Exponential Life

An Analysis of Exponential Technologies

Tyler Lund Hassan Safouhi*

February 9, 2019

Faculté Saint-Jean, University of Alberta
8406 Rue Marie-Anne Gaboury (91 Street)
Edmonton (AB) T6C 4G9, Canada

Abstract

Technology is ubiquitous and powerful in its effects. Despite the fact that many technologies follow a predictable exponential growth rate, due to our lack of exponential intuition, we often have the impression that technology appears unexpectedly. We term this subset of technology as exponential technology. Furthermore, the convergence of technology, which refers to the cumulation of several individual technologies to produce advancements greater than the original parts, can cause technology to improve at an even more impressive rate. Studying exponential technologies provides insight into how to best take advantage of these technologies, while minimizing their negative societal effects. In particular, educational practices should be reformed to involve new technology, and reasonable regulation frameworks should be established in order to minimize the risks without discouraging innovation.

*Corresponding author: hsafouhi@ualberta.ca
URL: www.safouhi.csj.ualberta.ca

Dr. Hassan Safouhi's research is primarily centered around designing new mathematical methods and models for solving large scale problems occurring in physics, engineering, and industry. His research career is singled out by its interdisciplinary nature and by the valuable impact of his theoretical contributions applied to practical industrial problems.

Dr. Safouhi has long expressed interest in technological advances and the concept of exponential and converging technologies, particularly due to their capacity to fundamentally transform our society by presenting new opportunities as well as unforeseen challenges.

Tyler Lund is a 2nd year Mechanical Engineering student in the biomedical option at the University of Alberta. He completed successfully his first year at Faculté Saint-Jean. His engineering aspirations and general appreciation of technology give him a keen interest in technological advances.

Working on this project under the supervisipon of Dr. Safouhi has given Tyler an understanding of how technology is able to make a difference, and a unique perspective on technology, his studies and his future career.

The objective of this work is to provide a better understanding of exponentially advancing technologies, in contrast to linear technologies whose growth is much more intuitive. The analysis also defines and promotes exponential thinking – a mindset that will be needed to solve many of the world's biggest impending challenges.

Acknowledgment

- ☞ *I am extremely grateful to my wife, Dre. Nisrine Mokri, for sharing her ideas and knowledge in the field of precision and molecular medicine.*

- ☞ *I would like to express my deepest gratitude to Jordan Lovrod for improving the quality of this text and providing valuable comments and suggestions. Jordan is presently in her 3rd year of the Specialization in Computational Mathematics program through the department of Mathematical & Statistical Sciences at the University of Alberta. She was a student at Faculté Saint-Jean for the 2016-2017 academic year, where she completed her first year in a General BSc. She is now studying in France at Université Paris Nanterre through an Exchange Program with the University of Alberta.*

- ☞ *I would also like to express my gratitude to Cindy Paul – Copyright Specialist, Copyright Office and Kim Kemmer – Copyright Permissions Assistant, for their assistance to ensure that copyrights for the images we used in the document are not infringed.*

- ☞ *Finally, thank you to everyone with whom I had the opportunity to discuss this subject and exchange ideas. Your contributions are greatly appreciated.*

Everything that's already in the world when you're born is just normal.

Anything that gets invented between then and before you turn thirty is incredibly exciting and creative and with any luck you can make a career out of it.

Anything that gets invented after you're thirty is against the natural order of things and the beginning of the end of civilization as we know it until it's been around for about ten years when it gradually turns out to be alright really.

Douglas Adams –
The Salmon of Doubt, *2005*

Contents

Executive Summary

Many technologies, from smartphones to complex computer systems, are said to be exponential, meaning that their capabilities grow exponentially as a function of time. As a result of our lack of intuition regarding exponential growth, these technologies can seem to appear and improve unexpectedly. Studying exponential growth and its ensuing disruptions is therefore essential in anticipation of the significant societal changes that will result from exponential technologies.

The convergence of technology refers to the merging of various technologies to cumulate in advancements more powerful than the individual pieces. Technologies comprised of several exponential technologies have the capacity to improve faster than the exponential technologies of which they are made. This is due to the fact that each of the building blocks is itself growing exponentially, thus making the whole improve at an even faster rate. In this work, we will examine the growth of the individual technologies of 3D printing, artificial intelligence and nanotechnology. We will then explore the potential of the convergence of these technologies in the context of health sciences, manufacturing, and education.

As society changes alongside technology, our approach to several different topics may need to be reconsidered. For instance, education should be revitalized to take advantage of new technologies and prepare students to use said technologies in future careers. We will also need to determine how to best regulate technology, which is particularly challenging due to the decentralized nature of technologies. One possible strategy would be to create international research groups dedicated to finding solutions to societal problems caused by technological advancements and establishing functional international regulations.

The impacts of exponential technologies will bring about a world far different from the one we know today. Provided we will be able to overcome the challenges that new technological developments will bring, exponential technologies will enable us to accomplish feats that were never before thought possible.

1 Introduction

Technology and society are inherently interconnected - our needs, which are in part determined by society, dictate the demand for and the role of technology, and in turn, technological advancements catalyze changes in society. These changes are vast and significant in nature. In particular, the prevalence of technology leads to the democratization of services as well as the transformation or elimination of many job sectors.

Consider, for instance, the world before the Internet. Information was gathered from books that were either owned or borrowed. Watching a new movie required going to the theater, and learning involved going to class, exchanging with others, or reading. Messages were sent through the post office, and international communication was done via expensive long-distance phone calls. Today, this can all be done over the Internet through online libraries, Netflix, YouTube, email, or video-chatting. The Internet is a manifestation of the democratization of information - any individual with an Internet connection is able to access more information than ever before, from wherever.

For another example, consider the modern smartphone. A smartphone functions as a GPS, music player, phone, camera, computer, radio, flashlight, and credit card, all in one. The small device has replaced a number of standalone technologies, all while occupying significantly less space. For comparison, a modern cellphone is approximately 1300 times more powerful than the computers that took Apollo 11 to the moon [1]. In spite of this, few people give thought to the power of the device in their pocket.

Rapid technological advances perpetuate a state of constant change in our society, and although modern technology has changed the world for the better in a number of ways, it is not without its drawbacks.

2 The exponential growth of technology

Many technologies can be gathered under the banner of exponential technology, referring to the type growth curve they follow. Outside the field of technology, exponential growth is uncommon, with linear growth being a predominant model for describing many phenomena. Technology, on the other hand, is especially prone to exponential growth curves, meaning that it often improves faster than we believe possible. Exponential growth is markedly fast, being, in particular, considerably more powerful than the better-known linear growth.

To illustrate the difference between linear and exponential growth, we will visit a popular example often cited by Peter Diamandis, Executive Chairman of the X Prize Foundation and cofounder and executive chairman of Singularity University. Suppose that a linear step is defined to be precisely one meter long. Then after 30 linear steps, a total of 30 meters has been traveled. Now if we define an exponential step to be twice as long as the previous step, then the distance that can be traveled in a small number of steps far exceeds what one might predict. Indeed, starting with a first exponential step of one meter long, the distance traveled after 30 exponential steps is 1,073,741,824 meters - approximately equivalent to walking the circumference of the Earth 26 times. This examples goes to show how poorly a typical human can conceptualize exponential growth. We tend to think linearly, and thus exponential growth seems unintuitive.

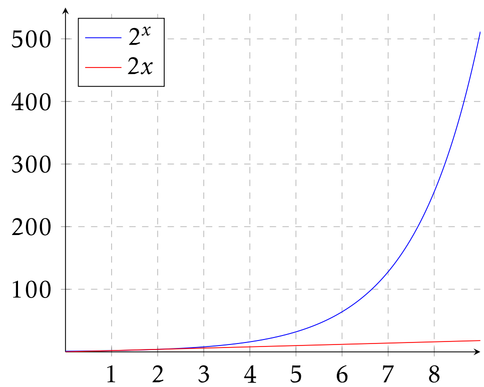

Figure 1: Comparing Linear and Exponential growth

One of the most renowned demonstrations of exponential technology was detailed by Gordon Moore in 1965 (Figure 2). Moore, a co-founder of Intel, predicted that the number of transistors on an integrated chip would double every two years, loosely signifying that computer performance would also double in that timespan [2].

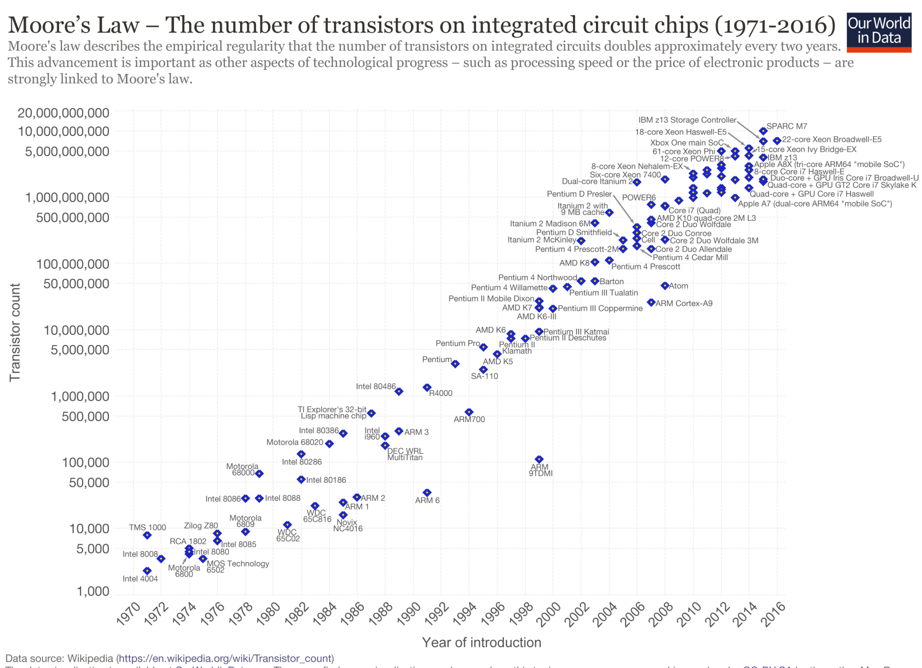

Figure 2: Moore's Law, Source [3]

However, Ray Kurzweil argues that Moore's Law is simply an example of "The Law of Accelerating Returns", which states that technology, in general, grows exponentially (Figure 3) [4]. In Kurzweil's Law, a positive feedback loop occurs as each new technological development is used to create the next, resulting in ever greater returns [5].

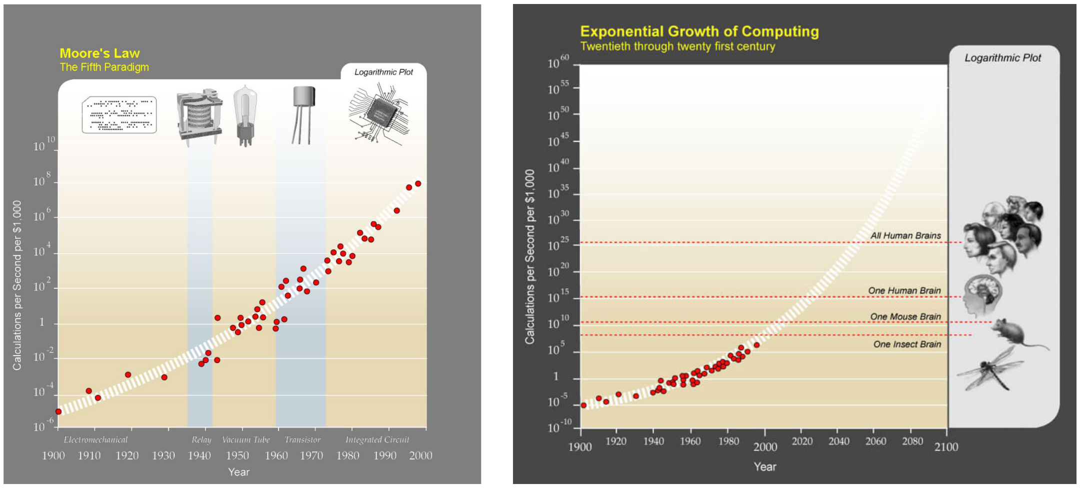

Figure 3: The Exponential Growth of Computing, Source [4]

Another exponential technological development is computer memory. Although less known than Moore's Law, Kryder's Law is equally important [6]. It states that the density of memory in magnetic drives doubles every 13 months, meaning that it grows even more quickly than that the number of transistors as stated by Moore's Law. Although it seems that the rate of improvement has slowed in recent years, it remains true that computer memory capabilities have improved tremendously (see Figure 4). The consequences of the improvements in memory technology have paved the way for the wealth of publicly available information. Consider that around the year 1990, it would have cost $10,000 to buy 1 GB of memory for a computer, but by 2014 the price had dropped to a mere $0.03 per GB. Although we rarely need the largest hard drives for personal use, businesses and governments must store large amounts of data, which could not be done efficiently without the recent developments in memory.

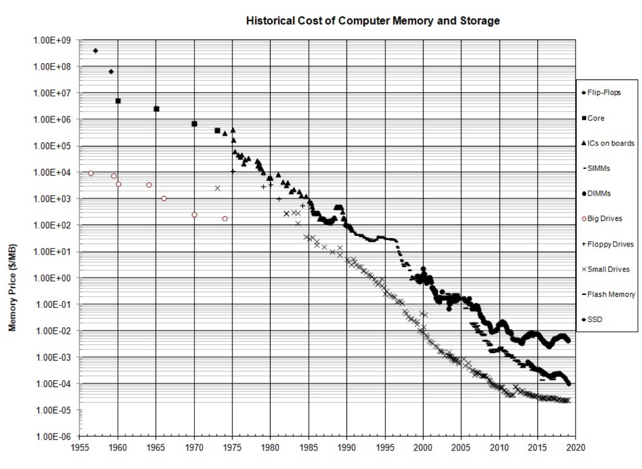

Figure 4: Kryder's Law, Source [7]

A standalone technology cannot improve exponentially indefinitely - each technology encounters a roadblock that is either impossible, or tremendously difficult and expensive to overcome. Even still, exponential growth can be experienced for impressive lengths of time simply by consistently making improvements to the underlying technology. For instance, computers were initially electromechanical. This technology was rendered obsolete by relays which were then replaced by faster and more reliable vacuum tubes. Vacuum tubes became increasingly small, making computers increasingly fast, until they were too small to preserve the vacuum [8]. The current challenge is that of the integrated circuit, which is characterized by Moore's Law. It has been long predicted that Moore's Law would eventually fail, and this point seems to be drawing near as the individual components of the integrated circuits are getting tremendously small, to the point where it is extremely difficult to keep the chip cool as the exponentially increasing number of transistors on each chip produces exponentially increasing quantities of heat. The future of computer chips could include using new compounds and materials in chips, or expanding chips into the third dimension [9]. Despite these challenges, because innovation drives profit, it is foreseeable chip companies will eventually develop a way to continue growing the computing capacity.

Because we have difficulty understanding exponential trends, exponential technologies seem to appear out of nowhere, leading to a lack of informed foresight to plan for new technologies. This is not due to the fact that technology is unpredictable. In fact, the growth of technology follows a highly predictable exponential growth curve. Rather, it is our persistent tendency to think in linear terms that gives the illusion of an unpredictable growth. It is therefore crucial to intentionally seek to understand the exponential growth of technologies such that we may be better prepared for future technologies.

3 The disruptive nature of exponential technologies

Due to the exponential growth rate of technology, there is a high turnover rate at which technology is replaced - technology becomes outdated extremely quickly, causing inconveniences for many individuals, companies, and industries as a whole. For this reason, exponential technology can be described as disruptive. One example of the disruptiveness of exponential technologies is found in the history of Kodak. In 1975, Steven Sasson invented the digital camera while working for Kodak as an engineer. Kodak already held a profitable monopoly on the film-photography market, so they were reluctant to embrace the new technology and instead continued with film photography alone [10]. The digital camera was, at the time, unable to produce high-quality images. Kodak failed to envision that the capacity of the digital camera would grow exponentially, quickly surpassing film-camera technology. At its peak in 1996, Kodak had 140,000 employees and a market cap of 28 billion dollars. However, their decision to not invest in the digital camera caused the company to lose their market share, and eventually file for bankruptcy in 2012. Kodak's choice to ignore the digital camera finished by effectively dismembering the company and pushing them out of the photography business [11]. As demonstrated by Kodak's bankruptcy, exponential technologies have the power to change an industry in a matter of a few years. A lack of exponential thinking can therefore jeopardize a company's long term-survival.

Unlike Kodak, Instagram effectively took advantage of an exponential technology (social media), leading to great success. In April of 2012 (the same year that Kodak went bankrupt), Instagram, which at the time had a mere 13 employees, was sold for one billion dollars. Both companies were in the photography business, but unlike Kodak, Instagram recognized and made use of a new exponential technology.

Exponential technologies enable new companies to satisfy demands more quickly than ever. This can cause new companies to take market share away from established companies, eventually causing them to be overpowered. In the case of exponential technology, established business often see the risks and experience the disruption of exponential technology, whereas the entrepreneurs often see the opportunity for investment and rapid success.

Many negative consequences emerge due to the prevalence of technology, one of which is its effect on employment. For one, the ubiquity of mobile phones and computers allows the workplace to invade the home, making it difficult for employees to disconnect from the stresses of the workplace. It has been found that employees with more work-related communication at home have increased conflict between work and family, and a higher risk of psychological distress and sleep problems [12]. This consequence, however, is not the extent of the negative impact of technology on employment. Two hundred years ago, about 70% of Americans lived on farms and worked in agriculture, yet today this percentage has dropped to a mere 1% [13]. This is mostly due to technological advancements in the farming industry, which has made each individual farmer significantly more efficient, therefore decreasing time needed to complete tasks, and thus the demand for farmers. This does not imply increased unemployment, but rather that people must now gain different qualifications to make a living.

The first jobs that computers replaced were those of the human computers: individuals who performed mathematical operations and approximations. When human computers could no longer compete with the speed and reliability of digital computers, many human computers turned to the new field of programming: as one profession died, another was born. Next to be replaced by computers were number-related and factory jobs. This pattern will continue repeating itself in various fields so long as technology continues to spread across disciplines and increase human efficiency.

It has been estimated that up to 47% of jobs in the United States could be computerized in the next 10-20 years [14]. Note that this is not a prediction of the number of jobs that will be lost, but rather a prediction of the number of jobs that computers and robots will have the ability to replace. These professions are unlikely to be eliminated until it is cost-effective to do so. Not all jobs will be equally affected by computerization. Studies have shown that jobs with higher wages and education requirements, for instance, are less susceptible to computerization. Another category of professions less likely to be computerized are jobs requiring high degrees of either social intelligence or creative thinking. Social intelligence is defined as the ability to understand others' feelings, persuade, or reconcile differences, and creativity is defined as the ability to innovate and problem solve. Jobs requiring high degrees of social intelligence or creative thinking include lawyers, social workers, politicians, artists, engineers, and musicians.

In contrast, many jobs will benefit from computerization. Technology can be developed to complement human workers' strengths - the social and create aspects of the job are left to be accomplished by the human. This can already be observed, as engineers and scientists use computers to perform intensive calculations, while results are left for human interpretation. In the near future, lawyers may have robotic paralegals or doctors artificially intelligent assistants.

Generally, we can expect that the time required to adopt modern technologies will be less than previous adoption periods, because in general, technology is being accepted more quickly [4].

Self-driving cars is another interesting example of modern technologies. A 2018 questionnaire found that at present, adults lack confidence in self-driving cars [15]. The study noted that most people trusted their own driving skills more than those of autonomous vehicles and indicated that they would like to have override controls if available. Although the respondents generally considered self-driving car as a convenience, most saw the technology as unnecessary and were unwilling to pay additional costs. Interestingly, age is not a significant factor in perspectives regarding autonomous vehicles. Only about half of existing studies conclude that younger populations are more likely to adopt the technology earlier, and some even conclude that older populations are more open to the technology due to the autonomy and independence it may grant.

4 Exponential technologies

4.1 3D Printing

3D printing, also known as additive manufacturing, was first conceived by Hideo Kodama in 1980 [16]. His technique, which would later be known as stereolithography, was to build the product using a liquid that hardened when exposed to ultraviolet light. Stereolithography was commercialized in 1984 by Charles W. Hull, and since then, the 3D printing industry has thrived. Apart from stereolithography, several additive manufacturing methods work by melting filaments (often made of plastic) in applied heat to form a solid object. Despite the fact that additive manufacturing has been in existence for the better part of four decades, it only entered general public awareness in the latter half of 2012. Recent improvements have not only increased the capabilities of additive manufacturing in industry, but also rendered the technology significantly more affordable for the general public [16].

3D printing has made a significant impact on the construction industry. The Chinese company WinSun alone has been able to print a 1100m^2 villa, a five-story apartment building, and 10 houses in just 24 hours [17]. That being said, additive manufacturing is not limited to the production of stationary structures. In 2014, Local Motors produced a 3D printed electric car, and in 2015, the United States Department of Energy unveiled a 3D printed Shelby Cobra (Figure 5).

Figure 5: 3D printed Shelby Cobra, Source: *U.S. Department of Energy*

Furthermore, engineers at General Electric Company were able to use additive manufacturing to create a small jet engine reaching up to 33,000 rotations per minute [18].

There are a number of advantages to industrial additive manufacturing. 3D printing improves worker safety since robots are used to complete the most dangerous tasks, enables companies to build more complex structures, and simplifies the task of prototyping and proofing. To expand on the latter, with 3D printing, testers can manufacture parts on-site as soon as modifications are made, meaning that external departments need not be involved, saving both time and labor costs. In fact, as much as 57% of 3D printing in industry is in the domain of new products development [19]. Furthermore, additive manufacturing minimizes waste by building structures in such a way that only the necessary materials are used. This allows construction companies to minimize costs and reduce environmental impact [16].

Desktop 3D printers are also gaining in popularity. Approximately 500,000 desktop 3D printers were sold in 2017, a 50% increase from 2016, and a 200% increase from 2015 [20]. This rapid growth in sales can be attributed to recent increases quality, usability, and affordability. In fact, in 2017, a 3D printer would sell for just 10% the price of a comparable printer sold in 2013 [20]. Although desktop printers are inferior in quality to industrial printers, they still have the capacity to print highly technical objects such as prosthetics.

Due to the democratization of manufacturing though 3D printing, more power will be placed in the hands of the consumers. As consumers gain the ability to produce their own goods, manufacturers will be faced with the challenge of offering highly customizable products in order to compete with the customized products that a consumer would be able to manufacture in their own home [16]. That being said, it is unlikely that 3D printing will eliminate the current manufacturing industry, as it allows for low-cost bulk creation of products that do not need to be custom-fit for each user [21].

4.2 Artificial Intelligence

Artificial intelligence (AI) has already been successfully applied to a variety of tasks, such as machines playing board games, robots displacing objects, and image recognition. Furthermore, AI is the key component in the present development of self-driving vehicles. Since AI is dependent on the computing power it can access, exponential increases in computing power will lead to exponential increases in AI capability. As a result, we can expect AI to soon become a practical and indispensable aspect of society.

One of AI's first major accomplishments was learning to play chess. Chess was considered an excellent project for AI, as the game has a well-defined set of rules (permitted moves and game structure) and clear winning conditions. These factors reduce the complexity of programs running on digital computers working in binary [22]. Furthermore, Chess seemed to be an appropriate level of difficulty - challenging, yet not impossible. Finally, because Chess requires a certain level of thought and planning, success on the project would lend credence to the idea that the machine was in fact intelligent.

Deep Blue, a chess-playing computer developed by IBM, was the first computer system to win a match against a reigning world champion. In short, the algorithm involved generating a series of possible moves in a tree-like structure, evaluating the resulting board positions, then "pruning" the tree by removing the branches that produced the least desirable results. It would then search deeper into the branches that provided desirable results, before eventually committing to the best move based on its evaluations [23, 24]. Deep Blue would evaluate each possible move based on a set of pre-defined rules that were hard coded into its system. In other words, Deep Blue did not learn to play chess - it was told how to play chess. Nonetheless, the algorithm proved to be effective as in 1997, Deep Blue beat Gerry Kasparov, the reigning world champion in a six-game match.

In 1950, Claude E. Shannon proposed that the greatest weakness of a hard coded program like Deep Blue is its inability to learn from mistakes [22]. Any improvement in performance would necessarily be due to adjustments made by the programmers. Shannon proposed that it would be better if the program was "self-improving", meaning that the program is able to modify the decision making process based on experience [22]. This idea eventually developed into the concept of neural networks and deep learning. Neural networks have allowed AI to move beyond hard coded systems such as Deep Blue - we can now program systems that can independently improve strategy through experience.

Go is a game played by placing a series of stones on a grid (usually 19x19), and eluded AI researchers until the development of neural networks. The number of possible moves made it impossible to program an efficient AI for Go with the same "brute force" approach as Deep Blue. It was long thought that human intuition was necessary for Go, unlike for Chess. However, in March 2016, an AI machine called AlphaGo used neural networks to teach itself how to play Go and was able to beat Lee Sedol (ranked fifteenth worldwide at the time) 4 games to 1 [25]. A neural network was able to replace and defeat human intuition, demonstrating how deep learning can effectively go beyond traditional AI. This is due to the fact that neural networks can generalize situations they have previously encountered, allowing them to make reasonable choices in new situations [26].

Neural networks are also gaining proficiency in image classification. Google's Inception-v3 has managed to correctly identify more than 80% of images in one guess and 96% within five guesses [27]. This is directly applicable to self-driving cars, which work by scanning the road with a number of sensors, building a model of its surroundings, and using the model to determine the best course of action. This entire process is repeated several times per second [28]. As image classification improves, these systems will be able to assess their surroundings more accurately and thus make safer decisions.

Another application of AI is IBM's Watson. Watson is an AI that can understand and answer natural-language questions. Watson acquired data by "reading" a number of sources including dictionaries, encyclopedias, thesauri, news articles, databases, and literary works [29]. Although Watson was not initially successful, it rapidly improved and was eventually given the challenge of a game of Jeopardy against the two greatest human players. Watson won the game, earning a score greater than the two other players combined [30]. Watson's victory symbolized significant progress in the field of Natural Language Processing (NLP), and AI as a whole.

Natural Language Processing poses a significant challenge for AI systems, principally because language is an unstructured data source [31]. Natural language is complex and nuanced, thus complicating the task of correct interpretation. For computers to interpret a natural language sentence, they need to understand all accents and dialects of the language and consider nuances such as tone, humour, and human error (such as misuses or misspellings of words). Computers process language by partitioning sentences into words, then identifying the roles each word might play (verb, noun, etc.). Next, the computers attempt to derive meaning from the sentence by piecing together each of the fragments. In particular, this is applicable to virtual assistants in smartphones and smart speakers, such as Google Home or Amazon Echo.

4.3 Nanotechnology

Nanotechnology is the subset of technology whose size is best measured in nanometers (each nanometer is equivalent to one billionth of a meter). Several issues arise with nanotechnology that are not encountered at larger scales. First, since there does not currently exist methods that allow us to assemble individual atoms or remove small groups of atoms from larger objects, traditional fabrication techniques cannot be applied. Moreover, the building blocks of nanotechnology must be thought of in terms of atoms and compounds instead of standard planks and beams, further complicating conceptualization, design, and fabrication. Despite these challenges, however, nanotechnologies have immense potential, especially in the field of computing.

One of the most prevalent nanoscale technologies is in computer chips, which have recently been fabricated with feature sizes as small as 10 nm [32]. At the lowest level, computing is accomplished by the movement of electrons. The distance that the electrons must travel is dependent on transistor size. In other words, the smaller the transistor, the faster the computer. Minimizing transistor size has therefore been a primary means of improving computer speed for at least the last 40 years. As such, nanotechnology is at the heart of the technological revolution sparked by computers and the Internet.

Transistors are often made by means of a process known as photolithography. Although expensive, so long as transistors continue to offer increasingly fast means of computing, the process of photolithography will likely persist. As mentioned earlier, however, no technology can improve indefinitely - at some point, this technology too will reach a limit in efficiency. This is in part due to the fact that as transistors become smaller, electron behavior can no longer be described by Newtonian mechanics, and rather begin to follow quantum mechanics. This produces seemingly unpredictable effects when viewed from a macro-scale perspective. These challenges have motivated the exploration of alternatives to transistors, one of which is carbon nanotubes [33].

Carbon nanotubes, which are hollow tubes made of graphene, have the potential to outperform transistors in computer chip applications. These tubes often range between 5nm and 50nm in diameter and can be up to 55cm in length [34]. Carbon nanotubes are possibly the strongest material known to date, being, in particular, considerably stronger than steel [35]. Furthermore, both heat and electricity can be efficiently conducted along the tube's length, however neither can be conducted along the tube's radius [36]. Like transistors therefore, these tubes can be used to effectively direct the movement of electrons. In fact, carbon nanotube technology promises to be of significant importance in future computing [33].

Carbon nanotubes can presently be found in sporting equipment, car parts, water filters and protective coatings, and are also used as an additive to provide other materials with strength, or electrical and thermal properties [37]. Beyond these applications, however, the use of carbon nanotubes has historically been prohibited by price; carbon nanotubes cost nearly nine dollars per gram on the consumer market. This price, however, pales in comparison to those seen before the turn of the century. On top of price, it has also proven to be difficult to fabricate nanotubes with consistent sizing. For many applications such as in computing, it is crucial that the nanotubes dimensions be consistent and precise. Researchers continue to work to develop new methods for production, including a recent technique that fabricates carbon nanotubes using atmospheric carbon dioxide. In this way, the repeated production of carbon nanotubes could potentially contribute to the reduction of greenhouse gases [38].

When Feynman was first exploring the concept of nanotechnology, he wondered whether it would be possible to fabricate objects from individual atoms through mechanical manipulation, and noted that the ability to image individual molecules could lead to significant breakthroughs in biology and chemistry [39]. Technologies that allow for direct measurements at the resolution of a single atom already exist and have been used in experiments for small-scale atom manipulation. These techniques are collectively known as Scanning Probe Microscopes (SPM) [40]. One method, known as the Scanning Tunnelling Microscope (STM), consists of a sharp metal stylus moving across a given surface, though not touching it. A voltage is applied, and when the STM is within one nanometer of the surface, electrons can tunnel into the surface. At any greater distance, however, the electrons do not transfer. Through this process, an image can be formed of the surface. This technique is not without fault - it only works for conductive surfaces, for instance. When the STM was first tested, it was noted that some atoms were moving, which confirmed that the STM has the ability to move individual atoms. Scientists and researchers have demonstrated this capacity with a variety of designs, among which include a quantum corral which traps a standing wave comprised of electrons [41], and a two-minute-long short film called "A Boy and His Atom" (a single frame is pictured in Figure 6) [42].

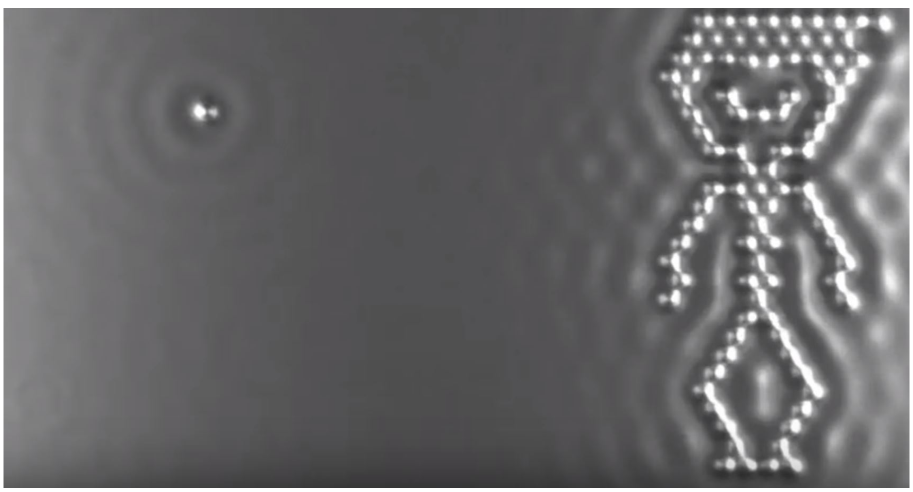

Figure 6: A Frame from "A Boy and His Atom", Source [42]

5 The convergence of technology and its societal impacts

In previous sections, we saw the power of 3D printing, AI, and nanotechnology as standalone advancements. We will now explore the impact that these technologies can have when considered in cumulation with other technologies. We define this as the convergence of technology: the combination of various technologies that results in even more influential advancements than might be expected from the individual parts.

5.1 The medical field

As we have seen, computer technology improves at an exponential rate, and Moore's Law states that the number of transistors in a computer chip doubles every two years, essentially doubling computing power. We have yet to discuss the effect of improvements in associated technologies, such as software programs, computer memory, and the Internet. It is the convergence of these technologies that gives us the ubiquitous computer which is central to many facets of modern life. For an in-depth example of the convergence of technology, we will examine how 3D printing, AI, and nanotechnology will, in conjunction with other select technologies, change the medical field.

In the medical field, organ transplants and prosthetics are often crucial, but it can be challenging to find an organ or prosthetic that matches the patient's needs. The 3D printer's ability to produce highly customized items using a number of different materials will begin to play a crucial role in the medical field. 3D printing an organ can be accomplished in several different ways [43]. One method involves printing a scaffold made of select polymers, after which the scaffold is manually seeded with the patient's cells [43]. The cells then multiply and eventually form the organ. Alternatively, a hydrogel can be printed and immediately covered in cells by the 3D printer to form the organ [44]. Both methods allow doctors to transplant organs grown out of the patient's own cells, thus decreasing the risk of rejection. 3D printed prosthetics are also becoming increasingly common [45]. For example, Enabling the Future is an online group creating new open-source 3D printed prosthetic hands and arms. By pairing individuals with 3D printers with patients in need of prosthetics, patients are able to receive customized prosthetics free of charge [46]. Furthermore, researchers are exploring how to create crowns, dentures, and orthodontic devices with 3D printers, which will in turn allow individuals to access dental treatment more quickly and for a lower cost [47].

The medical field is becoming increasingly reliant on data, which in turn creates demand for efficient means of screening large amounts of data. Herein lies the role of AI in healthcare. Health trackers are gaining popularity, and digital doctors are already in initial stages of development. The latter use sensors and AI to help individuals self-diagnose and make informed decisions on an appropriate course of action [48]. As these technologies improve, researchers will have access to invaluable information regarding health records, symptoms, and lifestyle. AI could then be used to investigate the data to identify patterns and suggest correlations between potential causes and effects. This could be used to further preventative medicine, allowing individuals to avoid certain risk factors and address others in order to preserve their health. The use of technology to record an individual's life comes with legitimate questions concerning privacy, but the benefits in the field of medicine could be worth the risk. Take for instance Watson, the natural language processing AI created by IBM. After winning Jeopardy, Watson transitioned to the medical field [49]. Watson's success in NPL is related to the development of artificial doctors - apps that, when given a description of symptoms, are able to provide personalized recommendations. Furthermore, Summit, a supercomputer completed in 2018 by Oak Ridge National Laboratories, was created with the purpose of efficient deep learning through neural networks [50]. One intended project for this computer is to analyze Americans' genetic code and health records to search for relationships between cancer and genetics. Speaking of cancer and AI, several studies have shown that when diagnosing certain types of cancer, a doctor working with AI identifies 7% more cases than the same doctor working alone [51].

Nanobots also have the potential to be significant medical forces in the near future. Robert A. Freitas Jr. has published a number of papers detailing designs for various nanobots, in particular the microbivore and the respirocyte. The microbivore is an artificial white blood cell small enough to fit through the smallest human capillaries. It works in virtually the same way as natural white blood cells - it locates unwanted intruders in the body and disassemble them, thus eliminating the threat. It is speculated that these nanobots could destroy an unwanted cell or bacterium every 30 seconds, regardless of antibiotic resistance. Similar nanobots could theoretically be directed to destroy cancer cells or circulatory obstructions, effectively eliminating disease and other health problems more reliably than natural white blood cells [50, 51]. Similarly, respirocytes are artificial red blood cells able to contain sufficient quantities of oxygen for an average individual to heavily exert themselves for more than ten minutes, or be immersed in water for up to 4 hours [52]. Both the microbivore and the respirocyte are theorized to be significantly more efficient than their natural counterparts. That being said, significant developments in nanoscale fabrication methods will need to happen before prototypes for these bots can be made. Additionally, Kurzweil suggests that nanobots could also be used for surgery or imaging. In essence, a nanobot could perform non-invasive surgery by being injected into the bloodstream, and reconstructing a given bone or organ using materials in the bloodstream. This would likely be accomplished using techniques borrowed from additive manufacturing [4]. He also suggests that nanobots could be used for high-resolution imaging, allowing scientists and doctors to better understand the inner workings of the body and brain, thus providing valuable information for understanding disease and developing treatment. All in all, it seems nanobots could bring significant improvements to our healthcare system, supposing we can find methods to create them.

5.2 Driverless Cars

For another example of the convergence of technology, we will study the effects that driverless cars could have on our society. Autonomous Vehicles (AVs) are predicted to emerge sometime between 2030 to 2050, although it remains unclear if these predictions are based on exponential or linear thinking [53, 54]. In any case, it will be a number of years yet before AVs outnumber traditional vehicles on the road. AVs vary in levels of autonomy; some AVs only perform minor tasks for the driver (such as parking), others perform all task but require human assistance in difficult conditions, while still others have no controls for drivers whatsoever [55]. When referring to AVs, we tend to picture the latter, but the partially self-driving vehicles must also be discussed, especially since their emergence will precede that of the fully self-driving vehicles.

It has been hypothesized that an average autonomous vehicles could replace up to 11 of its traditional counterparts [56]. This is due to the fact that while traditional vehicles often remain parked, AVs would be able to maximize their time in operation by transporting a number of different passengers. Correspondingly, the very concept car ownership may be remodeled. As AVs become more common, people may flock to ride-sharing companies that utilize driverless cars. Said companies would operate much like taxi companies today, the main difference being that a driver would no longer have to be paid, allowing for lower costs. This would lead to a reduced demand for standard taxis and ride-sharing platforms that use traditional vehicles. Another advantage AVs would provide is the opportunity for individuals to increase productivity by making use of their travel time for tasks other than operating the vehicle.

As implied above, jobs involving vehicle operation will be threatened when the cost of purchasing and maintaining an AV becomes less than the cost of paying a driver. The demand for workers in industries such as trucking, taxi driving, and bus driving will decrease as AI will supplant the need for drivers and undercut the associated costs. This will have repercussions across the economy as many jobs will be eliminated, leading to a risk of elevated unemployment and increased demand for retraining.

Another necessary consideration is safety. AVs are expected to become safer than traditional vehicles operated by the average driver, especially when taking into account drivers' inability to recognize their own weaknesses on the road. Studies have shown that despite the fact that roughly 80% of the driving population considers themselves above average, about 90% of crashes are caused by human error [57]. This failure to objectively self-asses driving skills prevents drivers from correcting themselves and maximizing road safety. Moreover, drivers can be distracted by cellphones, GPS devices, and other passengers, or can be impaired by alcohol, drugs, or lack of sleep. In contrast, Waymo, a leader in AVs and a subsidiary of Google, has used self-driving cars to drive over five million miles since 2009, with only one accident being caused by the self-driving technology [58]. Besides, the use of AVs would eliminate the safety risks associated with distracted and impaired driving. Since the performance of AVs is solely dependent on ability and hardware, driverless cars will contribute to a decrease in collisions and improve safety for everyone on the road.

Furthermore, driverless vehicles would grant increased independence to individuals who are unable to drive. In particular, this would include children and young adolescents, seniors, and people medically unable to drive. This would have a particularly poignant effect in smaller cities and towns where public transit is less developed. Access to AVs would therefore allow these individuals to gain in independence since they would no longer need to rely on other drivers for transportation.

A potential social side-effect of AVs may be a decreased prestige associated with driver's licenses. In many societies, a driver's license can serve as a significant milestone on an adolescent's journey to adulthood. That being said, as AVs become more commonplace, driver's licenses will no longer be as necessary or sought after. This milestone could be replaced by gaining the authorization to ride alone in an AV, which will likely be much more attainable than today's driver's licenses.

Several ethical questions arise when contemplating the use of AVs. As previously stated, about 90% of collisions are due to human error. Conversely, this suggests that approximately 10% of collisions cannot be avoided by the use of AVs. Non-human causes for collisions may include poor road or weather conditions, unpredictable animals, or mechanical failures. In such cases, the AI behind the car is responsible for making decisions. Consider a situation wherein a self-driving car must decide between harming a passenger and a pedestrian: whose safety should be prioritized? Now suppose another situation wherein a collision is bound to happen between two vehicles: which vehicle should suffer the greater damage? The answers to these questions must be thought through and coded into the AV. In contrast, with driver operated vehicles, reactions in such situations are often determined by sheer human reflex. Some people argue that such contrived scenarios ought not to be thought out at length because they are unlikely to occur [59]. Despite this criticism, however, answers to these questions must still be coded into the AV systems.

Given that human-driven cars will be more collision-prone than AVs, another ethical question worth scrutinizing is the banning of human drivers. This problem will become especially significant if there comes a point when all AVs are assuredly safer than all human drivers. From this ensues a dichotomy between the right to drive and the right to safety.

As it stands today, in the event of a collision, the responsible driver is held liable for harm to people and property. When the driver is an AI however, the question of liability becomes altogether more complex. Possible responsible parties include the manufacturer, the human occupant, or the individuals who programmed the AI in the first place. Depending on the answer, car companies may delay innovation in order to avoid potential liability surrounding AV collisions. This delay would ironically postpone the heightened safety that AVs will bring to our roads, and many human-error collisions that could be prevented by AVs will continue to occur [60]. Some studies suggest that a strict liability is the best approach to driverless vehicles, wherein the owner or operator of an autonomous car is always held liable, regardless of whether or not the harm was their fault or intent [61].

Another challenge to overcome is a settlement on international regulations surrounding AVs. If neighboring countries were to have different laws vis-a-vis self-driving car standards, then a legal AV on one side of the border may be illegal on the other side. As with all international regulations however, this process would likely be lengthy and complex, and could cause delays in the popular use of driverless vehicles.

Finally, computerized vehicles are at risk of hacking. This is analogous to the threat of hijacking with traditional cars. Hackers could potentially take control of AVs and harm the occupants of the vehicle, individuals outside the vehicle, or buildings and other property. The additional danger lies in the fact that an AV hacker need not be present in order to cause damage. This eliminates the risk of bodily harm for the perpetrator and can complicate the task of determining criminal intent and tracking those responsible for the crime.

The complex assortment of advantages and problems brought on by the AV refers back to the idea of disruptive innovation: an individual can choose to act based on the challenges or the opportunities that the new technology brings. So long as the obstacles surrounding AVs can be resolved, driverless vehicles will empower individuals with difficulties accessing transportation, improve road safety, and reduce the number of vehicles on the road. Ultimately, AVs highlight how the profound impact of exponential technology will not be limited to specialized scientific fields, but rather will reach society as a whole.

5.3 Education in a technology-centered world

As technology builds on existing technology and becomes integrated in society, we will witness several changes to the world we know today. Technology grows in an unintuitive, exponential fashion, meaning that the future is closer than we may think. This is of particular importance when regarding the education of young children. In preparation for these major societal changes, instead of focusing education around the world that is today, education should prepare the coming generations to be living in a world centered around technology.

Classroom teaching, even for young students, is often lecture style. Although this has long been thought the best means of education (short of a private tutor), it is no longer efficient nor pertinent given our new context of technology. In any classroom, there is much variation in students' style of learning, strengths, and ability to retain information. Efficient learning occurs when the students themselves are able to choose their method of learning, content, and pace. The reality of a traditional classroom setting is that only a select few students are provided with an optimal environment for learning, while some students get left behind, and still others are slowed by the group.

Modern technology delivers mechanisms for educators to provide more effective classroom settings. The Internet allows students to access a variety of recourses that complement their classroom-based learning. Sites such as Wikipedia provide repositories of information which can be accessed free of charge by anyone, anywhere. Specialized websites allow students to learn about specific subjects. For instance, sites such as Khan Academy provide lessons and practice material to teach students math from kindergarten through to post secondary level. Other specialized resources include entire online courses, especially for tech-based subjects like programming and computer science. Similarly, online language learning services allow people to learn languages of their choice at their own pace. The Internet allows children to acquire knowledge geared specifically towards their interests.

That being said, strictly online resources are not always better than traditional classroom. When Massive Online Open Courses (MOOCs) were first launched, some creators predicted that within a century, nearly all post-secondary institutions would be replaced by MOOCs [62]. However, the MOOCs, proved to be not nearly as successful as the creators thought. Although cutting edge in technology, the courses were behind the times in pedagogy. The MOOCs were lecture based, there was no peer interaction, and deadlines were not enforced, leaving students unmotivated. The MOOCs had a completion rate of less than 13%, and an even lower passing rate. Most MOOCs have been described as not meeting their target demographics. Typically, when these courses are designed it is to allow individuals without access to traditional schooling or with other disadvantages, but in fact individuals who successfully complete a MOOC tend to be those who would be learning about the subject regardless of whether or not they had easy access to a free course, whereas individuals who were reliant on the medium had even lower pass rates than the average [63]. It turns out that members of the target demographic generally do very poorly on these classes, earning lower grades than similar individuals enrolled in traditional classes covering the same material. Although MOOCs were an interesting concept, they will need to be significantly revamped to provide students with a modern, efficient method of education.

As a teenager, Jack Andraka, was able to take advantage of the Internet to develop a pancreatic cancer screening test in 2012 [64]. Although the merit of his work has been challenged by the scientific community [65, 66], it still serves as an example of how modern technology provides motivated individuals with the resources needed to push beyond what they would typically learn in a classroom. The education system should be fostering this type of motivation and ability to gather information, which would enable every student to study (and even contribute to) the fields of their interests.

The Center for Curriculum Redesign (CCR) published a book detailing their recommendations vis-a-vis the rapidly evolving educational needs. They posit that education in the 21st century should focus on teaching students creativity, critical thinking, communication and collaboration [67], claiming that these skills are among those which will be necessary for success in the future. It was also suggested that schools incorporate character education in order to prepare individuals to face the significant ethical issues that arise with modern technology. One way to incorporate these skills into the existing education system would be a transformation from lecture-based instruction to project-based discovery. As students work in groups to solve problems, they would need to use each of the required skills - creativity, critical thinking, communication, and collaboration in order to design a solution. To complete the project, students would need to propose new ideas, develop a plan, and then implement their solution. This would require significant time, thus meaning a re-prioritizing of subjects and a reallocation of classroom time would be in order. It is at this point that further study is necessary, as curriculum designers are not likely to entertain the idea of reducing content. With the Internet, however, memorized knowledge is becoming less important. We have access to information on all subject matters, even if we do not actively retain every bit. With this in mind, significant portions of knowledge-based curriculum could be supplanted by appropriately tailored problem-solving education.

This is not to say that students should be learning without guidance. Teachers should help students gain a solid understanding of the fundamentals in various subjects, especially those which are of growing in importance, such as numeracy, literacy and technological literacy. Long-term, AI teaching could replace the need for human teachers, leaving educators with tasks such as conflict resolution, encouraging focus, and keeping order in the classroom. AI would be able to analyze each student's strengths and weaknesses, as well as the manner in which they most effectively learn, in order to provide each student with a personalized education tailored to address their weaknesses and bolster their strengths. A system of this style could allow students to grasp the fundamentals of important disciplines more effectively.

Furthermore, the CCR suggests that the new curriculum should emphasize flexibility. Due to the exponential improvements of technology, the lifespan of a new curriculum centered around technology will itself exponentially decrease, meaning that it may not be possible for curricula to be overseen by the government in the same extent as it is today. Otherwise, curricula could refrain from the inclusion of specific knowledge, instead focusing on the general skills and attitudes that should be mastered upon students' entry into post-secondary or the workforce.

The exponential improvement of technology will most certainly have a significant impact on society and the skills required for success. The education system itself will therefore need to adapt in order to assist students in acquiring skills more adapted to a technology centered world.

5.4 The risks of modern technology

We have seen examples of how both opportunity and danger lie within each technological advancement. One of those dangers is that powerful individuals or companies may be able to exploit new technologies to further their agenda. Often, issues regarding abuses of power by means of technology is addressed by the adoption of laws, such as limits on cloning or restrictions on genetically modified/engineered food products [68,69]. That being said, regulations can take years to establish - much slower than the growth of technology itself. Being in an age of technological growth, governing bodies will be overwhelmed by the number of issues that need to be addressed with respect to technology. In this section, we will briefly examine a few of these issues.

Online privacy is a current topic of debate in many federal governments around the world. Websites are able to collect and sell data from site visitors, despite the fact that American citizens have demonstrated their opposition to online tracking and the selling of personal information [70]. Recent scandals involving Facebook and Cambridge Analytica have revealed the extent to which information collected online can be exploited. In short, Cambridge Analytica (CA) had an app on Facebook that purported to be a personality test, and collected the user's information along with the information of their friends [71]. Large portions of the subjects did not personally consent to the collection and use of their personal information. The collected data was then used in attempt to influence events like Brexit and the 2016 US presidential election [72]. Due to the lack of legislation, this issue fell into a legal gray area. A short while later, CA went bankrup. Although it seems that they are now operating under a different name [73]. The fact that a company was easily able to collect information on millions of people for the purpose of influencing political events caused significant discomfort and anger.

Privacy of health and lifestyle information is yet another significant issue. As discussed earlier, personal health trackers intended to be used as tool to promote healthy living are becoming increasingly common. This gives companies access to specific health information from large sample populations - information which has tremendous monetary value. On one hand, access to large amounts of information regarding lifestyle patterns and health could help researchers identify risk factors for many diseases. Ideally, the data should be anonymized, allowing for access to valuable information without invading user privacy. However, if there is a data breach, an individual's private health history could become public. There are currently laws in place regarding this privacy issue in both the United States and Canada [74, 75]. As the collection and transmission of healthcare-related information becomes commonplace, this risk will be amplified.

Part of the difficulty in regulating technology can be demonstrated by the attempts to restrict/ban Uber and other modern ride-sharing companies. Uber is a "platform that connects drivers and riders", and is therefore a direct competitor to taxis and public transportation [76]. As Uber branched into new cities, regulatory transportation boards attempted to ban the company, but they were often able to argue that they were not a taxi service, but rather a referral service, meaning that the established taxi rules did not apply to them [77]. Although rules were eventually established, it took several years to overcome this challenge. In the end, Uber is still in operation (although they have stopped serving some cities), and most of the regulatory threats have been avoided. Uber was not above the law, but they did consider themselves outside of the domain of applicable law. This is similar to modern technologies in that there are not pre-existing applicable laws.

Although exponential technology allows for the democratization of information, manufacturing, and other services, it threatens to increase economic inequality. Technology often enriches select individuals who possess the advantage of an exponential mindset, which comes through excellent education. Singularity University (SU) is a small university and research centre in Silicon Valley. It prides itself on teaching its students how to seize exponential opportunities and was, in fact, founded by Peter Diamandis and Ray Kurzweil, paragons of exponential thinking. Most interestingly, the school is sponsored by industry leaders such as Google, IBM Watson, and Deloitte. These companies sponsor the organization at the epicenter of exponential thinking, suggesting that they themselves embrace the concept. If these companies embrace the exponential mindset, then it seems reasonable to infer that exponential thinking can bring success, as these companies are leaders in their fields. In this way, exponential technology reinforces the troubling link between costly education and monetary success.

It is very difficult for a single organization, be it a city, state, province or even country, to restrict and regulate effectively. To efficiently address the exponential increase of technology and its associated risks, international committees should be formed. In the past, countries have collaborated to solve issues through meetings like the Paris Climate Conference or the G7 conferences. Such alliances address important issues such as pollution and security threats, but unfortunately there is not enough conversation surrounding issues regarding technology, which will only increase as technology itself grows exponentially.

6 Conclusion

Technology is a driving force behind the evolution of society, and as technology grows exponentially, societal changes will be expedited. We can already observe the capacity of technology, especially widespread computing and the Internet, to make cause societal change. Many of these technologies have opened up brand new worlds to explore, allowing some improvement of the human condition.

Much of technology can be said to be exponential, referring to the growth curve it follows. Humans tend to lack a strong exponential intuition, so in turn exponential technologies seem to have a habit of sneaking up on us and catching us unaware. Linear mindset. Adding to the rapid growth is the effect of the convergence of technologies, allowing for growth even faster than exponential.

These technologies are often a double-edged sword, allowing improvement but also challenging the existing structures and way of life. This is why they are considered disruptive technologies, as they disrupt the current paradigms and put in place new constructs. To survive in a world filled with disruptive technology, everyone, both companies and individuals, needs to be prepared to change how they act very quickly and either be willing to adopt the new way of life they present or lose relevancy.

Despite the benefits conferred on society by these technologies, they also bring with them previously unknown risks and vulnerabilities. Whether we consider regulation of technology, hacking computer systems, or data scandals, it is clear that the new technologies not only fix problems, but also create a few of their own. We want continued innovation; we do not want to live in a stagnant society. We must be able to protect everyone in the face of the new risks presented by technology, while still fostering innovation. To best regulate technology, we need to collaborate. We need public awareness of the benefits and risks presented by new technology so that the public at large understands what is at stake. Additionally, because of the decentralized nature of many new technologies, we need universal regulation, built by international or even supranational governing bodies, such that we may maximize the benefits while minimizing the negative impacts.

The exponential improvement of and convergence of technology will transform virtually all aspects of society in a short period of time. Although exponential technology brings a new set of risks that must be addressed, exponential technology will also enable humanity to reach a new level of discovery, health, efficiency, and safety that certainly makes the risks well worth facing.

References

[1] A. LoBaido. Apollo lunar missions' computing power vs. your iphone, 2016.

[2] G. E. Moore. Cramming More Components Onto Integrated Circuits. *Proc. IEEE*, 86:82–85, 1965.

[3] M. Roser and H. Ritchie. Technological Progress.
Data source: `https://en.wikipedia.org/wiki/Transistor_count`.

[4] R. Kurzweil. *The Singularity is Near: When Humans Transcend Biology*. Penguin Group, 2005. Sources for images: `https://en.wikipedia.org/wiki/The_Singularity_Is_Near#/media/File:PPTMooresLawai.jpg`, and `https://commons.wikimedia.org/wiki/File:PPTExponentialGrowthof_Computing.jpg`

[5] R. Kurzweil. *The Law of Accelerating Returns*.
http://www.kurzweilai.net/the-law-of-accelerating-returns, 2001.

[6] C. Walter. Kryder's Law. *Sci. Am.*, 293:32–33, 2005.

[7] Disk drive prices (1955-2013).
http://www.jcmit.com/diskprice.htm.

[8] W. Pottenger, P. Freiberger, D. Hemmendinger, and M. Swaine. Computer.
https://www.britannica.com/technology/computer.

[9] R. Courtland. Transistors could stop shrinking in 2021.
https://spectrum.ieee.org/semiconductors/devices/transistors-could-stop-shrinking-in-2021.

[10] C. Mui. How Kodak Failed.
https://www.forbes.com/sites/chunkamui/2012/01/18/how-kodak-failed/.

[11] J. Estrin. Kodak's First Digital Moment. *Lens Blog*, 2015.

[12] S. Schieman and M. C. Young. Are communications about work outside regular working hours associated with work-to-family conflict, psychological distress and sleep problems? *Work Stress*, 27:244–261, 2013.

[13] K. Kelly. Better Than Human. *Wired*, 21, 2013.

[14] C. B. Frey and M. A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technol. Forecast. Soc. Change*, 114:254–280, 2016.

[15] M. Kyriakidis B. van Arem S. Nordhoff, J. de Winter and R. Happee. Acceptance of Driverless Vehicles: Results from a Large Cross-National Questionnaire Study. *J. Adv. Transp.*, pages 1–22, 2018.

[16] The Free Beginner's Guide - 3D Printing Industry.
https://3dprintingindustry.com/3d-printing-basics-free-beginners-guide.

[17] M. Starr. World's first 3d-printed apartment building constructed in china. https://www.cnet.com/news/worlds-first-3d-printed-apartment-building-constructed-in-china/.

[18] GE Engineers 3D Print A Mini Jet Engine. https://www.ge.com/reports/post/118394013625/these-engineers-3d-printed-a-mini-jet-engine-then/.

[19] Sculpteo – the state of 3d printing. *Sculpteo's 4th annual report on 3D Printing and Digital Manufacturing*, 2018.

[20] S. Adams. Half million 3d printers sold in 2017 - on track for 100m sold in 2030. https://3dprintingindustry.com/news/half-million-3d-printers-sold-2017-track-100m-sold-2030-131642/, 2018.

[21] M. Timms. 3D printing cannot completely replace traditional manufacturing, say experts. https://www.worldfinance.com/infrastructure-investment/3d-printing-cannot-completely-replace-traditional-manufacturing-say-experts, 2014.

[22] C. E. Shannon. XXII. Programming a computer for playing chess. *Philosophical Magazine*, 41:256–275, 1950.

[23] M. Campbell, A. J. Hoane, and F. Hsu. Deep Blue. *Artif. Intell.*, 134:57–83, 2002.

[24] C. Schulze. Deep Blue.

[25] C. Metz. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/, 2016.

[26] J. Mahanta. Introduction to Neural Networks, Advantages and Applications. https://towardsdatascience.com/introduction-to-neural-networks-advantages-and-applications-96851bd1a207, 2017.

[27] S. Ioffe J. Shlens C. Szegedy, V. Vanhoucke and Z. Wojna. Rethinking the Inception Architecture for Computer Vision. *ArXiv151200567 Cs*, 2015.

[28] S. Rayej. How do self-driving cars work? *Robohub*, 2014.

[29] E. Brown J. Chu-Carroll J. Fan D. Gondek A. Kalyanpur A. Lally J. Murdock E. Nyberg J. Prager N. Schlaefer D. Ferrucci, D. Ferrucci and C. Welty. The AI Behind Watson - The Technical Article. http://www.aaai.org/Magazine/Watson/watson.php, 2010.

[30] S. Bagchi D. Gondek D. Ferrucci, A. Levas and E. T. Mueller. Watson: Beyond Jeopardy! *Artif. Intell.*, 199–200:93–105, 2013.

[31] J. Brownlee. What Is Natural Language Processing? *Machine Learning Mastery*, 2017.

[32] R. Courtland. Intel now packs 100 million transistors in each square millimeter. *IEEE Spectrum. https://spectrum.ieee.org/nanoclast/semiconductors/processors/intel-now-packs-100-million-transistors-in-each-square-millimeter*, 2017.

[33] M. M. Shulaker et al. Carbon nanotube computer. *Nature*, 501:526–530, 2013.

[34] Q. Zhang H. Xie-W. Qian R. Zhang, Y. Zhang and F. Wei. Growth of Half-Meter Long Carbon Nanotubes Based on Schulz-Flory Distribution. *ACS Nano*, 7:6156–6161, 2013.

[35] T. Johnson. What are carbon nanotubes. *ThoughtCo.*, 2017.

[36] Applications of Carbon Nanotubes. https://www.azonano.com/article.aspx?ArticleID=4842.

[37] R. H. Baughman M. F. L. D. Volder, S. H. Tawfick and A. J. Hart. Carbon Nanotubes: Present and Future Commercial Applications. *Science*, 339:535–539, 2013.

[38] M. Li A. Douglas, R. Carter and C. L. Pint. Toward Small-Diameter Carbon Nanotubes Synthesized from Captured Carbon Dioxide: Critical Role of Catalyst Coarsening. *ACS Appl. Mater. Interfaces*, 10:19010–19018, 2018.

[39] R. Feynman. There's Plenty of Room at the Bottom. http://www.zyvex.com/nanotech/feynman.html, 1960.

[40] J. Adams B. Rogers and S. Pennathur. Nanotechnology: The Whole Story. *Boca Raton: CRC Press, 1st ed.*, 2013.

[41] C. P. Lutz M. F. Crommie and D. M. Eigler. Confinement of Electrons to Quantum Corrals on a Metal Surface. *Science*, 262:218–220, 1993.

[42] A Boy And His Atom: The World's Smallest Movie. `https://en.wikipedia.org/wiki/A_Boy_and_His_Atom#/media/File:A_Boy_and_His_Atom_(still).jpg`.

[43] T. Trusk G. Forgacs V. Mironov, T. Boland and R. R. Markwald. Organ printing: computer-aided jet-based 3D tissue engineering. *Trends Biotechnol.*, 21:157–161, 2003.

[44] Z. Zheng et al. 3D Bioprinting of Self-Standing Silk-Based Bioink. *Adv. Healthc. Mater.*, 7, 2018.

[45] P. Najman J. Koprnicky and J. Safka. 3D printed bionic prosthetic hands. *IEEE International Workshop of ECMSM*, pages 1–6, 2017.

[46] About us. *Enabling The Future*, 2014.

[47] L. Gaget. Dental 3d printing: how does it impact the dental industry? https://www.sculpteo.com/blog/2017/12/18/dental-3d-printing-how-does-it-impact-the-dental-industry/, 2017.

[48] Underdog team wins millions in competition to make real-life tricorder . https://arstechnica.com/science/2017/04/underdog-team-wins-millions-in-competition-to-make-real-life-tricorder/.

[49] R. E. Gantenbein. Watson, come here! The role of intelligent systems in health care. *WAC*, pages 165–168, 2014.

[50] Summit. https://www.olcf.ornl.gov/summit/.

[51] K. Doi C. J. Vyborny R. A. Schmidt Y. Wu, M. L. Giger and C. E. Metz. Artificial neural networks in mammography: Application to decision making in the diagnosis of breast cancer. *Radiology*, 187:81–87, 1993.

[52] R. A. Freitas. Exploratory Design in Medical Nanotechnology: A Mechanical Artificial Red Cell. *Artif. Cells Blood Substit. Biotechnol.*, 26:411–430, 1998.

[53] C. Woudsma and L. Braun. Tomorrow Has Arrived: Cities and Autonomous Vehicles, 2017.

[54] G. Thibault and H. Moukanas. Autonomous Transportation Will Arrive Faster Than Predicted And Afford Bigger Business Opportunities, 2018. https://www.forbes.com/sites/oliverwyman/2018/02/05/autonomous-transportation-will-arrive-faster-than-predicted-and-afford-bigger-business-opportunities/2/.

[55] The 6 levels of self-driving car - and what they mean for motorists. TheJournal.ie, 2017. https://jrnl.ie/3603253.

[56] D. J. Fagnant and K. M. Kockelman. The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transp. Res. Part C Emerg. Technol.*, 40:1–13, 2014.

[57] M. DeBord. Americans are overconfident in their driving skills. http://www.businessinsider.com/americans-are-overconfident-in-their-driving-skills-2018-1.

[58] S. Captain. Humans were to blame in Google self-driving car crash, police say. https://www.fastcompany.com/40568609/humans-were-to-blame-in-google-self-driving-car-crash-police-say.

[59] J. Himmelreich. Never Mind the Trolley: The Ethics of Autonomous Vehicles in Mundane Situations. *Ethical Theory Moral Pract.*, pages 1–16, 2018.

[60] M. Schellekens. Self-driving cars and the chilling effect of liability law. *Comput. Law Secur. Rev.*, 31:506–517, 2015.

[61] Strict Liability. https://www.law.cornell.edu/wex/.

[62] R. Ubell. How the Pioneers of the MOOC Got It Wrong. https://spectrum.ieee.org/tech-talk/at-work/education/how-the-pioneers-of-the-mooc-got-it-wrong.

[63] M. Konnikova. Will MOOCs be Flukes? *The New Yorker*, 2014.

[64] A. Tucker. Jack Andraka, the Teen Prodigy of Pancreatic Cancer. https://www.smithsonianmag.com/science-nature/jack-andraka-the-teen-prodigy-of-pancreatic-cancer-135925809.

[65] M. Herper. Why Biotech Whiz Kid Jack Andraka Is Not On The Forbes 30 Under 30 List. https://www.forbes.com/sites/matthewherper/2014/01/08/why-biotech-whiz-kid-jack-andraka-is-not-on-the-forbes-30-under-30-list/.

[66] Z. W. Kendrick et al. Serum IGFBP2 and MSLN as diagnostic and prognostic biomarkers for pancreatic cancer. *HPB*, 16:670–676.

[67] B. Maya C. Fadel and B. Trilling. Four-Dimensional Education. *The Center for Curriculum Redesign*, 2015.

[68] Canada Closes Door on Cloning. https://www.wired.com/2004/03/canada-closes-door-on-cloning/.

[69] Health Canada. The Regulation of Genetically Modified Food. https://www.canada.ca/en/health-canada/services/science-research/reports-publications/biotechnology/regulation-genetically-modified-foods.html.

[70] C. J. Hoofnagle A. Bleakley J. Turow, J. King and M. Hennessy. Americans Reject Tailored Advertising and Three Activities that Enable It. *SSRN Electron. J.*, 2009.

[71] C. Cadwalladr and E. Graham-Harrison. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*, 2018.

[72] Cambridge Analytica, AggregateIQ and the Facebook scandal: A guide to who's accused of what., 2018.

[73] S. Nichols. Cambridge Analytica dismantled for good? Nope: It just changed its name to Emerdata. https://www.theregister.co.uk/2018/05/02/, 2018.

[74] Office of the Privacy Commissioner of Canada. Summary of privacy laws in Canada. https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/.

[75] Office for Civil Rights. Summary of the HIPAA Security Rule. https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html.

[76] Uber. How Does Uber Work? - Riding and Driving Basics. https://www.uber.com/en-CA/about/how-does-uber-work/.

[77] A. Armitage and A. K. Cordova. Design-Thinking: The Answer to the Impasse between Innovation and Regulation. *SSRN Electron. J.*, 2017.