPTs like the steam engine, electrification, and automobile, for example, will or can ever come again. In 2000, at the peak of the internet boom and bust, economist Robert Gordon (2000, pg. 12) reminded us that technological upheaval is not unique to any time:

> Internet surfing may be fun, but it represents a far smaller increment in the standard of living than what was achieved by the extension of the day into night achieved by electric light. The revolution in factory efficiency achieved by the electric motor, the flexibility and freedom achieved by the automobile, the saving of time and shrinking of the globe achieved by the airplane, the new materials achieved by the chemical industry, the first sense of live two-way communication achieved by the telephone, the arrival of live news and entertainment into the family parlor achieved by radio and then television, and the enormous improvements in life expectancy, health, and comfort achieved by urban sanitation and indoor plumbing.

To be clear, Gordon believes innovation will continue to deliver benefits, it's just t those benefits will never be relatively as great as past innovations. Only time will tell if he is right. Nevertheless, Gordon's comparison of a recent GPT (i.e., the Internet) to electrification, the automobile, and the airplane reinforces the complex nature of the process of innovation generally and GPTs specifically. Can we anticipate, with reasonable probability if a given innovation is a GPT? This research attempts to find out.

Robert Gordon shared his perspective at the dawn of the internet, and (Gordon, 2016) later asserted that TFP growth has been in a long-run decline since the 1970s, only to continue well into the future. In particular, he believes productivity gains from the computer revolution of the late 20th and early 21st century was a 'one-off,' never to be repeated. Other economists echo Gordon's view that the pace of technological progress and its impact has declined substantially, with little recovery to the levels of the 19th and early 20th century (David, 1991; Gordon, 2016; Cowen, 2011; Bloom et al., 2017; &

Nordhaus, 2015). How can there be so much innovation without a visible benefit, as some claim?

Various reasons have been offered for this seeming 'slow-growth-high-innovation' (SGHI) paradox, though a closer look reveals possible misunderstandings and mischaracterizations of the relationship between innovation and macroeconomic outcomes. For one, we too often fall prey to our all too human 'non-rational' tendencies when we consider choices, outcomes, and decisions, including, among them, our own innate bias toward overly optimistic expectations from innovations, discoveries, and inventions or at least claims as such (Kahneman, 2011*)*. Even Marvin Minsky, one of the pioneers of AI, predicted in 1967 that, "Within a generation, the problem of creating 'artificial intelligence' will be substantially solved." (Minsky, 1967, pg. 2).

Another explanation for this seeming paradox is that many new technologies are initially concentrated in a relatively small group of investors in that innovation and the (relatively small number of) 'early adopters' who can afford it (Andrews, et. al, 2016 & Furman & Orszag, 2015). The auto industry is an example, as its first decades in the late 1800s and early 1900s consisted of hundreds of small, custom manufacturers serving a relatively small number of wealthy customers. Mass diffusion only came once Henry Ford introduced an affordable and reliable automobile, the Model T, in 1908, which itself was made possible by the innovations of assembly-line manufacturing and (mass-scale) vertical integration (Lacey, 1986)

In fact, SGHI may be less a paradox and more a mismeasurement problem (Mokyr, 2010). Both the implicit economic (i.e., time and effort) and out-of-pocket (i.e.,

cash) costs of consumption often vary little relative to the innovation's utility, as, for example, anyone who has spent hours listening to podcasts or searching the internet can readily attest (Mokyr, 2010).

Finally, the simplest explanation may be that the greater the actual far-reaching nature of the innovation, the greater the underestimation of the time needed to adapt to and adopt the innovation! This may seem like circular reasoning but, as we will see later, it is the 'co-invention' of enabling technologies like, for example, airports, wireless communications, radar, and jet engines that allowed the innovation of mass commercial air travel to happen.

The Coming Macro-impact of AI

A 2022 survey by McKinsey Global Institute found fifty percent of public companies have already, or will soon, apply AI (Mayhew & Hall, 2022). S&P 500 listed companies acquired 52 AI startups in 2021 and early 2022, compared with 24 acquisitions five years previous (Burke, 2022). Recently, chatGPT, a free online Natural Language Processing (NLP) driven chatbot, set online records by attracting fifty-seven million monthly active users in the first month of introduction (Cerullo, 2023). In fact, we need look no further than our own smartphone, home (e.g., Nest), or car (e.g., Tesla) for AI's growing presence and relevance in our lives.

Our collective human experience of the changes wrought by successive 'industrial revolutions' from the mid-seventeenth century through to today's advances in AI, robotics, and medicine may have conditioned us to expect an ever-expanding and continuous stream of innovations (Schwab, 2017). Further, some innovations first appear

as ‘magical’ and anecdotal evidence suggests AI may have entered our collective consciousness (Cockburn, et al., 2019). Perhaps it is the mysterious way AI seems to work, so much so even its creators struggle to understand *precisely* how it produces usable results (Xiang, 2022). Maybe it is the speed, breadth, and depth with which AI is now part of our daily lives (Smith, 2021). Or perhaps it is just the fact that every innovation possesses the potential for good, bad, or even complicated mix of both (Moltzau, 2019).

Yet is AI all that special as an innovation within a greater historical context? Afterall, the ballpoint pen (Invented 1888, Giesbrecht, 2015), washing machine (Invented 1851, Brown, 2009), and cigarette (Invented 1843, Mishra & Mishra, 2013) were significant innovations too! Can they be usefully compared to AI? Yes and no. Pens, washing machines, and cigarettes continue to impact our lives (for good and/or bad). Yet, as the literature review will reveal, with the passage of time our ultimate experience of AI might turn out to be very different because of the vastly greater reach of its impact.

In fact, is AI already a GPT? This is one focus of our research, and our answer is ultimately *yes*. However, as the literature review will show, there is not yet a clear consensus among scholars, researchers, technologists, and policy makers supporting this view. We therefore explore empirical data providing early indications an innovation, AI/DL in our case, currently or on a clear path to being a GPT. I hope our data analysis will serve as a foundation for more thoroughgoing tools decision-makers can apply to

better anticipate whether and to what extent AI/DL or any innovation might be the next game changer.

Where is the consensus on AI as a GPT? For some academic researchers, industry analysts, and leading technologists, AI will have an economic and societal impact far greater than the railroad, electrification, automobile, or perhaps any innovation yet created (Agrawal, Gans, & Goldfarb, 2018a/b). In addition, recent empirical data suggests AI may already be on such path, as analysis of AI-related patents shows its diffusion has increased substantially over first two decades of the 20th century in industries as diverse as structural biology, transportation, and imaging (Cockburn, et al., 2019). Nevertheless, the empirical evidence as to whether AI is diffusing on a mass scale remains limited and, in some cases, mixed.

Countering the view of AI or any innovation being a GPT from this day forward is best exemplified by the views of Robert Gordon (2016), when we further asserted that between 1870 and 1970 the railroad, automobile, and electrification brought unprecedented income growth, improved health, increases in productivity, mass communications, and practical conveniences with such speed, scale, and impact that these will be impossible to repeat, ever.

His reasoning was based on the idea that the magnitude and suddenness of the macro-changes wrought by the better-known innovations of the late 19th through the mid 20th century (at least in the United States and Europe) were so great that any comparable gains would, at best, be marginal, no matter what improvements they might bring or, for that matter, even if replaced by newer innovations. Gordon offers the example of the

washing machine, which brought such significant change that, when compared to washing clothes by hand, no current advance in washing machine technology could ever match the degree of relative change in washing efficiency and effectiveness gains from hand washing to the first washing machines.

Innovation is a relentless force for change. Joseph Schumpeter first described it as a process of 'creative destruction' or a "…search for, and the discovery, experimentation, development, imitation, and adoption of new products, new production processes and new organizational setups (Dosi, 1988, pg. 222)." The fruits of innovation can sometimes be bittersweet, exemplified in the poignant examples of the automobile (i.e., unprecedented mass mobility along with urban sprawl, traffic fatalities, and pollution) and the internet (i.e., unprecedented mass communication along with reduced privacy, hacking, and polarization).

GPTs, by definition, have wide and deep impacts within and across market, industry, economic and geographic 'dimensions.' Given that, we would therefore expect to see early evidence of a possible GPT's pervasiveness deep within and across all these dimensions. This research proposes the 'completeness' of a GPT's pervasiveness is a key precursor of an innovation being or on a strong to becoming a GPT. To that end, we rely on empirical data exploring proxies for our pervasiveness dimensions, for AI/DL, in the form of substantial investments of time, money, and effort in AI/DL across startups, applications, patents, research papers, and publicly announced policies.

We sought data for this research that revealed the characteristics of pervasiveness of AI/DL (together and separately) for the key period of 2000 to 2020 for each

dimension of pervasiveness. We define pervasiveness as a tangible resource commitment within a given dimension. We further define 'complete pervasiveness' as the extent to which those resource commitments show across all dimensions.

Our five dimensions include: (1) products and services largely reliant on the innovation (2) novel applications of the innovation, (3) the transfer of knowledge of the innovation, (4) public policy initiatives specific to the innovation and (5) geographic dispersion of these four dimensions. Each is further defined in subsequent chapters. In the end, we hope this analysis will serve as the foundation for a decision tool that can inform our technology and technology policy choices.

CHAPTER 2

WHAT IS A GENERAL PURPOSE TECHNOLOGY?

According to Lipsey, et al, (2005, pg. 98) a general purpose technology (GPT) is a "…single generic technology, recognizable as such over its whole lifetime, that initially has much scope for improvement and eventually comes to be widely used, to have many uses, and to have many spillover effects." Early GPT innovations like the wheel and the written word have been with us for millennia while the steam engine, automobile, and electrification are more recent examples. All these and other notable innovations have had a similarly monumental impact on industries, markets, economies, and geographies. They are essential to the functioning of our world. Yet we are now in an age of robotics, artificial intelligence, and life sciences technologies that are changing our lives in yet-to-be-discovered ways and to a degree we are just beginning to understand.

As daunting as the possible changes from AI/DL appear, such changes are hardly without precedent in even more recent history. Steam power first drove large increases in industrial productivity in Britain, continental Europe, and, later, the United States, in the early 19th century on through to the early 20th century (Crafts & Mills, 2004). Though it was not until the latter half of the 19th century that it became *pervasive,* as it was deployed across coal mining, cotton textiles, and metal manufacturing (Kanefsky, 1979). It is this quality of pervasiveness – defined, discussed, and explained in Chapter 4 – that is the focus of our empirical analysis in Chapter 5. We further propose that GPT

characteristic of pervasiveness is largely what differentiates GPTs from all other innovations.

Understanding how, where, when, and why an innovation 'diffuses' is fundamental to understanding GPTs and innovation as a process. In the case of steam engines, global diffusion was slowed, however, by the limits of enabling technologies (i.e., vast amounts of coal needed to generate the steam and the limited capacity of early engines), though diffusion did pick up with the availability of more efficient, high-pressure engines (Kanefsky, 1979). We explore the theory of diffusion of innovations later in this chapter.

Another foundational aspect of our analysis is the relationship between innovation and economic disruption (Agrawal, Gans, & Goldfarb, 2018b & Jovanovic & Rousseau, 2005). GPT theory highlights the complex temporal connection between technological *cause* with macro effect. Decades can separate the two, though time lags can and appear to have collapsed with each technical advance in that innovation (Jovanovic & Rousseau, 2005; Bresnahan & Trajtenberg, 1995; & Lipsey, Bekar, & Carlaw, 1998). We therefore posit candidate GPTs should be revealed in a visible and consequential 'upswing' or cumulative growth in, in our case AI's, 'pervasiveness' over a substantial period time.

For some scholars, the number and variety of possible GPTs has grown beyond electrification, steam power, and the automobile to include cloud computing (Etro, 2009), blockchain (Filippova, 2019), nanotechnology (Forti, Munari, & Zhang 2019), 3D printing (Choi, 2018), and the Internet-of-Things (Edquist, Goodridge, & Haskel,

2019). 1 presents an extended list of possible GPTs. However, skeptics see these proposed GPTs as blurring the boundaries between 'run of the mill' innovation and GPTs (Mokyr, 2010), a topic we will turn to later. Figure 1 presents a list of possible GPTs through the ages.

Domestication of plants	Three-masted sailing ship	Motor vehicle
Domestication of animals	Printing	Airplane
Smelting of ore	Steam engine	Mass production
Wheel	Factory system	Continuous process
Writing	Railway	Lean production
Bronze	Iron steamship	Internet
Iron	Internal combustion engine	Biotechnology
Waterwheel	Electricity	Nanotechnology

(Lipsey, et. al, 2005, pg. 132)

Figure 1. List of possible GPTs

In response, Bekar, et. al., (2018, pg. 32) offered additional distinguishing characteristics:

- Complementarities with a cluster of technologies that define and support it.
- Complementarities with a cluster of technologies that it enables.
- Complementarities with a cluster of technologies that typically include those that are socially, politically, and economically transformative.
- [Having] no close substitutes.
- Having a wide array of applications.
- Initially crude but evolving in complexity.

However, it was Bresnahan & Trajtenberg (1995) who laid the foundation for much of GPT theory when they asserted some innovations standout as 'engines of growth' possessing a 'wide scope for technical improvement' with 'many uses across a wide range of products and services' while also 'complementing pre-existing

technologies' (Hall & Trajtenberg, 2004 & Petralia, 2020). One outcome of this research will be increased specificity in the definition of GPTs by elaborating on the characteristic of 'pervasiveness' which is currently defined in the research literature as the diffusion of an innovation across the depth and breadth of markets, industries, economies, and geographies, while retaining that innovation's distinctive features as it evolves beyond its original, first uses (Jovanovic & Rousseau, 2005).

Pervasiveness appears repeatedly across economics and innovation literature, although often under different names. One example is 'extended reach' which is the instance where an innovation has a substantial impact both across applications, products and/or services and disparate markets, industries, and geographies (Beaudry, et al., 2010 & Helpman, 1998).

Classifying innovations is difficult. The pencil, for example, has changed very little in form and function from the original graphite encased in wood invented in the 1600s. Similarly, the steam engine's functional elements have changed little from the time it first pumped water out of mines in the late 1700s to power generation on everything from the grid to nuclear submarines. In fact, contemporary steam turbines that convert various energy sources (e.g., coal, natural gas, hydro and nuclear power) to useable electricity and vary little in functional elements from its 19th century progenitors, despite dramatic increases and changes in their performance (e.g., efficiency, safety, etc.) and components (e.g., materials, tolerances, etc.), respectively.

The pencil offers an interesting contrast. Though an important innovation, it simply does not rival the steam engine in its macro-impact. This claim is not based on

the longevity of the innovation (both are still in use and the pencil predates the application of the steam engine by a century), differences in technical sophistication, or a waiting of intrinsic value. It is based on the extreme differences in the breadth and depth of impact between the two innovations across economies, markets, industries, and geographies.

This research hopes to make several contributions to both innovation and GPT theory and research. First, we hope to create an expanded, more useful definition of GPTs by adding an 'empirical grounding' to the concept of pervasiveness. Second, a more grounded definition of pervasiveness will provide tangible empirical precursors to the emergence of a possible GPT. This possibility will give decisionmakers a means to 'look ahead' to the implications of game-changing innovations. Third, this research could stimulate further research into the process of innovation and help explain how and why some innovations go on to become GPTs and, perhaps, others do not.

There is an increasing interest in AI as a likely GPT (Agrawal et al., 2019; Cockburn et al., 2019; & Brynjolfsson et al., 2021). Aghion et al., (2018), for one, claims AI is either at or approaching GPT status and transforming wide swaths of the global economy and society, including factory automation, autonomous driving, healthcare, entertainment, and communications. In fact, some are labelling AI a 'general purpose meta-technology' (Romer, 2008 & Agrawal et al., 2018b).

Some scholars have tracked AI's diffusion, and we review the more common approaches to doing so in the next chapter. However, one significant challenge with assessing AI/DL's macro-impact is its limited history as a large-scale application,

particularly when compared to longevity of prior GPTs (Stuart & Norvig, 2003). Though AI/DL's conceptual beginnings go back 75 years, it is only in the first two decades of the 21st century that it has been applied on any scale. This research will take the history of AI/DL as it is and attempt to identify useful empirical 'dimensions.'

One of the founders of AI, Jeffrey Hinton (2021) believes it is a disruptive technology but one with a future that is as hard to predict as any technology yet created. From its earliest days in the late 1940s AI showed promise, initially as logic-driven, rules-based 'software' for playing checkers and, later with early expert systems applied to medical diagnoses (Domingos, 2015). Despite this demonstrated potential, early AI research failed to sustain funding through the 1960's and 70's, owing partly to the complexity of the dominant rules-based approach to early AI systems and their limited utility outside of controlled lab settings and highly constrained problem domains (Markov, 2015). Further, extremely limited compute power (by today's standards) and limited funding brought on the AI 'winters' of the late 1970s and early 1990s (Stuart & Norvig, 2003).

Nevertheless, AI's (first) 'resurrection' came in the 1980s with the simultaneous emergence of four factors: (1) an increasing accessible supply of varied data; (2) increases in processing power and efficiency (owing to the inevitability of Moore's Law); (3) renewed academic, governmental, and corporate research, and (4) entrepreneurs testing the limits of AI's capabilities and applying it in wide variety of ways in varied settings – learning what worked and what didn't (Agrawal, Gans, & Goldfarb, 2018a).

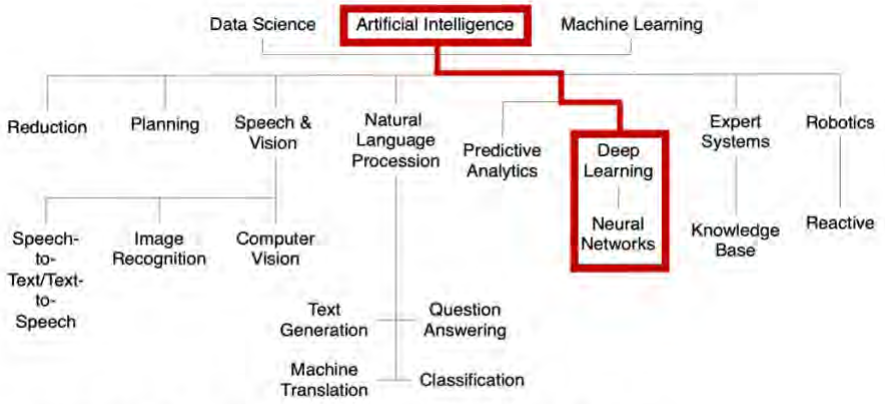

Figure 2. The AI family is defined by Deep Learning

AI's second resurgence came with the growth of the first large-scale real-world applications in the late 1990s and early 2000s such as with Google Search. These early deployments spawned the growth and evolution of AI's 'family tree' of complementary technologies and techniques, as shown in Figure 2.

Specifically, DL emerged as critical enabling 'branch' technology that informed or enhanced the capability and capacity of AI across this family tree. AI pioneer, Geoffrey Hinton (2021) would go on to assert that: "…take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than existing techniques." This is possible because DL uniquely possesses, "…the ability to take known information to generate new information. Better prediction is likely to have widespread consequences because prediction is fundamental to decision-making (Agrawal, et al., 2018a, pg. 5)."

At a fundamental level, DL permits AI applications to recognize subtle and novel patterns in massive data sets, resulting in a continuous, self-reinforcing loop of

prediction and learning. However, just as it is the case for most modern technological innovations, understanding AI/DL's unique capabilities can be a daunting challenge without deep technical training in and knowledge of computer science.

Nevertheless, a conceptual understanding of those capabilities may be sufficient to understand this research. For example, we can readily understand what Google Maps and Uber ride apps do for us and, to some extent, how these applications work. Yet we are not quite sure at a deeper technical level how they really work. Yet, what are the defining features of a given AI application? This research assumes that a core, definitive capability of a growing number of AI applications is an ability that we human's take for granted, which is the capacity to perceive complex, dynamic patterns in the real-world (i.e., unstructured) data and improve through continuous feedback. In a sense, DL moves AI and many AI-driven applications such as robotics and autonomous driving closer to overcoming the challenge of Moravec's Paradox, (Moravec, 1988, pgs. 15-16):

> Encoded in the large, highly evolved sensory and motor portions of the human brain is a billion years of experience about the nature of the world and how to survive in it. The deliberate process we call reasoning is, I believe, the thinnest veneer of human thought, effective only because it is supported by this much older and much more powerful, though usually unconscious, sensorimotor knowledge. We are all prodigious olympians in perceptual and motor areas, so good that we make the difficult look easy. Abstract thought, though, is a new trick, perhaps less than 100 thousand years old. We have not yet mastered it. It is not all that intrinsically difficult; it just seems so when we do it.

Thus, deep technical knowledge of the algorithms, complex software, hardware, or other technical details is not essential to appreciating what is vital about AI/DL for this research. Fortunately, we are also helped in analysis of AI/DL as a GPT by a wealth of theory and research across a range of disciplines including innovation theory

(Bresnahan & Trajtenberg, 1995 & Andergassen, et al., 2017, macroeconomic policy (Van Zon & Kronenberg, 2003), economic history (Crafts & Mills, 2004), and the organization of markets and industries (Lindmark, 2005). Gaps in our understanding of complex innovations like GPTs remain and further research is needed (Lipsey, Carlaw, & Bekar, 2005). As we will see in the next section, determining what is and is not a GPT systematically and consistently is a challenge, though hardly an impossible one (Moser & Nicholas, 2004).

Critiques of GPT Theory

Not everyone supports GPT theory. Some skeptics simply do not believe a 'general purpose' technology can alone have a massive impact, at least as defined by current GPT theory (Field, 2011; Crafts & Mills, 2004; Cantner & Vannuccini, 2012; Ristuccia & Solomos, 2014; & Gordon, 2016). For the skeptics, innovations are complex interconnected webs of dependencies that make precise prediction of their evolution nearly impossible. Further, the process of innovation is far too dependent on time, place, or circumstance to draw generalizable laws or principles. We address these concerns and critiques next.

As Lipsey, et al. notes (2005), "…if the concept of a GPT is to be useful, then GPTs must be identifiable." Some take Lipsey's caution further – perhaps further than he intended – and claim the criteria for what is and is not a GPT are too subjective to be usefully applied and thus GPT theory lacks empirical validity and relevance (Field, 2011). For example, Field (2011) asks if the railroad's distinguishing feature is the track, locomotive, freight car, or some combination of the above? Field further argues that

innovation is not a measurable thing, but a complex, continuous, and evolving *process* of combining and recombining ideas, knowledge, and experience applied to real-world problems (2011, pg. 4):"…many single-purpose innovations…with arguably large implications for TFP…" and that "…there can be 'general-purpose' innovations that represent relatively marginal improvements over previously existing technology and have low social saving/contribution to TFP (e.g., the felt-tipped pen)."

Still, much of available empirical research supports a positive association between a GPT 'diffusion' or adoption across markets, industries, applications, and geographies in ways predicted by theoretical models (Aghion & Howitt, 2000 & Helpman & Trajtenberg, 1998a; 1998b). While it is true innovations 'stand on the shoulders' of earlier innovations, some uniquely stand out. Field's critiques do not appear admit this.

Crafts & Mills (2004) extend Field's critique and assert the steam engine fails to qualify as a GPT according to the theory's own assumptions. For instance, one of the earliest applications of steam power was the Newcomen engine invented in the early 1700s and first used for pumping water from mines. Later versions of the steam engine powered ships, locomotives, and factories from the early 1800s through the early 20th century (Crafts & Mills, 2004).

Of course, the diffusion of steam power depended on advances in, and the availability of, 'enabling' innovations, such as the condenser and practical pressure gauges, rather than singularly on the steam engine (Crafts & Mills, 2004). Crafts & Mills (2004) found two-thirds of steam power was utilized across coal mining, cotton textiles,

and metal fabrication by 1870, and yet steam power was still slow to diffuse and not sufficiently pervasive (at that time) to meet the key defining criteria of a GPT. It was only later advances that brought enabling technologies, processes, and skills that made mass diffusion possible.

Most innovations trace their origin to earlier innovations. Further, most innovations are possible only because other technologies 'enabled' the birth of and evolution of a technology. Popcorn (Invented circa 4700 BCE, Kummer, 1988), one outcome of the agriculture revolution of ten thousand years ago, would not itself have been possible if not for the innovation of the two enabling technologies of fire (First applied circa 300,000 BCE, Johnson, 2022) and iron smelting (Invented circa 5000 BCE, Wertime & Muhly, 1979)!

But one enabling technology (i.e., fire) came much earlier than agriculture while another (i.e., iron) came much later. Therefore, asserting steam engine's diffusion was 'slow' is a relative statement without providing a context for GPTs and how they might diffuse. As shown later, empirical research has connected cause (i.e., GPT) and effect (i.e., macro-impact), though separated by decades (Rosenberg, 1976; Jovanovic & Rousseau, 2005; & Bresnahan & Trajtenberg, 1995).

The diffusion of innovations is a social phenomenon accelerated or slowed by different cultural, societal, economic, and political contexts and factors (Jovanovic & Rousseau, 2005; & Lipsey, Carlaw & Bekar, 2005). For example, diffusion of, say, an agricultural innovation in 18th century northern Europe would differ a great deal from the diffusion of that same innovation across 17th century China or 19th century Kansas or

Iowa in North America. In the case of AI/DL, contextual factors such as complementary enabling algorithmic, software, hardware, data storage innovations; new processes; organizational changes; and new skills matter a great deal (Brynjolfsson et al. 2017; David 1991).

Another example of why context matters is the transistor (Invented 1947, Zorpette, 2022). It enabled the diffusion of early computers and other electronic devices, including the hearing aid. Yet a major industry at the time, the telecommunications sector – a major beneficiary of the technology – took far longer to adopt it than was the case in other industries, despite other advances in technology within the telecommunications industry (Bresnahan & Trajtenberg 1995).

Another example may be helpful. The first industrial dynamo was introduced in the 1870s, yet its impact on manufacturing would not be felt until the 1920s (Bresnahan & Trajtenberg 1995). Some factory owners and managers adapted their organization and management practices, like changes to factory layout and the addition of complementary innovations, while utilizing steam turbine for centralized power generation through a utility, and some did not and, consequently, failed (Smil 2005 & David 1991).

For GPT skeptics, the differences between GPTs and other innovations are not often distinct enough to be usefully applied (Cantner & Vannuccini, 2012). David & Wright (1999, pg. 3) criticized the growing number of technologies labeled as GPTs, but also allowing for GPT 'sub-types:'

> One has only to consider the length of such proposed lists of GPTs to begin to worry that the concept may be getting out of hand. History may not have been long enough to contain this many separate and distinct revolutionary changes. On closer inspection, it may be that some of these sweeping innovations should be

> better viewed as sub-categories of deeper conceptual breakthroughs in a hierarchical structure. Alternatively, […] historical episodes may be fruitfully understood in terms of interactions between one or more GPTs on previously separate historical paths.

Ristuccia & Solomos (2014) questioned the macroeconomic impact of GPTs in their study of European electrification in the early 20th century, where they found no direct relationship between economic growth in Europe and its electrification. In fact, they were heavily critical of the concept of a GPT theory (Ristuccia & Solomos, 2014):

> …we might need to demote GPT theory to a specific case of a more general and mundane relationship between technology and long-term growth according to which a big (not necessarily a pervasive) technological change will have big effects.

The lag between a GPT's early diffusion and its macro-impact is often cited as undermining the case for GPT theory. Electrification and the automobile took time to diffuse, in part, *because* they are so far reaching. And much of GPT research, including that of the theory's skeptics, supports this assertion (Rosenberg, 1976; Jovanovic & Rousseau, 2005; & Bresnahan & Trajtenberg, 1995). However, the weight of research from our literature review reveals that, though substantial time lags exist (e.g., decades, or less) between cause (i.e., critical mass diffusion of the GPT) and effect (impact on TFP), these lags do not eliminate the possibility of that innovation being a GPT.

Understanding Diffusion

Innovations diffuse much like a virus across and through industries, markets, economies, and geographies. That is why diffusion theory is important to understand, as it focuses less on what is diffused (e.g., ideas, techniques, products, know-how, etc.) and more on how, when, where, and why an innovation diffused (e.g., what mediums were

used to communicate, how often they were communicated, where they were communicated, etc.) (Greenhalgh, et al., 2005). New ideas bring technological, social, economic, and even political change which becomes the tangible evidence of that diffusion. Diffusion can be viewed as a 'socio-communicative' process influencing and/or controlling the communication of the innovation and its diffusion (Rogers, 2003).

For Rogers (2003), 'adoption' occurs when a person, group, and/or organization does something differently after encountering the innovation than before. The process takes time as individuals with differing interests, resources, abilities, and locations are exposed to and choose to adopt the idea, product, or concept. This process is frequently rendered as a simple logistic progression or 'S' curve as presented earlier in Figure 3 and it directly corresponds with a (Gaussian) adoption curve.

So-called 'early adopters' are the adventurous few who first try out the innovation and, with success, are followed by more pragmatic 'late' adopters, and on to mass adoption (Bass, 1969). It must be kept in mind not all innovations diffuse. In fact, many don't. Schumpeter (1911), for one, knew ideas inevitably become new products and services, as any experienced entrepreneur or investor can attest to! For him, this causal chain from idea to prototype to product and service is long, complex, and, consequently, very hard to predict, as every technology researcher, entrepreneur, investor, and executive knows (Gottfredson & Aspinall, 2005).

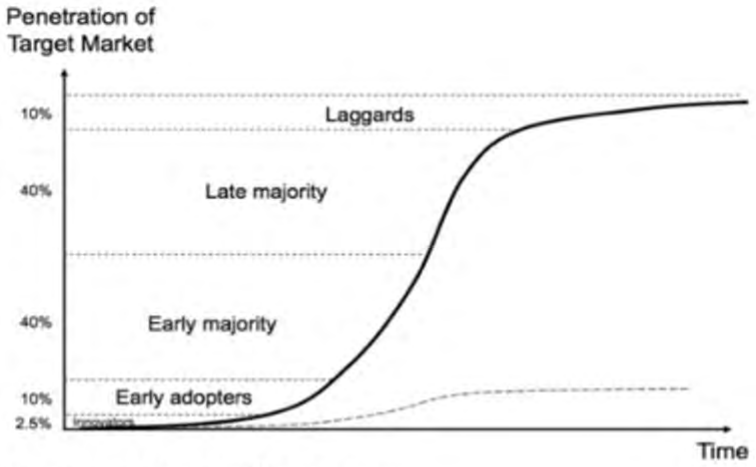

Figure 3. A classic diffusion curve

Often, the caution of potential users and the need for changed processes and skills can slow or even block diffusion (Lee, et. al.., 2010). In fact, growth in the awareness and adoption of an innovation across an organization, market, economy, group (e.g., engineers, dentists, teenagers, etc.), or geography can be thought of as our working definition of diffusion and pervasiveness being the outcome of diffusion (Peres, Muller, & Mahajan, 2010).

Diffusion theory originated with E. M. Rogers in the early 1960s and his seminal study, “Diffusion of Innovations” (Rogers, 2003). In it, he traced the flow of ideas, products, and concepts through populations, organizations, or society. Diffusion is multi-stage process of awareness, interest, evaluation, trial, and adoption (Rogers, 1995). Diffusion, and within it, adoption, are complex processes and their progress differs in unpredictable ways, even when the innovations and their respective contexts are very similar (Jovanovic & Rousseau, 2005).

For example, Jovanovic & Rousseau (2005) found, when comparing electrification and information technology, each diffused at different rates resulting in differing levels of disruption. Though both diffused widely, electrification diffused at a faster rate than IT, while IT improved faster than electrification and, ultimately, both birthed more innovation IT may ultimately have the larger impact (Jovanovic & Rousseau, 2005).

Precise and accurate prediction of the path of diffusion is very, very difficult, in part, because of the number and complexity of interacting variables (Furman 2018; Goolsbee 2018; & Naimat, 2016). David (1991, pg. 356) highlights this complexity:

> [the]…extended trajectory of incremental technical improvements, the gradual and protracted process of diffusion into widespread use, and the confluence with other streams of technological innovation, all of which are interdependent features of the dynamic process through which a general-purpose engine acquires a broad domain of specific applications.

These challenges informed our effort to reduce 'definitional confusion (discussed in depth later).' Afterall, AI is a now a label applied to so many things, so often that it may have reached 'semantic satiation' or the point where it has lost useful meaning (Jakobovits, 1962). As a result, overlapping data was utilized for our data analysis to mitigate, to the extent possible, any latent definitional ambiguity inherent in the data sets (López-Cobo et al., 2019).

It would not be the first, or last, time diffusion studies were challenged by definitional ambiguity, particularly for the study of AI (Davis & Chouinard, 2016). For example, attempts to quantify AI's diffusion often relied on absolute and relative investments as a percentage of total expenditures on seemingly AI-related and/or

dependent technologies, which may under or over-represent the true amount of investment (IDC, 2017 & Orinox, 2017).

Third, R&D expenditures and venture capital investments have been used as proxies for the diffusion of AI, but definitional ambiguity sometimes resulted in misleading, confusing, and non-comparable results (MGI, 2017; OECD, 2018; & Science-Business, 2018). Further, patent data has also been used as an empirical measure of AI's diffusion (Eaton & Kortum, 1999). We mitigated the impacts of these issues on our measurements of pervasiveness by using overlapping data sets in our analysis.

Adding to the complexity of our analysis is the fact that diffusion is not a single moment in time but rather a series of successive waves (Rogers, 2003; Helpman & Trajtenberg; 1994; Crafts, 2003; Jovanovic & Rousseau; 2005; Bresnahan, 2010; & David & Wright, 2003) with the first waves often generating negative productivity growth (Helpman & Trajtenberg, 1994; Crafts, 2003; Jovanovic & Rousseau, 2005; Bresnahan, 2010, & David & Wright, 2003). This negative economic impact, paradoxically, slows the initial growth in the innovation that ultimately brings growth. Why might this be? A simple example helps.

Anyone promoted into a new professional role knows that being productive in a new job means learning new routines, new social connections, terminology, and skills. This learning (i.e., adapting) takes time and effort away from doing the core work of the new role (i.e., production). This is especially true when the employer makes a major

departure from the current way doing things. In fact, researchers have found it can take eight months for a newly promoted employee to reach full productivity (Ferrazzi, 2015).

The same logic applies to innovations, as the initial productivity losses will be (hopefully!) more than made up for once it is fully adopted. This process is helped by supplementary, and complementary enabling technologies (Crafts, 2003; Jovanovic & Rousseau, 2005; Bresnahan, 2012; & David & Wright, 2003).

This 'change over' from innovate–diffuse–adopt–adapt creates a 'friction' as progressive abandonment of current tools is nearly simultaneous with an adoption of new tools (Lipsey, Carlaw & Bekar, 2005). An example is found with the transition from the horse to the tractor in early 20th century agriculture in the United States as in Figure 4. The final cost of this friction, which is unknown when the choice is made to adopt the innovation, is often the difference between success or failure of individuals, organizations, and even nations (Kennedy & Funk, 2016).

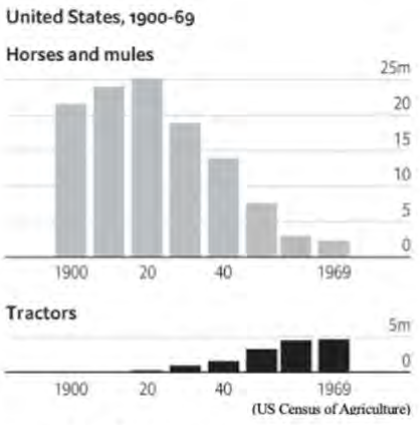

Figure 4. The transition from horses & mules to tractors

As has been the case, we combined AI and DL into a meta-concept by using 'AI/DL' for convenience's sake and have differentiated these concepts when necessary.

For now, AI/DL refers to any machine or algorithm capable of observing its environment, learning, and, based on the knowledge and experience gained, taking intelligent actions and decisions (Craglia et al., 2018). Nilsson (2009) offers a similar definition: "…activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment."

Measuring diffusion is difficult but not impossible (Valdes & Rudyk, 2017; Webb, 2019; Alderucci et al., 2020; & Schumpeter, 1939). Some, however, believe measuring AI/DL's diffusion is not yet practical because it is still in is earliest stages of that diffusion (Brynjolfsson et al., 2021; Valdes & Rudyk, 2017; Webb, 2019; & Alderucci, et al., 2020). Others assert the opposite and point to empirical research asserting AI/DL's diffusion is well underway and may be in its second diffusion wave (Trajtenberg, 2019). Gries and Naude (2018) and Acemoglu et al., (2020) make separate but similar cases that AI triggered an expansion of output and changing occupational task structures, while stimulating long-run employment. The weight of the available research suggests that AI/DL diffusion is underway and has been since roughly 2000, at least applied form.

Yet another complexity with understanding diffusion is the role of public policy in facilitating and/or impeding diffusion. For example, economists and policy analysts have intellectual property law, which though intended to protect the integrity of an innovation as it diffuses, it can also impede its diffusion (Mokyr, 2010; Boldrin & Levine, 2008; & Merges & Nelson, 1994). Mokyr (1994) claims the steam engine would

have diffused at a faster rate but for the fact that, "… [the inventor Watt] held a wide-ranging patent, he succeeded in blocking [the development of high-pressure steam engines] for many years". Boldrin and Levine (2008) argue our GPT example steam engine that: "…keeping prices high and prevented others from producing cheaper or better steam engines, Boulton and Watt hampered capital accumulation and slowed economic growth." The airplane would have diffused faster, according to Merges and Nelson (1990), because "…there is good reason to believe that the Wright [brother's] patent [on an efficient stabilizing and steering system] significantly held back the pace of aircraft development in the United States." Patents are tangible evidence of the diffusion of AI/DL and therefore, utilized in our later data analysis.

Finding GPTs

Predicting an innovation's evolutionary path is nearly impossible. Electric motors, first introduced in the late 19th century, reached mass diffusion in the early 20th, and replaced the centrally located steam power machines over decades. Yet the most forward-thinking managers and factory owners of the day originally saw electric motors as providing only incremental and marginal improvements in efficiency (David, 1991). In the end, the benefits of reduced energy costs over steam-driven power turned out to be only one of the electric motor's many unanticipated benefits, as the productivity gains from new factory designs, automation, and the flexibility to locate production virtually anywhere were even greater (David, 1991).

Predicting an innovation's path is so difficult partly because multiple, countervailing forces slow diffusion and in unexpected, hard-to-predict ways. For

example, ‘lock-in’ is an instance when innovation adoption and the resources used for that adoption become self-defeating barriers to further adoption (Klitkou, et. al., 2015). In that sense, innovation can become an ‘escalator’ from which it is too costly to get on (e.g., opportunity cost) or off (e.g., lock-in).

Investment in new routines, new policies, and enabling technologies are often an essential component of rapid diffusion (Fichman, 1992). But these investments too can also be self-limiting (Bresnahan, et al., 1996). On an individual level, we experience this when we buy something new – from smartphone to air-fryer – and need to acquire and/or understand a bewildering array of supporting and complementary accessories and new routines.

Yet another factor greatly complicating any precise and accurate prediction is that innovations evolve far from their original purpose. Fire for cooking meat, for example, is the progenitor to the steam engine, which itself is a progenitor of power generation and so on. This evolutionary path is not unlike the ‘natural selection’ process in evolution and just as inherently difficult to connect an innovation’s origin to its current version, let alone its theoretical future self (Afeyan & Pisano, 2021 & Liu, et al., 2020). Who could possibly have imagined the evolutionary path of the telegraph (Invented 1843, Coe, 2003) would evolve into iPhone and FaceTime?

Despite diffusion’s atypical path, GPTs characteristically retain their essential ‘form’ or distinguishing features and functionality for years and even decades later (Bresnahan & Trajtenberg, 1995). The modern steam power (i.e., the conversion of thermal energy to produce mechanical energy); the grid (i.e., the large-scale generation,

transmission, and distribution of electricity); and the automobile (i.e., internal combustion engine powered vehicle) would be recognized by their respective inventors (i.e., Thomas Savery, Thomas Edison, and Karl Benz). This is true despite advances in the innovation's components (e.g., new materials, electronics, computer technology, and machining technology, etc.).

Could AI/DL retain its form or distinguishing capabilities of pattern recognition, self-learning, and prediction as it diffuses across applications, market sectors and geographies? Again, time will tell.

GPT Sub-types

Bekar, et. al. (2018) believes GPTs cannot be assessed solely on their macro effects, as some fall into (relatively large) niches. He offers the laser as an example. Lasers would qualify as a GPT because they are: (a) seemingly pervasive in everything from health care, music, communications, and the military; (b) improving with time, and (c) enabling other innovations (Bekar, et. al., 2018). Nevertheless, lasers have failed to have the disruptive macro-impact equal to electrification, the automobile, or railroads (Bekar, et. al., 2018).

Lipsey, Carlaw, and Bekar (2005) pointed out the laser's diffusion was helped by being 'best fitted' to the structure of the economy at that time, which is an instance when an innovation requires far fewer 'up-front' or enabling investments to diffuse (i.e., pervade) on a mass scale. To that end, the internet, software engineering, along with the advance of Moore's law in hardware, may have provided a readily accessible 'best fit' for the last two decade's diffusion of AI/DL. Could it be asserted that the current

structure of the economy is a best fit to enable the pervasiveness of AI/DL? This is a question for future research.

For Frischmann (2012), GPTs are less 'singular innovations' like the steam engine or automobile and more 'sub-types' or clusters of innovations like, (a) transportation systems, such as highway systems, railway systems, airline systems, and ports; (b) communication systems, such as telephone networks and postal services; (c) governance systems, such as court systems; and (d) basic public services and facilities, such as schools, sewers, and water systems. Frischmann (2012) believes these GPT sub-types are best framed as 'infrastructure' because:

- GPTs can be used in a way that is 'non-rivalrous' over a large range of demand.
- Where demand is influenced by downstream productive activities that require the resource as an input.
- GPTs can be an input into a wide range of goods and services which may include private goods, public goods, and social goods.

Bekar, et. al (2018) echoes Frischmann's sub-types by asserting GPTs are distinguishable from other innovations by their 'internal nature,' 'external interaction with other technologies,' and 'efficiency of the technology in delivering its services.' Ultimately, Bekar, et al. (2018) believes GPTs are:

> … a single technology, or closely related group of technologies, that is widely used across most of the economy, is technologically dynamic in the sense that it evolves in efficiency and range of use in its own right, and complementary with many downstream sectors where those uses enable a cascade of further inventions and innovations.

This chapter provided the foundations for what is and is not a GPT, why they matter, how they matter, and what should be considered when assessing an innovation as a GPT. We now turn to the essential question of whether AI/DL is a GPT.

CHAPTER 3

ESSAY 1 – IS AI A GENERAL PURPOSE TECHNOLOGY?

AI's conceptual birth goes back centuries with the work of Reverend Thomas Bayes, whose Bayesian Probability provides the reasoning behind statistics, decision sciences theory, operations research, and prediction algorithms (Friston, 2012). However, Applied AI's birth goes back to the 1940s with the 'programming' electro-mechanical computers to perform basic reasoning and, for the time, complex problem-solving (Haigh, 2016). One important milestone was Edmund Berkeley's (1949) *"*Giant Brains: Or Machines That Think" and his view of the future of computer science*:*

> Recently there have been a good deal of news about strange giant machines that can handle information with vast speed and skill… These machines are similar to what a brain would be if it were made of hardware and wire instead of flesh and nerves… A machine can handle information; it can calculate, conclude, and choose; it can perform reasonable operations with information. A machine, therefore, can think.

Yet it was Berkeley's contemporary, mathematician Alan Turing, author of "Computing Machinery and Intelligence (Turing, 1950*,*)*,*" who laid the foundation for AI and computer science. In his landmark paper, the 'Turing Test' where an individual poses random questions to someone (or something) on the other side of a wall, which generates responses back to the questioner. The 'test' is whether the questioner can tell if the 'respondent' is person or machine. According to Turing, if the questioners cannot determine the true source of the responses, then we have AI.

AI advanced further with the Dartmouth Summer Research Project on Artificial Intelligence workshop in 1956 (Haenlein & Kaplan, 2019). Attendees included future

Nobel Laureates and some of the founders of the disciplines of computer science and information theory (Haenlein & Kaplan, 2019). Their goal was ambitious (McCarthy, et. al., 1955, pg. 2):

> …to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

From those early days, AI evolved through cycles of growth and contraction (Menzies, 2003). Early on, computer science pioneer Edward Feigenbaum introduced so-called 'expert systems', which mimicked aspects of human (i.e., expert) decision-making. These and similar clever applications briefly sustained interest in AI. However, this brief AI 'boom' was followed by a long 'bust' of nearly a decade, as users found existing applications limited in their real-world utility by rigid rules and narrow application (Haenlein, et. al., 2019).

The 1980s through the early 2000s brought on a second boom, as AI was applied to complex and dynamic real-world problems (Haenlein, et. al., 2019 & Menzies, 2003). In the last decade of the 20th century John Hopefield and David Rumelhart introduced the concept of 'deep learning,' which allowed computers to find distinctive, useful patterns in massive, disparate structured and unstructured data sets and which also allowed AI applications to substantially diffused (Haenlein, et. al., 2019). Importantly, AI began to reveal its distinctive capacity for continuous learning and complex pattern detection from both structured and unstructured data – an essential capacity enabled by DL.

The algorithmic capability (e.g., 'reinforcement' learning, supervised, and unsupervised learning, and back propagation, etc.) infused much of the AI family of applications and technologies that, in turn, accelerating 'AI technology' diffusion across the breadth and depth of markets, applications, economies, and geographies (Haenlein, et. al., 2019). AI's diffusion was also preceded by the growth in its capacity to process large volumes of data through the development of complementary (i.e., enabling) hardware advances (e.g., increased computing speed and power, lower processing costs, etc.) (Haenlein, et. al., 2019). Simultaneous advances in AI/DL's capability and capacity to both process and use ever larger data sets accelerated its application (and pervasiveness) within and across a growing variety of complex real-world problems within and across different markets, applications, economies, and geographies (Schmidhuber, 2015).

The critical time frame for AI/DL's mass diffusion was 2000 to 2020 (Cockburn et al., 2019; McKendrick, 2022; Argawal, 2019: Trajtenberg, 2019; & Goldfarb, et al, 2021) These two decades are therefore the focus of our data analysis. This is not to suggest the more recent time-period from 2020 to 2024 is unimportant. Rather, the global impact of COVID is only now being understood (as of early 2024) and it would therefore be difficult to account for AI's impact across markets, industries, applications, economies, and geographies in our data analysis (Roman, S., et. al., 2022).

GPTs possess an 'emergent' quality in that some innovations become GPTs if, and when they cross a 'threshold' (e.g., diffusion, 'technology trajectories' and 'technology clusters', etc.) (Korzinov, & Savin, 2016). As stated previously, almost all

innovations owe their existence to one or more earlier innovations, directly and/or indirectly, and this is particularly the case with AI/DL, which can be said to exist within the AI 'family tree' in Figure 5.

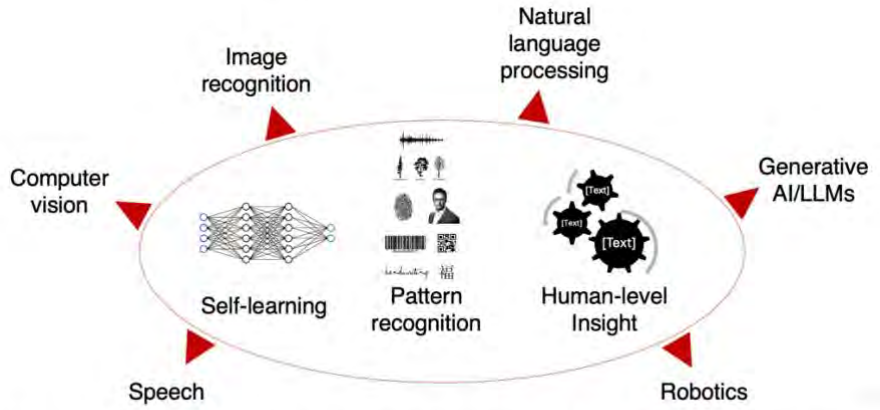

Figure 5. Deep Learning is the enabler of many classic AI technologies

Figure 5 encompasses the disparate techniques, technologies, and disciplines (Klinger et al., 2020 & Wheatley, 2020), most of which overlap and/or intersect with still other techniques, technologies, and disciplines not presented in the Figure, including statistics, linguistics, logic, various mathematics, and neuroscience. As a result, definitional boundaries between what is and is not AI and DL can become blurred (Krafft et al., 2020).

One way to address these fuzzy boundaries is to make distinctions based on the AI family member's 'higher order' capabilities and distinctive features (Krafft et al., 2020; Klinger et al., 2020; & Wheatley, 2020) and might also include AI's 'persistent' and 'capacitive" characteristics (McCarthy, 2007):

> It [AI] is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

To that end, a GPT's 'generalizability' is closely related to pervasiveness put forward in our research. For example, when electrification is compared to, say, nuclear power or fMRI technology the concept of generalizability comes into sharper focus. It is the capacity of the innovation to drive advances in across a disparate array of industries, markets, and applications with very little modification to its basic, distinctive features (Agrawal et al., 2018a/b & Brynjolfsson et al., 2019). As we highlighted earlier, it is DL's capacity to infer patterns from structured and, importantly, unstructured data that makes it 'generalizable' and making mass diffusion possible across markets, industries, applications, geographies, and knowledge domains (LeCun, et al., 2015).

DL's evolution began with first attempts to build AI systems, including using symbolic logic and, later, software systems 'embedded' with human expert-sourced knowledge and heuristics (Wooldridge, 2020). It was only after decades of trial and error that machine learning (ML) and later DL were developed and embedded within software that AI could infer patterns from labelled and, later, unlabeled datasets (LeCun et al., 2015). DL's evolution was furthered helped by newer enabling technologies, such as backpropagation, which linked digital 'neurons' in the DL-derived network to minimize prediction error (Chauvin & Rumelhart, 1995).

AI/DL continues to evolve with recent advances in enabling algorithmic and hardware advances (Wooldridge, 2020 & Hooker, 2020). Hooker (2020), for one, believes AI/DL's diffusion was helped by a 'hardware lottery,' with massive private investments, including primarily venture capital, in enabling technologies. This theme

echoed by Korzinov and Savin's (2016) who emphasize 'technology trajectories' and 'technology clusters' as necessary conditions for the diffusion of GPTs.

AI/DL's diffusion was furthered by advances in other complementary technologies such as computer vision, image processing, video generation, speech recognition/synthesis, language translation, question answering, robotics, and gaming drove (Agrawal et al., 2018a). DL became tightly woven into the fabric of a growing number of everyday products and services like search engines, social networking sites, translation systems, and digital personal assistants (Agrawal et al., 2018a/b).

The consensus among researchers and scholars of innovation as well as technology, and industry experts is AI/DL has moved beyond a promising technology in the late 20th century to practical tool applicable across a growing spectrum of real-world problems. To that end, scholars are modeling AI as a GPT in the hopes of improved prediction and pattern recognition and applied across a range of complex and non-linear contexts (Agrawal et al., 2018b; Crafts & Mills, 2021). We propose 'complete' pervasiveness is synonymous with diffusion across markets, industries, applications, geographies, and knowledge domains could be predictive and/or indicative of an innovation being a possible GPT.

Factors Slowing Diffusion

A logistic 'S' curve is an idealized view of diffusion, but it does not do justice to the true nature of innovation in the real-world. Importantly, innovations rarely go through single and complete adoption cycle but rather adoption cycles beget more branching adoption cycles, creating an evolving 'discontinuous' path of newer branching

diffusion paths (Petrick & Echols, 2004). The 'digital music player' gave birth to an adoption cycle and diffusion path that to the Apple iPod which, in turn, birthed an adoption cycle and leading to the first Apple iPhone. This complex evolutionary path is one reason game-changing innovations are never so easily identifiable, even to those (i.e., executives, technologists, industry analysts, researchers, and venture capitalists) paid to spot and exploit these cycles (Petrick & Echols, 2004).

Despite DL's generalizability, several factors impeded its diffusion. First, DL was constrained by the generalizability of the datasets on which it was trained (D'Amour, et al., 2020; & Marcus, 2018).[4] The result in the case of cybersecurity and autonomous vehicles is that DL was confounded by malicious users and/or unforeseen situations, with sometimes disastrous results (Brundage et al., 2018). DL also requires masses of data and corresponding large-scale computational infrastructure to train and deploy, which is not always readily available and may further slow DL's diffusion (Thompson et al., 2021 & Strubell, et al., 2020). Nevertheless, recent advances in 'Generative AI' and 'Large Language Models' (LLMs), though originating from limited raw data sets, may be capable of accelerating the generalizability of DL applications (Meyer, J.G., *et al.*, 2023).

Another barrier to diffusion is that some datasets are (presently) too large to be 'curated' or filtered to remove, say, personal information and discriminatory and inflammatory content. The result is unforeseen privacy issues and the risk that the

[4] Though this may soon become less of a constraint with recent advances in generative AI. Time will tell.

resulting models will be biased and unfair towards minorities and disadvantaged groups and further slowing DL's diffusion (Raji et al., 2021 & Bender et al., 2021). It is possible continued advances in processing power will mitigate this barrier.

Last, the concept of 'best fit,' discussed previously, can help to explain a slowing of diffusion. For example, Thompson and Spanuth (2021) assert early mainframe computers served as a 'platform' model but fragmented once 'component (i.e., enabling?) technologies' within the computer itself (i.e., processing power) evolved in speed and scale. The result was a rapid evolution of enabling and complementary technologies as the core processors diffused across applications and markets, outstripping, undermining the viability of the mainframe market. For Thompson, & Spanuth (2021), the technical advances inherent in a GPT may slow its diffusion, until and unless other, newer best-fit applications can be found.

What Makes DL Special?

It is important to set DL set apart from other AI technologies. As we discussed, it is its self-learning capacity that makes a level of prediction when, in the case of neural networks, specific inputs are available to the application. Neural networks evolved from earlier work on 'perceptrons' in the 1950s and 1960s, which attempted to show that digital technology could mimic the function of neurons on the human brain by learning as they are fed inputs (Rosenblatt, 1958, 1962). This happens through: (1) combinations of weights and thresholds that translate inputs into outputs, (2) measure the 'closeness' of these outputs to reality, and (3) adjust these weights to narrow the distance between

outputs and reality. The result is an ability to 'perceive' patterns within unstructured data.

Through the 1980s, AI pioneer Jeffrey Hinton (2006) advanced neural network techniques through 'back-propagation,' which created a practical self-reinforcing feedback loop and, consequently, significantly improved the practical application of 'supervised' learning (i.e., using curated data to 'train' an application) (Rumelhart, Hinton, & Williams, 1986). Advances in algorithmic approaches continue to this day, revealing the potential of DL's predictive power from larger and larger data sets (Hinton & Salakhutdinov, 2006).

Back propagation makes possible continuous self-learning from real-world problems. In a sense, DL's learning capacity results in virtuous spiral of 'learning-by-doing,' which further sets it apart from traditional 'static' software applications. This gives AI/DL a self-sustaining and self-reinforcing capability. Instead of requiring, say, a developer to code each process step to be automated, ML algorithms 'discover' their own function(s) and connect set(s) of inputs to set(s) of output(s), assuming large enough datasets exist for input/output mapping (Brynjolfsson & Mitchell 2017).

As with the other 'AI family' of technologies, neural networks are enabled by DL, as Figure 5 reflects and assume the current trend of greater reliance on DL across many AI applications to continue.

DL Informs Human Decisions

Prediction is at the heart of AI/DL's real-world relevance (Agrawal, et. al. 2018a; Lindsay & Serra, 2022) and, in the 2nd decade of the 21st century, it is oriented toward

informing human predictions and decision-making (Agrawal, et. al. 2018a; & Lindsay & Serra, 2022). Examples abound, from finding (i.e., predicting) the shortest route to the nearest Starbucks to safely and autonomous robots moving (i.e., predicting) the best path through a busy warehouse.

The founders of AI sought to augment or even replace human intelligence with machine intelligence, not to create the dystopian world of The Terminator but to use AI to move beyond the highest levels of human intelligence and decision-making to the benefit the world (Kurzweil, 2012*).* Though that goal is yet to be achieved, progress in the application of AI to real-world problems has been significant and is accelerating at an accelerating rate, particularly in the first two decades of the 21st century (Liu et al., 2018 & Boucher, 2020).

Our literature review revealed the first two decades of the 20th century are the time when AI/DL appears to be undergoing a mass adoption across and within markets, industries, economies, and geographies. Further as we will see, DL's unique capabilities appear to be accelerating AI's adoption.

Challenges With AI/DL as a GPT

Some academic researchers believe AI is a possible GPT, even though it has not yet reached mass diffusion despite seventy years of history (Agrawal et al., 2019; Cockburn et al., 2019; & Brynjolfsson et al., 2021). Goldfarb (2020) puts the challenge of determining if AI/DL is a possible GPT succinctly:

> …resolving this speculation using existing empirical methods for identifying GPTs requires many years of hindsight. By that point, the managerial relevance of knowing whether ML, blockchain, cloud

computing, and other technologies are GPTs is limited.

We saw how enabling technologies make diffusion possible. To that end, some researchers assert investment in AI/DL's complementary and enabling technologies is well underway and, consequently, so is AI/DL's mass diffusion (Brynjolfsson et al., 2018). Yet, Brynjolfsson et al., (2018) believes that a transition to an AI-driven economy will require more and substantial investments in complementary, intangible assets including data, skills, and organizational capital at least 2.7 to 4.1 times bigger than the investments in tangible assets themselves.

Every innovation comes with explicit (i.e., cash) and implicit (i.e., time, effort, opportunity, etc.) adoption costs and AI/DL is no different (Brynjolfsson et al., 2018). Recently, the recognition of these costs, from anecdotal reports, is motivating a growing 'rethink' of many long-established organizational practices (Brynjolfsson, Hitt, & Yang, 2002). For example, DL enabled AI technologies are only now emerging and, consequently, so is an accounting of the costs of this diffusion (Teece, 2018 & Rathje & Katila, 2021). A tally of the total implicit and explicit costs AI/DL's adoption awaits the passage of time and great deal of effort, but they are expected to be massive (Jovanovic & Rousseau, 2005; Greenwood & Yorukoglu, 1997; Aral, Brynjolfsson, & Wu 2012; & Tambe, Hitt, & Brynjolfsson, 2012)

The experience of electrification is telling, as half of US manufacturing establishments remained unelectrified until 1919, about thirty years after the shift to polyphase alternating current began (David, 1991). One reason for this is Henderson's (1993, 2006) belief that manufacturing plants, at the advent of electrification, were in the

thrall of the 'curse of knowledge' or constrained by the designs of the time. Even today, the costs of adoption of electrification across the United States and Europe during this first adoption phase has never been calculated.

The different ways researchers have characterized GPTs calls for a consistent method for characterizing GPTs (Moser & Nicholas, 2004; Hall & Trajtenberg, 2006; Feldman & Yoon, 2012; Graham & Iacopetta, 2014; & Petralia, 2020). Our research uses disparate and overlapping data sets from the first two decades of the 20th century to support our proposition that AI/DL is a possible GPT. We hope to enhance the consistency and reliability with which GPTs can be characterized.

The lack of a common framework has the following consequences: (1) GPT researchers are less consistent in their judgements, (2) various data sources remain unexplored before more useful tools can be found, and (3) GPT theory must be more forward-looking to be useful decisionmakers. Therefore, our research method must address these three concerns to be useful and turn to that in the next chapter.

Levels of Analysis and Identifying a GPT

To understand AI/DL as GPT we must understand what is its 'essential' defining character as a GPT. The case of the railroad is illustrative (Donaldson & Hornbeck, 2016). The railroad brought about dramatic changes in the US, European and, ultimately, global society and economy as masses of people and goods could now be moved far faster, far greater distances, at a far lower cost, and on an unprecedently massive scale. In addition, the second and third order changes from the railroad,

including changes in population distribution, the built environment, the military, and communications were similarly profound.

The defining characteristic that made the locomotive possible was the ‘portable’ motive force of the steam engine. The other components comprising early and contemporary railroads, including rails, carriage, and other features existed before or well before the invention of the railroad (Donaldson & Hornbeck, 2016). That same motive force (i.e., steam engine) also transformed sea travel, manufacturing, and a host of other industries, as we discussed previously. In essence, the steam engine is the GPT that enabled another great innovation (i.e., the railroad). In the same vein, DL is the essential ‘motive force,’ and therefore GPT, within the larger innovation of AI.

This foregoing provides a clear example of the need for multiple levels of analysis – from broad to narrow – associated with all GPTs. In the above case, the broad theoretical concept of 'the physics of steam power' is the knowledge base, the 'steam engine' is the (general purpose) technology, and the 'railroad' is just one application. The chronological unfolding seen in the historical record follows these levels, proceeding from broad to narrow. Similarly, for this dissertation, and the case of AI, the broad theoretical concept of Large Language Models (i.e., LLMs) are the knowledge base, the DL is the GPT, and, for example, ChatGPT is one application.

CHAPTER 4

ESSAY 2 – THE IMPORTANCE OF PERVASIVENESS

GPT theory is focused on prior innovations and their historical and current macroeconomic impact. This work has contributed to our understanding of technology and its impact on our world. Can GPT theory help build a forward-looking tool capable of assessing *current* innovations and their prospective macro-impact? At a minimum, such a tool would give decision-makers a range of outcomes and inform the creation, development, management, and regulation of potentially game-changing technologies (Agrawal, Gans, & Goldfarb, 2019)? This would be an exercise in technology forecasting, which is not a new idea, as this chapter shows. Our research hopes to contribute to such a tool.

Predicting an innovation's path is very, very difficult, yet looking to what *is on the horizon* is the job of every executive, venture capitalist, research scientist, policymaker, and entrepreneur (Davenport, 2022 & Buchannan, 2013). And nowhere is prediction more difficult than in new products and services (Weller, 2015), as ninety percent of roughly thirty thousand new products introduced every year fail (Emmer, 2018). Without the ability to anticipate what may be coming, economic, and societal advance would be nearly impossible (Davenport, 2022). Anticipating what *might come next* has been and remains the focus of technology forecasting.

Trying to Forecast Innovation

Management theorist John Brian Quinn (1967) framed technology forecasting's main purpose*:*

> [technology forecasts] …do not necessarily need to predict the precise form technology will take in a given application at some specific future date. Like any other forecasts, their purpose is simply to help evaluate the probability and significance of various possible future developments so that managers can make better decisions.

Our research seeks to know what forecasting tools, systems, and techniques have and have not proved useful and in different contexts. We did not seek, nor did we find, a 'crystal ball.' We turn to Peter Drucker (1969) again, who reminds us that doing so would be a waste of time:

> …the impacts of technology are almost never capable of being anticipated very far in advance. And the "expert's" predictions are no better than anyone else's, and often worse. In the early fifties the best experts predicted … a total worldwide market of no more than a hundred computers altogether but also "automation unemployment" of six, if not of ten, million people by 1970. Their projections were completely valid, their samples scientific, and their techniques rigorously quantitative. It was only their premises that were wrong.

We examined a wide variety of forecasting techniques to identify those that could contribute to our research. Fortunately, the discipline of technology forecasting has a rich history going back a century with President Franklin Roosevelt's establishment of the National Resources Committee, whose mission was to look forward and assess the possible impacts of advances in science technology on employment in the United States (Schnaars, 1989 & National Resources Committee, 1937).

The outbreak of World War II drove the need to exploit science and technology for the war effort, culminating in the seminal work of Vannevar Bush, science adviser to

President Roosevelt, of *Science: The Endless Frontier* (Bush, 1990). In it, Bush presciently spoke of the need to anticipate the growing societal, economic, and cultural impact of science and technology. Later, the Cold War brought further increases in government funded science and engineering R&D (Roussel, 1991), which drove the evolution of tools like systems analysis and the Delphi method, both still in use today (Murdock, 1974 & Gordon & Helmer, 1964).

The Cold War also birthed public, private, and non-profit organizations dedicated to supporting the needs of military, policymakers, private investors, researchers, and industry executives to anticipate advances in science and engineering. Organizations like Rand Corporation, MITRE, The Institute for Defense Analysis, The White House Office of Technology Assessment, DARPA, and the Office of Technology Assessment were created for that purpose (Jantsch, 1967).

Despite the spread of these tools, disillusionment grew as their relevance to real-world problems like the Vietnam War and the United States energy crisis of the 1970s began to be questioned (Coates, 1994). According to Coates (1994):

> Technology forecasting was reduced in practice to a set of tools and methods; forecasts produced between 1975 and the early 1990s were relatively few, generally poorly defined, and executed without much attention to formal assumptions, time horizons, or limitations.

Nevertheless, technology forecasting has contributed much, including to the evolution of business strategy and strategic management. So much so it is now taken for granted that *every* organization must anticipate abandoning maturing technologies and invest in newer ones to stay competitive (Christensen, 2013 & Huang, 2014). Not surprisingly, forecasting is a global industry today (Feng, et. al. 2022). Contemporary

management theory rests on an appreciation for and an assumption of a growing collective economic and structural interconnectedness in combination with a growing dependence on technology (Esmaelian, et al, 2017 & Mudambi & Ricketts, 1998).

What Can We Know About the Future?

Technology foresight is a variant of forecasting that, according to Fergnani (2020, p. 7) helps decisionmakers, "…interpret changes in the environment, outline and evaluate possible futures based on these changes, and use information from these futures for the organization's competitive advantage." Martin (1995) adds nuance:

> [technology foresight] …is the systematic study of scientific, technological, economic, and social developments in the longer term. Its goal is to identify strategic research areas and select the technology clusters that will make the greatest contributions to economic and social interests.

Firat, Madnick, & Woon (2008, p.7) provide additional nuance:

> [technology foresight] …applies to all purposeful and systematic attempts to anticipate and understand the potential direction, rate, characteristics, and effects of technological change, especially invention, innovation, adoption, and use.

Forecasting cannot yet and likely never will consistently yield precise predictions of an innovation's path, but it has provided useful alternative evolutionary paths (Day & Shoemaker, 2005 & Porter, A. L. et al. 2004). Precise and accurate prediction is likely never possible because of the nearly infinite contingent contextual variables along the way from idea to prototype to mass diffusion (Bresnahan & Trajtenberg, 1995; Lipsey, Bekar, & Carlaw, 1998; & Jovanovic & Rousseau, 2005).

And it is 'contingent complexity' that animates forecasting skeptics as they point to the fact that empirical evidence to date shows that the best forecasts have limited

accuracy in the middle and long-term, poor repeatability, and limited adaptability (Kucharavy & De Guio, 2005). Philip Tetlock, Professor at the Wharton School, for one, agrees with that critique in his popular book Super Forecasters: The Art and Science of Prediction (Tetlock & Gardner, 2015).

In this fifteen-year study of expert prediction, he assessed the accuracy of 28,000 economic and political predictions of 284 financial executives, journalists, economists, and policymakers. An expert, which he dubbed a 'Hedgehog,' possesses deep knowledge of a specific domain (e.g., economics, financial markets, international relations, politics, etc.). Yet despite the depth of knowledge, Hedgehog, or expert predictions were no better than random chance, even when their predictions were well within the expert's field.

Tetlock also ran tests on the predictions of random non-experts or 'Foxes, which are people who know a 'little bit about a lot of things.' He found 'Fox' or non-expert predictions were more accurate and, in some cases, far more accurate than the experts. Importantly, the non-expert or Fox predictions depended not on their tools or biases but on how they thought about the prediction process and their assumptions therein.

Non-experts consciously included an array of seemingly unrelated considerations and data, according to Tetlock, greatly enhanced the quality of their predictions. Though not explicitly about forecasting technology, Tetlock's work points to the need to include a sufficient breadth of disparate data.

Predictions about the process of innovation is not to be confused with prediction in the natural sciences, of which the latter relies on established theory, data, and

experiment. In contrast, technology forecasting – like predictions across the social sciences – must account for a large degree of messy real-world complexity which is not inherent in the natural sciences (Buchanan, 2013). Too often, according to Klosterman (2013),"…technology forecasts (too often) rely on information that is out of date, estimated, or collected at the wrong level of aggregation."

One way to reduce this 'messiness' is to extrapolate from a single variable, yet this approach is constrained by the data's timeframe, along with a dependence on the context of the variable of interest (Wissema, 1982). Further, though predictors strive for accuracy it often must come at the expense of decision-informing efficacy (Klosterman, 2013*)*. Likewise, other tools may possess greater efficacy (depending on context), such as the case with substitution analysis, the Delphi method and complex multivariable modeling (discussed later), but even these tools demand rare expert judgement, vast amounts of data, and are heavily context dependent (Wissema, 1982). Alas, there is no free lunch with prediction.

We return to Peter Drucker (1969), who commented on the complexity of predicting a technology's possible impacts which:

> …tend to be synergetic and the result of several developments, each independent in its origins and the outgrowth of a separate discipline with its own, …are often quite indirect, and by-products rather than main products &…tend to be threshold phenomena; until they reach a certain level, there is no impact at all.

Highlighting technology forecasting's limitations does not, however, undermine its utility. It orients our expectations about our analysis and, ultimately, our conclusions, however tentative.

Understanding Forecasting Tools

We live in an era of growing mountains of data. Yet for some innovation scholars, the world may already be awash in enough data to enhance our predictions of tomorrow's innovations (Li, et. al., 2019). If true, could AI/DL one day be used to predict its future self? Though an interesting question, we return to understanding the strengths and limitations of forecasting tools to identify those most relevant to our data analysis (Daim, et al., 2006 & Magruk, 2011).

Forecasting tools vary widely in their approach and uses of data (Coates et al., 2001). All, however, rely heavily on prior and current data to inform their assumptions about the future (Klosterman, 2013 & Ascher, 1981). Therefore, high quality forecasting depends on the accurate and precise tracing of an innovation's previous and recent history to generate useful alternative future paths of the innovation (Klosterman, 2013). These 'possible futures' help decisionmakers better anticipate and plan for what could come.

Slocum & Lundberg (2001) provide a review of forecasting tools, together with our comments on the potential relevance for our data analysis and forward-looking tool:

- Intuitive Forecasts: An 'ask the expert'' approach akin to consulting Tetlock's Hedgehogs. Not surprisingly, this approach is readily undermined by inevitable unexpected occurrences, invalid not-so-expert 'expert' opinion, and latent forecaster bias (Harteis & Gruber, 2008).

- Consensus View: These methods are characterized as 'wisdom of the crowds' (Surowiecki, 2005) approach to prediction. Examples include so-called 'prediction markets' that used voting or consensus forecasts to determine the probability of certain future outcomes (Rhode & Strumpf 2004; Wolfers & Zitzewitz 2004; Snowberg, Wolfers, & Zitzewitz 2005; & Leigh & Wolfers 2007).

- Delphi Method: Convening expert panels to arrive at a consensus forecast while controlling for the qualifications of the experts and their biases. However, there are recurring questions around human bias and the tendency for 'group think' impacting the efficacy of this technique for technology forecasting (Woundenberg, 1991).

- Statistical Modeling and Simulation: Relying on empirical observations of the phenomenon being predicted while identifying patterns of association between the variables accounting for that phenomenon. Patent citations and linking them to technological road mapping and forecasting is a common forecasting approach (Érdi, et al., 2013; Fleming, 2001 & Seung-Wook, et al., 2014).

- Application of a Causal Model: This method is like statistical modeling, while also identifying candidate 'causative agent(s)' of the predicted phenomenon. Econometric modeling is a common application of this approach (Érdi, et al., 2013).

- Analogy Method: Analogies allow for a systematic comparison of the forecasted technology to an earlier technology that appears similar in key respects. However, this approach is not without its limitations that, according to Martino (1983) and echoing Drucker's (1969) concerns, "…are based on the assumption that there is a 'normal' way for people to behave and that given similar situations, they will act in similar ways. However, there is no guarantee that people today will act as people did in the model situation. Hence the forecast is at most probable, never certain".

- Structural Modeling: Again, econometric modeling is one common example of structural modeling, which is the application of mathematical or analytical models to technological innovation. It highlights key elements of the evolution of an innovation, identify its functional aspects, and express them symbolically or mathematically, while also possessing some of the limitations of the other approaches describe here (Slocum & Lundberg, 2001).

- Scenarios: Scenarios are alternative narratives built on differing, well-defined assumptions about the future. Each narrative tests the implications of each of the key assumptions within that narrative. Technology 'roadmaps' are an example and can be run for generations (e.g., 2 to 10 years) (Kostoff & Schaller, 1999). Yet a recent McKinsey study of industry road mapping found it failed to, "identify 'white space' in technology development and provide a holistic view on how technologies and markets would evolve and what these changes suggest about future technologies (Heim, et al., 2017)."

- Extrapolation: This method extends an existing trend, often with a 'cone' or range of probable paths. Graphically, this often appears as a logistic or ''S'' curve. The main critique is the inductive assumption 'tomorrow will just be a little bit different from today, give or take' (Sahal, 1977 & Sahal & Yee, 1975). More sophisticated versions feature the progression of a given innovation along the hypothetical S curve super-imposed on bell-shaped life-cycle curve, which highlight its evolution through adoption, growth, maturity, and decline (Coates, et. al, 2001 & Firat, et. al., 2008).

Extrapolation assumes 'path dependence' (Mahoney & Schensul, 2006) or the idea that an innovation's previous 'path' through to the present constrains its possible alternative further paths. Path dependence underlies predictions in: (1) weather forecasting & search traffic (Jun et al., 2017), (2) R&D spending (Taegu, Kim, & Lee, 2014), (3) technology or product life cycle evolution (Park et al., 2015), (4) averaged analogy method based on parameters of multiple products, (5) technology road mapping, and (6) bibliometrics (Daim, et. al., 2006). Therefore, extrapolation, augmented by data analysis, appears to be the most appropriate method for our analysis, even though a most scholars assert there may not yet be enough historical data of AI/DL available to plot its disruptive path (Goldfarb, et.al., 2021 & Christensen, 2000).

Interestingly, AI/DL's limited history as a practical application may also limit its path dependence. A shorter history may translate into a greater range of alternative futures. Is there an inherent path dependence in AI/DL's diffusion to this point in time. Yet again, only time will tell. We are not making precise predictions. In essence, we propose a given technological innovation (i.e., AI/DL) fits into a particular class of innovations (i.e., GPT). We now turn to understanding and defining what variable we are extrapolating: the GPT characteristic of pervasiveness and complete pervasiveness.

What Are We Looking For?

For decades, scholars have tried to distinguish GPTs from other innovations but with mixed results. Nevertheless, a consensus appears to have formed around the following three distinct characteristics of GPTs (Bresnahan & Trajtenberg, 1995; Lipsey, Bekar, & Carlaw, 1998; Jovanovic & Rousseau, 2005 & Ristuccia & Solomou, 2014):

(1) They are 'pervasive' or diffuse across the breadth and depth of markets, industries, economies, applications, and geographies.

(2) They show 'improvement' or increase in efficiency, effectiveness, capability, and capacity with use over time.

(3) They 'spawn' or stimulate other innovations, some of which may become other GPTs.

These characteristics make GPTs distinctively the 'drivers of seismic macroeconomic shifts' (Strohmaier & Rainer, 2016). We assert that pervasiveness is the most definitive of the GPT characteristics. After all, many innovations could meet the criteria of 'spawning innovations' and 'improving over timing,' but not pervasiveness. An example is functional magnetic resonance imaging (fMRI), which measures small changes in blood flow that come with brain activity.

Figure 6. Common GPT characteristics & pervasiveness

Figure 6 presents a comparison of fMRI with oft-cited GPTs along with the three GPT characteristics. fMRI has saved and/or improved countless lives and spawned innovation in healthcare since its introduction in 1984, while also continuously improving its own performance (Raichle, 2009). Nevertheless, fMRI is a healthcare application and thus lacks the quality of complete pervasiveness of the GPTs cited in Figure 6, which we define next.

What is Complete Pervasiveness?

The word pervasiveness is used differently in our research than how it is sometimes used in discussions of technology (Kourouthanassis, et. al., 2008; Efanov, & Roschin, 2018). For example, the National Institutes of Standards and Technology (NIST), in fact, has a program called The Pervasive Information Technologies Program that "facilitates the creation of standards for sensor communication, networking interoperability, and sensor information security..." and defines the term as "the trend towards increasingly ubiquitous connected computing sensors, devices, and networks that monitor and respond transparently to human needs." (Marbukh, et. al., 2010).

In other instances, 'pervasiveness' refers to the insidious quality of technology to (Mitroff, 2019):

> [make] …our lives incomparably better in numerous ways, now constitutes one of our biggest threats. It threatens not only our physical, but our mental well-being It forces us to confront some of the most important questions facing the human condition: What is the extent and the depth to which we are willing to allow technology to intrude into our lives? In other words, how pervasive and invasive are we willing to let technology be?

Complete pervasiveness, however, is defined in our research as the extent to which AI/DL is tangibly present within and across all markets, industries, economies,

and/or geographies. Electrification is an exemplar, as it has changed everything from manufacturing, transportation, agriculture, and communications and is considered one of the greatest engineering feats of all time (Holden, 2000 & Gordon, 2000).

In manufacturing alone, it generated massive TFP gains across the United States and Europe from the late 19th century to the mid 20th century, as factories were redesigned to accommodate modular electric motor driven manufacturing equipment (David, 1991 & Gordon 2000). From its birth in the US, the 'grid' developed by Thomas Edison in New York in 1882, electrification's impact extends to this day.

Pervasiveness and AI/DL

Anecdotal data suggests AI/DL is diffusing broadly and deeply, particularly over the last twenty years. However, prior empirical work presents a mixed picture. For example, some analysts predicted forty seven percent of current jobs in the US would be obsolete by as soon as the end of 2030 (Frey & Osborne, 2017), while others predicted many occupations are or will be replaced by AI (Pethokoukis, 2019). Still another study of the maturity levels of AI deployments across half a million businesses in the US found less than 1,500 were making substantial investments in AI (Naimat, 2016).

AI/DL is the product of massive, initial resource commitments made over decades in the form public and private R&D, public, corporate, as well as entrepreneurial trial and error (Trajtenberg, 2019; Argawal, 2019; & Goldfarb, et. al, 2021). Consequently, there should be tangible indicators of resource commitments of AI/DL's pervasiveness, within and across markets, industries, economies, and geographies, through the crucial period of mass diffusion from 2000 to 2020.

Our central proposition, therefore, holds that complete pervasiveness is true when all the following criteria are met:

1. Accelerating cumulative growth of AI/DI across different markets and industries.
2. Accelerating cumulative growth of AI/DI across disparate, novel applications.
3. Accelerating cumulative growth of AI/DI knowledge, specifically within and across academic research.
4. Accelerating cumulative growth of national public policy initiatives in support of AI/DI.
5. Accelerating cumulative growth of geographic dispersion AI/DI globally.

A challenge in multi-disciplinary research is the consistency in the application of terminology. For example, in GPT theory the concept of 'generalizability' matches our definition of complete pervasiveness closely. The term has been used by innovation scholars to describe the versatility of a given innovation to be applied across applications, markets, and industries (Agrawal et al., 2018 & Brynjolfsson et al., 2019). Nuclear power held the promise of 'free' electricity but failed to have the impact across industries, markets, and geographies like mass transportation, communications, healthcare, and agriculture as originally predicted (Rhodes, 2018) because, in part, of its lack of 'generality' (Agrawal et al., 2018b & Brynjolfsson et al., 2019).

Another term is 'knowledge density' or the idea that the "structure of our knowledge should have a sufficient density for a GPT to become pervasive, whereby, structure we mean interconnectedness of our ideas and by density interchangeability of our knowledge among various applications *(*Korzinov & Savin, 2016)." An extension of

knowledge density is Brynjolfsson's (2019) notion of 'connective tissue' or the idea that investments in technologies and skills are a necessary precondition of diffusion.

As previously highlighted, diffusion is wave-like and each successive wave brings second and third order economic, technological, and societal effects (Geroski, 2000). Yet, each wave can take years or even decades to be observable (Bresnahan & Trajtenberg, 1995 & Jovanovic & Rousseau, 2005). This wave-like characteristic gives diffusion its distinctive logistic or 'S' curve, as in Figure 7, which is the cumulative result of many individuals, groups, and organizations communicating, considering, and adopting a 'new thing' (Bresnahan & Trajtenberg, 1995 & Geroski, 2000). Thought the possible patterns of diffusion of a GPT, are like all other innovations, they are *very* different in the depth and breadth of the innovation's 'reach' across market, industries, economies, and geographies.

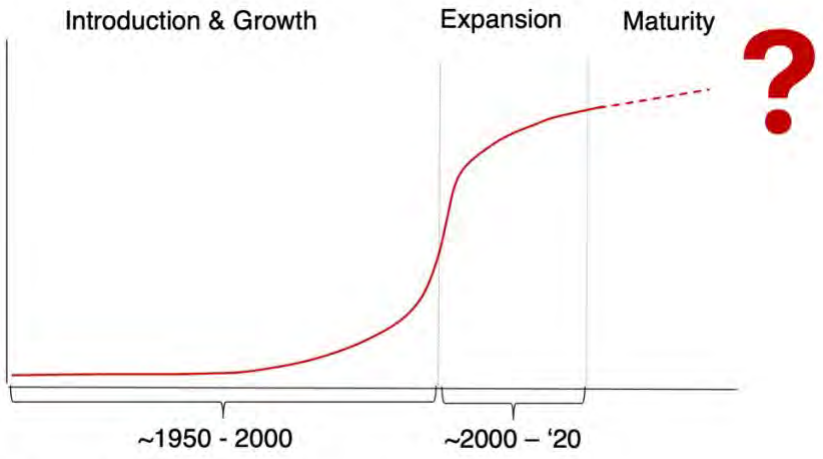

Figure 7. AI's diffusion curve

Different innovations impact us differently. Sliced bread was exceptional when it was invented 1928 (Latson, 2015) but its impact was not comparable to the automobile

or electric grid of that day through to today. As we have seen, some scholars remain skeptical of single innovations having macro-scale impact (David, 1991). Recalling the GPT skeptics, diffusion is a vast, complex web of interconnections and interacting variables, making definitive cause and effect connections nearly impossible (David, 1991).

For GPT skeptics, the scale of impact is irrelevant. For them, the only difference between, say, sliced bread and the automobile is the latter differs only in the number and complexity of interactions between and among variables and not from a singular force. Yet, the weight of reviewed research supports the view that GPTs do have a discernable, disproportionate, and direct macro-impact within and across industries, markets, applications, economies, and geographies (Jovanovic & Rosseau, 2005). For GPTs, quantity (i.e., size of macro-impact) has a quality (i.e., defining characteristic) all its own.

Bresnahan and Trajtenberg (1995) introduced another helpful term, 'complementarity,' which was later expanded on by Petralia (2020). That is, GPTs are enabled by 'upstream' innovations (e.g., machine learning, reinforcement learning, deep learning, etc. for AI/DL) while simultaneously enabling 'downstream' innovations (e.g., cloud computing, 'big data' technologies, etc. for AI/DL) (Brynjolfsson, et al., 2017). This concept assumes AI/DL possesses underlying 'structures' very similar to other information technologies (Brynjolfsson, et al., 2017).

Complete Pervasiveness

It could be said that pervasiveness is the latent potential existing within every innovation, manifesting only if, and when some 'diffusion threshold' is reached. Such a threshold can exist within a single 'dimension' (e.g., market, industry, economy, or geography) or cumulatively across multiple or, in the case of complete pervasiveness, across all dimensions. For example, it might be said that pervasiveness in the market dimension is achieved when a certain level of investment and/or intellectual capital exists (Strebulaev & Gornall, 2015). We would also expect to see an increase in the cumulative growth and acceleration in the rate of growth of investment and patents as a prelude to the achievement of pervasiveness for, say, the product and services dimension. But growth across one or two dimensions is not sufficient to qualify as complete pervasiveness and, therefore, a GPT. It must show across all dimensions.

Pervasiveness overlaps with the concept of 'saturation' from economics, which is the point at which the rate of diffusion of a product or service plateaus (Popp, 2002). Saturation occurs deep in the life cycle of a product's diffusion, when diffusion is at the top of the (diffusion) S curve, discussed earlier. Goldfarb. et. al, (2020) asserts the wide use, ongoing improvement, and follow-on innovations are difficult to measure without the benefit of hindsight, which is why most GPT studies use data (e.g., evidence of widespread adoption, whether patents, patent citations and/or some evidence of productivity impact, etc.) a decade or more after their introduction.

Pervasiveness within a single dimension cannot alone predict a GPT, as we have seen many factors can slow or accelerate diffusion. As innovation diffuses across more

and more dimensions toward complete pervasiveness then there is a strong possibility that innovation is a GPT. These 'pre hoc' dimensions are the foundations for a forward-looking tool. Our central proposition is therefore that 'complete' pervasiveness is the instance when all five dimensions, as indicated in Figure 8, are present for an innovation.

Figure 8. The dimensions of complete pervasiveness

Though some scholars believe GPTs sit on a 'innovation spectrum' and not on threshold of adoption (Teece, 2018; Gambardella et al., 2021; Rathje & Katila, 2021; & Bojovic, 2022), our proposition is that complete pervasiveness is a very strong indicator that an innovation is a GPT. Complete pervasiveness can also be thought of as a 'phase transition' of innovation from one state (e.g., interesting innovation) to another state (e.g., potential macro-impact innovation or possible GPT) as in Figure 9.

A phase transition is a concept borrowed from thermodynamics that describes a physical process where matter (e.g., solid ice) transforms into another form of matter (e.g., liquid water) or gas (e.g., gas steam). Such a phase transition in our research is a

distinct ‘changeover’ from an innovation being ‘normal and novel’ (i.e., pervasive) to a ‘supra-normal’ with macro-impact potential (i.e. complete pervasiveness).

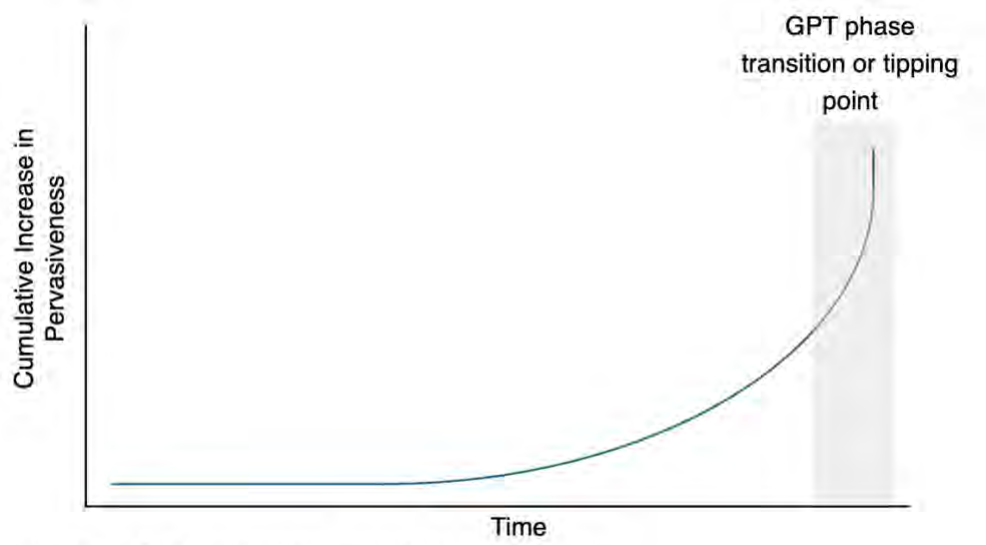

Figure 9. A characteristic growth in pervasiveness

The idea of a GPT is a phase transition from an innovation to a supra-innovation can be applied to everything from the first Model Ts was offered for sale (1908[5]) to the first instance of the electric grid (1882[6]). It is possible that AI/DL may have entered its own ‘phase transition’ the latter part of the second decade of the 21st century. The next chapter describes the methodology supporting that proposition.

Research Goals and Methods

Our research *goal* of a forward-looking tool is comprised of three research objectives. The first is to deepen our understanding of the innovation process generally and GPTs specifically. As our literature review revealed, innovation is not an ‘event’ but

[5]From the Ford Motor Company website (www.ford.com) (2024).
[6]IEEE.

a complex, largely social, process of diffusion from concept to invention to application to adoption (McLaughlin & McLaughlin, 2021).

The second objective is to assess AI/DL as a GPT. AI/DL became our focus, in part, because it satisfied our third goal: identifying a recent innovation that appears to be an economic, technological, and social phenomenon. The literature review drew from an array of research disciplines and across academic, industrial, and public research from 1950 to the present.

At some level, our research seeks a more applied definition of GPTs. Scholars and decisionmakers have tried to do so, as they explored interconnected and overlapping networks of technologies that shared 'underlying principles and mutual dependencies' (Petralia, 2020); assessed patent growth rates that capture the 'elaboration of technologies like GPTs' (Moser & Nicholas, 2004; Jovanovic & Rousseau, 2005; & Petralia, 2020) and assessed patent growth rates for 'unbiased generality measures' (Pavitt, 1988; Griliches, 1990; & Hall & Trajtenberg, 2006).

In fact, patent research has been especially fruitful, but challenging as a single source of insight (George & Walsh, 2022):

> …if [the] AI become(s) more knowledgeable and skilled than all people in a field, it is unclear how a human patent examiner could assess whether an AI's invention was obvious. An AI system built to review all information published about an area of technology before it invents would possess a much larger body of knowledge than any human could. Assessed against all knowledge, almost everything would seem obvious…

Other quantitative and qualitative approaches to defining GPTs have included:

- Using the Herfindahl–Hirschman (HHI) index to calculate a 'technical commonality index' (Moser & Nicholas, 2004).

- Qualitative measurement of 'technology complementarity' and 'technical applicability' (Feldman & Yoon, 2011).
- Qualitative measurement of 'universality,' 'innovation,' and a 'wide range of influences' (Youtie, et. al., 2008).
- Qualitative measurement of various firm-level measures (Clarke, et. al., 2015).

All these methods helped guide for our data analysis as have industry reports (Chui, 2023 & Lynch, 2023). Our review of the available data revealed several useful sources, and these are discussed in the next section.

The Oxford Dictionary defines pervasiveness, "…as the fact of being present and noticeable in every part of a thing or place." We sought data reflecting pervasiveness of AI and DL (together and separately) for the key period of 2000 to 2020 for each of five dimensions. We define pervasiveness as a rapidly growing identity or tangible presence within a given dimension and complete pervasiveness as the same across all five dimensions.

Our pervasiveness dimensions are the following:

- Products and Services: AI/DL-driven new startups and corporate investments,
- Novel Applications: AI/DL-focused patents,
- Knowledge Transfer*:* AI/DL-themed research papers and conferences,
- Public Policy Initiatives*:* AI/DL-focused policies, and
- Geographic Dispersion: The previous indicators globally dispersed.

For each dimension we found public and/or proprietary data sources that supplied multiple empirical indicators for each pervasiveness dimension from 2000 to

2020. We assessed the progression for each data source for each of the five dimensions seeking to identify a characteristic growth curve comparable to Figure 9.

Our overall research effort is propositional, not hypothetical. That is, we're not offering testable hypotheses. Our propositions are intended to stimulate further research. Propositional research is an approach articulated by Alfred North Whitehead (1929/1978) who held propositional research was a vital, speculative, creative act intended to motivate further questions, questioning and, ideally, more research. We propose a tool out of which future research can identify testable hypotheses.

We were unable to obtain data for every year from 2000 to 2020 across all five pervasiveness dimensions. In some cases, we included data beyond 2020. In other cases, sources commenced in 2010 (or later) while others were available for the entire period. Nevertheless, this was not thought to be a problem, as some pervasiveness dimensions did not manifest significantly until the beginning of the second decade of the 21st century.

Our research is not immune to definitional ambiguity and is to be expected with over-used terms like AI. It is the reason we used varied and overlapping data sets for each of the five pervasiveness dimensions. Additional data sets beyond those used in this research, admittedly, may have brought further clarity to our analysis. However, some of the data sets examined reviewed but not utilized in our research, while interesting, were either not available to the researcher because of cost considerations or were thought to be of marginal informational value. Future research may warrant consideration and analysis of these additional data.

Data Analysis

Our central proposition holds complete pervasiveness, or substantial increases in measures across all our proposed pervasiveness dimensions is strongly indicative of a GPT. On the other hand, pervasiveness, or substantial increases across one or a few dimensions is partially suggestive of innovation that may be on a path to GPT status. Ultimately, complete pervasiveness is most likely with a rapid absolute growth and rate of growth across all five pervasiveness dimensions.

We do not claim any pervasiveness dimension is of greater or lesser importance than any other. Rather, in the instance that fewer than all five dimensions reflect pervasiveness then it possible to make one of more of the following inferences:

- Our definition of pervasiveness does not capture AI/DL's actual or potential impact: We consistently attempted, throughout the literature review and in the crafting of this research, to define terms and apply them in a consistent way.
- The data underlying one or more pervasiveness dimensions do not reflect AI/DL's our definition of pervasiveness: Each data source was carefully vetted. However, this was not original data collected by the researcher and some inconsistencies may therefore be inherent in the data. That is the primary reason multiple data sources were used for each dimension.
- The data underlying one or more pervasiveness dimensions partially reflects our definition of AI/DL's pervasiveness: Again, each source was carefully vetted and since the research had to make do with largely public sources some precision may have been unintentionally sacrificed along the way.

Measuring Pervasiveness Dimensions

As stressed previously, multiple data sources were used to measure each dimension. The data sources including the following:

- Products and Services: We used the founding date of AI/DL-driven new startups and announced corporate investments in AI/DL. Our data sources for these were Crunchbase, The AI Index 2023 Annual Report, and CB Insights.

- Novel Applications: We used the number of new AI/DL patent filings from the World Intellectual Property Organization.

- Knowledge Transfer: We used the number of AI/DL-themed research papers presented at conferences and our data source was the Center for Security and Emerging Technology.

- Public Policy Initiatives: We identified instances when national governments announced or published a comprehensive AI/DL-focused policy or policies. For that we used Center for Security and Emerging Technology and The AI Index 2023 Annual Report.

- Geographic Dispersion: The geographic dispersion of each dimension was tracked.

A wide variety of possible data sources were explored. We focused on three criteria when comparing alternative data sources. First, we needed to data sources that provided sufficiently clear definition of AI and DL to avoid definitional ambiguity. Second, we sought data sources that matched our pervasiveness dimensions as closely as possible. Third, we sought data from credible, recognized sources. In the process, we chose the following public and proprietary data sources:

- Crunchbase: Founded in 2007, it is proprietary database of over three million records of startups globally. Cumulative totals of AI and DL startups for the period from 2000 to 2020 were exclusive of firms the ceased operations (either through closure, acquisition, or otherwise). In addition, only those startups that received venture capital funding and had made explicit claims of utilizing AI, DL, or both in their products and services were included in our analysis as well.

- The AI Index 2023 Annual Report: This data set is well-regarded aggregate annual data base AI-related investment, research, and activity administered by Stanford University Institute for Human-Centered Artificial Intelligence.

Changes in the dollar value of investments for all data used are in real dollars utilizing the Consumer Price Index (CPI).

- World Intellectual Property Organization (WIPO): This data set includes more than 43 million patent documents, including international patent applications submitted under the International Patent System.

- Center for Security and Emerging Technology (CSET): This think tank is an extension of the Georgetown University's School of Foreign Service. The CSET project identifies the universe of research publications relating to the application and development of artificial intelligence and is, itself, a project of The Allen Institute for Artificial Intelligence.

- CB Insights: This is a proprietary database of machine learning and advanced analytics to targeted venture capital, startups, patents, and other data to predict market trends.

We now describe the intent of each pervasiveness dimension:

Pervasiveness Dimension: Products and Services

As we saw, there is both anecdotal and empirical support showing the AI/DL is in ever more products and services across markets, economies, and geographies. However, a precise and through accounting of AI/DL's impact on each product or service, though immensely valuable, is beyond the capacity of this research, given the vast number and variety of new of products and services introduced everyday all over the world for 2000 to 2020.

However, there are useful proxies for this dimension in the *source* of new products and services, which are the startups and established companies providing and investing in those AI/DL-infused products and services. We therefore assessed (1) startups utilizing AI/DL as an essential component of their offering and (2) announced investments in AI/DL by public companies.

The challenges with measuring this dimension in this way was three-fold. First, commercial claims of the extent to which AI/DL were employed versus 'marketing spin' is an open question. Further, only startups founded within 2000 to 2020 were included in the analysis. A distinction was also made between startups utilizing AI generally and those that specifically claimed to utilize DL to add granularity to our analysis.

Similarly, the veracity of announcements by publicly funded companies utilizing AI/DL could also be questioned as marketing spin. However, the fact that they are public companies, and thus liable for extreme claims, does lend credibility to their announced claims. An additional aspect of pervasiveness for products and services is the extent to which those goods and services are being created or offered across industrial and/or market sectors as well as geographic regions. Again, we used startups and private investment as proxies for the pervasiveness of AI/DL across goods and services, in this case across market sector and geographies.

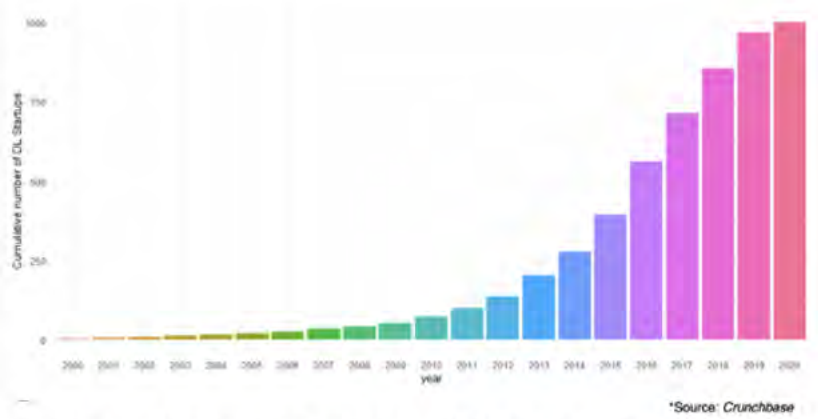

Figure 10. Cumulative number of DL startups by year

Figure 10 shows the cumulative number of DL startups by year globally. These numbers are net of startups that were founded within the 2000 to 2020 timeframe and ceased operations. The data shows a rapid diffusion of DL across startups, although 2020 appears to show a slow-down in the rate of growth. Figure 11 shows that same growth of DL but broken out across five market verticals including education, finance, healthcare, IT, and others. The category of 'others' includes a wide array of verticals from communications, manufacturing, and services, to name a few. The figures show a consistent diffusion and equitable dispersion across verticals.

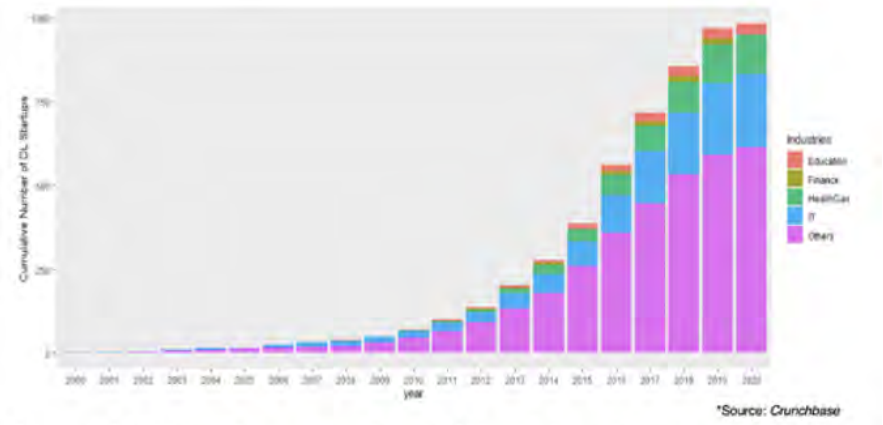

Figure 11. Cumulative number of DL startups by year & industry

Figure 12 shows the cumulative corporate investment by type investment of globally in AI for the time-period 2013 partially through the pandemic in 2022. Though this period is inside of our 2000 to 2020 timeframe, 2013 is the first year that investment became substantial. Again, we see both the characteristic rapid relative diffusion and fairly equitable dispersion across investment type. Figure 13 shows the annual investment in AI by industry sector for the time-period 2017 to 2022. This data shows a

substantial breadth of application across a very wide array of industry sectors for the second decade of our time-period.

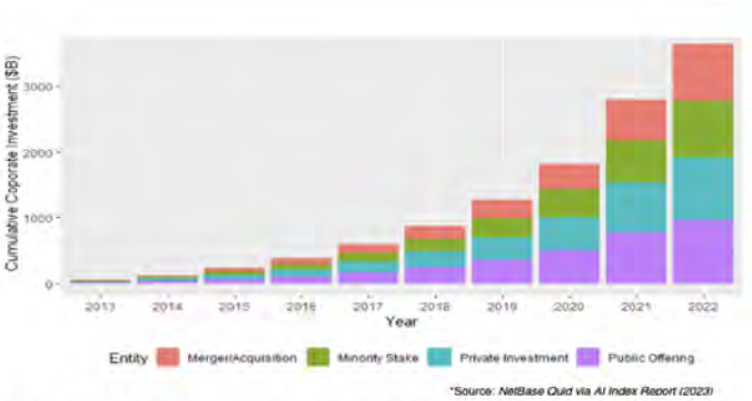

Figure 12. Cumulative corporate investment by year & investment type

Pervasiveness Dimension 2: Novel Applications

As the literature review revealed, prior studies examined AI's diffusion by counting the growth of filed and/or granted patents. We applied a similar analysis to measure the extent to which AI/DL are reflected in novel applications.

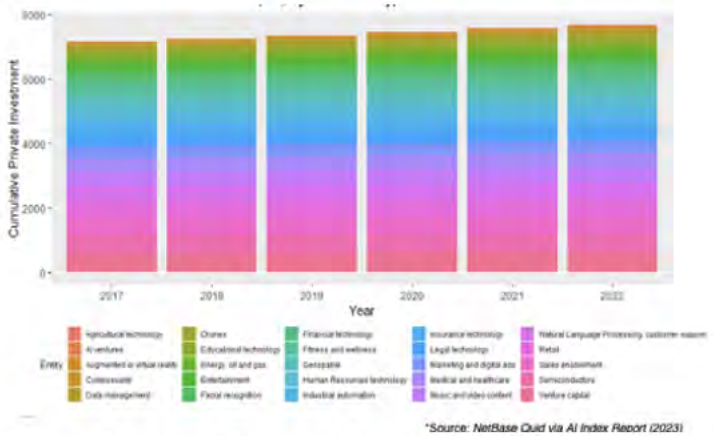

Figure 13. Cumulative private investment by year & type

Specifically, we counted new patent filings and not patents granted is the time between filings and issuance likely varies widely across all patents. We also closely evaluated the claims within the filings to ensure that the patent conveyed that AI, DL or both were fundamental to the technology's application. Figure 14 shows the cumulative number of AI-infused patents globally across market sectors for the time-period of 2010 to 2020. The data shows a rapid growth in the diffusion of AI through patents and an equitable dispersion across industry sectors.

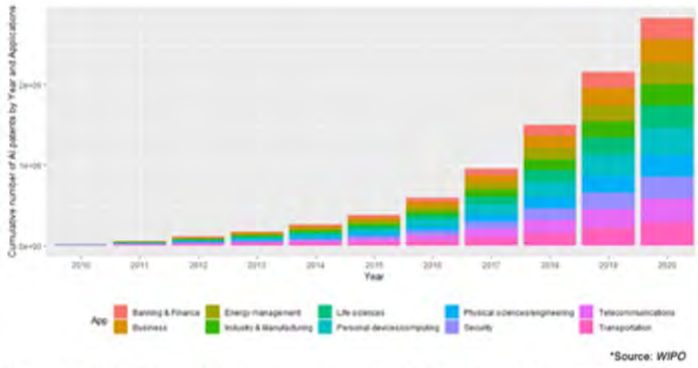

Figure 14. Cumulative number of AI patents by year & application

Pervasiveness Dimension 3: Knowledge Transfer

It is essential to measure to the extent to which AI/DL is under consideration within academic research, as it is an important source of concepts and new ideas. We therefore counted published papers, subject conferences, and presentations are research conferences in all cases where the focus was AI/DL.

While far from comprehensively measuring the transfer of AI/DL knowledge, they do, nevertheless, provide a meaningful indicator of the relative level and relative

changes in the intensity of that knowledge transfer. Figure 15 shows the cumulative number of AI publications by year and region from the time-period 2010 to 2020. The data shows a rapid diffusion and an equitable dispersion across regions.

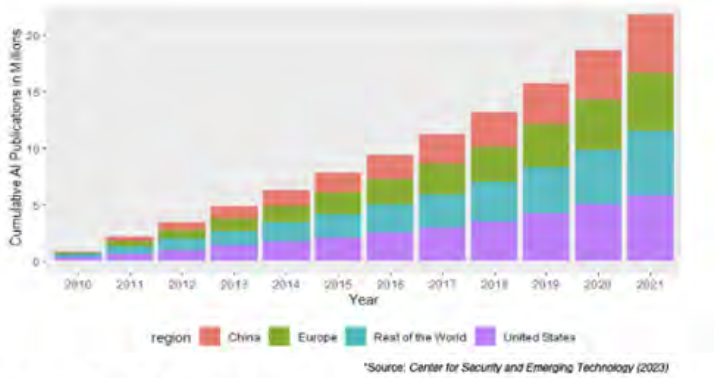

Figure 15. Cumulative AI publications by year & region

Pervasiveness Dimension 4: Public Policy Initiatives

Public policy exists to drive changes in society, but public policy initiatives are also often responses (i.e., signals) to perceived changes within and across a given nation (Mueller, 2020). Therefore, as AI diffuses and begins to have a macro-impact within a region (e.g., the European Union, Middle East North Africa, etc.) or national governments and regional governing bodies would likely have announced or are undertaking AI-focused public initiatives in response.

The nature of the initiatives will vary as widely as there are numbers of countries and regions. Some may view the perceived changes wrought by AI as a threat, others as an opportunity, and still others as a little bit of both. Such indicators would best be found in public pronouncements of national policy initiatives. As Figure 16 shows those

nations and regions who have announced but not released a national AI strategy initiative (red) and those nations and regions who have released their national AI strategy. The data shows a breadth of national strategies globally.

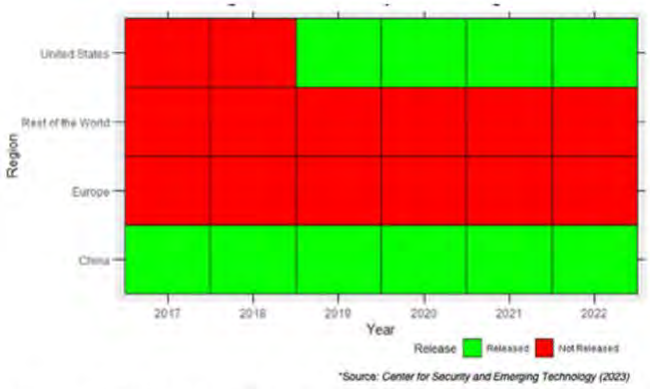

Figure 16. National strategies in AI released by year & region

Pervasiveness Dimension 5: Geographic Dispersion

The oft-cited GPTs of electrification, the steam engine, and the automobile became global phenomena. For these GPTs, diffusion started in the industrialized countries of Europe, North America, and parts of Asia of the mid to late 19th century through the early 20th century. Ultimately, these GPTs quickly spanned the globe, as today 92 percent of world population has direct access to electricity (World Bank, 2023) with 1.4 billion automobiles today (Saja, 2020). The extent to which AI/DL reaches across regions or countries within each dimension is an aspect of pervasiveness.

Figure 17 shows the cumulative number of AI startups founded (net of startup closures) between 2013 and 2022 by region (i.e., China, EU/UK, US, and Rest of

World). The data shows both a rapid growth in diffusion globally and an equitable dispersion across regions.

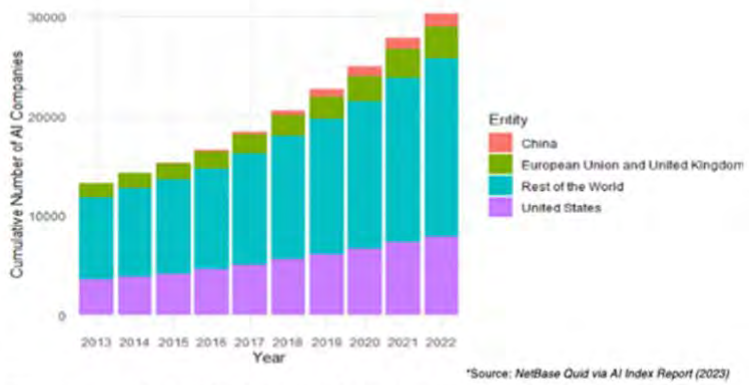

Figure 17. Cumulative number of AI startups by year & region

As Figure 18 shows the cumulative number of DL companies founded (net of startup closures) between 2000 and 2022 by region (i.e., China, EU/UK, US, and Rest of World). The data shows both a rapid growth in diffusion globally and an equitable dispersion across regions.

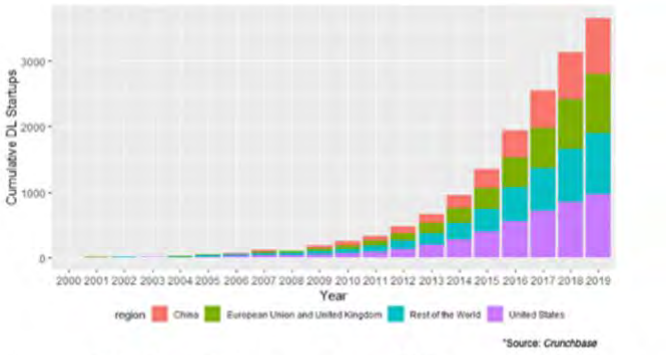

Figure 18. Cumulative number of DL startups by year & region

Figure 19 shows the cumulative number of AI-infused patents filed between 2020 and 2022 by region (i.e., China, EU/UK, US, and Rest of World). The data shows both a rapid growth in diffusion globally and an equitable dispersion across regions.

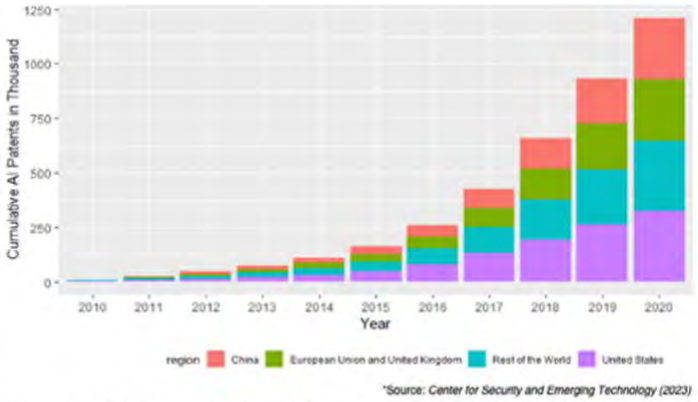

Figure 19. Cumulative AI patents by year & region

CHAPTER 5

CONCLUSIONS

Our focus on AI as a GPT served several purposes. First, AI proved to be an excellent candidate for GPT status! The most recent advances in, for example, Generative AI and Large Language Models, alone hold the potential for a macro-scale impact. Second, AI/DL provides a real-time window into how science can birth 'impractical' ideas holding little or no promise only to have them, with patient, persistent effort over decades, emerge as a GPT impacting billions of lives for decades to come – for better, worse, or maybe a little of both. Finally, our literature review of historic GPTs, along with our analysis of AI, showed how technologies separated by decades and even centuries can have a great deal in common and yield useful insights for the future, if one can ignore the hype and confusion and look closely at what is happening.

The Limitations of Our Approach

First, there is no way of knowing (now) if the cumulative growth curve presented in Figure 11 for each dimension is the beginning of a continuous asymptotic inflection in the coming years ahead or a temporary inflection that will soon crest and decline. That is why this research included as wide a variety of overlapping data sets as possible supporting our central proposition. Yet again, time will tell!

Second, as we mentioned we applied a 'propositional approach' to our research, mainly because of the scale and temporal challenges cited in our research constructs (i.e., GPTs, pervasiveness, and AI/DL assessed over time to the present). From the

outset, we sought to identify patterns *possibly* connecting our 'pre-hoc' constructs (e.g., the slow and then very rapid growth of pervasiveness for each dimension) and not testable hypotheses.

It is also the reason we stressed the conditional nature of this research effort. We are predicting, though across a range of possible outcomes, and as Nobel Laureate physicist Niels Bohr reminds us, "Prediction is very difficult, especially if it's about the future!" Nevertheless, our approach is no less worthy or useful.

Last, we drew from multiple disciplines. The literature review captures, to the fullest extent, the nuanced differences in meaning behind different terms used by different researchers within and, in some cases, across their respective disciplines. As a result, some meaning may have been unintentionally lost along the way. In effect, we faced challenges common in multi-disciplinary research (Shanableh, et. al., 2022) but were inspired to press on by philosopher and scientist Johann Wolfgang von Goethe when he said, "Whatever you can do, or dream you can, begin it. Boldness has genius, power, and magic in it." We next turn to our data analysis.

AI and the Physical World

The automobile changed our world in so many ways. At the advent of the Model T anyone could now afford a car that could allow one to travel nearly anywhere at any time and as often as they chose. Yet the GPT that is the automobile also brought second and third order impacts to our *physical* world. For one, the automobile motivated the construction of millions of miles of paved roads, highways, bridges, filling stations, hotels, to name a few (Hugill, 1982). The massive number of refineries, manufacturing

plants, asphalt, and concrete plants, show rooms, parking lots and garages would also need to be included to get a greater sense of the full transformative physical impact of the automobile on our planet. And, of course, the vast environmental consequences of over a century of the automobile would also need to be considered.

The GPT that is the electric grid also brought massive physical changes to our world and they are only now being assessed and tallied (Popovich, N. & Plumer, B., 2023). One need only consider the countless power stations, power lines, nuclear power plants, hydroelectric dams, wind, and solar farms, along with vast expanse of consumers of electric power. In a sense, the interstate highways, and beltways along with power lines and coal-fired power plants are only some of the physical enabling, ongoing costs of adoption of these two GPTs.

However, the advance of the digital world has not encroached on our physical world to the degree, breadth, and depth of the automobile and the electric grid, or at least not yet. However, AI/DL may change that. One possible reason is that AI and DL have given machines the ability to perceive the physical world – overcoming Moravec's Paradox – function within it, and ultimately change it, directly and indirectly, and for better or worse.

Robots, autonomous vehicles, and an ever-growing number of IOT devices are likely just the beginning of machines taking on virtually all the tasks that are or even could be done by humans. The possibilities are nearly infinite as the consequences for how our built world will likely change. It is possible that the pace of technological

innovation and change will make the current versions of these advanced technologies look like crude instruments.

It is one thing to give machines the ability to perceive the world around them and quite another to also add the ability to 'reason.' That is, machines are increasingly able to function independent of direct human supervision on increasingly complex tasks. Autonomous driving is one obvious example, and it may hold the promise and challenge of massive changes in our built environment. For example, would parking lots and garages be needed in a world of autonomous vehicles? Indeed, would there even be a need for nearly as many vehicles or, for that matter, would car ownership itself become a thing of the past? In that event, will automobile manufacturers start to look like Uber? With far fewer cars will the road network look the same?

There are increasing concerns about the ultimate impact of AI on our world. Some see a dystopian world, others foresee a world where humans are relieved of the burden of work, and still others who see a new evolutionary cycle where human and machine symbiotically evolve. 'Technophobia' is now a treatable condition, suffered mainly older adults as they attempt to coup with newer digital technologies (Di Giacomo D. et al., 2019).

Informing Decisions

It is hoped this research informs researcher and decisionmaker alike and anticipate the next possible GPT. For example, investors and executives will understand the role and impact of complementary and supporting technologies and, consequently, the impact of investments in those technologies. Policymakers too are often under

tremendous pressure to find and invest in the next GPT (Hidalgo & Hausmann, 2009 & Delgado et al., 2014). Further, they're increasingly called to 'regulate' technologies that are just emerging and possessing extraordinary potential (Trajtenberg, 2019; Beraja et al., 2021; & Austin & Macauley, 2001). Our research can inform those choices before the mass diffusion and mass impact of the GPT is felt.

We must acknowledge prediction is very hard. However, a forward-looking tool, however imprecise, can guide R&D decisions and assist in forming innovation alliances (Cassiman & Veugelers, 2006 & Cohen 2010) and the purchase and/or outsourcing of third-party services (Attewell 1992 & Bresnahan, et al., 1996). Further, such a tool can allow the innovators to better capture the profits from developing complementary technologies (Gambardella, et al., 2021). Last, they can prepare for what is over the horizon, by helping with the explanation and prediction of economic change and growth (Helpman, 2004).

A deeper understanding of GPTs means a deeper understanding of the process of innovation, which is vital to improving our world. Nobel laureate Kenneth Arrow (2012) put it best:

> …it is hopeless to develop a model which will genuinely predict innovations… and in claiming that those models and the considerations behind them will provide some useful idea of the average rate of technological change, of the degree of fluctuations and the kinds of surprise that we may find in the future. We cannot, of course, predict a surprise; that is a contradiction in terms. But we can predict the kind of surprises that might occur.

The hope of this research is that it will inform the decisions of a variety of decisionmakers. For example, this research is a starting point for improved private investment decisions in 'transformative' technologies by giving investors a more

discerning instrument for identifying and differentiating among innovations of all sorts. Such a tool would then allow for more thoughtful and focused long and near-term investment decisions.

Another potential benefit of our research is the ability to guide innovators *when* to capture the potential benefits of their GPT. That is, knowing what kind of GPT will point the way to the innovators as to what types of complementary assets to invest in. In fact, a GPT may focus public and private investing on the appropriate complementary technologies and supporting skills, as complementary assets play a less central role in value capture if the technology is not general purpose (Gambardella, et al. 2021).

Last, timing is everything and innovation within and across a given industry needs time to diffuse and have an impact. This research may help guide the timing of public and private investment in innovation as horizons are likely to be long term. A GPT is a long, big game and investors, public or private, need to know what game they are playing. This research hopes to inform those investors as to which game they're in and how to play it.

REFERENCES

Agrawal, A., Gans, J., & Goldfarb, A. eds. (2018a). *The Economics of Artificial Intelligence: An Agenda*. Chicago: University of Chicago Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2018b) *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, Boston, MA.

Agrawal, A., Gans, J., & Goldfarb, A. (2019) *The Economics of Artificial Intelligence: An Agenda, National Bureau of Economic Research*, The University of Chicago Press.

Alderucci, D. & Ashley. K. (2020) Using AI to Analyze Patent Claim Indefiniteness. *IP Theory* 9.1:2.

Andergassen, R., Nardini, F., & Ricottilli, M. (2017) Innovation diffusion, general purpose technologies and economic growth." *Structural Change and Economic Dynamics* 40: 72-80. http://dx.doi.org/10.1016/j.strueco.2016.12.003

Arrow, K. J. (2012) *Social Choice and Individual Values*. Yale University Press. http://www.jstor.org/stable/j.ctt1nqb90

Ascher, W. (1981) The forecasting potential of complex models. *Policy Sciences*, 13(3), 247–267. http://dx.doi.org/10.1007/BF00138485

Bass, F.M. (1969) A new product growth model for consumer durables. *Management Science*, 15: 215–27. http://dx.doi.org/10.1287/mnsc.15.5.215

Beaudry, P., Doms, M., & Lewis, E. (2010) Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas. *Journal of Political Economy,* 118(5), 988–1036. https://doi.org/10.1086/658371

Bekar, C., Carlaw, K., & Lipsey, R. (2018) General purpose technologies in theory, application, and controversy: A review. J*ournal of Evolutionary Economics*, 28, p. 1005–1033. 10.1007/s00191-017-0546-0

Bengio, Y., Lecun, Y., & Hinton, G. (2021). Deep learning for AI. *Communications of the ACM*, 64(7), 58-65. http://dx.doi.org/10.1145/3448250

Berkeley E. (1961) *Giant Brains: Or Machines That Think, Science* Editions, Inc. New York, New York. http://dx.doi.org/10.1063/1.3066893

Boldrin, M. & Levine, D. (2008) *Against intellectual monopoly*. Vol. 62. Cambridge: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511510854

Bresnahan, T. & Trajtenberg, M. (1995) General purpose technologies 'Engines of growth'? *Journal of Econometics*, 65, 1, p. 83–108. https://doi.org/10.1016/0304-4076(94)01598-T

Bresnahan, T., Greenstein, S., Brownstone, D. & Flamm, K. (1996) Technical Progress and Co-Invention in Computing and in the Uses of Computers, *Brookings Papers on Economic Activity*. Microeconomics, Vol. 1996, pp. 1-83. http://dx.doi.org/10.2307/2534746

Bresnahan, T. (2012), *Generality, Recombination and Re-Use in J. Lerner and S. Stern(eds) The Rate and Direction of Inventive Activity Revisited*, University Chicago Press, p. 611-656. https://press.uchicago.edu/ucp/books/book/chicago/R/bo12947031.html

Bresnahan, T. & Gambardella, A. (1998) *The Division of Inventive Labor and the Extent of the Market* in *General Purpose Technologies and Economic Growth*, ed. Elhanan Helpman, Cambridge, MA: MIT Press, p. 253-282.

Bresnahan, T. & Trajtenberg, M. (1995) General Purpose Technologies: 'Engines of Growth?' *Journal of Econometrics*, 95, p. 83-108. http://dx.doi.org/10.1016/0304-4076(94)01598-T

Brown, K. (2009) *Foul Bodies: Cleanliness in Early America*, New Haven: Yale University Press. http://dx.doi.org/10.12987/9780300160277

Brynjolfsson E., & McAfee, A. (2017) "The Business of Artificial Intelligence," *Harvard Business Review*, p. 1-20.

Brynjolfsson, E., Rock, D. & Syverson, C. (2019) *Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics*. In Agrawal, Gans, Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, p. 23-57.

Buchanan, M. (2013) *Forecast: What Physics, Meteorology, And the Natural Sciences Can Teach Us About Economics*, Bloomsbury, New York, New York.

Bush, V. (1990) Science: The Endless Frontier, National Science Foundation (NSF 90-8), Washington, DC, (Orig. ed. 1945).

http://dx.doi.org/10.7208/chicago/9780226613475.003.0001

Cantner, U. & Vannuccini, S. (2012) A new view of general purpose technologies, Jena *Economic Research Papers*, No. 2012,054, Friedrich Schiller University, Jena and Max Planck Institute of Economics, Jena Carlaw, Kenneth I. and Richard G. Lipsey. 2002. "Externalities, Technological Complementarities and Sustained Economic Growth." *Research Policy* 31: p. 1305-1315.

Carlaw, K. & Lipsey, R. (2006) GPT-Driven, Endogenous Growth, *The Economic Journal*, 116(508), p. 155-174. http://dx.doi.org/10.1111/j.1468-0297.2006.01051.x

Cerullo, M. (2023) "ChatGPT user base is growing faster than TikTok," MoneyWatch.

Choi, S. W., You, Y. Y., & Na, K. S. (2014). Forecasting promising technology using analysis of patent information: Focused on next generation mobile communications. *Journal of Central South University*, 21, 4303-4310. http://dx.doi.org/10.1007/s11771-014-2429-y

Christensen, C.M. (2000) *The Innovator's Dilemma*, HarperCollins Publishers, New York. http://dx.doi.org/10.15358/9783800642816

Christensen, C.M. (2013) *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Review Press.

Chui, M., (2023) "The state of AI in 2023: Generative AI's breakout year," McKinsey & Co. online survey.

Coates, V., Faroque, M., Klavins, R., Lapid, K., Linstone, H.A., Pistorius, C., Porter, A.L. (2001) On the future of technological forecasting. *Technological Forecasting and Social Change*, 67 (1), 1–17. http://dx.doi.org/10.1016/S0040-1625(00)00122-0

Coates, J. F., Mahaffie, J., & Hines, A. (1994) Technological forecasting: 1970–1993, *Technological Forecasting and Social Change*, 47, 23–33. http://dx.doi.org/10.1016/0040-1625(94)90037-X

Cockburn, I., Henderson, R. & Stern, S. (2019) *The Impact of Artificial Intelligence on Innovation*. In Agrawal, Gans, Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, p. 115-146. http://dx.doi.org/10.7208/chicago/9780226613475.003.0004

Coe, L. (2003) *The Telegraph: A History of Morse's Invention and Its Predecessors in the United States*, McFarland & Company.

Crafts, N.F.R. (2021) Artificial intelligence as a general-purpose technology: an historical perspective, *Oxford Review of Economic Policy*, Volume 37, p. 521–536. http://dx.doi.org/10.1093/oxrep/grab012

Crafts, N.F.R. & Mills, T.C. (2004) Was 19th century British growth steam powered? the climacteric revisited. *Explorations in Economic History,* 41. http://dx.doi.org/10.1016/j.eeh.2003.10.001

Daim, T.U., Rueda, G., Martin, H., & Gerdsri, P. (2006) Forecasting emerging technologies: use of bibliometrics and patent analysis. *Technological Forecasting and Social Change*, 73 (8), 981–1012 Oct. http://dx.doi.org/10.1016/j.techfore.2006.04.004

D'Amour, A., et al. (2022). Under-specification presents challenges for credibility in modern machine learning. *The Journal of Machine Learning Research*, 23(1), 10237-10297.

David, P. A. (1991) 'Computer and dynamo: the modern productivity paradox in a not-too-distant mirror', in *Technology and Productivity: The Challenge for Economic Policy*, p. 315–348, Paris: OECD.

Davenport, D. (2022) "Predictably Bad Investments: Evidence from Venture Capitalists" http://dx.doi.org/10.2139/ssrn.4135861

David, P. (1991) Computer and dynamo: The modern productivity paradox in a not-too-distant mirror. In *Technology and Productivity: The Challenge for Economic Policy*, p. 315-347 (OECD, Paris).

David, P. A. & G. Wright (2003) *General Purpose Technologies and Surges in Productivity: Historical Reflections on the future of the ICT Revolution*, In: P. A. David and M. Thomas (eds.), *The Economic Future in Historical Perspective*, Oxford University Press. http://dx.doi.org/10.5871/bacad/9780197263471.003.0005

David, P.A. & Wright, G. (1999) "General purpose technologies and surges in productivity: Historical reflections on the future of the ICT revolution." http://dx.doi.org/10.5871/bacad/9780197263471.003.0005

David, P. A. (1990) The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox, *American Economic Review,* 80, p. 355-61.

Davis, J. L., & Chouinard, J. B. (2016) Theorizing Affordances: From Request to Refuse. Bulletin of Science, *Technology & Society*, 36(4), 241–248. https://doi.org/10.1177/0270467617714944

Day, G. & Schoemaker, P. (2005) "Scanning the periphery." *Harvard Business Review* (pp. 1–12).

Di Giacomo D, et al., (2019) Psychological Barriers to Digital Living in Older Adults: Computer Anxiety as Predictive Mechanism for Technophobia. *Behavioral Science* (Basel).;9(9):96. doi: 10.3390/bs9090096.

Domingos, P., (2015) *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books.

Donaldson, D. & Hornbeck, R., (2016) Railroads and American Economic Growth: A "Market Access" Approach , *The Quarterly Journal of Economics*, Volume 131, Issue 2, Pages 799–858, https://doi.org/10.1093/qje/qjw002

Dosi, G. (1988) *The Nature of the Innovation Process, Technical Change and Economic Theory*, Pinter, London.

Drezner, Daniel W. (2019) "Economic statecraft in the age of Trump." *The Washington Quarterly* 42.3: 7-24. http://dx.doi.org/10.1080/0163660X.2019.1663072

Drucker, P. F. (1969) Comment: Is Technology Predictable? *Technology and Culture*, 10(4), 522–527. http://dx.doi.org/10.1353/tech.1969.a892397

Eaton, J., & Kortum, S. (1999) International Technology Diffusion: Theory and Measurement. *International Economic Review*, 40(3), 537–570. http://dx.doi.org/10.1111/1468-2354.00028

Efanov, D. & Roschin, P., (2018) The All-Pervasiveness of the Blockchain Technology, *Procedia Computer Science*, Volume 123, Pages 116- http://dx.doi.org/10.1016/j.procs.2018.01.019

Emmer, M., (2018) "95 Percent of New Products Fail." *Inc. Magazine*.

Érdi, P., Makovi, K., Somogyvári, Z. et al. (2013) Prediction of emerging technologies based on analysis of the US patent citation network. *Scientometrics*, 95, 225–242.

Erikson, R. S., and C. Wlezien. (2008) Are political markets really superior to polls as election predictors? *Public Opinion Quarterly*, 72(2): 190-215. Green, K. C, J. S. Arms. http://dx.doi.org/10.1093/poq/nfn010

Esmaelian, M., Tavana, M., Di Caprio, D., Ansari, R. (2017) A multiple correspondence analysis model for evaluating technology foresight methods. *Technological Forecasting and Social Change,* 125, 188–205. http://dx.doi.org/10.1016/j.techfore.2017.07.022

Feldman, M. P. & Yoon, J. W. (2011) An empirical test for general purpose technology: An examination of the Cohen–Boyer RDNA technology. *Industrial and Corporate Change*, 21(2), p. 249–275. http://dx.doi.org/10.1093/icc/dtr040

Feng, L, Wang, Q, Wang J, Lin, KY. (2022) *A Review of Technological Forecasting from the Perspective of Complex Systems*. Entropy (Basel). 4;24(6):787. http://dx.doi.org/10.3390/e24060787

Fergnani, A. (2022). Corporate foresight: A new frontier for strategy and management. *Academy of Management Perspectives*, 36(2), 820-844. http://dx.doi.org/10.5465/amp.2018.0178

Ferrazzi, K., (2015) Technology Can Save Onboarding from Itself, *Harvard Business Review*, online.

Fichman, R. G. (1992) Information technology diffusion: A review of empirical research. *In ICIS* (Vol. 195, p. 206)

Field, A. (2011) *A Great Leap Forward: 1930s Depression and U.S. Economic Growth*. New Haven: Yale University Press. http://dx.doi.org/10.12987/9780300168754

Firat, A.K., Madnick, S., & Woon, W.L. (2008) "Technology forecasting: A review." In: *Working Paper CISL# 2008–15. MIT*.

Fleming, L. (2001) Recombinant uncertainty in technological search. *Management Science*, 47,117–132. http://dx.doi.org/10.1287/mnsc.47.1.117.10671.

Freeman, C. & Soete, L. (1995) *The economics of industrial innovation*. 3rd ed. London: Routledge. http://dx.doi.org/10.4324/9780203064474

Frey, C. & Osborne, M. (2017) "The future of employment: How susceptible are jobs to computerization?" *Technological Forecasting and Social Change*. 114: 254-280. http://dx.doi.org/10.1016/j.techfore.2016.08.019

Friedrich, S., Antes, G., & Behr, S. (2022) Is there a role for statistics in artificial intelligence? *Advanced Data Analysis Classification*, 16, 823–846 https://doi.org/10.1007/s11634-021-00455-6

Frischmann, B. M. (2012) *Infrastructure: The social value of shared resources*. Oxford: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780199895656.001.0001

Friston, C. (2012) The history of the future of the Bayesian brain, *NeuroImage*, v 62, Issue 2, p. 1230-1233. http://dx.doi.org/10.1016/j.neuroimage.2011.10.004

Gambardella, A.; Heaton, S., Novelli, E. & Teece, D. (2021) Profiting from Enabling Technologies? *Strategy Science*, 6/1: p. 75-90. http://dx.doi.org/10.1287/stsc.2020.0119

Garimella, S.V., Persoons, T., Weibel, J., & Yeh, L.T. (2013) Technological drivers in data centers and telecom systems: Multiscale thermal, electrical, and energy management. *Applied Energy*, 107, 66–80. http://dx.doi.org/10.1016/j.apenergy.2013.02.047

Giesbrecht, J. (2015) "How the Ballpoint Pen Killed Cursive," The Atlantic.

Goldfarb, A., Taska, B., & Teodoridis, F. (2023). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy*, 52(1), 104653. http://dx.doi.org/10.1016/j.respol.2022.104653

Gordon, R. (2000) "Does the "New Economy" Measure up to the Great Inventions of the Past?" *NBER Working Paper* No. w7833. http://dx.doi.org/10.1017/CBO9780511616587.005

Gordon, R. (2016) *The Rise and Fall of American Growth*. Princeton University Press. http://dx.doi.org/10.1515/9781400888955

Gordon, T. J., & Helmer, O. (1964) *Report on a Long-Range Forecasting Study*. RAND Corporation, Santa Monica, CA.

Gottfredson, M. & Aspinall, K (2005) "Innovation versus Complexity: What is Too Much of Good Thing," Harvard Business Review.

Geroski, A. (2000) Models of technology diffusion, *Research Policy*, Volume 29, Issues 4–5, Pages 603-625, http://dx.doi.org/10.1016/S0048-7333(99)00092-X

Green, K. C, J. S. Armstrong, and A. Graefe. (2007) Methods to elicit forecasts from groups: Delphi and prediction markets compared. Foresight: *The International Journal of Applied Forecasting*, 8:17-20.

Greenhalgh, T., Robert, G., Macfarlane, F., Bate, P., Kyriakidou, O., & Peacock, R. (2005). Storylines of research in diffusion of innovation: a meta-narrative approach to systematic review. *Social science & medicine*, 61(2), 417-430. http://dx.doi.org/10.1016/j.socscimed.2004.12.001

Hall, B. & Trajtenberg, M. (2004) Uncovering GPTS with Patent Data. *The Journal of Economic History*, 64(1), p. 61-99.

Haenlein, M. & Kaplan, A. (2019) A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*. 61. 000812561986492. 10.1177/0008125619864925.

Haigh, T., Priestley, P. M., & Rope, C. (2016). *ENIAC in action: Making and remaking the modern computer*. MIT Press. http://dx.doi.org/10.7551/mitpress/9780262033985.001.0001

Harteis, C. & Gruber, H. (2008) Intuition and Professional Competence: Intuitive Versus Rational Forecasting of the Stock Market. *Vocations and Learning*, 1, 71–85. http://dx.doi.org/10.1007/s12186-007-9000-z

Heim, U., Heuss, R., & Katzir, T. (2017) "Building an Integrated Technology Road Map to Drive Successful Innovation," McKinsey.

Helpman, E. (2004) *The Mystery of Economic Growth*. Cambridge, MA: Harvard. http://dx.doi.org/10.2307/j.ctv1c5cx4v

Helpman, E. (1998) *General Purpose Technologies and Economic Growth*, Cambridge, MA: MIT Press.

Helpman, E. & Trajtenberg, M. (1998) "*A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies*," Helpman, E. (ed.), *General Purpose Technologies and Economic Growth*, Cambridge: MIT Press, p. 85-119.

Hogendorn, C., & Frischmann, B. (2020). Infrastructure and general purpose technologies: A technology flow framework. *European Journal of Law and Economics*, 50(3), 469-488. http://dx.doi.org/10.1007/s10657-020-09642-w

Holden, C. (2000) "Engineering's Greatest Hits," *Science*. doi: 10.1126/article.37790

Houdé, O., Kayser, D., Koenig, O., Proust, J., & Rastier, F. (Eds.). (2004). *Dictionary of cognitive science: Neuroscience, psychology, artificial intelligence, linguistics, and philosophy*. Routledge. https://doi.org/10.4324/9780203486030

Hötte, K., Tarannum, T., Verendel, V., & Bennett, L. (2022). Exploring Artificial Intelligence as a General Purpose Technology with Patent Data--A Systematic Comparison of Four Classification Approaches. arXiv preprint arXiv:2204.10304.

Huang, L., Zhang, Y., Guo, Y., Zhu, D. & Porter, A.L. (2014) "Four-dimensional science and technology planning: a new approach based on bibliometrics and technology road mapping", *Technological Forecasting and Social Change*, Vol. 81, pp. 39-48. http://dx.doi.org/10.1016/j.techfore.2012.09.010

Hugill, P. J. (1982). Good Roads and the Automobile in the United States 1880-1929. *Geographical Review*, *72*(3), 327–349. https://doi.org/10.2307/214531 IEEE, Pearl Street Station, IEEE *Global History Network*.

Jantsch, E. (1967) *Technological Forecasting in Perspective*, OECD, Paris

Jakobovits, L. A. (1962) *"Effects of Repeated Stimulation on Cognitive Aspects of Behavior: Some Experiments on the Phenomenon of Semantic Satiation"* (PhD), McGill University,

Johnson, M. (2022) "Ancient Relatives Used Fire, Surprising Discovery Suggests," Washington Post.

Jovanovic, B. & Rousseau, P. (2005) General Purpose Technologies, Working Paper 11093, *National Bureau of Economic Research*, Cambridge, MA http://www.nber.org/papers/w11093

Jovanovic, B. & Rousseau, P. L. (2005) *General Purpose Technologies*, in Anghion P., Durlauf S., (eds) *Handbook of Economic Growth*, Elsevier, p. 1182-1224. http://dx.doi.org/10.1016/S1574-0684(05)01018-X

Jun, S.P., Sung, T.E., & Park, H.W. (2017) Forecasting by analogy using the web search traffic. *Technology Forecasting and Social Change*, 115, 37–51. http://dx.doi.org/10.1016/j.techfore.2016.09.014

Kahneman, Daniel, (2011) *Thinking, fast and slow*. New York, Farrar, Straus, and Giroux.

Kennedy, B., & Funk, C. (2016) "28% of Americans are 'strong' early adopters of technology," Pew Research Center (online).

Kent, M.L., Saffer & A.J. (2014) A Delphi study of the future of new technology research in public relations. *Public Relations Review*, 40 (3), 568–576. http://dx.doi.org/10.1016/j.pubrev.2014.02.008

Kim, T., et al., (2014) Determining the scale of R&D investment for renewable energy in Korea using a comparative analogy approach. *Renewable & Sustainable Energy Review*, 37 (37), 307–317.c http://dx.doi.org/10.1016/j.rser.2014.05.025

Klitkou, A., Bolwig, S., Hansen, T., & Wessberg, N. (2015) The role of lock-in mechanisms in transition processes: The case of energy for road transport, *Environmental Innovation and Societal Transitions*, Volume 16, 2015, Pages 22 37. http://dx.doi.org/10.1016/j.eist.2015.07.005

Klosterman, R. (2013) Lessons Learned About Planning, *Journal of the American Planning Association*, 79:2, 161-169. http://dx.doi.org/10.1080/01944363.2013.882647

Korzinov, V. & Savin, I. (2016) Pervasive enough? General purpose technologies as an emergent property, *KIT Working Paper Series in Economics*, No. 95, https://doi.org/10.5445/IR/1000062013

Kostoff, R. N., & Schaller, R. R. (1999) Science and Technology Roadmaps. *George Mason University School of Public Policy Working Paper* No. 99:3. http://dx.doi.org/10.1109/17.922473

Kourouthanassis, P., et.al, (2008) "Delineating the Degree of 'Pervasiveness' in Pervasive Information Systems: An Assessment Framework and Design Implications," *Panhellenic Conference on Informatics*, Samos, Greece, pp. 251-255, doi: 10.1109/PCI.2008.28. http://dx.doi.org/10.1109/PCI.2008.28

Kucharavy, D., & De Guio, R. (2005) Problems of forecast. *En Proceeding of conference ETRIA TRIZ Future 2005*, p. 219-235.

Kummer, C. (1988) “Hot Popcorn: The First Popcorn Was Made by Accident. Now There Are Better Ways,” The Atlantic.

Kurzweil, R., (2012). How to Create a Mind: The Secret of Human Thought Revealed. New York, New York, Viking Books.

Lacey, R., (1986) Ford, the Men, and the Machine. Boston: Little, Brown.

Latson, J., (2015) “How Sliced Bread Became the “Greatest Thing,” Time magazine. July 7, 2015.

Lee, W. J., & Sohn, S. Y. (2014). Patent analysis to identify shale gas development in China and the United States. *Energy policy*, 74, 111-115. http://dx.doi.org/10.1016/j.enpol.2014.08.009

Leigh, A., & Wolfers, J. (2007) Prediction markets for business and public policy. *The Melbourne Review,* 3(1):7 15.

Li, X.; Xie, Q.Q.; Daim, T.; Huang, L.C. (2019) Forecasting technology trends using text mining of the gaps between science and technology: The case of perovskite solar cell technology. *Technology Forecasting and Social Change,* 146, 432–449. http://dx.doi.org/10.1016/j.techfore.2019.01.012

Lindsay, G. & Serre, T. (2022) *Deep Learning Networks and Visual Perception*, Oxford Research Encyclopedia of Psychology. http://dx.doi.org/10.1093/acrefore/9780190236557.013.841

Lipsey, R., Carlaw, K. & Bekar, C. (2005) *General purpose technologies and long-term economic growth*. Oxford: Oxford University Press.

Liu, H. & Chen, Z., Tang, J. et al. (2020) Mapping the technology evolution path: a novel model for dynamic topic detection and tracking. *Scientometrics* 125, 2043–2090 http://dx.doi.org/10.1007/s11192-020-03700-5

Liu, J., et al., (2018) "Artificial Intelligence in the 21st Century," in *IEEE Access*, vol. 6, pp. 34403-34421, http://dx.doi.org/10.1109/ACCESS.2018.2819688
Lynch. S. (2023) "State of AI in 14 Charts," *Stanford HAI online*.

Machiavelli, N. (2003) *The prince* (G. Bull, Translation). Penguin Classics. (Original work published 1532). http://dx.doi.org/10.1522/cla.man.pri

Magruk, A. (2011) Innovative classification of technology foresight methods. *Technological and Economic Development of Economy,* 17 (4), 700–715.

Mahoney, J & Schensul, D. (2006) *Historical Context and Path Dependence*, Oxford University Press. pp. 454–471. http://dx.doi.org/10.1093/oxfordhb/9780199270439.003.0024

Marbukh, V., et. al. (2010) "Towards Evolutionary-Pricing Framework for Mobile Sensor Network Self-organization", *2010 IEEE World Congress on Computational Intelligence*, Barcelona. http://dx.doi.org/10.1109/CEC.2010.5586175

Martin, B.R. (1995) Foresight in Science and Technology. *Technology Analysis and Strategic Management*. 7, 139–168. http://dx.doi.org/10.1080/09537329508524202

Martino, J (1983) *Technological Forecasting for Decision Making*, 2nd ed. New York, NY: North-Holland, pp: 39, 40.

Martino, J. P. (Ed.). (1972). *An introduction to technological forecasting* (v1). Taylor & Francis. http://dx.doi.org/10.4324/9781351106450

McCarthy, J., Minsky, M., Rochester, N., & Shannon, C.E. (1955) *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, pg. 2.

McCarthy, J. (2007) *What is Artificial Intelligence*? Stanford University, School of Computer Science, http://www-formal.stanford.edu/jmc/

McKendrick, J. (2021) "AI Adoption Skyrocketed Over the Last 18 Months'" *Harvard Business Review* (online) https://hbr.org/2021/09/ai-adoption-skyrocketed-over-the-last-18-months

McLaughlin, L. and McLaughlin, J.F. (2021) "Do we innovate as we believe we do?", *Journal of Research in Innovative Teaching & Learning,* Vol. 14 No. 2, pp. 218-231. http://dx.doi.org/10.1108/JRIT-10-2019-0073

Menzies, T. (2003) "21st-Century AI: Proud, Not Smug," *IEEE Computer Society* 1094-7167, p. 18-24. http://dx.doi.org/10.1109/MIS.2003.1200723

Merges, R. & Nelson, R. (1994) "On limiting or encouraging rivalry in technical progress: The effect of patent scope decisions*", Journal of Economic Behavior and Organization*, 25 (1): p. 1–24. http://dx.doi.org/10.1016/0167-2681(94)90083-3

Meyer, J.G., et al. (2023) ChatGPT and large language models in academia: opportunities and challenges. *BioData Mining,* 16, 20. http://dx.doi.org/10.1186/s13040-023-00339-9

Miles, I., Popper, R. (2008) *The handbook of technology foresight*. MPG Books.

Miles, M., Ryman, T.K., Dietz, V., Zell, E., Luman, E.T. (2013) Validity of vaccination cards and parental recall to estimate vaccination coverage: A systematic review of the literature. *Vaccine* 31 (12), 1560–1568. http://dx.doi.org/10.1016/j.vaccine.2012.10.089

Mishra S, & Mishra MB (2013) Tobacco: Its historical, cultural, oral, and periodontal health association, *Journal of International Society of Preventive and Community Dentistry* 2013 Jan;3(1):12-8. http://dx.doi.org/10.4103/2231-0762.115708

Mitroff, I. (2019) "Pushing Back on Pervasive Technology," *Sloan Management Review*.

Mokyr, J. (2010) *The Contribution of Economic History to the Study of Innovation and Technical Change: 1750-1914*, Handbook of the Economics of Innovation 1, 11–50.

Moltzau, A. (2019) “Artificial Intelligence – The Good, The Bad and The Ugly,” *ODSC Journal*.

Moravec, H. (1988), *Mind Children*, Harvard University Press, pgs. 15 - 16.

Moser, P, & Nicholas, T. (2004) "Was Electricity a General Purpose Technology? Evidence from Historical Patent Citations." *American Economic Review*, 94 (2): 388-394. http://dx.doi.org/10.1257/0002828041301407

Mudambi, R. & Ricketts, M. (1998) *The Organization of the Firm: International Business Perspectives*. London, Routledge.

Mueller, B. (2020) Why public policies fail: Policymaking under complexity, *EconomiA*, Volume 21, Issue 2, Pages 311-323, http://dx.doi.org/10.1016/j.econ.2019.11.002

Murdock, C. A. (1974) Defense policy formation: a comparative analysis of the McNamara era. Suny Press.

Myhrvold, N., (2018) “What the History of Math Can Teach Us about the Future of AI” Observations, *Scientific American*.

National Resources Committee: Technological Trends and National Policy, (1937) Including the Social Implications of New Inventions. *Report of the National Resources Committee on the Social Implications of New Inventions*, Washington, DC.

Nilsson N.J. (2009) *The Quest for Artificial Intelligence*. Cambridge University Press.

Ott I, Papilloud C, & Zuelsdorf T. (2009) What drives innovation? Causes of and consequences for nanotechnologies. *Managing Global Transitions,* 7(1):5 p. 26.

Park, S., Lee, S.J. & Jun, S. (2015) A network analysis model for selecting sustainable technology. *Sustain*, 7 (10), 13126–13141. http://dx.doi.org/10.3390/su71013126

Peres, R., Muller, E. & Mahajan, V. (2010) Innovation diffusion and new product growth models: A critical review and research directions. *International Journal of Research in Marketing*, 27: 91–106. http://dx.doi.org/10.1016/j.ijresmar.2009.12.012

Petrick, I. & Echols, A. (2004) Technology roadmapping in review: A tool for making sustainable new product development decisions, *Technological Forecasting and Social Change*, v71, Issues 1–2, Pages 81-100. http://dx.doi.org/10.1016/S0040-1625(03)00064-7

Popovich, N. & Plumer, B. (2023) "*Why the U.S. Electric Grid Isn't Ready for the Energy Transition*," The New York Times.

Popp, D. (2002) "Induced Innovation and Energy Prices." *American Economic Review*, 92 (1): 160-180. http://dx.doi.org/10.1257/000282802760015658

Porter, A. L. et al. (2004) Technology Future Analysis: Toward integration of the field and new methods. *Technological Forecasting and Social Change*, 2004. v71, p. 287-303.

Quinn, J.B., (1967) "Technological Forecasting," Harvard Business Review (online).

Raichle, M. (2009) A brief history of human brain mapping, *Trends in Neurosciences*, v32, Issue 2, Pages 118-126, ISSN 0166-2236, https://doi.org/10.1016/j.tins.2008.11.001.

Rhode, P. W, & Strumpf, K.S. (2004) Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2): 127-41. http://dx.doi.org/10.1257/0895330041371277

Rhodes, R. (2018) *Energy: A Human History*. New York: Simon & Schuster, 2018.

Ristuccia, C., & Solomos S., (2014) "Can general purpose technology theory explain economic growth? Electrical power as a case study." *European Review of Economic History*, 18.3: p. 227-247.v http://dx.doi.org/10.1093/ereh/heu008

Rogers, E.M. (2003) *Diffusion of Innovati*ons, 5th ed., Simon & Schuster. http://dx.doi.org/10.4324/9780203710753-35

Roman, S., et. al. (2022) "The Coronavirus Pandemic's Economic Impact," US Census Bureau.

Rosenberg, N. (1982) Inside the black box, Technology and economics, Cambridge: Cambridge University Press.

Rosenberg, N. (1976) *Technological Change in the Machine Tool Industry, 1840-1910, Perspectives in Technology*, Cambridge: Cambridge University Press, p. 9-31.

Rosenberg, N. & Trajtenberg, M. (2004) A General-purpose Technology at Work: The Corliss Steam Engine in the Late-nineteenth-century United States, *The Journal of Economic history*, 64(1), 61-99. http://dx.doi.org/10.1017/CBO9780511561313.002

Roussel, P.A., Saad, K.N. & Erickson, T.J. (1991) *Third Generation R&D: Managing the Link to Corporate Strategy*, Arthur D. Little, Boston, MA.

Routley, M. & Phaal, R. (2013) Exploring industry dynamics and interactions. *Technological Forecasting and Social Change*, 80 (6), 1147–1161. http://dx.doi.org/10.1016/j.techfore.2012.04.015

Russo, A., (2020) "Recession and Automation Changes Our Future of Work, But There are Jobs Coming," Report Says, *World Economic Forum*.

Russo, A. (2002) "Recession and Automation Changes Our Future of Work, But There are Jobs Coming, Report Says," *World Economic Forum*.

Sahal, D., (1977) A theory of measurement of technological change, *International Journal. Systems Science*. 8 671–682. http://dx.doi.org/10.1080/00207727708942073

Sahal, D., & K. Yee, (1975) Delphi: An investigation from a Bayesian viewpoint, *Technology Forecasting and Social Change*. 7, 165–178. http://dx.doi.org/10.1016/0040-1625(75)90056-6

Saja, F. (April 2020) "How Many Cars Are There in the World," Drive Tribe.

Schumpeter, J. (1911) The Theory of Economic Development. Harvard University Press, Cambridge. http://dx.doi.org/10.4324/9781315135564

Schumpeter, J. (1939) Business cycles. New York: McGraw-Hill Book Company. http://dx.doi.org/10.1522/030021081

Schnaars, S.P. (1989) Megamistakes: Forecasting and the Myth of Rapid Technological Change, Free Press, New York and London, p. 202.

Schmidhuber, J. (2015) Deep Learning in Neural Networks: An Overview. Neural Networks, Volume 61, Pages 85-117. http://dx.doi.org/10.1016/j.neunet.2014.09.003

Schwab, K. (2017) The Fourth Industrial Revolution, New York, Crown Business.

Shanableh, A., Aderibigbe, S., Omar, M., Shabib, A. (2022) *Challenges and Opportunities of Multi-Disciplinary, Inter-Disciplinary and Trans-Disciplinary*

Research. In: Badran, A., Baydoun, E., Hillman, J.R. (eds) *Higher Education in the Arab World: Research and Development*. Springer, Cham. http://dx.doi.org/10.1007/978-3-030-80122-9_18

Sickles, R., & Zelenyuk, V. (2019) *Measurement of Productivity and Efficiency: Theory and Practice*. Cambridge: Cambridge University Press. http://dx.doi.org/10.1017/9781139565981

Slocum, M.S. & Lundberg, C.O. (2001) Technology Forecasting: From Emotional to Empirical. *Creativity and Innovation Management*, 10: 139-152. http://dx.doi.org/10.1111/1467-8691.00215

Smil, V. (2015) *Power density: a key to understanding energy sources and uses*. MIT Press. http://dx.doi.org/10.7551/mitpress/10046.001.0001

Smith, C. (2021) "A.I. Here, There, Everywhere," The New York Times.

Snowberg, E., Wolfers, J. & Zitzewitz, E. (2005) *Information (in)efficiency in prediction markets*. In *Financial and betting markets*, edited by L. Vaughan Williams. Cambridge, MA: Cambridge University Press, pp. 366-86.

Sood A. & Tellis G. (2005) Technological evolution and radical innovation. Journal of Marketing, 69, 152–168. http://dx.doi.org/10.1509/jmkg.69.3.152.66361

Surowiecki, J. (2005) The Wisdom of Crowds. United States: Knopf Doubleday Publishing Group.

Steinert, M. & Leifer, L. (2010) "Scrutinizing Gartner's hype cycle approach," PICMET 2010 *Technology Management for Economic Growth*, pp. 1-13.

Stevenson, B., (2018) "*AI, Income, Employment, and Meaning*." In *The Economics of Artificial Intelligence: An Agenda*, ed. Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press.

Strebulaev, I. & Gornall, W. (2015) "How Much Does Venture Capital Drive the U.S. Economy?" *Insights by Stanford Busines*s, Stanford Business School.

Strohmaier R, & Rainer A. (2016) *Studying general purpose technologies in a multi-sector framework: The case of ICT* in *Denmark. Structural Change and Economic Dynamics*, 36: p. 34-49. http://dx.doi.org/10.1016/j.strueco.2015.10.003

Tetlock, P. E., & Gardner, D. (2015) *Superforecasting: The art and science of prediction*. Crown Publishers/Random House.

Trajtenberg, Manuel. (2019) *AI as the next GPT: A Political Economy Perspective. In Agrawal, Gans, Goldfarb Eds. The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, p. 175-186. http://dx.doi.org/10.7208/chicago/9780226613475.003.0006

Weller, C., (2015) "9 'world-changing' innovations that ended up as failures," Business Insider.

Wertime, T. & Muhly, J.D, eds. (1979) *The Coming of the Age of Iron*, New Haven.

Wheatley, A.D, (2020) "The AI Family Tree," *The LibrAIry Blog*.

Whitehead, A. (1978) *Process and Reality: An essay in cosmology* (corrected edition). The Free Press.

Wissema, J.G. (1982) Trends in technology forecasting. *R&D Management*, 12: 27-36.

Wolfers, J. & E. Zitzewitz. (2004) Prediction markets. *Journal of Economic Perspectives*, 18(2): 107-26. http://dx.doi.org/10.1257/0895330041371321

Wolfers, J. & Leigh, A. (2002) Three tools for forecasting federal elections: Lessons from 2001. *Australian Journal of Political Science*, 37(2):223-40. http://dx.doi.org/10.1080/10361140220148115

World Bank (2023) *The Energy Progress Report*, World Bank, Washington, DC.

Woundenberg, F. (1991) An evaluation of Delphi, *Technology Forecasting and Social Change,* 40 (1991) 131 – 150.

Van Zon, A., Fortune, E., & Kronenberg, T. (2003) "How to Sow and Reap as You Go: a Simple Model of Cyclical Endogenous Growth," *MERIT Research Memoranda*

Xiang, C. (2022) "Scientists Increasingly Can't Explain How AI Works." *Vice*.

Zorpette, G. (2022) "How the First Transistor Worked," *IEEE Spetrum (online)*.

DATA SOURCES

Abdulla, S. (2023) "Leading the Charge: A Look at the Top-Producing AI Programs in U.S. Colleges and Universities" (*Center for Security and Emerging Technology*).

Abdulla, S., (2022) "Examining patent data in PARAT: Patent Applications" (*Center for Security and Emerging Technology*).

CRA Taulbee Survey (2023), *Computing Research Association*.

Epoch, (2023) 'Parameter, Compute and Data Trends in Machine Learning'. Published online at epochai.org. Retrieved from: 'https://epochai.org/mlinputs/visualization' [online resource]

Kinoshita, R., (2023) "Views of AI PhD Recipients on Resources to Build the Domestic Talent Pool" (*Center for Security and Emerging Technology*).

Maslej, N., Fattorini, L., Brynjolfsson, E., et. al., (2023) "The AI Index 2023 Annual Report," *AI Index Steering Committee, Institute for Human-Centered AI, Stanford University*, Stanford, CA.

McKinsey Report, (2022) The state of AI in 2022—and a half decade in review.

NetBase Quid (2022) via "The AI Index 2023 Annual Report," *AI Index Steering Committee, Institute for Human-Centered AI, Stanford University*, Stanford, CA, April 2023

Roser, M. (2023) - "*Artificial intelligence has advanced despite having few resources dedicated to its development – now investments have increased substantially*" Published online at *OurWorldInData.org*. https://ourworldindata.org/ai-investments

US Census of Agriculture, 1900 – 1969.

Data provided by the author upon request.