THIS IS A PREPRINT VERSION 29.10.2024

First Author: Dr. Veli Virmajoki, Senior Researcher, Finland Futures Research Centre, University of Turku Finland
vevirm@utu.fi

Second Author: Dr. Mika-Petri Laakkonen, Lead/Principal Research Scientist, Oulu University of Applied Sciences
mika-petri.laakkonen@oamk.fi

Scenarios of AI and Scenarios of the Society. Marriage of Love, Not Convenience

Abstract

This paper presents a new way to think about the future of artificial intelligence by looking at two types of scenarios together: how AI might develop, and what kind of world it might develop in. Using the history of atomic technology as an example, we show that the effects of a new technology depend heavily on when and where it emerges. Our argument is that we cannot understand AI's impact by looking at the technology alone - we must also consider the social setting it will operate in, and how timing shapes its development. This combined approach to analyzing future scenarios helps us better plan for AI's development and offers useful insights for policy makers and those concerned with AI's ethical implications.

Keywords: Artificial intelligence, Scenario planning, Technological development, Counterfactual history, Futures research, Though Experiments

1. Introduction

Let us begin by traveling back in time to a possible world. Imagine we are somewhere in the 1930s, and we have increasingly begun to understand that the theories in physics tell us that an atomic bomb is possible. Imagine also that we need to know where this technology could lead us. How could we have foreseen that it would lead to the use of atomic bombs at the end of a world war and then to the geopolitics of the Cold War, which shaped the world throughout the century and still has (sad) ramifications in our era? Probably we could not have seen this. But we might have had a chance if we focused on the possible developments of the technology and the timeline of that development, and if we paid attention to possible scenarios for societies and to the fact that a worldwide war would reach its horrible peak in the 1940s.[1]

To strengthen this thought experiment, let us assume that the theoretical background and, thereby, the technology for the atomic bomb were ready earlier. Countries would have developed such bombs in, let us say, the early 1930s – countries with sufficient resources. What would have followed? Would the unstable and belligerent societies of the 1930s have used those weapons earlier? Who would have had access to them? No matter what the answers are, we can be fairly confident that the history of the late 1930s to the mid-1940s would have looked drastically different, and so would the history of the rest of the century.[2]

This thought experiment reveals two important issues. First, we cannot understand a technology's impact without understanding how the world surrounding that technology develops. Second, the timing and context of a technology's emergence profoundly influence its use and effects. Current AI technologies can be thought – within some limits (see the next section) – in analogy with the historical scenario in our thought experiment. New technologies appear in a world of societal and geopolitical shifts and developments. When we look towards the future, it is necessary to consider not only the technological trajectories of AI but also the diverse contexts in which it may be employed impact the world.

This conceptual paper presents a broad approach to analyze AI's future within wider sets of future changes. Rather than attempting to predict specific outcomes, we aim to construct a framework that can serve as a foundation for more focused foresight activities. Our intention is to stimulate reflection and provide a tool for those engaged in more detailed work in their relevant context. This method emphasizes the interactions between technological advancement and societal shifts, and we argue that this is a fruitful way to understand how AI will impact us through the world surrounding us.

---

[1] While finishing this paper, John Hopfield and Geoffrey Hinton received the Nobel Prize in *physics* for developing key methods in artificial neural networks and laying the foundation for modern machine learning. This coincidence strengthens the utility thought experiment in our paper.

[2] As the history we present here is rather commonly known, we do not use references throughout. However, we used the following works as our materials: For atomic technology: Cirincione, 2007; Gosling, 2005; Herrera, 2006; Rhodes, 1988; Rhodes, 1995. For 1930s context: Brendon, 2002; Overy & Wheatcroft, 2000.

To be more precise, we propose “a dual approach” to exploring AI futures. This dual approach consists of developing scenarios of the future of AI and scenarios of contexts, and, most importantly, connecting the two sets. In this framework, we define "context" broadly. It refers to any environment in which AI might operate, ranging from global scenarios to small group settings. This conceptual framework is designed to be adaptable and enable any group of actors to apply the dual approach within *their* relevant context. By examining both technological and contextual scenarios together, it is possible to gain understanding of impacts of AI across various domains. Our approach recognizes the complex and sometimes contingent relationship between innovation and society. At its core, the dual approach is a general approach with rather universal application.

We proceed as follows. Section 2 explores the use of historical thought experiments in informing the development of future scenarios and how futures researchers have analyzed the connection. Section 3 presents our *dual approach* in detail and introduces the concept of "injecting" AI scenarios into scenarios of different contexts. Section 4 illustrates our approach through an example focusing on the future of interdisciplinary research in universities. Section 5 analyzes the broader implications of our approach for considerations of time, scenarios, and ethics in futures research. Section 6 concludes the paper by summarizing our key arguments and contributions. Throughout the paper, we emphasize how historical parallels and thought experiments can make approaching the development and impact of transformative technologies like AI more manageable and provide a sense of scale of possible trajectories.

2. Why Use Historical Thought Experiments?

The relationships between history, historiography, and futures research have been a topic in futures research. In the discussion about the issue, the connections between history, historiography, and futures research have been seen. Many researchers recognize the value of historical thinking in understanding possible futures. This connection is particularly relevant to our approach of where we integrate scenarios of futures of AI with scenarios of future contexts. Bradfield et al. (2016) emphasize that we can gain insights from the past even if historical trajectories do not repeat themselves. They suggest that we can compare, contrast and debate possible future changes against the causal framework of the past. In a bit more provocative take, Virmajoki (2023a) argues that historical understanding always generates future possibilities because both require generalizations, while Virmajoki (2023b) argues that causal explanations in historiography can guide our approach to future scenarios. Moreover, Staley (2010) suggests that historical thinking is important for studying the future because both require contextual analysis. He argues that historians of the future make broad, non-demonstrative inferences based on available evidence.

The concept of counterfactual histories is particularly relevant to our methodology. Staley (2002) emphasizes that addressing key historiographical questions often requires the study and tracking of alternative historical outcomes. Similarly, Booth et al. (2009) propose that the theoretical literature on counterfactuals and modal narratives can help clarify theoretical issues related to foresight methods. Green (2012) has also examined the similarities between the cognitive approaches used in historical study and strategic foresight and argued for the similarity between counterfactual histories and future scenarios.

These lines of reasoning provide a fertile ground for using historical thought experiments, such as exploring alternate paths in atomic technology, to inform our approach to AI future. We need historical analogies and counterfactuals to understand the scale and determinants of a major issues for the future, such as AI.

Of course, counterfactual histories have often been criticized as mere speculation. What counts as a plausible or relevant counterfactual scenario is a genuine question. The problem is that we cannot have direct evidence of counterfactual scenarios because, by definition, they did not happen. One could argue that if we are not able to distinguish between plausible and far-fetched counterfactual scenarios, there is little hope that counterfactual considerations provide any helpful insight in futures research – it would be speculation after speculation. However, Virmajoki (2024) suggests that, when approached carefully, historical counterfactuals are legitimate and valuable analytical tools. While there are challenges in constructing and evaluating these thought experiments, they are not insurmountable. There are issues that we can, according to Virmajoki, control in our thought experiments so that they do not get out of hand and turn into mere irrelevant speculation that has not historical of futures research value. Thought experiments work in many cases and for many reasons. This perspective underpins our use of historical thinking while analyzing the potential trajectories of AI in society.

To show the value of historical thought experiments, let’s put some flesh on the bones of above consideration through a historical case.

A Thought Experiment

To simplify the history a bit, we can provide the following scheme of how atomic weapons became to affect the world fundamentally.

Leo Szilard's work in the mid-1930s introduced the concept of a self-sustaining nuclear chain reaction that could release vast amounts of energy. This laid the groundwork for nuclear energy and atomic weapons. In 1934, Szilard filed a patent on the idea. Concerned about the potential military applications, he sought to keep the patent secret and offered it to the British government for safekeeping, but the War Office rejected his proposal. In December 1938, Otto Hahn and Fritz Strassmann discovered nuclear fission. By early 1939, Lise Meitner and Otto Frisch explained that uranium fission released large amounts of energy and additional neutrons, and this could cause a self-sustaining reaction.

As noted, scientists like Leo Szilard understood relatively early the potential military applications of what was, at this point, theoretical work. During the course of the events, they helped draft a letter for Albert Einstein to send to President Roosevelt in August 1939. The letter warned of Germany's possible exploration of atomic weapons and this prompted the United States to accelerate nuclear research. This led to the Manhattan Project. After years of intense work, the project culminated in the Trinity test on July 16, 1945, where the first atomic bomb (or ‘nuclear device’) was successfully detonated. This was followed by the atomic bombings of Hiroshima on August 6 and Nagasaki on August 9, 1945. Once Soviet Union achieved the capability to create nuclear weapons, the outcome was a nuclear arms race between two geopolitical blocks with vast political differences. The arm race led United States and Soviet Union pursuing even more powerful thermonuclear weapons which culminated in the creation of the hydrogen bomb in the early 1950s.

Let us consider an alternative scenario: what if these nuclear science breakthroughs had occurred earlier, perhaps in the early 1930s? The world at that time faced economic instability and rising political extremism. The situation was extremely different from the one where Europe was being cut between to different geopolitical systems at 1945. One thing is clear: earlier development of nuclear technology would have influenced events in various ways.

Many countries in the early 1930s struggled with the Great Depression and political upheaval. Totalitarian regimes gained power and geopolitical tensions intensified. The development of atomic weapons during this period presents several possibilities. Germany, under Hitler's leadership, might have pursued nuclear weapons as part of its military rearmament. Japan, with its imperial ambitions in Asia, could have sought such weapons to secure regional dominance.

What could have happened in these conditions? Would the countries have used the weapons and, if so, how and with what consequences and chain reactions in world politics? These questions lack definitive answers. However, imagining these possibilities allows us to consider alternative historical paths. Nuclear weapons in the 1930s might have altered the balance of power in unexpected ways. Earlier and more severe conflicts could have resulted.

The Great Depression had created widespread desperation and discontent. Rising unemployment, economic hardship, and political instability drove the growth of extremist movements in many countries. Hitler's Nazi regime had begun to militarize Germany, violating the Treaty of Versailles. The availability of nuclear weapons in this context could have profoundly altered the geopolitical landscape. Germany's aggressive foreign policy, driven by Hitler's ambition for notorious 'Lebensraum', might have become far more dangerous with access to nuclear weapons. Given the horrible nature of the actual history, the mere introduction of the idea of 'more dangerous' here already gives us an understanding of how open historical paths are. Japan, pursuing imperial ambitions in China and Southeast Asia, could have viewed nuclear weapons as the ultimate tool for securing regional dominance. In this alternative timeline, an arms race between nations might have included rapid nuclear weapons development, potentially leading to early use in regional conflicts or as tools of intimidation.

The outcome of such a scenario remains unpredictable and subject to many factors beyond technology. The mere presence of atomic bombs might have deterred conflict, or the desperation and ideological fervor of the time could have led to their early use. These open questions highlight the complexity of predicting new technologies' impact in volatile times.

The broader context of the 1930s plays an important role in this thought experiment. Diplomatic failures, like the League of Nations' inability to prevent aggression and Western powers' policy of appeasement, set the stage for conflict. The addition of atomic bombs to this mix raises questions about potential shifts in diplomatic strategies or further escalation of tensions.

The scientific environment of the time also warrants consideration. Scientists fleeing Nazi Germany in the 1930s contributed to intellectual migration that shaped scientific progress in the West. Accelerated development of nuclear physics and engineering might have led to an earlier "arms race" among world powers, with nuclear weapons pursuit intensifying in the mid-1930s rather than the 1940s.

This alternative scenario opens up different pathways for how World War II and its aftermath might have unfolded. Brief but devastating wars involving atomic bombs could have occurred in the 1930s. Alternatively, nations might have recognized these weapons' catastrophic potential and sought new diplomatic strategies, potentially averting the scale of conflict that characterized World War II.

Atomic Technology and AI: A Systematic Comparison

To better understand the potential trajectories and impacts of AI, we can draw valuable lessons from the development of atomic technology. Both represent transformative technologies with far-reaching implications. The following table outlines key parallels (Table 1):

Factor	**Atomic Technology**	**Artificial Intelligence**	**Lesson for AI Development**
Pace of development	Rapid progression from theoretical possibility to practical application during WWII	Accelerating advancements, with significant breakthroughs occurring in short timeframes	The speed of AI development may outpace societal and regulatory adaptation, requiring proactive governance
Potential for dual-use	Used for both energy production and weapons, raising complex security issues	Applicable across various fields, from healthcare to warfare, presenting similar dual-use challenges	Early consideration of potential applications is crucial to guide development towards beneficial uses
Global impact and geopolitical implications	Reshaped international relations, led to arms races and new power dynamics	Has the potential to alter economic structures, military capabilities, and global power balances	International cooperation and shared governance frameworks are essential to manage global impacts
Ethical challenges	Raised questions about the morality of mass destruction and	Presents issues related to privacy, autonomy, bias, and potential existential risks	Ethical considerations must be integrated into the development process from the outset

	scientists' responsibilities		
Public perception and fear	Induced both awe at human achievement and fear of annihilation	Generates excitement about possibilities and anxiety about job displacement and loss of control	Clear communication and public engagement are vital to build understanding and trust

Table 1

The timing of atomic technology development, coinciding with World War II, profoundly shaped its initial applications and subsequent impact. Similarly, the trajectory of AI development will be influenced by current global challenges, economic conditions, and geopolitical tensions. For instance, just as the wartime context accelerated atomic weapon development, today's competitive international environment might drive rapid AI advancements in areas like autonomous weapons or surveillance technologies.

This comparison underscores the critical importance of considering both technological capabilities and societal contexts when envisioning possible AI futures. By learning from the atomic age, we can better anticipate and shape the development of AI in ways that maximize benefits while mitigating potential risks.

Exploring counterfactual scenarios for atomic technology development can further illuminate potential paths for AI. For instance, what if atomic technology had been developed a decade earlier, during the Great Depression and the rise of totalitarian regimes? Or what if international cooperation had led to shared control of the technology from the outset? These "what-if" scenarios highlight how different societal contexts could have dramatically altered the trajectory and impact of atomic technology. Similarly, for AI, we must consider various potential futures. What if a major AI breakthrough occurs during a global economic crisis? What if AI development becomes concentrated in a single country or corporation? By systematically exploring these alternative scenarios, we can better prepare for a range of possible AI futures and work towards shaping the most beneficial outcomes.

What this teaches?

This thought experiment explores how timing and context shape a technology's global impact. The 1930s world was key to nuclear technology's eventual use. Similarly, current conditions will influence AI's future. Looking at this alternate history does not claim to know what might have happened. Instead, it sparks imagination and shows that technology does not develop in isolation. It is tied to its time. AI today faces challenges like those of 1930s nuclear technology. This exercise makes us think about how technology and society interact to shape the world. Imagining different futures helps us understand possible paths and what might influence them.

We want to balance learning from history and staying open to new ideas when thinking about AI futures. By mixing AI development scenarios with future social contexts, we hope to make a framework that is based in history but flexible enough for new possibilities. Using thought experiments, like imagining different timelines for atomic technology, serves many purposes. It shows how technology development depends on society, helps identify key factors in AI's growth, and encourages thinking about many possible futures, even unlikely ones. This approach gives insights into how technology and society change together, helping us deal with the uncertainty in AI's fast-changing field.

3. The Dual Approach and Injections

Our conceptual framework ('the dual approach') rests on the insight that, in order to analyze the future of AI, we must employ two distinct yet interconnected sets of scenarios. This dual approach stems from our understanding of how technologies, such as the atomic bomb, profoundly influenced, and were influenced by, their contexts and timing.

The first set of scenarios is about possible future contexts where technology might be used. Contexts cover the possible states of worlds, societies, or situations in which AI might be developed and utilized. For example, the impact of atomic technology was significantly shaped by the global political landscape of the 1940s. Our thought experiment demonstrates that the timing of the development of a technology can be as critical as its intrinsic capabilities. Had nuclear weapons been developed in the early 1930s with the instability and escalating tensions characterizing that era, the course of history might have been different – both the world and the use and development of nuclear technology. Similarly, the future of AI will be greatly affected by the social, economic, and political environments in which it evolves. These contextual scenarios may concern a wide range of contexts from the whole world to a specific organization or even a small group. Given the extensive scenario work already conducted in different domains, we will not go into the details of the scenarios of different contexts. However, the next section provides an example of contextual scenarios and illustrates how an AI scenario can be integrated into it.

It is crucial to recognize that, in our dual scenario approach, one must study multiple possible future scenarios of contexts in order to understand different what-ifs of AI. One can employ various techniques to generate these future scenarios depending on the topic, the participants involved, and the intended application of the scenarios. We refrain from prescribing a specific method for scenario creation, as different fields and organizations may have their preferred approaches. For example, some might utilize Delphi techniques to gather expert opinions, while others might employ horizon scanning to identify emerging trends. What matters in the dual approach is that there is a heterogenous set of scenarios of any given context of interest. Our approach is adaptable and compatible with various scenario-building methods, provided they produce multiple, diverse scenarios rather than a single projection. The core idea is that, regardless of how the future scenarios are developed, they can be systematically combined with AI development scenarios to understand AI and its effects better.

The second set of scenarios concerns the possible directions AI technologies could take, focusing on advancements, breakthroughs, and enduring challenges. AI's development can be understood through four main models: neural networks, genetics, genealogy, and physical phenomena. Neural networks, inspired by brain structure, use connected nodes to simulate neurons, establishing the groundwork for machine learning and deep learning. The genetics model introduced evolutionary algorithms, applying natural selection principles to refine solutions. Genealogy extended this by using historical data and lineage to improve performance further. The physical phenomena model, still emerging, takes inspiration from physics, aiming to integrate physical principles into computational frameworks for AI. (Darwin & Huxley 2003; Fausett 1994; Goldberg 1998; Kramer 2017; Lee 2019; Nie et al. 2020; Humphreys 2009; Coeckelbergh 2021; Laakkonen 2024).

For the reasons above, we structure our discussion about the issue, we can identify four primary paradigms[3]

1. *Modeling Neural Networks*: Based on the human brain and memory, using historical data for pattern recognition and decision-making.

[3] Later, we discuss the developmental arcs of AI. As 'Darwin and Huxley' reference shows, very long trajectories in the history of science are relevant to understanding AI technologies fully. However, that is a topic or story for another paper and we cannot go to that length in this paper.

2. *Genetic Modeling*: Rooted in biology and genetics, focusing on evolutionary processes to solve complex problems.
3. *Genealogical Modeling*: Considering long-term adaptive processes based on population dynamics.
4. *Physical Phenomena Modeling*: Inspired by real-world physics, viewing AI as a dynamic, self-organizing system.

These paradigms offer a framework to discuss and understand the future capabilities of AI. Each paradigm implies a distinct approach to AI development with its own assumptions, methodologies, and potential trajectories.

The Modeling Neural Networks paradigm is inspired by the human brain's function and has proven crucial for pattern recognition and decision-making tasks. When we go back in time a bit, Gasteiger and Zupan's (1993) work is an example of an early application of this paradigm (in this case, in chemistry). Their research showed how neural networks could analyze chemical structures and predict properties. This work can be seen as an example of a situation where the approach's potential in complex scientific domains was proven through application, and the work established a foundation for decades of subsequent development.

Genetic Modeling takes is inspired by biological evolution and has for a relatively long time provided robust tools for optimization and problem-solving. For example, Hambardzumyan et al.'s (2011) study demonstrated the application of genetic modeling to complex medical research topics such as cancer research, specifically gliomas. This work, again, indicated the paradigm's potential in medicine and, thereby, was a sign of expanding body of research in this area.

The Genealogical Modeling paradigm concentrates on long-term adaptive processes and has provided a framework to understand and predict complex population dynamics for several decades. For example, Degnan and Rosenberg's 2009 paper "Gene tree discordance, phylogenetic inference and the multispecies coalescent" is an application of this paradigm in evolutionary ecology. Their work showed how genealogical modeling could help to understand complex evolutionary processes. It built upon and extended earlier research in the field.

Physical Phenomena Modeling derives from our understanding of the physical world and views AI as a dynamic, self-organizing system. For example, Rai and Sahu's 2020 review "Driven by Data or Derived Through Physics? A Review of Hybrid Physics Guided Machine Learning Techniques

With Cyber-Physical System (CPS) Focus" offers a recent perspective. Their work focuses on how physics-based modeling can bridge data-driven and physics-based approaches in cyber-physical systems. It reflects on the evolution of this paradigm over time.

These paradigms have undergone significant evolution since their birth and inception. Each in their own way shows how steps in the development of AI can be taken to multiple directions and in multiple ways from past to present to future. The paradigms above serve as an analytic tool but also as a pointer towards the multiplicity of the thing we often simplify by talking about “AI”. As we attempt to analyze the future trajectories of AI development, we must understand both the historical foundations and ongoing advancements in each of the areas presented here through the paradigms. For example, recent progress in natural language processing, which falls under the Modeling Neural Networks paradigm, has provided systems that can analyze multiple documents and generate concise summaries. These developments demonstrate substantial advancement in AI's ability to comprehend and generate human-like text (Chen et al., 2022). Researchers have also developed novel approaches to utilize AI for Earth system modeling. They work with partial or incomplete data to understand and predict environmental changes (Ayed et al., 2022). This application illustrates how AI adapts to and addresses pressing global challenges.

It is important to note that researchers are also focusing on making AI systems safer and more reliable – even if we need to take this with a pinch of salt sometimes. For example, the development of SAMBA (SAfe Model-Based and Active reinforcement learning) aims to make reinforcement learning - a key technique in AI - safer and more efficient (Cowen-Rivers et al., 2022). This is an example of research that highlights that, in the dual approach, we need to consider safety and ethical implications alongside technological advancements in our scenario planning. There are sometimes ethics and safety *within* AI technologies and not only ethics and safety *about* the technology.

Injecting Scenarios

We have now taken a glance on the diversity in the notion of AI. Now, by mixing these ideas with different scenarios of contexts, we can better understand how AI technologies might develop in various situations. This resembles how atomic technology was shaped by what was happening in its context – the geopolitics of the world – at 1940’s. In what follows, we discuss the examples above and show how they build on certain assumptions about how the world is and might be – on scenarios.

As we saw, Gasteiger and Zupan (1993) analyzed the use of neural networks in chemistry. Their work assumes computers will keep getting more powerful and allow for more complex chemical analysis. The work also assumes scientists will rely and trust AI more when its usefulness is established in certain field. In a future where grand environmental problems dominate us, this kind of AI might help solve these issues and thereby expand radically. But just like how nuclear technology changed unexpectedly due to world events, AI in chemistry could change a lot if global priorities shift or if we make unexpected breakthroughs in other technologies. What will happen to particular technology depends on how well its core assumptions and principles fit the development of different contexts.

Hambardzumyan et al. (2011) studied genetic modeling in mice. Their work might be viewed as a sign of a future where medicine is personalized based on AI analysis of genetic data. But, in order to understand the different trajectories of the technology, we need to also think about how this future might change if, for example, strict rules about genetic data were agreed worldwide – similar to how international agreements have shaped how we use nuclear technology.

Degnan and Rosenberg (2009) worked on genealogical modeling in ecology. Their approach might be useful in a scenario where we face rapid environmental changes. But if we imagine a world where early efforts to protect nature had been more successful, this technology might have developed differently, maybe focusing more on keeping ecosystems healthy rather than managing their decline. What exact implications this would have had is a question that cannot be discussed here.

Rai and Sahu (2020) analyzed at how machine learning can be guided by physics in systems that combine physical and digital elements. Their work assumes that *Internet of Things* technologies and smart infrastructure will keep growing. But if we imagine a world where people rejected these technologies because of privacy concerns - like how some countries have rejected nuclear power because of safety worries - the development of these AI techniques might, again, look very different.

As we think about these possible futures, it becomes clear that AI's development will be shaped by what the technology can do, what society do and wish to do, and what problems the world faces. The path of AI is not set in stone – it can change based on what we as a society decide to do. Here we face difficult ethical questions too. As AI gets more involved in research and decision-making, we need to figure out who is responsible for this or that aspect of the technology. We also need to make sure we can understand how these AI systems work. The story of nuclear technology shows us that powerful technologies can have big, unexpected effects. The way nuclear weapons worked and the scale in which they were created made a great impact on the history of the world. Moreover, as many of the issues AI is entangled with are global, countries and other major players should work together on AI in order to make it really sustainable. But just like how nuclear technology was affected by world politics, AI will probably be influenced by which countries and other entities have power and what different countries and entities want.

To sum up, understanding how AI might develop in different future scenarios enables us to see what problems and opportunities might come up. This shows that we need to be flexible in how we plan for AI's future – we need to understand the diversity of technology, its assumptions, and the world where it is planted. It is also important that different experts to work together - not just technology experts, but also philosophers, sociologists, ecologists, policymakers, and so on. Only by thinking about all the possible futures for AI and its contexts, we can guide AI's development in desirable way. The future of AI is not decided yet; it will be affected by unexpected changes in contexts *at the time* when the AI is planted in that context.

Scenario-Injection and Time-Horizons

Our approach examines AI technologies across three time frames:

Short-term: 0-5 years.

Medium-term: 5-15 years.

Long-term: 15+ years. Looking further into the future, things become more uncertain.

In the short term (0-5 years), we can see clearer paths, though clarity does not imply comfort. Artificial intelligence is expected to prioritize responsible AI practices, focusing on ethical use and transparency in applications. Notable advancements are anticipated in generative AI, particularly with multimodal models that process and generate text, images, and video. Operational efficiency in AI platforms will improve, enabling easier integration into business processes. The adoption of small and wide data approaches will enhance AI's accessibility, making it effective even with limited data. (Sandeep Singh Sengar et al., 2024; Siau & Wang, 2020; Mucha & Seppala, 2020). For example, large language models have greatly advanced natural language processing and generation, reshaping how we interact with AI. This rapid progress mirrors the inevitable development of atomic weapons following the scientific breakthroughs of the 1930s and 1940s. Even within this short time frame, small changes could lead to significant deviations. Had Germany not lost in 1945, what might have happened? If you were in 1943 and knew about the atomic bomb, could you have foreseen the Cold War? In the short term, things can shift quickly.

In the medium term (5-15 years), AI and contextual scenarios grow more complex, leading to a broader range of outcomes. Multimodal AI in healthcare and urban mobility has driven transformative changes. Leveraging biomedical data from biobanks, health records, medical imaging, and wearable biosensors, AI has led to advanced solutions reflecting the complexity of human health. Personalized medicine now tailors treatments to genetic profiles, while digital clinical trials speed up new therapy development. AI-enabled remote monitoring improves chronic disease management and early detection. Pandemic surveillance systems predict and mitigate disease spread in real time. Autonomous vehicles, meanwhile, are reshaping urban mobility, challenging policymakers to manage the resulting changes. Yet, critical issues around data privacy, security, and ethics demand robust regulatory frameworks for responsible AI use. (Acosta et al., 2022; Noorbakhsh-Sabet et al., 2019; Faisal et al., 2019) For instance, Rai and Sahu (2020) explored hybrid AI models combining physics and machine learning, allowing AI to handle more complex scenarios. This resembles the Cold War's impact, where atomic technology affected world politics. What if the 1950s political situation had differed? Could atomic energy have been used more peacefully? These "what ifs" show that medium-term AI developments could vary widely based on societal shifts. Medium-term scenarios help us understand diverse AI trajectories and how society might address emerging ethical challenges.

In the long term (15+ years), AI's trajectory is far more open and speculative. Artificial General Intelligence (AGI) may reach unprecedented levels. The Tong test, emphasizing dynamic embodied physical and social interactions (DEPSI), has become the gold standard for AGI assessment. AGI systems embedded in daily life excel in personalized education, advanced healthcare diagnostics, and treatment planning. In smart cities, AGI optimizes infrastructure and public safety, while advanced robotics manage complex tasks in industrial and home settings. AGI also drives major scientific discoveries in fields like medicine, energy, and environmental science, contributing to sustainable solutions for global challenges. Ethical and regulatory frameworks will focus on transparency, accountability, and fairness. (Peng et al., 2024; Goertzel, 2014.) For example, Bonnet et al. (2022) focused on deep learning in computational fluid dynamics, but their research could inspire broader long-term impacts. Similar to how nuclear technology shifted from military to energy and medicine, AI’s future over 15+ years may hold unexpected changes. Potential long-term scenarios include breakthroughs in regenerative medicine, new energy sources, or reshaped human-machine interaction. While speculative, these outcomes are essential to understanding AI’s long-term impact on society, even if they go beyond current research boundaries.

However, no matter the issue, the interplay between AI development scenarios and societal contexts evolves across these timeframes (Table 2):

TIME HORIZON	AI DEVELOPMENT SCENARIO	SOCIETAL CONTEXT SCENARIO	KEY CONSIDERATIONS
SHORT-TERM (0-5 YEARS)	Incremental AI improvements	Reactive policy adjustments	Immediate ethical concerns, public perception
MEDIUM-TERM (5-15 YEARS)	Potential AI breakthroughs	Shifting social norms and structures	Emerging ethical challenges, adaptation strategies
LONG-TERM (15+ YEARS)	Transformative AI capabilities	Fundamental societal changes	Speculative scenarios, long-term human values

Table 2

As we think about the developments of AI over time, we can see that looking further ahead opens up more possibilities. Short-term scenarios might seem limited by today's technology while medium and long-term scenarios are become more unclear. This shows we need to think flexibly and have sets of scenarios where AI is injected in both in short, medium, and long-term.

AI Technologies Have Long Development Arcs

It is here that we need to recognize that AI technologies, like other major and diverse technologies, have long histories rooted in decades of research. Just as nuclear technology was built on early atomic physics discoveries, AI has grown over time and been shaped by various historical and social situations – and so it will continue to be.

The historical development of AI technologies, much like nuclear technology, illustrates the complex interplay between scientific progress and societal context. Rai and Sahu's (2020) work on integrating physics with machine learning and Zupan and Gasteiger's (1993) application of neural networks in chemistry demonstrate how AI has evolved by merging with traditional scientific disciplines. These examples remind us of the development of nuclear technology, where early research could have taken various paths, given the context, even though the theoretical basis was already established. This historical analogy enables us to consider alternative trajectories for AI: For example, what if early AI research had prioritized different approaches over neural networks? Such "what if" scenarios could show us the critical role of societal context and historical contingencies in shaping technological development. At least, they should remind us that the path of innovation is neither linear nor predetermined.

AI is not only some contemporary phenomenon and we should, thereby, be aware of its history. This is the main point of this paper: to show that AI technologies, like those before them, are part of a longer story and deeply connected to their historical and social contexts. By recognizing these long development paths, we can better understand how AI got to where it is now and estimate where it might go in the future. By acknowledging these long development paths and the potential for different outcomes, we move beyond simple views of AI. Instead, we get a richer understanding of how AI grows along with society, much like how nuclear technology shaped and was shaped by 20th-century world politics.

Challenges and Limitations in AI Development

While our approach helps us to understand possible futures of AI, one need to talk about the grand challenges shaping AI development currently. While this outside the scope of our paper, there are some issues that can be mentioned as exemplary cases so that the point is not lost out of sight.

For example, Mohamed et al. (2020) discuss data quality and availability issues. AI systems need data, but this data often reflects society's biases and limits. So, the very data powering AI also limits what it can do in different situations and how it affects the context surrounding it. Zednik (2021) points out that AI computation complexity is another challenge. This limits AI's ability to scale up and creates energy demands that are a major issue when global sustainability is the concern. Another great challenge is making AI systems understandable, especially as they are used more in societal decision-making. For example, Erasmus et al. (2021) argue that if we cannot understand how AI makes its decisions, it is difficult to trust it – at least in the ethical sense – in important areas like healthcare or criminal justice. This shows how the technical complexity of AI is tied to society's need for accountability and ethical oversight. Moreover, Buckner (2018) analyzes AI's struggle to apply what it learns to new situations. This limits AI's usefulness in changing environments – or ‘contexts’ from the perspective of this paper and scenarios.

Our scenario-based approach gives us some, admittedly humble, tools to navigate these interconnected challenges. It offers a framework to analyze how AI might develop in different context and delimit the set of issues to those that are concerns in that particular context; banks and research institutions, for example, share different worries. By considering both the technological possibilities and social contexts, we can better prepare for the various ways AI might evolve and impact our world.

4. An Illustration of the Dual Approach: Universities, Interdisciplinarity, and AI

To illustrate our dual approach in a simple manner, in this section, we consider on specific technology in specific context: unsupervised clustering in the context of interdisciplinarity of universities. The future of interdisciplinary research is currently a major issue for universities, as interdisciplinarity is seen as a way to address grand challenges of our time such as climate change and social inequalities (Bursztyn & Drummond, 2014; Pimentel et al., 2023). However, achieving interdisciplinary collaboration is not straightforward and can happen in different ways depending on institutional structures, funding priorities, and technological capabilities.

To show how our dual scenario method works, let's look at unsupervised clustered AI models, such as self-organizing maps, and how they might be related to interdisciplinary research. This example shows how injecting scenarios of AI development with broader university future scenarios can give us valuable insights.

In one possible scenario of the future of universities, there might be a strong push towards interdisciplinarity and breaking down disciplinary silos (Steger, 2019; Millar, 2016; Tjörnbo & McGowan, 2022; Bursztyn & Drummond, 2014; Salmela et al., 2021; Bromham et al., 2016). Here, advanced unsupervised clustering AI models might play a role. These AI tools could analyze vast amount of research data, including publications, citations, and keywords, and, in this way, suggest novel connections between seemingly unrelated fields. By grouping related research areas across disciplines, even when they use different terminology, AI could suggest collaborative projects that people might not be aware of.

But how these AI tools could impact interdisciplinarity would not depend solely on their technological capabilities, but also on the broader institutional and societal contexts in which they are used (or not used). In a scenario where universities stay mostly siloed and resist change, even most capable AI clustering tools would not be much of a help. On the other hand, in a future where universities have already taken steps toward in breaking down disciplinary silos, these AI tools could speed up and deepen interdisciplinary collaboration.

Our dual scenario approach lets us systematically explore these interactions (see Table 3 for summary). By considering multiple possible futures for university structures and research priorities (scenarios of the future of interdisciplinarity) alongside various possible paths for AI development (scenarios of AI technologies), we can get a more nuanced understanding of how interdisciplinary research might evolve.

Consider, as an illustration, a more detailed example: Salmela et al. (2021) study internally incentivized interdisciplinarity in universities. Their research examines a case where a university reallocated internal funds to "research platforms" requiring participation from multiple departments or faculties. This approach led to specific tensions in (i) resource allocation, (ii) division of labor, and (iii) scientific output.

When it comes to resource allocation, unsupervised clustering AI could analyze past project outcomes, publication patterns, and collaboration networks to suggest certain funding distribution across platforms. In a possible future context of universities with increased competition for resources, this AI could provide a basis for decisions. However, it cannot resolve the fundamental issue of limited resources or account for the political dynamics within institutions. So in a scenario where unsupervised AI is used in platforms, more efficient resource allocation might be achieved, but the fundamental issue of limited resources remains unresolved and political dynamics within institutions may still influence decisions. This situation could potentially drive the development of AI systems specifically designed to navigate academic politics and resource allocation. This would be a situation similar to how the unique organizational challenges of the Manhattan Project led to innovations in project management and resource coordination during World War II.

When it comes to the division of labor tensions, AI could cluster researchers based on complementary skills and past collaborations and, thereby, identify effective interdisciplinary teams. In a possible future context of a university where boundaries between disciplines become more fluid, this could facilitate even more collaborations. However, the AI might struggle to capture the qualitative aspects of research contributions and, thereby, reinforce existing hierarchies or overlook valuable but less quantifiable inputs. So in a scenario where unsupervised AI is used in platforms, team formation might be optimized, but the AI might struggle to capture qualitative aspects of research contributions and, in this way, reinforce existing hierarchies or overlook valuable but less quantifiable inputs.

For scientific output tensions, unsupervised AI could cluster publication venues and citation patterns across disciplines by possibly identifying valuable interdisciplinary publication opportunities. In a possible future context of a university that where traditional disciplinary metrics are esteemed, this could help researchers balance interdisciplinary work with career advancement. However, the AI might not be able to resolve the underlying conflict between interdisciplinary research and discipline-specific career progression structures. So in a scenario where unsupervised AI is used in platforms, researchers might find it easier to identify suitable publication venues for interdisciplinary work, but the underlying conflict between interdisciplinary research and discipline-specific career progression structures remains unresolved and may continue to influence researchers' choices.

The discussion above shows how our dual scenario approach captures the complex interplay between technological possibilities and institutional contexts. It allows us to move beyond simplistic, linear projections of either AI development or university futures. Instead, it encourages us to consider how these factors might co-evolve and influence each other over time. In Table 3, we summarize this simplified scenario exercise.

	No interdisciplinary push	**Interdisciplinary push**
Advanced Unsupervised Clustering AI	Limited adoption of AI tools. Minimal impact on interdisciplinary research. Potential tension between AI-driven insights and institutional structures	Accelerated interdisciplinary collaboration. Emergence of new interdisciplinary fields. AI-driven identification of novel collaborative projects. Risks bias and over-reliance on AI recommendations, potentially overlooking human intuition
No Advanced Unsupervised Clustering AI	Reinforcement of existing disciplinary boundaries Minimal change in research practices	Slower progress in interdisciplinary research. Reliance on traditional methods for identifying collaborative opportunities. Potential mismatch between

	Missed opportunities for interdisciplinary breakthroughs	institutional goals and technological capabilities.

Table 3

Moreover, this approach illustrates the inherent ethical dimensions of the futures of AI. For example, how might the use of AI tools to guide research collaborations affect academic freedom and human-driven discovery? What might need to be taken to ensure that AI-driven interdisciplinary initiatives don not exacerbate existing power imbalances within and outside academia? (Salmela et al., 2021; Pelacho et al., 2021) In general, understanding the future of AI through its possible uses in specific contexts enables us to ask targeted ethical questions that are difficult but manageable.

5. Time, Knowledge, and Ethics of AI futures

As we have seen, the dual approach to AI futures connects scenarios of AI development with scenarios of its contexts of use. In this way, it not only builds on but also expands traditional scenario planning methods. This approach is not only a practical tool for AI governance and development but it also tackles some core questions about time, knowledge, and ethics in futures research. By looking closer at these areas, we can better understand what our method brings and what it means. Inayatullah (2008) emphasizes the importance of scenarios in futures research, and our method takes scenarios seriously by building two sets of them and then connecting the two sets through injection.

At the heart of our dual scenario framework is the idea that time and context are intertwined in an intimate way, especially when it comes to the future of AI. This matches the main idea of scenario planning, which tries to spot a range of possible futures – each with its own risks and chances (Amer et al., 2013). Our approach takes this idea further by looking at AI development within different contexts and by recognizing that the timing of some AI development is not separate from where they happen and, moreover, the timing and contexts are tangled up with the capabilities of the AI technology as hand. Our thought experiment about atomic technology being developed earlier is a way to get a grasp of the idea (see Section 2). By thinking about how atomic weapons might have affected the world and been developed further if they had been created in the 1930s instead of the 1940s, we can understand how technological capabilities and social, political, and economic contexts mix together. This experiment highlights how timing is crucial in shaping how technologies unfold and impact things. The development of atomic weapons during World War II and the Cold War that followed fundamentally shaped how they were used and the whole geopolitical scene of the 20th century. In the same way, how AI develops in various possible future contexts - each with its own economic, political, and social challenges – will influence where different AI technologies go and how they impact things. By thinking about these contextual factors through our dual scenario method, we can understand the complex nature of the futures of AI.

As we already noted, AI technologies have developmental arcs similar to atomic science's historical path. This particular arc is only one example of how crucial timing is in technology development and impact. Just like nuclear physics had a rich history before major technological breakthroughs like atomic weapons, AI paradigms and technologies have evolved over decades. For example, the neural network models behind today's deep learning breakthroughs have roots in cybernetics research from the 1940s and 1950s. Why they are on the state they are now is something that only future historians can explain to us; but as futures researchers we need to have techniques to build modal understanding of the future.

Our method of "injecting" AI development scenarios into scenarios of contexts is a novel extension on existing scenario planning techniques. While it is similar to The Intuitive Logics approach explores key driving forces shaping the future to develop plausible scenarios and improve strategic decision making, (Amer et al., 2013), our dual scenario method goes further by systematically analyzing the scenarios of technology and scenarios of contexts and, only after this is done, understand the whole picture through injecting the scenarios into each other. This approach lets us generate new knowledge about possible futures by focusing on the relevant context (from the point of view of some actors) and then focusing on the effect of general trajectories of AI in that context.

By analyzing how different AI capabilities might show up in various contexts, we can uncover cause-and-effect pathways and interactions that might be missed in more traditional, linear projections – by linear we mean here trajectories without possible turning points that underly contingency (see Beatty 2016 for similar perspective in historical sciences). Our atomic technology thought experiment shows this idea. By thinking about how the existence of atomic weapons in the unstable 1930s might have changed history, we get insights into the complex relationships between technological capabilities, geopolitical tensions, and societal values. This same approach, applied to AI futures through our dual scenario method, can help us analyze how AI capabilities contextual factors might interact in ways we might not see if we only looked at tech development by itself.

When it comes to selecting and validating scenarios, our dual scenario approach fits with best practices in the field while bringing in new ideas. Our approach, which looks several possible AI trajectories and social contexts, lets us explore several of possibilities while still being manageable. The key is selecting scenarios that capture the most important uncertainties without overwhelming ourselves. This aligns with Swart et al.'s (2004) suggestion that scenarios should be coherent and plausible. We validate our scenarios using criteria similar to those used in traditional scenario planning, like plausibility and internal consistency. But we take the idea further by considering how plausible and consistent the interactions between AI development paths and social contexts are. This is achieved by creating two sets of scenarios and injecting them into each other. BThis aligns with Inayatullah's (2008) view of scenarios as the "tool par excellence of futures research" – and we marry together two sets of scenarios here.

Importantly, the dual approach, as described in this paper, teaches us that we are not mere receivers of AI’s impact. We live in the contexts we create. By understanding the contexts, we understand ourselves. And by understanding ourselves, we can affect the contexts. Given that AI develops only in a context, it follows that we can affect how AI develops through knowing ourselves. This aligns with what Masini (2006) describes as “learning to live *with* the future” through systematic futures thinking that is directed at visioning and consequent acting. This already shows that the dual approach forces us to take ethical responsibility about the future of AI because we create the contexts.

At the ethical level, generally speaking, our dual scenario approach builds value considerations in the foresight process itself. The approach covers the normative side of futures research, where the focus is not on what could happen but what would we prefer (desire) to happen (Bell 1997). By thinking about how AI might develop and be used in different social contexts, we naturally raise questions about values, priorities, and the kind of future we want to create. This ethical dimension is especially crucial for AI futures, given the technology's potential to fundamentally reshape social, economic, and political structures. The atomic technology thought experiment highlights the ethical implications of technological development. The decision to use atomic weapons at the end of World War II and the arms race during the Cold War raised ethical questions beyond comprehension about using technology, deterrence, and the responsibilities of scientists and policymakers. Similarly, the development of AI raises critical ethical questions about privacy, autonomy, fairness, and the future of work, among others (see the previous sections). Our method also addresses concerns about inclusivity and power dynamics in futures thinking. By systematically analyzing multiple possible future contexts and trajectories, our dual scenario approach helps democratize the process of imagining AI futures. It allows us to think about the AI through the contexts different sets of actors (ranging from one actor to global population) may find themselves living in.

The atomic technology thought experiment illustrates the importance of considering diverse perspectives in futures thinking. The development and use of atomic weapons were shaped by a particular set of voices and values. Our dual scenario approach to AI futures seeks to broaden the range of perspectives considered, recognizing that the development and deployment of AI will have global implications that affect diverse communities in different ways.

Also, our dual scenario approach helps with the challenge of imagining futures in futures research. Virmajoki (2022) argues truly transformative futures might be beyond what we can currently imagine. This is especially true for AI, where there is potential for big breakthroughs and paradigm shifts. A systematic analysis of multiple contexts and AI trajectories can help push the boundaries of what we can imagine and consider. This fits with Bendor et al.'s (2021) call to expand the futurological imagination and open it up to new possibilities for knowledge and action. Moreover, the dual approach echoes with Ramírez and Wilkinson's (2016) emphasis on the importance of reframing in scenario planning. By considering alternative imageries of futures, we can challenge existing assumptions and create more robust scenarios. In our approach reframing is done through the scenarios of contexts and injections – this shows the trajectories of AI from *a* perspective. Reframing allows us to look at AI development from different perspectives and thus break free from conventional thinking patterns that might limit our understanding of possible futures.

Using historical analogies and "what-ifs" in our approach, like our atomic technology thought experiment, further helps us imagine alternative futures. By thinking about how a major technology like atomic bomb might have developed in different historical contexts, we can get insights into how AI might unfold in various future contexts. This method aligns with van Notten et al.'s (2003) characterization of scenarios as descriptions of possible futures that reflect different perspectives on *the past, present, and future*.

Interestingly, our method also resonates with Dator's (2009) claim that all scenarios fall into one of four archetypes: growth, collapse, discipline, and transformation. Our dual scenario approach allows us to analyze how AI might develop in terms of the archetypes. For example, we can examine how AI might evolve in a growth scenario where technological progress accelerates rapidly, or in a collapse scenario where societal breakdowns impact AI development. We can also consider discipline scenarios where strict regulations shape AI trajectories, and transformation scenarios where AI fundamentally alters social structures and human-machine relationships. What is important here is to notice that, when looked from different contexts, the same AI developments may be associated with different archetype. For example, rapid advancements in AI capabilities could be seen as growth from a technological perspective, but as potential collapse from a labor market or social stability viewpoint. We need scenarios of contexts to truly understand the archetypes of the scenarios of AI. Nothing is a collapse in itself; that something is always a collapse from the point of view of the actors of a context. Consider, for example, the year of 1945 and what it meant to different groups throughout the globe.

The openness and contingency of the future, a core principle in futures research, is covered in our dual scenario approach. As Bell (1997) argues, following Amara, "The future is a domain of liberty not simply because we cannot know the future in any certain sense. It is also because the future itself is contingent, not only of our knowing of it" (p. 151). Our approach acknowledges this fundamental uncertainty while still giving tools to manage this uncertainty through foresight. By analyzing several scenarios of AI and contexts, we can tame, at least a bit, the grand phenomenon of AI development. The atomic bomb thought experiment illustrates this fundamental openness and contingency. The course of history might have been radically different if atomic weapons had been developed earlier or in a different geopolitical context – and if different decisions had been made in early 1940s. Similarly, the future of AI is not predetermined but will be shaped by the interaction of technological developments and societal choices. This is exactly why we need two sets of scenarios: The scenarios of the contexts enable us to see where we might live in and where we want to live in. Then we inject scenarios of AI to these scenarios and attempt to make best of it.

In fact, we can tame it through contingency. Once we see that the future of AI depends on its context and timing, we can make it something that does not only happen to us but allows us to live with it through our own identity: Once we know who we are and what we want, we can adapt the trajectories of AI to suit ourselves. This resonates with Sardar's (2010) second law of futures studies, MAD (Mutually Assured Diversity). This law requires that human diversity is assured and thriving, rather than merely surviving, in any future scenario. Futures research must account for various cultural systems and ensure that all possible paths remain open for the recognition and appreciation of diverse ways of being human.

It follows that our approach also recognizes how futures thinking can shape the future, especially in technology development. Popper (1957) famously pointed out that our ideas about the future can influence how it actually unfolds. This is important for AI, where public views and policy choices can greatly affect the trajectory of the technologies. By thinking about how different contexts might shape AI development, our dual scenario method looks at this back-and-forth between foresight and reality.

To sum up, our dual scenario approach offers a strong and ethically-based method for approaching AI futures. It builds on established scenario planning ideas while adding new elements that fit the complex, heterogenous, and fast-changing field of AI. Our approach helps shape the future of AI in more informed, inclusive, and responsible ways. It gives the tools to use the results of more predictive and forecasting works in a systematic way that takes into account contexts and human actors in those contexts.

6. Conclusion

The main idea of this paper is that, in order to understand the future of AI in its full heterogeneity, we must consider both the possible trajectories of AI development and the various societal contexts where these technologies might emerge and be used. This dual scenario approach, informed by historical analogies and thought experiments like the one concerning the development of atomic technology, provides a framework for understanding and influencing the interplay between AI capabilities and societal factors across different time horizons. This approach allows us to see beyond simplistic and deterministic views of technological progress and instead recognize the interactions between technological capabilities, societal factors, and the element of timing.

Our dual scenario method, with its emphasis on context and timing, offers some advantages. First, it encourages a more holistic view of AI development as it is built on the notion that the impact of these technologies will depend not only on their capabilities, but also on where and when they emerge and the social, economic, and political environments of that time and context. The thought experiment about the earlier development of atomic weapons vividly illustrates how timing can dramatically alter the course of history for centuries. Second, our approach provides a framework for analyzing possible challenges and opportunities that might arise at the intersection of AI capabilities, societal contexts, and specific time horizons. This temporal dimension enables us to understand how AI might evolve and impact society over different time scales. Finally, the dual approach places ethical considerations into the very process of studying the possible futures of AI by encouraging us to consider not just what is technologically possible, but also who we are, who we will be, and what we desire.

REFERENCES

Acosta, J. N., Falcone, G. J., Rajpurkar, P., et al. (2022). Multimodal biomedical AI. *Nature Medicine*, *28*, 1773-1784.

Amer, M., Daim, T. U., & Jetter, A. (2013). A review of scenario planning. *Futures*, *46*, 23-40.

Ayed, I., de Bézenac, E., Pajot, A., Brajard, J., & Gallinari, P. (2022). Modelling spatiotemporal dynamics from Earth observation data with neural differential equations. *Machine Learning*, *111*(6), 2349-2380.

Beatty, J. (2016). What are narratives good for? *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *58*, 33-40.

Bell, W. (1997). *Foundations of Futures Studies Volume 1*. Transaction Publishers.

Bendor, R., Eriksson, E., & Pargman, D. (2021). Looking backward to the future: On past-facing approaches to futuring. *Futures*, *125*, Article 102677.

Bonnet, F., Mazari, J. A., Munzer, T., Yser, P., & Gallinari, P. (2022). An extensible benchmarking graph-mesh dataset for studying steady-state incompressible Navier-Stokes equations. In *Proceedings of ICLR 2022 Workshop on Geometrical and Topological Representation Learning*.

Booth, C., Rowlinson, M., Clark, P., Delahaye, A., & Procter, P. (2009). Scenarios and counterfactuals as modal narratives. *Futures*, *41*(2), 87-95.

Bradfield, R., Derbyshire, J., & Wright, G. (2016). The critical role of history in scenario thinking: Augmenting causal analysis within the intuitive logics scenario development methodology. *Futures*, *77*, 56-66.

Brendon, P. (2002). *The Dark Valley: A Panorama of the 1930s*. Vintage.

Bromham, L., Dinnage, R., & Hua, X. (2016). Interdisciplinary research has consistently lower funding success. *Nature*, *534*(7609), 684-687.

Buckner, C. (2018). Empiricism without magic: Transformational abstraction in deep convolutional neural networks. *Synthese*, *195*(12), 5339-5372.

Bursztyn, M., & Drummond, J. (2014). Sustainability science and the university: Pitfalls and bridges to interdisciplinarity. *Environmental Education Research*, *20*(3), 313-332.

Chen, Z., Xu, J., Liao, M., Xue, T., & He, K. (2022). Two-phase Multi-document Event Summarization on Core Event Graphs. *Journal of Artificial Intelligence Research*, *74*, 1037-1057.

Cirincione, J. (2007). *Bomb Scare: The History and Future of Nuclear Weapons*. Columbia University Press.

Coeckelbergh, M. (2021). Time Machines: Artificial Intelligence, Process, and Narrative. *Philosophy & Technology*, *34*, 1623-1628.

Cowen-Rivers, A. I., Palenicek, D., Moens, V., Mohammed, M., Sornette, A., & Wang, J. (2022). SAMBA: safe model-based and active reinforcement learning. *Machine Learning*, *111*(1), 173-203.

Darwin, C., & Huxley, T. H. (2003). *The Origin of Species: 150th Anniversary Edition*. Signet.

Dator, J. (2009). Alternative futures at the Manoa School. *Journal of Futures Studies, 14*(2), 1–18.

Degnan, J. H., & Rosenberg, N. A. (2009). Gene tree discordance, phylogenetic inference, and the multispecies coalescent. *Trends in Ecology & Evolution*, *24*(6), 332-340.

Erasmus, A., Brunet, T. D. P., & Fisher, E. (2021). What is Interpretability? *Philosophy & Technology*, *34*(4), 833-862.

Faisal, A., Kamruzzaman, M., Yigitcanlar, T., & Currie, G. (2019). Understanding autonomous vehicles: A systematic literature review on capability, impact, planning and policy. *Journal of Transport and Land Use*, *12*(1), 45-72.

Fausett, L. (1994). *Fundamentals of Neural Networks: Architectures, Algorithms and Applications*. Prentice Hall.

Gasteiger, J., & Zupan, J. (1993). Neural Networks in Chemistry. *Angewandte Chemie International Edition in English*, *32*(4), 503-527.

Goertzel, B. (2014). Artificial General Intelligence: Concept, State of the Art, and Future Prospects. *Journal of Artificial General Intelligence, 5*(1), 1-46.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Professional.

Gosling, F. G. (2005). *The Manhattan Project: Making the Atomic Bomb*. University Press of the Pacific.

Green, A. (2012). Continuity, contingency and context: Bringing the historian's cognitive toolkit into university futures and public policy development. *Futures, 44*(2), 174-180.

Hambardzumyan, D., Parada, L. F., Holland, E. C., & Charest, A. (2011). Genetic modeling of gliomas in mice: New tools to tackle old problems. *Glia*, *59*(8), 1155-1168.

Herrera, G. L. (2006). *Technology and International Transformation: The Railroad, the Atom Bomb, and the Politics of Technological Change*. State University of New York Press.

Humphreys, P. (2009). The philosophical novelty of computer simulation methods. *Synthese*, *169*(3), 615-626.

Inayatullah, S. (2008). Six pillars: futures thinking for transforming. *Foresight*, *10*(1), 4-21.

Kramer, O. (2017). *Genetic Algorithm Essentials*. Studies in Computational Intelligence. Springer.

Laakkonen, M-P., Kivivirta, V., & Mazari, A. (2024). *Artificial Intelligence Ultimate Nature of Being: The Background Assumptions of Modelling AI Algorithms*. Cambridge Scholars Publishing.

Lee, R. S. T. (2019). AI Powerful Tools in Quantum Finance. In *Quantum Finance* (pp. 159-208). Springer.

Masini, E. (2006). Rethinking futures research. *Futures*, *38*(10), 1158-1168.

Millar, V. (2016). Interdisciplinary curriculum reform in the changing university. *Teaching in Higher Education*, *21*(4), 471-483.

Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology*, *33*(4), 659-684.

Mucha, T., & Seppala, T. (2020). *Artificial Intelligence Platforms – A New Research Agenda for Digital Platform Economy*. ETLA Working Papers No. 76.

Nie, Y., Bo, H., Zhang, W., & Zhang, H. (2020). Research on Hybrid Wind Speed Prediction System Based on Artificial Intelligence and Double Prediction Scheme. *Scientific Programming, 2020,* Article ID 9601763.

Noorbakhsh-Sabet, N., Zand, R., Zhang, Y., & Abedi, V. (2019). Artificial Intelligence Transforms the Future of Health Care. *The American Journal of Medicine*, *132*(7), 795-801.

Overy, R., & Wheatcroft, A. (2000). *The Road to War* (Rev. ed.). Penguin.

Pelacho, M., Rodríguez, H., Broncano, F., Kubus, R., Sanz García, F., Gavete, B., & Lafuente, A. (2021). Science as a commons: Improving the governance of knowledge through citizen science. In K. Vohland et al. (Eds.), *The Science of Citizen Science* (pp. 79-94). Springer.

Peng, Y., Han, J., Zhang, Z., Fan, L., Liu, T., Qi, S., Feng, X., Ma, Y., Wang, Y., & Zhu, S. C. (2024). The Tong Test: Evaluating Artificial General Intelligence Through Dynamic Embodied Physical and Social Interactions. *Engineering*, *34,* 12-22.

Pimentel, E., Cho, C. H., & Bothello, J. (2023). The blind spots of interdisciplinarity in addressing grand challenges. *Critical Perspectives on Accounting*, *93,* Article 102475.

Popper, K. (1957). *The Poverty of Historicism*. Routledge.

Rai, R., & Sahu, C. K. (2020). Driven by data or derived through physics? A review of hybrid physics-guided machine learning techniques with cyber-physical system (CPS) focus. *IEEE Access*, *8*, 71050-71073.

Ramírez, R., & Wilkinson, A. (2016). *Strategic reframing: The Oxford scenario planning approach.* Oxford University Press.

Rhodes, R. (1988). *The Making of the Atomic Bomb*. Simon & Schuster.

Rhodes, R. (1995). *Dark Sun: The Making of the Hydrogen Bomb*. Simon & Schuster.

Salmela, M., MacLeod, M., & Munck Af Rosenschöld, J. (2021). Internally incentivized interdisciplinarity: Organizational restructuring of research and emerging tensions. *Minerva*, *59*(3), 355-377.

Sandeep Singh Sengar, Affan Bin Hasan, Sanjay Kumar & Fiona Carroll (2024). Generative artificial intelligence: a systematic review and applications. *Multimedia Tools and Applications*.

Sardar, Z. (2010). The namesake: Futures; futures studies; futurology; futuristic; foresight—What's in a name? *Futures*, *42*(3), 177-184.

Siau, K., & Wang, W. (2020). Artificial Intelligence (AI) Ethics: Ethics of AI and Ethical AI. *Journal of Database Management*, *31*(2), 74-87.

Staley, D. J. (2002). A History of Future. *History and Theory*, *41*(Theme Issue), 72-89.

Staley, D. J. (2010). *History and Future: Using Historical Thinking to Imagine the Future*. Lexington Books.

Steger, M. B. (2019). Committing to cultures of creativity: The significance of transdisciplinarity. *Globalizations*, *16*(5), 763-769.

Swart, R. J., Raskin, P., & Robinson, J. (2004). The problem of the future: Sustainability science and scenario analysis. *Global Environmental Change*, *14*(2), 137-146.

Tjörnbo, O., & McGowan, K. (2022). A complex-systems perspective on the role of universities in social innovation. *Technological Forecasting and Social Change*, *174*, Article 121247.

van Notten, P. W., Rotmans, J., Van Asselt, M. B., & Rothman, D. S. (2003). An updated scenario typology. *Futures*, *35*(5), 423-443.

Virmajoki, V. (2022). Limits of conceivability in the study of the future: Lessons from philosophy of science. *Futures*, *142*, Article 102993.

Virmajoki, V. (2023a). Frameworks in Historiography: Explanation, Scenarios, and Futures. *Journal of the Philosophy of History*, *17*(2), 167-193.

Virmajoki, V. (2023b). *Causal Explanation in Historiography*. Palgrave Macmillan.

Virmajoki, V. (2024). On the Function and Nature of Historical Counterfactuals: Clarifying Confusions. *Journal of the Philosophy of History*. [Add DOI when available]

Wright, G., Bradfield, R., & Cairns, G. (2013). Does the intuitive logics method – and its recent enhancements – produce "effective" scenarios? *Technological Forecasting and Social Change*, *80*(4), 631-642.

Zednik, C. (2021). Solving the Black Box Problem: A Normative Framework for Explainable Artificial Intelligence. *Philosophy & Technology*, *34*(2), 265-288.

Zupan, J., & Gasteiger, J. (1995). Neural networks in chemistry. *Journal of Chemical Physics*, *1*(1), 503-521.