

This Time Is Different: A comparison of three waves of AI Adoption in the Sciences

January 23, 2026

Abstract

This study examines the diffusion of artificial intelligence within the scientific enterprise to understand whether the recent wave of generative AI represents a fundamental departure from previous technological paradigms. We undertake a large-scale quantitative analysis of 4.8 million academic publications from 1950 to 2024, systematically mapping the adoption patterns of three distinct AI waves: deep learning, natural language processing, and GenAI. Our analysis reveals that GenAI has diffused with unprecedented speed and breadth across scientific disciplines. We find that GenAI adoption is associated with markedly lower concentration levels and is correlated with the formation of smaller research teams, suggesting a reversal of long-standing trends toward greater specialization and larger team sizes in science. These findings indicate that GenAI constitutes a significant shift in the relationship between AI and scientific discovery.

1 Introduction

Artificial intelligence (AI) appears to be reshaping the practice of science at a scale and speed that few anticipated. From astronomy to zoology, from theoretical physics to clinical medicine, artificial intelligence has become an increasingly important tool for hypothesis generation, data analysis, pattern recognition, and knowledge discovery in science (Gao and Wang, 2024). The appearance of rapid adoption by scientists raises several foundational questions about the trajectory of AI in science. Has the relationship between AI and scientific discovery undergone any profound transformations in the last few years, upon the spread of generative artificial intelligence? To what extent has AI evolved from a niche computational approach into a pervasive force across the scientific enterprise? Has it become a force capable of reshaping every corner of science? Answering these questions requires a systematic, evidence-based approach that goes beyond anecdotal observations.

The broader scientific community's engagement with artificial intelligence has historically been shaped by the technology's demanding requirements for specialized knowledge and computational infrastructure. AI applications in science have mostly necessitated substantial complementary investments—particularly in technical expertise and collaborative partnerships with computer scientists. This pattern has created a dynamic where the benefits of AI adoption have been concentrated

among institutions and researchers with access to these complementary resources and AI-engaged research spreads across fields but does not fully integrate with traditional disciplinary approaches (Duede et al., 2024). Moreover, science as an institution presents two unique challenges for AI adoption: the need for and emphasis on transparency and reproducibility can create tensions with "black box" technologies like AI, while established disciplinary boundaries may resist cross-cutting technological change, where established ways of working, knowing, and credit allocation may be upended.

Our paper addresses these questions by undertaking a large-scale, quantitative analysis of the diffusion of AI in science from 1950 to the present day. We systematically map the adoption of AI across all scientific fields, providing comprehensive empirical evidence of its trajectory. We describe the long-term trends, examining the patterns of use and adoption over time and across disciplines. We further stratify our analysis by three distinct technological waves in AI—deep learning, natural language processing, and generative AI—to better understand the forces governing this diffusion process. By doing so, we aim to move beyond speculation and provide a descriptive account of AI’s evolving role in science, revealing whether this time is, indeed, different.

Our analysis of 4.8 million papers comprising of 6.7 million authors across all scientific fields) shows that AI has indeed undergone a notable transformation in its relationship with scientific discovery, with this transformation being particularly pronounced since 2015, when advances in deep learning and computational power converged to create opportunities for scientific applications (e.g., Chen et al. (2018)). The integration of AI into scientific workflows has enabled researchers to tackle previously intractable problems, from predicting protein structures to analyzing vast astronomical datasets, from modeling climate systems to accelerating drug discovery . Our analysis reveals that this remarkable diffusion of AI across scientific disciplines has followed patterns that are characteristic of general-purpose technologies —requiring substantial complementary investments, exhibiting uneven adoption across fields, and demanding significant organizational and methodological adaptations. Typically, AI applications in science have necessitated "complementary assets": the human capital, organizational capabilities, and institutional arrangements required to effectively deploy and benefit from technological innovations (Teece 1986).

However, the recent widespread adoption of generative artificial intelligence (GenAI) represents a potential inflection point in this evolutionary trajectory (Noy and Zhang, 2023; Brynjolfsson et al., 2025; Dell’Acqua et al., 2023). The introduction of large language models, diffusion models, and other GenAI systems—particularly following the public release of ChatGPT and similar tools since late 2022—has altered the accessibility and usability of AI capabilities for scientific research. Unlike previous waves of AI adoption that required extensive technical expertise and collaborative arrangements with computer science departments, GenAI tools offer what might be characterized as "plug-and-play" functionality, enabling researchers across disciplines to leverage sophisticated AI

capabilities with minimal additional training or infrastructure investment. In essence, this is gaining access to complementary expertise through the AI models instead of relying on other human collaborators (Dell’Acqua et al., 2025a)

GenAI has important implications for knowledge work tasks in the scientific process. Scientists stand to be significantly influenced by the advent of new AI capabilities: LLMs have demonstrated ability to automate intricate scientific tasks with remarkable efficiency (e.g., Huang et al. (2023)) and have scored well on tests of scientific reasoning (e.g., West (2023)). They also show promise in being able to contribute significantly to fields such as mathematics (NASEM, 2023), biology (Tong and Zhang 2023), and astronomy (Medeiros et al. 2023). Their potential extends to applied domains, such as medicine (Lind Plesner et al. 2023).

We present early evidence that the diffusion of GenAI in science represents a fundamental departure from established patterns of technological adoption—one that challenges core assumptions about how transformative technologies spread through knowledge-intensive institutions. Our analysis reveals three crucial dimensions along which GenAI differs from previous waves of AI adoption in science:

First, GenAI is diffusing across scientific disciplines at an unprecedented pace. While previous AI technologies exhibited the gradual adoption curves characteristic of GPTs, GenAI tools have achieved widespread adoption within months rather than years of their introduction. This acceleration reflects not merely improved marketing or reduced costs, but a clear change in the accessibility and immediate utility of AI capabilities for scientific research. Second, the diffusion of GenAI exhibits markedly greater equality across scientific fields compared to previous AI waves. Our analysis of adoption patterns reveals significantly lower concentration coefficients for GenAI usage compared to deep learning and NLP applications, suggesting that the benefits of this technology are being distributed more evenly across disciplines. This pattern contrasts sharply with the concentrated adoption observed in earlier AI waves, where fields with stronger computational traditions and closer ties to computer science departments captured disproportionate benefits.

Third, and most disruptively, GenAI appears to break the traditional requirement for complementary assets that has characterized previous GPTs. Unlike deep learning and NLP applications, which typically require increased collaboration with computer science experts and substantial investments in specialized human capital, GenAI tools can be deployed effectively by researchers without extensive additional training or institutional support. This "plug-and-play" characteristic alters the economics of AI adoption and may explain both the speed and equality of its diffusion. Strikingly, our data reveals that teams using GenAI are systematically smaller than teams using traditional AI algorithms—suggesting that GenAI may be reversing decades-long trends toward larger teams and greater specialization in science (Wuchty et al., 2007). If sustained, this shift could fundamen-

tally alter the human capital requirements for frontier research and lower longstanding barriers to scientific innovation.

Our findings have implications for our understanding of technological diffusion in science and raise important questions about the institutional and methodological challenges that GenAI poses for the scientific enterprise. If GenAI indeed has severed the traditional link between complementary investments and productivity gains, then existing theoretical frameworks may be inadequate for predicting its long-term impacts on scientific practice, institutional arrangements, and knowledge production processes. Moreover, the rapid and widespread adoption of tools that operate as "black boxes" creates potential tensions with fundamental scientific values of transparency, reproducibility, and explainability that merit careful examination (Lebovitz et al. 2022). This tension is further underscored by the possibility that AI may not simply augment the existing scientific process but fundamentally reorganize it, shifting the goal from discovering human-interpretable theories to building and probing complex, learned representations of the world (Mullainathan and Rambachan, 2025).

Here we situate our inquiry within the canonical view of general-purpose technologies (GPTs). Established accounts emphasize that transformative technologies diffuse only after organizations undertake sizable complementary investments—in skills, tools, workflows, and infrastructure (Bresnahan and Trajtenberg, 1995). By contrast, for the present wave of generative AI high-capability systems are accessible through simple interfaces and can be embedded into existing cognitive workflows with comparatively modest upfront cost.

A second implication concerns who benefits and how quickly. Long-standing trends in science—rising specialization, larger teams, and the burden of knowledge—suggest that frontier work typically concentrates among highly skilled, well-resourced actors (Wuchty et al. 2007; Jones 2009, 2010). GenAI may relax some of these constraints by lowering fixed costs for advanced tasks, allowing earlier or peripheral participants to contribute more readily and enabling established teams to reallocate effort toward higher-value activities. Overall, while it is too early to draw firm conclusions, GenAI may ultimately reverse some long-standing trends in science.

The empirical evidence presented in this paper suggests that GenAI's diffusion patterns represent a genuine departure from established norms, with implications that extend far beyond the adoption of a new set of research tools. We are picking up early signals of patterns that may define the future of scientific research. While GenAI's widespread adoption began only in late 2022, the unprecedented speed of diffusion, the equality of adoption across fields, and the reduction in team sizes we document are already visible. These are not speculative projections but empirical observations of a transformation in progress. There is no reason to expect these patterns are temporary artifacts. Rather, we may be witnessing the opening stages of a fundamental shift in how scientific

knowledge is produced and by whom.

2 Theoretical Framework

To understand the significance of GenAI’s distinctive diffusion patterns, it is essential to situate this analysis within the broader theoretical framework of GPTs. The concept of GPTs provides a powerful lens for understanding how certain technologies come to drive "whole eras of technical progress and economic growth" (Bresnahan and Trajtenberg, 1995). The historical examples like steam power and electricity shared a common pattern of diffusion characterized by initially slow adoption followed by rapid acceleration as complementary assets and supporting infrastructure developed (Bresnahan and Trajtenberg 1995). The steam engine, for instance, required not only mechanical improvements but also the development of new manufacturing processes, transportation systems, and organizational forms before its transformative potential could be realized.

Crucially, the GPT framework emphasizes that the economic and social impacts of these technologies depend heavily on the development of complementary assets—the specialized knowledge, organizational capabilities, and institutional arrangements necessary to effectively deploy and benefit from the technology (Hall and Trajtenberg 2004). This insight suggests that the distribution of GPT benefits is often determined not by access to the technology itself, but by the ability to develop and control the complementary resources required for its effective utilization (Teece, 1986)

The application of GPT theory to artificial intelligence has gained considerable traction among economists and innovation scholars. In particular, Agrawal, Gans, and Goldfarb posit that AI functions as a general purpose technology primarily by notably reducing the cost of prediction—a crucial input to decision-making across virtually all domains of human activity (Agrawal et al. 2018). This perspective reframes AI not as a collection of specific applications, but as a transformative capability that can enhance productivity and enable new forms of value creation across diverse sectors and disciplines.

The GPT framework also suggests that AI adoption should exhibit the characteristic pattern of initially slow diffusion followed by acceleration as complementary assets develop. This pattern has indeed been observed in previous waves of AI adoption in science, specifically deep learning and natural language processing (NLP), which we use as benchmarks in our analyses. The requirement for such complementary assets has created barriers to adoption that have shaped the uneven distribution of AI benefits across scientific disciplines and institutions.

Despite the technology’s obvious potential, measurable productivity gains have been slower to materialize than many observers expected, a pattern that some attribute to the time and investment required to develop the complementary assets necessary for effective AI deployment and define it as

“AI productivity paradox” (Brynjolfsson et al., 2019). This analysis suggests that the full economic and scientific benefits of AI may take decades to realize, as organizations gradually adapt their processes, train their personnel, and develop new forms of human-AI collaboration.

Recent work by Agrawal et al. (2025) develops a task-based model that explores how AI augments different stages of the scientific process, with particular attention to the "jagged frontier" of AI capabilities—the uneven landscape where AI excels at some tasks while remaining limited in others (Dell’Acqua et al., 2023). Drawing on frameworks from the economics of automation and task-based models (Acemoglu and Restrepo, 2018), they examine how AI can aid hypothesis generation through pattern recognition in complex data (Ludwig and Mullainathan 2024) and enhance understanding through computational tools that reveal structure in high-dimensional spaces (Krenn et al., 2022). Central to their framework is the role of specialized human capital, distinguishing between "AI-expert scientists," who are positioned to reap significant productivity gains, and "ordinary scientists." Their analysis focuses primarily on the need for complementary investments in AI expertise to realize productivity benefits from the technology. In contrast, Mullainathan and Rambachan (2025) offer a more radical vision, arguing that AI may fundamentally reorganize the scientific process itself: shifting from the discovery of human-interpretable theories to the construction and interrogation of learned, high-dimensional representations of the world. This vision draws inspiration from historical analogies, like the adoption of electricity in manufacturing (David 1990). Our empirical analysis provides early evidence on these contrasting perspectives: we demonstrate that the accessibility of GenAI fundamentally alters traditional diffusion patterns, suggesting that ease-of-use is not merely a moderating factor but a transformative force that reshapes the economics of AI adoption in science.

AI Waves

To understand the distinctive diffusion patterns of GenAI in scientific research, we focus on two major waves of artificial intelligence that have shaped the contemporary research landscape: deep learning and NLP, and use these as a comparison for GenAI. Deep learning resembles GenAI in its technical sophistication and widespread use, while NLP resembles the most widespread family of GenAI algorithms, large language models, in their domain. Each wave represents not merely a technological advancement, but a shift in how researchers approach computational problems, with distinct algorithmic foundations and practical implementations that have evolved from specialized research tools to accessible platforms for scientific inquiry.

Deep Learning

Deep learning emerged as a transformative force in artificial intelligence through the development of multi-layered neural networks capable of learning hierarchical representations from data.

The algorithmic foundations of deep learning trace back to the perceptron model of the 1950s and the backpropagation algorithm developed in the 1980s, but the field experienced its renaissance in the 2000s when computational advances and large datasets converged to make training deep neural networks feasible (LeCun et al., 2015). The core algorithmic innovation lies in the ability of these networks to automatically learn feature representations through multiple layers of nonlinear transformations, eliminating the need for manual feature engineering that had constrained earlier machine learning approaches.

The foundational architecture of deep learning systems consists of interconnected layers of artificial neurons, each applying learned weights and biases to input data before passing the transformed information to subsequent layers. Convolutional Neural Networks (CNNs), introduced by LeCun and colleagues, revolutionized image processing by incorporating spatial locality and translation invariance through convolutional operations and pooling layers (LeCun and Bengio, 1998). Recurrent Neural Networks (RNNs) and their more sophisticated variants, Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), addressed sequential data processing by maintaining internal memory states that allow the network to process temporal dependencies (Hochreiter and Schmidhuber, 1997). These architectural innovations enabled deep learning systems to achieve higher performance on tasks ranging from image classification to speech recognition.

The algorithmic sophistication of deep learning systems required substantial computational resources and technical expertise for implementation. Training deep neural networks involves solving high-dimensional optimization problems through gradient descent algorithms, requiring careful initialization strategies, regularization techniques, and hyperparameter tuning to achieve convergence. The computational demands of these algorithms necessitated specialized hardware, particularly Graphics Processing Units (GPUs), which could perform the parallel matrix operations essential for efficient neural network training (Raina et al., 2009). This technical complexity created significant barriers to adoption, limiting deep learning applications to research groups with access to computational resources and specialized knowledge.

The transition from algorithms to accessible tools in deep learning occurred gradually through the development of software frameworks that abstracted much of the underlying complexity. TensorFlow, released by Google in 2015, provided a comprehensive platform for building and deploying deep learning models, offering high-level application programming interfaces (APIs) that simplified common operations while maintaining flexibility for advanced users (Abadi et al., 2016). PyTorch, developed by Facebook's AI Research lab, introduced dynamic computational graphs and a more intuitive programming interface that accelerated research and experimentation (Paszke et al., 2019). These frameworks, along with others like Keras and Caffe, democratized access to deep learning capabilities by providing pre-built architectures, optimization algorithms, and training procedures that researchers could adapt to their specific domains. Despite these advances in tooling, deep

learning applications in scientific research continued to require substantial complementary assets. Researchers needed to understand the mathematical foundations of neural networks, the principles of gradient-based optimization, and the intricacies of AI model architecture design. The complexity of hyperparameter tuning, model validation, and interpretation of results created additional barriers that limited adoption to institutions and research groups with appropriate technical capabilities. Using these tools required at least basic knowledge of neural network architecture. There were numerous discretionary decisions the user of these libraries had to make. Successful implementations typically required collaboration with computer science experts or significant investment in technical training.

Natural Language Processing (NLP)

Natural Language Processing represents a distinct wave of AI development focused on enabling computers to understand, interpret, and generate human language. The field has evolved through several paradigmatic shifts, from rule-based systems of the 1960s and 1970s to statistical approaches of the 1990s and 2000s, culminating in the neural revolution that began in the 2010s (Jurafsky and Martin, 2019). Each phase brought algorithmic innovations that progressively improved the ability of computational systems to process linguistic data, but also introduced new complexities and requirements for successful implementation.

The early algorithmic foundations of NLP relied heavily on linguistic theory and hand-crafted rules. Parsing algorithms like the Cocke-Younger-Kasami (CYK) algorithm and Earley's algorithm provided systematic approaches to syntactic analysis, while semantic processing relied on knowledge representation schemes and logical inference systems. These rule-based approaches achieved impressive results in constrained domains but struggled with the ambiguity, variability, and contextual dependencies that characterize natural language use. The brittleness of rule-based systems and the enormous effort required to encode linguistic knowledge for new domains limited their scalability and practical applicability.

The statistical revolution in NLP, beginning in the 1990s, introduced probabilistic models that could learn patterns from large text corpora rather than relying solely on hand-crafted rules. Hidden Markov Models (HMMs) became the foundation for part-of-speech tagging and named entity recognition, while n-gram language models enabled probabilistic prediction of word sequences (Manning and Schütze, 1999). The introduction of the Expectation-Maximization (EM) algorithm and other unsupervised learning techniques allowed these models to discover latent linguistic structures from unlabeled text data. Statistical machine translation systems, exemplified by IBM's alignment models and phrase-based translation systems, demonstrated the power of data-driven approaches to complex linguistic tasks (Brown et al., 1993).

The algorithmic sophistication of statistical NLP systems required expertise in probability theory, information theory, and computational linguistics. Researchers needed to understand concepts like maximum likelihood estimation, Bayesian inference, and various smoothing techniques to handle data sparsity. The development of effective features for statistical models required deep linguistic knowledge combined with computational skills. Training and evaluation of these systems demanded large annotated corpora and sophisticated experimental methodologies that were not readily available to researchers outside the NLP community.

The neural revolution in NLP, beginning around 2010, introduced distributed representations and deep learning architectures specifically designed for language processing. Word embeddings, particularly the Word2Vec and GloVe algorithms, provided dense vector representations that captured semantic relationships between words through their distributional properties in large text corpora (Mikolov et al. 2013). These representations enabled neural networks to process linguistic data more effectively by encoding semantic similarity in geometric relationships within high-dimensional vector spaces.

Tools for NLP development evolved to provide higher-level abstractions, but still required substantial technical knowledge for effective use. Libraries like NLTK and spaCy provided implementations of common NLP algorithms and preprocessing tools, while frameworks like Hugging Face Transformers offered pre-trained models and fine-tuning capabilities (Wof et al. 2020). However, successful application of these tools to scientific research problems typically required an understanding of linguistic concepts, familiarity with machine learning principles, and the ability to adapt existing models to domain-specific requirements.

Generative Artificial Intelligence (GenAI)

Generative artificial intelligence represents a paradigmatic shift in AI development, characterized by systems capable of creating new content—text, images, code, and other media—that exhibits human-like quality and creativity. While the algorithmic foundations of GenAI build upon decades of research in machine learning and neural networks, the emergence of large-scale generative models has altered the accessibility and applicability of AI technologies across diverse domains, including scientific research.

The algorithmic core of modern GenAI rests on several key innovations that enable the modeling of complex, high-dimensional probability distributions. Generative Adversarial Networks (GANs) established a framework for training generative models through adversarial competition between a generator network that creates synthetic data and a discriminator network that attempts to distinguish real from generated samples (Goodfellow et al. 2014). This adversarial training paradigm enabled the generation of remarkably realistic images and other media, but suffered from training

instability and mode collapse issues that limited their practical applicability.

The transformer architecture, developed in a seminal 2017 paper by a majority of Google researchers titled "Attention is All You Need", which was originally designed to work with textual data, became the foundation for the most impactful GenAI systems. Large Language Models (LLMs) demonstrate the power of scaling transformer architectures to billions or trillions of parameters, trained on vast text corpora using self-supervised learning objectives. These models learn to predict the next token in a sequence, but this simple objective enables them to capture complex patterns of language use, world knowledge, and reasoning capabilities that emerge from the scale of training data and model parameters.

Diffusion models represent another crucial algorithmic innovation in GenAI, particularly for image and multimedia generation. These models learn to reverse a gradual noise corruption process, starting from pure noise and iteratively refining the output to produce high-quality samples (Ho et al. 2020). Stable Diffusion and similar systems demonstrate how diffusion models can generate photorealistic images from text descriptions, opening new possibilities for creative and scientific applications.

The algorithmic sophistication of GenAI systems represents the culmination of decades of research in machine learning, requiring expertise in deep learning, optimization theory, and large-scale distributed computing. Training large language models involves solving optimization problems with billions of parameters using massive datasets and computational resources that are accessible only to major technology companies and well-funded research institutions. The technical complexity of implementing attention mechanisms, handling numerical stability in deep networks, and managing the computational requirements of large-scale training creates substantial barriers for researchers seeking to develop their own GenAI systems.

However, the transformation of GenAI from algorithms to accessible tools represents a clear departure from previous AI waves. The release of ChatGPT by OpenAI in November 2022 marked a watershed moment in AI accessibility, providing a conversational interface that allows users to leverage sophisticated language generation capabilities without any technical knowledge of the underlying algorithms. Unlike previous AI tools that required programming skills, understanding of machine learning concepts, or familiarity with specialized software frameworks, ChatGPT and similar systems can be accessed through simple web interfaces using natural language prompts.

The ecosystem of GenAI tools has rapidly expanded to encompass diverse modalities and applications. GPT-4 and its variants provide advanced text generation, analysis, and reasoning capabilities through conversational interfaces. Claude, developed by Anthropic, offers similar capabilities with an emphasis on safety and helpfulness. Specialized tools like GitHub Copilot integrate GenAI

directly into software development environments, providing code completion and generation capabilities. Image generation tools like DALL-E, Midjourney, and Stable Diffusion enable users to create sophisticated visual content from text descriptions. Audio generation systems like Whisper for speech recognition and various text-to-speech systems extend generative capabilities to multimedia content.

The "plug-and-play" nature of these tools alters the requirements for AI adoption in scientific research. Researchers can now access state-of-the-art AI capabilities without investing in computational infrastructure, learning programming languages, or developing expertise in machine learning algorithms. The natural language interfaces of these systems allow researchers to express their needs in familiar terms rather than translating their problems into technical specifications. This accessibility has enabled rapid adoption across diverse scientific disciplines, from humanities scholars using ChatGPT for literature analysis to physicists employing GenAI for hypothesis generation and experimental design .

The democratization of AI access through generative tools has notable implications for the distribution of AI benefits in scientific research. Previous AI waves concentrated benefits among institutions and researchers with access to computational resources, technical expertise, and collaborative relationships with computer science departments. GenAI tools, by contrast, can be effectively utilized by any researcher with internet access and basic digital literacy skills. This shift from specialized technical tools to accessible consumer applications represents a notable change in the economics of AI adoption, potentially explaining the rapid and egalitarian diffusion patterns observed in GenAI usage across scientific fields.

The distinction between GenAI algorithms and tools also reveals important differences in adoption patterns and requirements. While the development and training of GenAI algorithms continues to require substantial technical expertise and computational resources, the deployment of these algorithms as accessible tools has created new pathways for AI adoption that bypass traditional barriers. This separation between algorithm development and tool usage represents a novel pattern in technology diffusion, where the benefits of sophisticated AI systems can be captured without direct engagement with their technical complexity.

The implications of this algorithmic-tool distinction extend beyond mere accessibility to encompass fundamental questions about the nature of scientific expertise and the role of technical knowledge in research. As GenAI tools become increasingly sophisticated and accessible, they may reshape not only how research is conducted but also what kinds of knowledge and skills are valued in scientific communities. Understanding these transformations requires careful analysis of how different types of GenAI applications—algorithmic versus tool-based—are being adopted across scientific disciplines and what this reveals about the changing nature of scientific practice in the age

of artificial intelligence.

3 Data

We have used Dimensions.ai as the source of our publications datasets. It is a proprietary dataset that records the universe of publications from various sources. It includes conference proceedings and preprints as well as published, peer-reviewed articles. Dimensions includes 159M publications as of August 2025. We accessed the publications and downloaded the relevant metadata using both the API and Google BigQuery access options. To compile each of the datasets discussed in this paper, we compiled relevant vocabularies and downloaded each of the entries on the Dimensions publications dataset that included at least one of the terms in the relevant vocabulary in either the title column or the abstract column. Including a technical term related to AI or one of its waves in the title or abstract of a paper signifies a considerable engagement with that concept, tool, or algorithm. Dimensions offers more than 70 columns of data about the publications. These include information such as title, abstract, date published, arXiv ID, if any, and authors. We only downloaded 71 of these columns. The full text of the publications is not available for download, even though the function of full-text search is offered. We did not use this option, as it would lead to the inclusion of many papers that mentioned AI terms somewhere other than the body of the publication. We compiled a list of 267 AI terms and used those to identify AI papers.

Dimensions gives each paper field tags. Each publication can belong to one or more academic fields, and sometimes, subfields are specified as well. The fields used by Dimensions come from the Australian and New Zealand Standard Research Classification (ANZSRC), 2020, which is the most recent version of this classification. According to ANZSRC, there are 23 academic fields. These fields are as follows: Agricultural, Veterinary, and Food Sciences, Biological Sciences, Biomedical and Clinical Sciences, Built Environment and Design, Chemical Sciences, Commerce, Management, Tourism and Services, Creative Arts and Writing, Earth Sciences, Economics, Education, Engineering, Environmental Sciences, Health Sciences, History, Heritage and Archeology, Human Society, Indigenous Studies, Information and Computing Sciences, Language, Communication and Culture, Law and Legal Studies, Mathematical Sciences, Philosophy and Religious Studies, Physical Sciences, and Psychology.

Identifying AI Papers

The first step in compiling a dataset of AI publications was building a comprehensive list of major AI terms. We started with an initial list of 56 AI terms compiled by data science experts. We tried to make sure the major waves of AI and areas of application are represented in this seed lexicon so that we ensure the resultant vocabulary does not miss any major subset of AI publica-

tions. We then prompted the GPT4 API to give ten other AI terms, including tools, concepts, and algorithms closely associated with each one of the 56 original ones. This process was repeated 100 times to minimize the impact of stochasticity in the GPT-4 model’s output, and all of the returned terms were collected. After deleting duplicates and ambiguous terms that have usage outside of AI, we were left with a final list comprised of 271 AI terms. Finally, to obtain a comprehensive dataset of AI publications, we downloaded the metadata of all of the papers that include at least one of these terms in their title or abstract, yielding a dataset of more than .3 million publications from 1950 until the end of 2024. This dataset is primarily used to demonstrate the overall trends of AI adoption in academic disciplines and count the total number of AI papers in a given year.

To construct this dataset, we made sure to include all of the major concepts and algorithms in AI. This is ensured by the careful crafting of the seed lexicon and the iterative process we employed when constructing our final AI vocabulary. We do not claim that we have captured all AI terms, as that seems an impossible undertaking. We are not bothered by the absence of less common AI terms from our vocabulary, as they will not introduce a substantial bias in our data. As our data shows, 95% of AI publications can be explained by only 95 AI terms.

Deep Learning, Natural Language Processing, and Generative AI Papers

The deep learning (DL), natural language processing (NLP), and Generative AI vocabularies were obtained using ChatGPT’s deep research tool. We started by isolating 19 AI terms from our main AI vocabulary, which were NLP terms. These included bag-of-words and text embedding. We then used these to teach ChatGPT’s deep research tool what we mean by NLP AI terms. Then we asked it to return an expanded list of NLP terms, and this process was repeated for deep learning and generative AI as well. The returned lists were then cleaned of duplicates and general terms. This resulted in 137 NLP terms, 116 generative AI terms, and 124 deep learning terms.

Our generative AI lexicon includes tools, concepts, and algorithms that have the potential to generate new content. Any researcher on AI knows that the spaces of NLP, deep learning, and generative AI are not mutually exclusive. Nor are the boundaries between them well-defined. However, for the most part, there are well-established algorithms, concepts, and tools that are agreed upon by the scientific community to reside in each of these spaces. We aimed at capturing such terms.

Initially, ‘transformer’ was one of the terms in the generative AI lexicon. When we downloaded GenAI publications from Dimensions, upon inspection, “transformer” proved to be a problematic term, as it yielded many false positive papers in the field of engineering, so we went back to the deep research tool and constructed more granular collocations with “transformer” that specifically referred to GenAI, such as “vision transformer”. This resulted in a vocabulary of 152 GenAI terms including various GenAI terms including various AI-related collocations of transformers.

Although we included many more terms when constructing the vocabulary for the three generations of AI than we did with the main AI dataset, we should note that only one deep learning term (‘neural networks’) determines 95% of the deep learning dataset; for NLP, 16 terms determine 95% of the dataset, and for generative AI, 7 terms cover 95% of the dataset. Of the 16 NLP terms and 7 generative AI terms, 8 and 3 are present in the main AI vocabulary. DL, NLP, and generative AI datasets have 1.19 million unique publications. In the results of this paper, we merge these four datasets and use it as our main AI dataset.

Unpacking Generative AI

We divide up GenAI vocabulary into two main subsets: those relating to tools, and those relating to algorithms. ChatGPT was released in November 2022, and since then, many papers in various fields have used it to study the impact of generative AI on various phenomena. Slightly before and after the release of ChatGPT, other tools, especially chatbots, had been released that are now being used in research. This is a fundamental distinction: GenAI can be used for scientific research with a plug and play approach, without any adjustment to the technology itself. This phenomenon has democratized access to state-of-the-art AI models for various tasks. Hence, we decided to divide up the the GenAI dataset, into two parts: tools, for papers using the plug-and-play tools, and algorithms, for papers leveraging different aspects of the technology. Out of the 113,450 papers in the generative AI dataset, 16,674 of them used some sort of generative AI tool and another 84,367 of them used algorithms.¹ Examples of tools are “ChatGPT” and “Stable Diffusion”, and examples of algorithms are “transformer” and “self-attention”.

4 Results

Our empirical analysis reveals striking patterns in the diffusion of artificial intelligence across scientific disciplines, with GenAI exhibiting distinctly different adoption dynamics compared to previous AI waves. The results presented here demonstrate three key dimensions along which GenAI diverges from established patterns of technological diffusion in science: distinctive speed of adoption, greater equality of distribution across fields, and reduced dependence on complementary assets.

The Exponential Growth of AI in Scientific Research

Figure 1 presents a comprehensive view of AI adoption in scientific research from 1950 to 2024, revealing the acceleration in AI-engaged publications across all waves of artificial intelligence devel-

¹We used a vocabulary approach for this categorization. Out of the 152 generative AI terms, we assigned 31 to the tools category and 80 to the algorithms category. Five terms, corresponding to 3,131 papers (2.76% of the generative AI dataset), were classified as neither tools nor algorithms.

opment. The time series analysis shows that while AI research maintained relatively modest growth rates for several decades, the period since 2015 has witnessed an explosion in AI applications across scientific disciplines.

Figures 2 and 3 show that in all of the academic fields present in our data, this exponential increasing trend can be observed. This is true of both the absolute numbers and percentages of papers using AI. As noted before, the growth looks particularly pronounced after 2015.

Figure 4 shows the breakdown by AI waves. The overall trajectory of AI research publications demonstrates the characteristic exponential growth pattern associated with transformative technologies. From fewer than 440 AI-engaged publications annually in the 1950s, the field has grown to over 828,434 publications per year by 2024, representing more than a 1,882.8-fold increase over seven decades. However, this aggregate growth masks important differences in the adoption patterns of distinct AI technologies, with GenAI showing particular acceleration in recent years.

Deep learning publications, represented by the orange line, show steady growth beginning around 2005, with acceleration becoming pronounced after 2015. This pattern reflects the gradual maturation of neural network architectures and the increasing availability of computational resources necessary for training deep learning models. The growth trajectory of deep learning follows the classic S-curve pattern predicted by diffusion theory (Griliches, 1957; Rogers, 1962; Mansfield, 1983), with an initial period of slow adoption followed by rapid acceleration as the technology matured and supporting infrastructure developed. NLP publications, shown in blue, exhibit a similar pattern with much more limited growth.

The most striking feature of Figure 4 is the sharp growth of GenAI publications, represented by the green line, beginning around 2018 and accelerating after 2022. The GenAI trajectory shows an ascent that dwarfs the growth rates of previous AI waves, reaching over 40,000 publications annually by 2024 despite being the most recent addition to the AI landscape. This unprecedented growth rate suggests that GenAI is diffusing through scientific communities at a pace that exceeds historical patterns of technology adoption.

To show these trends numerically, we fit an exponential function in the form of

$$\text{number of papers in year } t = ab^t$$

to each of the three datasets separately, where a shows the number of papers in year zero and b is the multiplicative factor. For the start year within the range of 1995-2005 for DL, we have $b_{DL} \in [1.125, 1.181]$, which gives a 12.5% – 18.1% increase in DL papers annually. With 2010-2018 range as the start year for GenAI, we get $b_{GenAI} \in [1.333, 1.692]$, which gives a 33.3% – 69.2% increase in GenAI papers annually. For NLP, we again experiment with start years in the 1995-2005 range and get $b_{NLP} \in [1.194, 1.208]$, giving us an annual growth rate of 19.4% – 20.8% in NLP papers. These make GenAI the fastest-growing, followed by NLP. Deep Learning grows slower than

any other waves.

The temporal dynamics revealed in Figures 1 and 4 provide crucial context for understanding the distinctive nature of GenAI diffusion. While deep learning and NLP required decades to achieve widespread adoption, GenAI has accomplished similar penetration in a matter of years. This acceleration cannot be explained solely by the cumulative nature of technological development or improvements in computational infrastructure, suggesting that GenAI possesses characteristics that alter the dynamics of technology adoption in scientific contexts.

Unprecedented Growth Rates Across AI Waves

Figures 5 and 6 provide a more detailed analysis of growth rates across the three AI waves, revealing the exceptional nature of GenAI’s diffusion trajectory. Figure 5 presents three-year annual growth rates for each AI technology from 2000 to 2024, while Figure 6 shows the distribution of AI research across six representative academic disciplines, illustrating how different fields have engaged with each wave of AI development.

The growth rate analysis reveals that GenAI has achieved peak annual growth rates of about 90% in 2023, far surpassing the maximum growth rates observed for deep learning (approximately 50%) and NLP (approximately 40%). These growth rates are particularly remarkable given that they represent percentage increases from increasingly large base numbers, indicating that GenAI adoption is accelerating even as the absolute number of publications grows rapidly.

The temporal pattern of growth rates also differs significantly across AI waves. Deep learning and NLP show the characteristic pattern of early rapid growth followed by gradual deceleration as the technologies mature and adoption saturates. GenAI, by contrast, shows sustained high growth rates that have not yet begun to decelerate, suggesting that the technology is still in the early phases of its diffusion trajectory despite its rapid initial adoption.

The disciplinary analysis in Figure 6 reveals how GenAI has achieved broad penetration across diverse academic fields with striking speed. The six-panel display shows publication trends for Computer Science, Engineering, Physical Sciences, Life Sciences, Social Sciences, and Humanities, representing the full spectrum of scientific inquiry from highly computational fields to traditionally non-quantitative disciplines.

Computer Science shows early and sustained growth across all three AI waves, reflecting the field’s role as the primary source of AI innovation. Yet even here, GenAI represents a qualitatively different phenomenon with the same growth pattern observed in aggregate data. Engineering disciplines demonstrate rapid deep learning adoption beginning around 2010, consistent with the field’s emphasis on computational methods. NLP adoption remained more modest due to limited language processing applications in engineering contexts. Life Sciences show strong deep learning

adoption for analyzing biological data, with more limited NLP use reflecting specialized biological text processing needs.

Social Sciences present an example of GenAI’s distinctive diffusion pattern. While deep learning and NLP adoption remained modest, GenAI showed notable growth beginning in 2022, suggesting these tools have democratized AI access in ways previous technologies could not achieve. The Humanities provide the most compelling evidence for GenAI’s unique accessibility. Traditional AI technologies showed minimal adoption, consistent with limited computational engagement. Yet GenAI adoption shows remarkable growth, with 2024 usage levels exceeding deep learning and NLP combined—representing a fundamental shift in the relationship between computational methods and humanistic inquiry.

Democratization of AI Access Across Scientific Disciplines

The disciplinary analysis reveals a crucial insight about GenAI’s diffusion pattern: it has achieved significant penetration in fields that showed limited engagement with previous AI technologies. This pattern suggests that GenAI has overcome barriers that previously limited AI adoption to computationally oriented disciplines, potentially democratizing access to AI capabilities across the full spectrum of scientific inquiry.

The comparison between the Humanities and other disciplines is particularly illuminating. In 2024, Humanities researchers were using GenAI at levels that exceeded their historical usage of deep learning and NLP combined, representing a shift in the computational engagement of traditionally non-quantitative fields. This pattern cannot be explained by a sudden methodological shift in entire fields or increasing computational literacy in humanities disciplines within the span of a few years. The reason behind this fast adoption of GenAI by the Humanities and, to a lesser extent, the Social Sciences must be sought in the nature of GenAI, suggesting that GenAI possesses characteristics that alter the accessibility of AI technologies.

The relatively new penetration of GenAI in the Social Sciences is led by GenAI tools, whereas in Computer Science, Engineering, Physical Sciences, and the Life Sciences, these are GenAI algorithms that are mostly being used by researchers. This can be seen in Figure [7](#)

Greater Equality in AI Diffusion Patterns

Figures [8](#) and [9](#) present a comprehensive analysis of distributional equality across AI technologies, providing quantitative evidence for our claim that GenAI exhibits more egalitarian diffusion patterns than previous AI waves. Figure [8](#) shows Gini coefficients calculated across the 22 academic fields represented in our dataset, while Figure [9](#) presents a detailed comparison of median author

counts for GenAI tools versus algorithms across all academic disciplines.

The Gini coefficient analysis reveals striking differences in the distributional equality of different AI technologies. Traditional measures of inequality show that deep learning and NLP exhibit high levels of concentration across academic fields, indicating that the benefits of these technologies have been captured disproportionately by a small number of fields. GenAI algorithms (Gini = 0.656) show even higher concentration, consistent with the technical complexity and resource requirements associated with developing and implementing GenAI systems. The early wave of GenAI was in many ways comparable to prior AI waves.

However, GenAI tools present a distinctly different pattern, with a Gini coefficient of 0.451 that represents significantly greater distributional equality than any other AI technology in our analysis. This finding provides quantitative confirmation of our qualitative observations about the democratizing effects of GenAI tools. The lower Gini coefficient indicates that the benefits of GenAI tools are being distributed more evenly across academic disciplines, rather than being concentrated in computationally oriented fields.

The distinction between GenAI algorithms and tools is crucial for understanding these distributional patterns. While GenAI algorithms show the highest concentration of any AI technology (consistent with the substantial technical expertise and computational resources required for their utilization), GenAI tools show the lowest concentration. This pattern suggests that the transformation of sophisticated AI algorithms into accessible tools has altered the economics of AI adoption, enabling broader participation across diverse academic disciplines.

The median author analysis in Figure 9 provides additional insight into the collaborative patterns associated with different AI technologies. The horizontal bar chart shows median author counts for GenAI tools and algorithms across all 22 academic fields in our dataset. The consistent pattern across academic fields shows that papers using GenAI tools have weakly fewer authors than papers using GenAI algorithms, suggesting that tools enable more independent research while algorithms require larger collaborative teams.

This pattern has important implications for understanding the complementary asset requirements of different AI technologies. The larger team sizes associated with GenAI algorithms likely reflect the need for diverse expertise spanning domain knowledge, computational skills, and AI technical knowledge. The smaller team sizes for GenAI tools suggest that these applications can be successfully implemented by researchers without extensive collaborative arrangements, reducing one of the traditional barriers to AI adoption in scientific research.

Dependence on Complementary Assets

Figure 10 provides a focused analysis of team composition differences between GenAI tools and algorithms, offering direct evidence for our claim that tools reduce the complementary asset requirements that have traditionally characterized AI adoption in scientific research. The simple bar chart compares mean author counts for papers using GenAI algorithms (4.32 authors) versus tools (3.05 authors), representing a 29% reduction in team size associated with tool usage. This difference in team sizes provides quantitative evidence for the reduced collaborative requirements of GenAI tools compared to algorithms. The smaller team sizes associated with tools suggest that researchers can successfully implement these technologies without assembling the diverse expertise traditionally required for AI applications. This finding supports our theoretical argument that GenAI tools have internalized many of the complementary assets that previously created barriers to AI adoption.

Computer Science Collaboration Patterns

Figure 11 presents additional evidence for our argument about reduced complementary asset requirements in GenAI adoption. The bar chart shows the percentage of papers in each AI category that include at least one author with a Computer Science affiliation, providing a direct measure of the extent to which different AI technologies require collaboration with technical experts.

The results reveal a clear hierarchy in Computer Science collaboration requirements across AI technologies. NLP shows the highest collaboration rates (47.4%), reflecting the technical complexity of NLP algorithms and the specialized knowledge required for their implementation. Deep Learning shows slightly lower but still substantial collaboration rates (37.4%), consistent with the broad applicability of deep learning methods across disciplines but continued need for technical expertise.

GenAI algorithms show collaboration rates (37.6%) similar to deep learning, indicating that the development and implementation of GenAI algorithms continues to require substantial technical expertise. However, GenAI tools show far lower collaboration rates (19.8%), representing a 47% reduction compared to GenAI algorithms and even larger reductions compared to other AI technologies. This pattern provides direct evidence for our claim that GenAI tools have reduced the complementary asset requirements that traditionally characterized AI adoption in scientific research. The lower collaboration rates with Computer Science experts suggest that researchers across disciplines can successfully implement GenAI tools without the extensive technical partnerships that were necessary for previous AI technologies.

The implications of this finding extend beyond mere convenience to encompass critical questions about the distribution of AI benefits in scientific research. Previous AI waves concentrated benefits among institutions and researchers with access to Computer Science expertise, creating systematic advantages for well-connected research groups and institutions with strong computational programs.

The reduced collaboration requirements of GenAI tools suggest that these systematic advantages may be diminishing, potentially leading to more equitable distribution of AI benefits across the scientific community.

5 Discussion

The empirical evidence presented in these five figures provides compelling support for our central thesis that GenAI represents a notable departure from established patterns of technological diffusion in scientific research. The unprecedented growth rates, broad disciplinary penetration, greater distributional equality, reduced team size requirements, and lower collaboration rates with technical experts all point to a technology that operates according to different economic and institutional dynamics than previous AI waves.

The distinction between GenAI algorithms and tools emerges as a crucial factor in understanding these diffusion patterns. While GenAI algorithms continue to exhibit the characteristics of traditional GPTs, GenAI tools have altered the accessibility landscape by internalizing many of the complementary assets that previously created barriers to adoption.

These findings have important implications for our understanding of technological diffusion in scientific contexts and raise critical questions about the future trajectory of AI adoption in research. If GenAI tools continue to evolve in ways that further reduce barriers to adoption while expanding their capabilities, they may represent the emergence of a new class of technologies that challenge established theoretical frameworks for understanding innovation diffusion.

Previous research has documented how the increasing size of teams is transforming the nature of knowledge production, with larger teams becoming necessary to tackle complex scientific problems (Wuchty et al., 2007). Our results reveal that GenAI tools may be reversing this trend. In each field, the median size of teams using GenAI tools is smaller than the median size of teams using traditional AI algorithms, suggesting that GenAI may reduce the human capital required as input for a unit of innovation. Experimental evidence shows that while AI substitution can reduce team coordination in some contexts, AI augmentation can effectively replicate certain benefits of human collaboration and enable more efficient team structures, particularly when AI serves as a "cybernetic teammate" that enhances rather than replaces human expertise (Dell'Acqua et al., 2025b a). As this new wave of AI diffuses further and becomes a standard tool in the innovation process, its positive impact on lowering the burden of knowledge may become more pronounced in the data. Treating the number of innovators as fixed, this could lead to an increase in output per innovator, thereby boosting the efficiency of scientific innovation activities and potentially reversing long-standing trends toward larger teams and even later onset of innovation by innovators.

At the same time, this rapid and widespread adoption also raises important questions about the quality and reliability of GenAI applications in scientific contexts. The ease of use that facilitates broad adoption may also enable inappropriate or ineffective applications, potentially creating new forms of research quality concerns. The tension between accessibility and expertise represents a challenge for the scientific community as it grapples with the implications of democratized AI access.

The rapid diffusion of GenAI in scientific research has created unusual tensions with the fundamental institutions, values, and practices that have governed scientific inquiry since the Enlightenment. While the technology's accessibility and immediate utility have driven widespread adoption, its "black box" nature and probabilistic outputs challenge core scientific principles of transparency, reproducibility, and explainability that have long served as the foundation of reliable knowledge production. These tensions suggest that the GenAI revolution may represent not merely a technological shift, but a potential paradigm challenge to established scientific epistemology.

GenAI poses unique challenges to this institutional framework because its core operational characteristics appear to conflict with foundational scientific values. The principle of explainability—the requirement that scientific methods and reasoning be transparent and comprehensible to qualified peers—is directly challenged by the opacity of large language models and other GenAI systems. While researchers can observe the inputs and outputs of these systems, the internal processes by which they generate results remain largely inscrutable, even to their creators.

The hypothesis testing framework that underlies much of empirical science also faces challenges from GenAI applications. Traditional scientific methodology emphasizes the formulation of explicit, testable hypotheses prior to data collection and analysis, with clear protocols for evaluating evidence and drawing conclusions. GenAI tools, however, often operate through pattern recognition and statistical inference processes that may identify relationships or generate insights without explicit hypothesis formulation. While this capability can be valuable for exploratory research and hypothesis generation, it creates tensions with established standards for scientific inference and evidence evaluation.

The resolution of these tensions will likely require the development of new institutional frameworks, methodological standards, and quality control mechanisms specifically designed for the age of GenAI. This process of institutional adaptation represents one of the most significant challenges facing the scientific community as it grapples with the implications of widespread AI adoption. The outcome of this adaptation process will have profound implications not only for the practice of science but also for public trust in scientific institutions and the broader social role of scientific knowledge.

6 Conclusion

Our paper makes several important contributions to our understanding of technological diffusion in science and the distinctive characteristics of the GenAI revolution. We provide the first systematic empirical analysis comparing the diffusion patterns of GenAI with those of previous AI waves, demonstrating quantitatively that GenAI exhibits distinctly different adoption dynamics. Through analysis of publication data, collaboration networks, and institutional surveys, we document the striking speed, equality, and accessibility that characterize GenAI diffusion in scientific contexts.

Our analysis demonstrates that the plug-and-play nature of GenAI tools alters the economics of AI adoption by reducing barriers to entry and democratizing access to advanced capabilities. This finding has important implications for innovation theory more broadly, suggesting that certain technological characteristics can disrupt established diffusion patterns in ways that existing theoretical frameworks do not adequately capture.

The central argument of this paper—that GenAI represents a clear departure from established patterns of technological diffusion in science—has implications that extend far beyond the specific case of AI adoption. If our analysis is correct, then the GenAI phenomenon may signal the emergence of a new class of technologies characterized by rapid, egalitarian diffusion and reduced dependence on complementary assets. Understanding these new diffusion dynamics is crucial for predicting the impacts of future technological innovations and developing appropriate institutional responses to rapid technological change.

The stakes of these questions are particularly high given the potential of GenAI to transform not only how science is conducted but also what kinds of knowledge can be produced and by whom. If GenAI indeed democratizes access to advanced research capabilities, it may alter the distribution of scientific authority and the social organization of knowledge production. Understanding these transformations is essential for ensuring that the benefits of AI-enabled research are realized while maintaining the institutional foundations that have made science a reliable source of knowledge about the natural and social world.

References

- Abadi, M., Barham, P., Chen, J., and Chen, Z. (2016). Tensorflow: A system for large-scale machine learning. *OSDI*.
- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*.
- Agrawal, A., Gans, J., and Goldfarb, A. (2018). *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.
- Agrawal, A., McHale, J., and Oettl, A. (2025). AI in science. Working Paper.
- Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies “engines of growth”? *Journal of Econometrics*, 65(1):83–108.
- Brown, P. F., Della Pietra, S. A., Della Pietra, V. J., and Mercer, R. L. (1993). The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*.
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative ai at work. *The Quarterly Journal of Economics*.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2019). Artificial intelligence and the modern productivity paradox. In *The economics of artificial intelligence: An agenda*, volume 23, pages 23–57. University of Chicago Press.
- Chen, H., Engkvist, O., Wang, Y., Olivecrona, M., and Blaschke, T. (2018). The rise of deep learning in drug discovery. *Drug Discovery Today*, 23(6):1241–1250.
- David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *The American Economic Review*, 80(2):355–361.
- Dell’Acqua, F., Ayoubi, C., Lifshitz-Assaf, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., et al. (2025a). The cybernetic teammate: A field experiment on generative ai reshaping teamwork and expertise. Working Paper 33641, National Bureau of Economic Research.
- Dell’Acqua, F., Kogut, B., and Perkowski, P. (2025b). Super mario meets ai: Experimental effects of automation and skills on team performance and coordination. *The Review of Economics and Statistics*.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. Working Paper 24-013, Harvard Business School Technology & Operations Mgt. Unit.

- Duede, E., Dolan, W., Bauer, A., Foster, I., and Lakhani, K. (2024). Oil & water? diffusion of AI within and across scientific fields. *arXiv preprint arXiv:2405.15828*.
- Gao, J. and Wang, D. (2024). Quantifying the use and potential benefits of artificial intelligence in scientific research. *Nature Human Behaviour*, 8(12):2281–2292.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., and Warde-Farley, D. (2014). Generative adversarial networks. *Advances in neural information processing systems*.
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4):501–522.
- Hall, B. H. and Trajtenberg, M. (2004). Uncovering gpts with patent data. *NBER working paper*.
- Ho, J., Jain, A., and Abbeel, P. (Advances in neural information processing systems). Denoising diffusion probabilistic models. *Advances in neural information processing systems*.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*.
- Huang, Y., Goma, A., Weissmann, T., Grigo, J., Ben Tkhayat, H., Frey, B., Gaipl, U. S., et al. (2023). Benchmarking chatgpt-4 on acr radiation oncology in-training exam (txit): Potentials and challenges for ai-assisted medical education and decision making in radiation oncology. *arXiv preprint arXiv:2304.11957*.
- Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317.
- Jones, B. F. (2010). Age and great invention. *The Review of Economics and Statistics*, 92(1):1–14.
- Jurafsky, D. and Martin, J. H. (2019). Speech and language processing. Stanford University.
- Krenn, M., Pollice, R., Guo, S. Y., Aldeghi, M., Cervera-Lierta, A., Friederich, P., Gomes, G. d. P., et al. (2022). On scientific understanding with artificial intelligence. *Nature Reviews Physics*, 4(12):761–769.
- Lebovitz, S., Lifshitz-Assaf, H., and Levina, N. (2022). To engage or not to engage with ai for critical judgments: How professionals deal with opacity when using ai for medical diagnosis. *Organization Science*, 33(1):126–148.
- LeCun, Y. and Bengio, Y. (1998). Convolutional networks for images, speech, and time series. In *The handbook of brain theory and neural networks*.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*.
- Lind Plesner, L., Müller, F. C., Brejnebo, M. W., Laustrop, L. C., Rasmussen, F., Nielsen, O. W., Boesen, M., and Andersen, M. B. (2023). Commercially available chest radiograph ai tools for detecting airspace disease, pneumothorax, and pleural effusion. *Radiology*, 308(3):e231236.

- Ludwig, J. and Mullainathan, S. (2024). Machine learning as a tool for hypothesis generation. *The Quarterly Journal of Economics*, 139(2):751–827.
- Manning, C. and Schütze, H. (1999). *Foundations of Statistical Natural Language Processing*. MIT Press.
- Mansfield, E. (1983). Technological change and market structure: An empirical study. *The American Economic Review*, 73(2):205–209.
- Medeiros, L., Psaltis, D., Lauer, T. R., and Özel, F. (2023). The image of the m87 black hole reconstructed with primo. *The Astrophysical Journal Letters*, 947(1):L7.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mullainathan, S. and Rambachan, A. (2025). Science in the age of algorithms. In *The Economics of Transformative AI*.
- NASEM (2023). *Artificial Intelligence to Assist Mathematical Reasoning: Proceedings of a Workshop*. The National Academies Press, Washington, DC.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6684).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, volume 32.
- Raina, R., Madhavan, A., and Ng, A. Y. (2009). Large-scale deep unsupervised learning using graphics processors. In *Proceedings of the 26th annual international conference on machine learning*, pages 873–880.
- Rogers, E. M. (1962). *Diffusion of innovations*. Free Press.
- Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6):285–305.
- Tong, Y. and Zhang, L. (2023). Discovering the next decade’s synthetic biology research trends with chatgpt. *Synthetic and Systems Biotechnology*, 8(2):220.
- West, C. G. (2023). Advances in apparent conceptual physics reasoning in gpt-4. *arXiv e-prints*, pages arXiv–2303.
- Wof, T., Debut, L., Sanh, V., Chaumond, J., and Delangue, C. (2020). Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*.

Wuchty, S., Jones, B. F., and Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039.

Figure 1

AI Research Papers Published Over Time (All AI)




Figure 2

AI Publications by Field (Absolute Counts, 1995-2022)

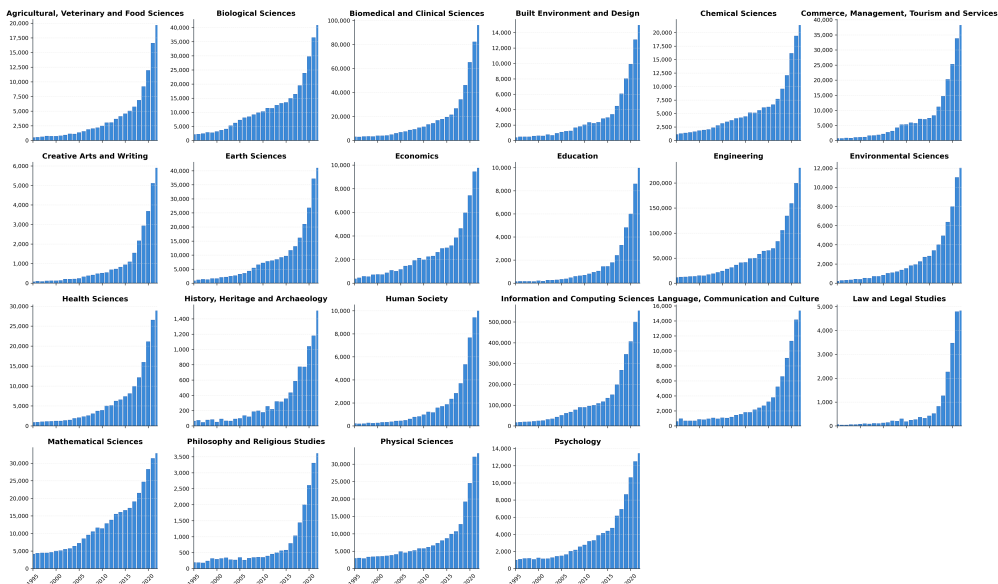


Figure 3



Figure 4




Figure 5

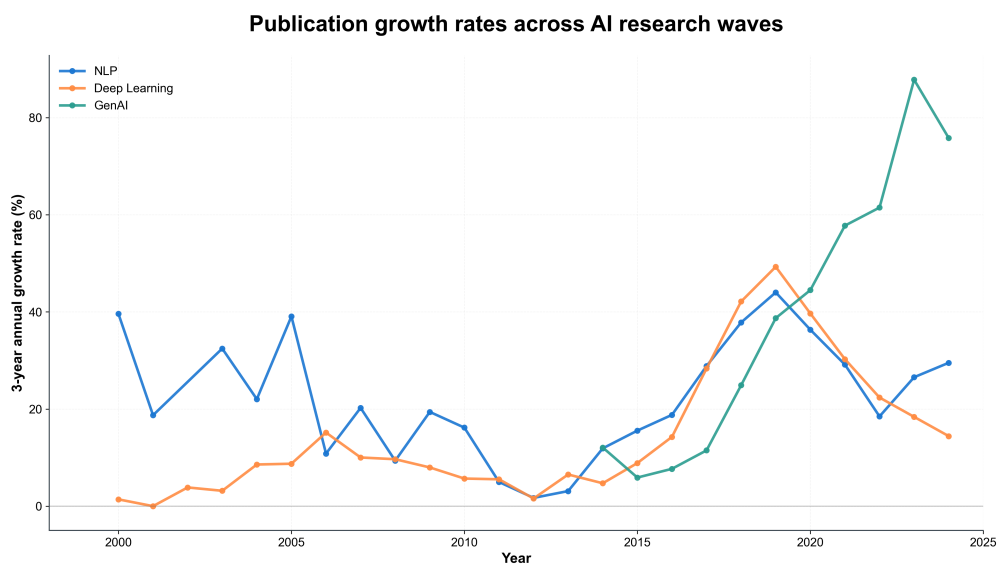


Figure 6

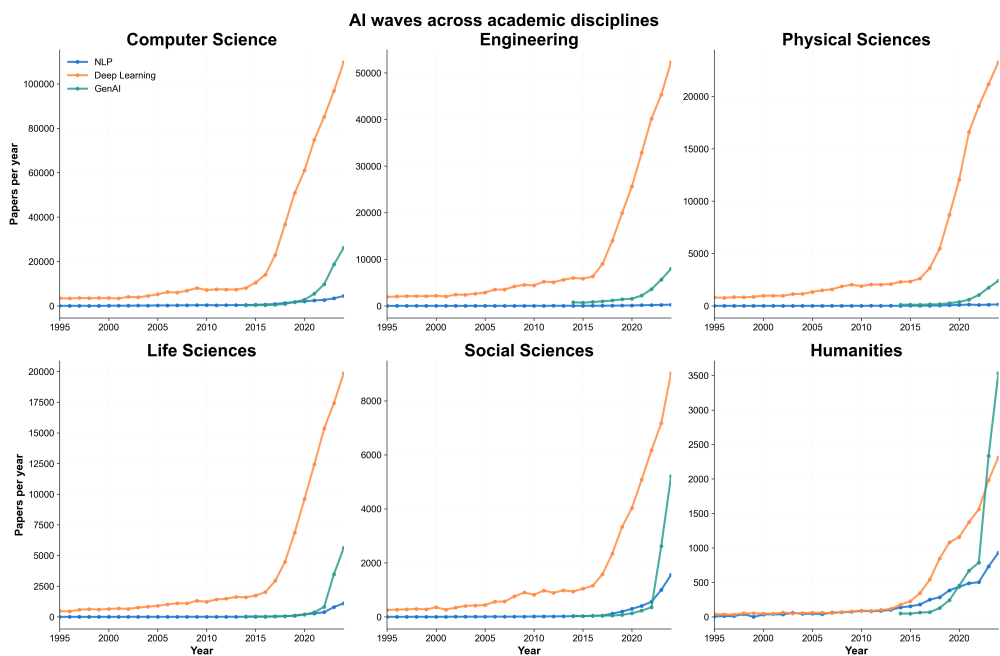


Figure 7

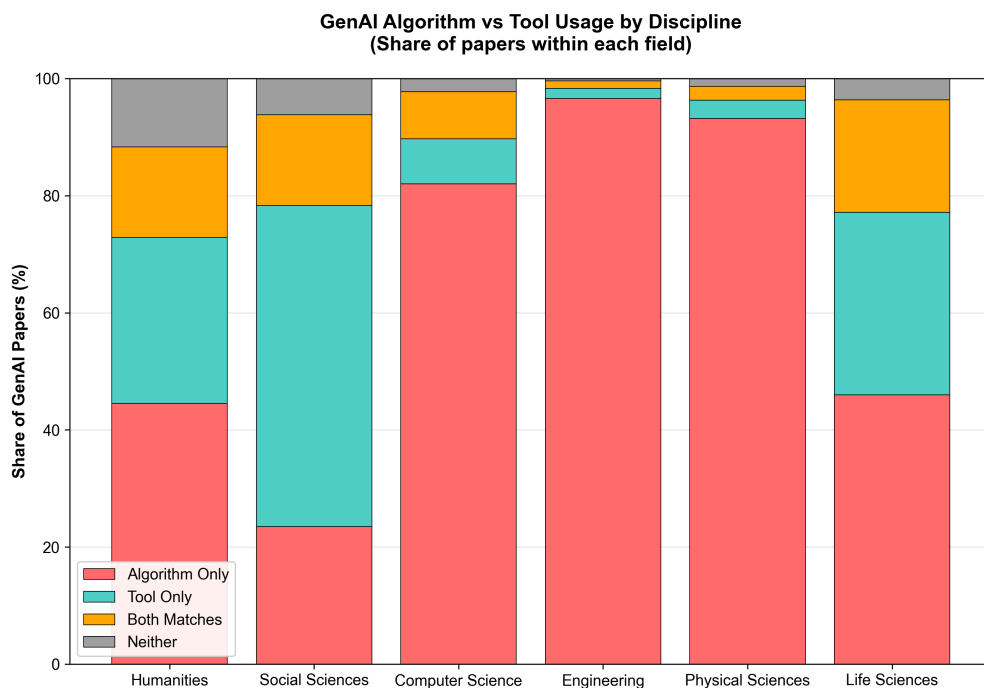


Figure 8

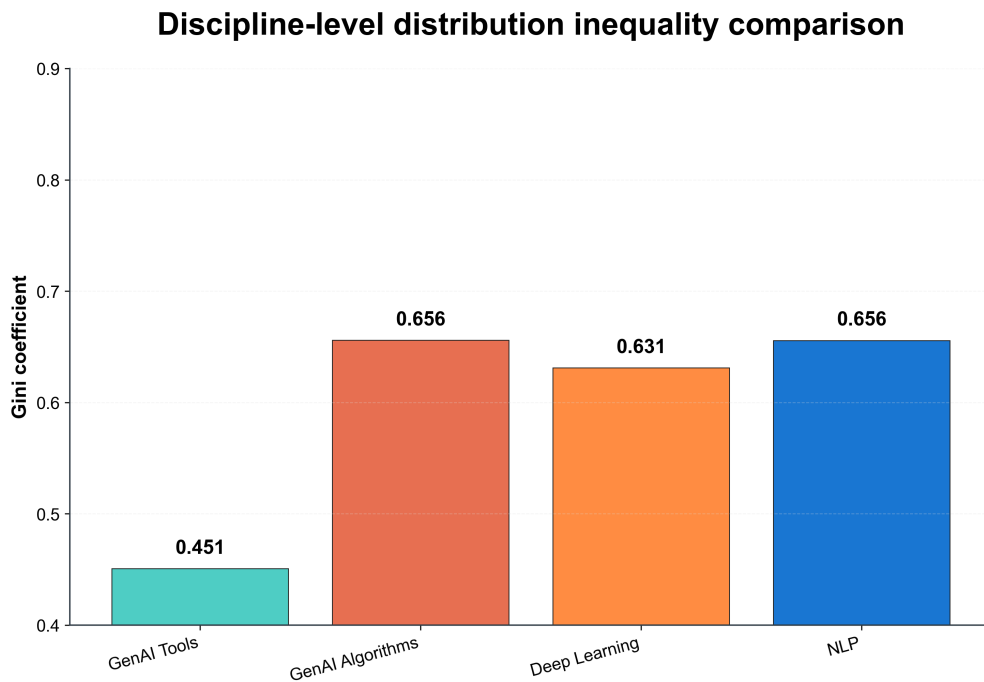


Figure 9

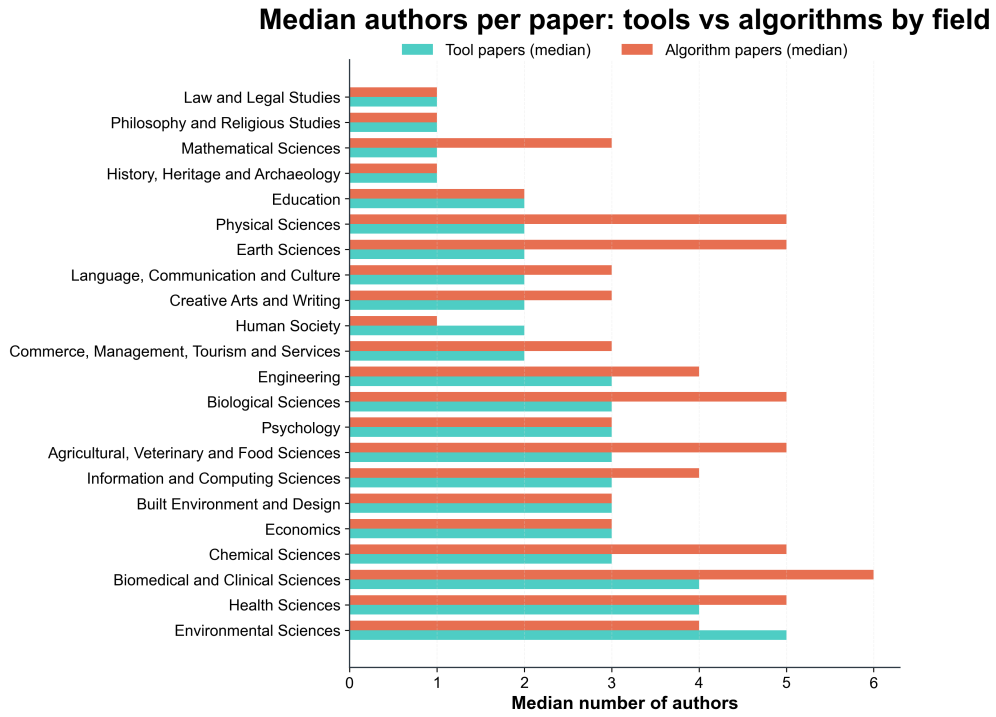


Figure 10

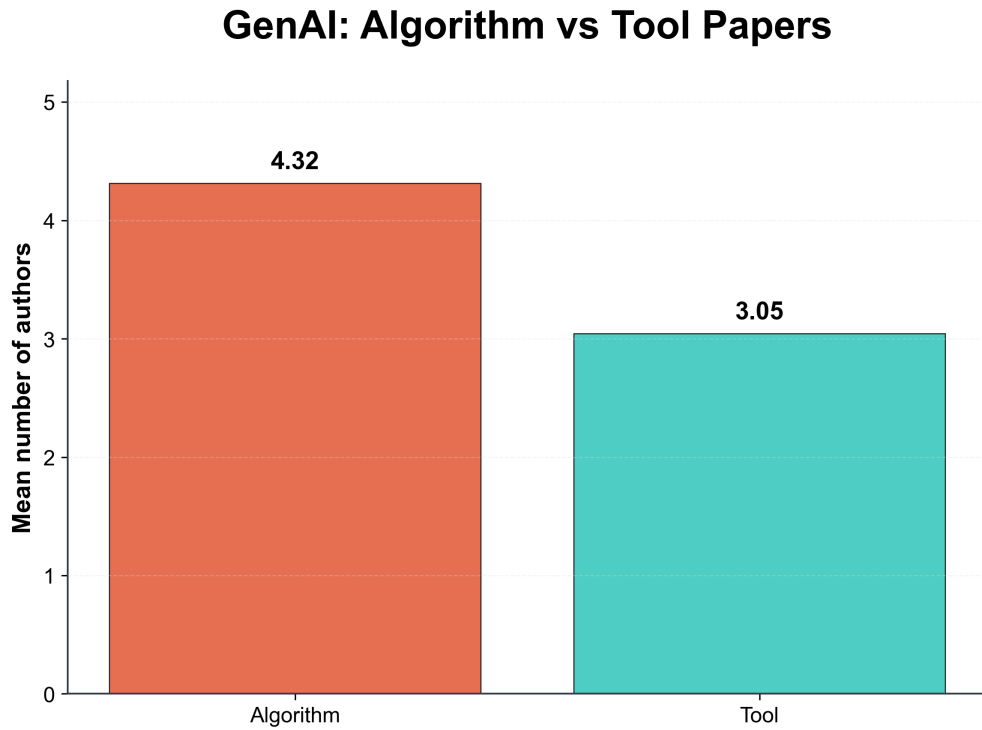


Figure 11

Computer science collaboration rates across AI research areas

