

AI Governance and Liability in Europe

A Primer

Edited by

Ceyhun Necati Pehlivan
Nikolaus Forgó
Peggy Valcke



Published by:

Kluwer Law International B.V.
PO Box 316
2400 AH Alphen aan den Rijn
The Netherlands
E-mail: lrs-sales@wolterskluwer.com
Website: www.wolterskluwer.com/en/solutions/kluwerlawinternational

Sold and distributed by:

Wolters Kluwer Legal & Regulatory U.S.
920 Links Avenue
Landisville, PA 17538
United States of America
E-mail: customer.service@wolterskluwer.com

ISBN 978-94-035-0549-7

e-Book: ISBN 978-94-035-0559-6
web-PDF: ISBN 978-94-035-0569-5

© 2025 Ceyhun Necati Pehlivan, Nikolaus Forgó & Peggy Valcke

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher.

Permission to use this content must be obtained from the copyright owner. More information can be found at: www.wolterskluwer.com/en/solutions/legal-regulatory/permissions-reprints-and-licensing

Printed in the Netherlands.

CHAPTER 4

Genesis of the Artificial Intelligence Act: When AI Beats the Imitation Game

Christof Koolen

1 INTRODUCTION

Only a few phenomena have triggered legal and ethical questions quite like the development of artificial intelligence (AI). What began as simple constructs designed to solve number-crunching problems has evolved into complex and intricate systems capable of interacting with speech, language, and vision. From generating images to writing code and solving logic puzzles: the current sentiment surrounding AI models is a blend of both excitement and fear.¹ On the positive end of the spectrum, AI promises to profoundly impact our lives, whether by aiding in drug development,² designing new materials,³ or helping to combat climate change.⁴ On the negative end of the spectrum, however, these very same AI systems can just as easily contribute to the spread of misinformation, aid in the design of deadly pathogens, and play a role in military applications.

If anything, AI is poised to transform our relationship with technology. Its recent resurgence therefore raises a pertinent question: *how to ensure that AI neither evolves into a tool of wrongdoing nor poses a threat to humanity?* While doomsday scenarios may be exaggerated, the question should nonetheless be taken seriously, as the capabilities of AI models are clearly outrunning those of their human creators. With AI increasingly mimicking human skills to an eerie degree, and with fundamental human

1. The Economist 2023, 9-10.

2. Mock et al. 2023, 468.

3. Henshall 2023.

4. Cows et al. 2023, 284-285; Stein 2020, 898-900.

rights and values hanging in the balance, a serious discussion about the regulation of AI is necessary. While calls for regulatory action were voiced across the globe, the governance of AI has proven challenging, in part due to the speed of AI developments, ambivalence about the precise nature of AI, and uncertainty about how AI should be regulated. As a result, questions linger as to how lawmakers should address AI, its uses, and its consequences.

In Europe, efforts were launched as early as 2018 to find appropriate answers to these questions. Six years later, on 13 March 2024, these questions were finally answered with the adoption of the AI Act. Together with, among others, the General Data Protection Regulation (GDPR), the Digital Markets Act (DMA) and Digital Services Act (DSA), the Data Governance Act (DGA) and Data Act (DA), the adoption of the AI Act is expected to further cement the European Union's (EU's) position as a global leader in setting rules for the digital world.

It is against this backdrop that this book sets out to contribute to a better understanding of the regulation of AI in Europe. By way of introduction, this chapter seeks to contextualize the topic of AI, provide further background information on some of the main AI developments and, more importantly, relate the AI Act to several other ongoing or completed legislative initiatives in the EU.

To attain these objectives, the chapter is divided into six sections. Following this introductory section, section 2 takes a brief look back on the development of AI and sheds light on the advances in computer power and the availability of data. Section 3 then touches on the legislative history of the AI Act. Section 4 subsequently expands on the core underpinnings and main principles of the AI Act. Section 5 then frames the AI Act in relation to several other legal instruments, notably the GDPR (subsection 5.1), the DGA (subsection 5.2), the DA (subsection 5.3), the DMA (subsection 5.4), and the DSA (subsection 5.5). Section 6 concludes the chapter.

2 A BRIEF LOOK BACKWARDS ON AI

2.1 Beating the Imitation Game

The origins of AI date back to the 1950s and the research of British mathematician Alan Turing. In the opening sentences of his seminal article entitled *Computing Machinery and Intelligence*, he poses the following fundamental question: 'Can machines think?'⁵ Throughout the article, Turing elaborated on a methodological thought framework that could be used to determine whether AI is distinguishable from human intelligence. Fundamentally, he proposed playing a game with three different players placed in different rooms: a human, a machine, and an interrogator.⁶ The goal of the interrogator is to pose questions to the other two players in order to determine who is human and

5. Turing 1950, 433.

6. *Ibid.*, at 434.

who is machine. At the same time, both the human and the machine are tasked with trying to convince the interrogator that they are the human respondent. This test became known as the Turing Test, or more colloquially, the *Imitation Game*.⁷

Anno 2024, AI is capable of convincingly beating this Imitation Game. In many respects, the output of AI models is completely indistinguishable from human-generated output. What is more, the capabilities of AI models have even surpassed those of humans in several domains, notably in relation to reading comprehension, language understanding,⁸ and image recognition.⁹ In other words, in terms of quality, AI is capable of producing output that is at least as good as human output. More importantly, AI is currently also capable of engaging in behaviour that attracts legal consequences. A befitting example in this respect is the conclusion of contracts.¹⁰ More macabre are situations where AI systems cause personal injury or even death.¹¹

2.2 From AI Winter to AI Summer

While Alan Turing fairly accurately predicted that, in due time, machine intelligence would be able to convince its interrogator of its human-esque capabilities,¹² the journey of AI has not been without its problems. In fact, the history of AI has been defined by ups and downs, and by booms and busts: periods marked by high levels of optimism, followed by the realization that much-anticipated progress did not materialize.¹³ Two particular periods – the 1974-1980 era¹⁴ and the 1987-1994 era¹⁵ – have been described in literature as *AI winters*, representing stretches of time when AI research went into hibernation for the better half of a decade (see Figure 4.1).¹⁶

7. In 2022, discussions about the Turing Test resurfaced in the press following (unsubstantiated) claims that Google's LaMDA (Language Model for Dialog Applications) chatbot had achieved 'sentience': Dave 2022; Reuters 2022; for an academic analysis of the Turing Test, consider Saygin et al. 2000.

8. Regarding AI benchmarks for language understanding, see Wang et al. 2018; Wang et al. 2019.

9. Kiela et al. 2024; Kiela et al. 2023.

10. See, for example, Herbosch 2023, 66-67, para. 79; and the references cited therein.

11. Consider, for instance, Atkinson 2023; Tesla 2016.


12. Turing 1950, at 442: 'I believe that in about fifty years' time it will be possible to programme computers ... to make them play the imitation game so well that an average interrogator will not have more than 70 per cent, chance of making the right identification after five minutes of questioning. The original question, "Can machines think?" I believe to be too meaningless to deserve discussion.'

13. Duan et al. 2019, 63.

14. After an initial wave of excitement, it became clear that many of the problems that AI promised to solve could not yet be adequately framed from a computational point of view. The resulting pullback in funding for AI research was termed the first AI winter, which stretched from 1974 to 1980.

15. Interest in AI was renewed following the emergence of so-called expert systems, which replicated human decision-making in an 'if ... then ... else' format. However, the limitations of this type of conditional reasoning soon became apparent, leading to the second AI winter from 1987 to 1994.

16. Nilsson 2009, 408-409.

Figure 4.1 Evolution of AI Popularity over Time¹⁷


Over time, however, it became clear that problems previously thought impossible to solve with AI – whether due to a lack of computational power or due to challenges of computationally framing logic and reasoning – were slowly but surely coming within our grasp. As a result, the excitement surrounding AI rekindled during the late 1990s, with the famous defeat of Garry Kasparov by the computer chess-playing system Deep Blue in 1997 serving as one of the key highlights.¹⁸

During the 2010s, AI research branched out exponentially. Long gone are the times when AI systems rely on rule-based logic and the simple indexing and retrieval of knowledge. AI systems now have the ability to learn and improve their performance over time for certain tasks, they can interact with the intricate nuances of language, and convert complex inputs – such as speech and vision – into numerical representations. Currently, the state-of-the-art is focused on *deep learning*, a class of machine learning algorithms that uses artificial neural networks (ANNs) with multiple layers of ‘depth’ to extract high-level features from unstructured input data (see Figure 4.2).¹⁹

17. The visual is the author’s own work, but was based on EPRS 2020, 5. Note, however, that the visual is primarily intended to illustrate how interest in AI has fluctuated over time, rather than to provide an accurate measure of AI’s popularity.

18. Campbell et al. 2002, 57; readers interested in learning more about the development of Deep Blue may also consider Hsu 2002.

19. Foster & Friston 2023, 23; Kelleher 2019, 68; Goodfellow et al. 2016, 5-6.

Figure 4.2 Artificial Neural Network Consisting of Three Hidden Layers²⁰

A key characteristic of these neural networks is their ability to independently learn which features are most relevant based on a given set of training data. The learning done by the neural network essentially involves optimizing its internal parameters (or *weights*) to maximize the probability of correctly treating new, unseen input data (e.g., classifying pictures of cats and dogs).²¹ Over time, the field of deep learning has evolved into assorted subdisciplines, each of which focuses on a specific type of problem or application.²²

20. The visual is the author's own work. When classifying images of cats and dogs using an ANN, the network automatically learns to identify features important for distinguishing between these two animals. In the early stages of training, the initial hidden layers (e.g., hidden layer 1) learn to recognize simple visual elements (e.g., colours, edges, lines). As the network passes through deeper hidden layers (e.g., hidden layer 3), these features are combined into more complex patterns specific to cats and dogs (e.g., shape and placement of the eyes and ears, texture of the hair coat, proportions of the animal's body). Ultimately, the network synthesizes all the extracted features from the different hidden layers to accurately classify an image as either a cat or a dog.

21. Kelleher 2019, 49-51.

22. For example, *autoencoders* are a particular type of ANN used to process unlabelled data, where the network learns to represent input data in a compressed manner and with fewer dimensions (i.e., encode) and then reconstruct the data as close as possible to the original input data (i.e.,

2.3 Advances in Computing Power and Availability of Data

Before delving into the legal developments, it is fair to ask what has caused the sudden resurgence of interest in AI. In this respect, the Commission points to two key drivers fuelling the current AI momentum: exponential increases in computing power and the increased availability of data.²³ Combined, these two dynamics create an ideal climate for AI to flourish and for AI systems to be applied to various facets of our lives.²⁴ To better understand this, it is worth taking a closer look at each of these factors individually.

2.3.1 Computing Power

When assessing computing power, the discussion quickly shifts to so-called *graphics processing units* (GPUs). While originally designed to render computer graphics, these chips are now increasingly being used for their general-purpose computing capabilities.²⁵ In essence, GPUs are architecturally optimized to be extremely efficient at performing calculations. When describing the computational capabilities of a GPU, the unit of measurement typically used is floating point operations per second (FLOPs). Fundamentally, floating point arithmetic involves applying basic arithmetic functions – such as additions, subtractions, multiplications, or divisions – to *floating point numbers*. Consider, for instance, the number 481, which can be re-expressed as a floating point number by rewriting it as a combination of an integer that is scaled by an exponent of a fixed base:²⁶

$$481 = 48.1 \times 10^1 = 4.81 \times 10^2 = 0.481 \times 10^3$$

A single floating point operation could then look as follows:

$$0.481 \times 10^3 + 7.947 \times 10^3 = 8.428 \times 10^3$$

decode). *Generative adversarial networks* (GANs) are another type of ANN, where two different neural networks (a generator and a discriminator) are pitted against each other in an effort to generate new data with the same statistical patterns as the training data. Also well-known are *convolutional neural networks* (CNNs), which are suited to grid-like data such as images, and *recurrent neural networks* (RNNs), which are used to process sequential data where the directional flow of information is important, such as natural language. *Transformers*, a deep learning architecture where text is represented in a numerical format called ‘tokens’, have also received a lot of attention recently and have seen widespread application in the context of training LLMs.

23. White Paper on AI, 2.

24. AIA, Rec. 3: ‘AI systems can be easily deployed in a large variety of sectors of the economy and many parts of society, including across borders, and can easily circulate throughout the Union’.


25. The term general-purpose graphics processing unit (GPGPU) is often mentioned in this context.

26. In this sense, ‘floating-point’ refers to the simple fact that a number can be expressed in a decimal floating point form and that the decimal can quite literally float by changing the exponent of the fixed base.

The advantage of using floating point numbers instead of integers is that very small numbers and very large numbers can be expressed in a very compact and consistent format. As such, floating point numbers represent a good balance between, on the one hand, numerical precision and, on the other hand, numerical stability when dealing with a wide range of values.²⁷

When looking at the evolution of GPUs, the number of FLOPs that a GPU chip can perform has increased exponentially over time (see Figure 4.3).

Figure 4.3 Evolution of Single GPU Computing Power over Time in Consumer-Grade GPUs²⁸



Consider, for example, the information taken up in Table 4.1: the most advanced enterprise-level chips currently available on the market – Nvidia’s H100 GPUs – are capable of performing a staggering 51,220,000,000 floating point operations *every second*. What is even more remarkable is that this represents a $\sim 53,354\times$ increase in computing power between 1999 and 2023.


27. Floating point arithmetic typically uses a limited number of ‘significant digits’ compared to integer arithmetic. For example, the integer value 849,058 has six significant digits, whereas the floating point value 0.849×10^6 has only three significant digits. Note, therefore, that these two different notation methods represent a trade-off between precision and range.

28. The visual is the author’s own work, but was based on numbers from TechPowerUp (undated) and Nvidia GPU specification sheets. See also Table 4.3 in s. 7 (Appendix).

Table 4.1 Evolution of GPU Computing Power over Time²⁹

	Nvidia GeForce 256	Nvidia Tesla K20X	Nvidia H100
Release date	1999	2012	2023
Transistor count	17 million (0.24%)	7,080 million (100%)	80,000 million (1,130%)
Memory size	0.032 GB (0.53%)	6 GB (100%)	80 GB (1,333%)
Memory bus	128-bit (33%)	384-bit (100%)	5,120-bit (1,333%)
Bandwidth	4.8 GB/s (0.48%)	1,008 GB/s (100%)	2,039 GB/s (202%)
32-bit Floating Point	0.00096 teraFLOPs (0.01%)	9,935 teraFLOPs (100%)	51.22 teraFLOPs (516%)
MSRP at release	USD 249	USD 7,699	USD 30,000

Figure 4.4 also illustrates that the cost of acquiring a computing capacity of one teraFLOP per second (i.e., 1,000,000,000,000 floating point operations per second) has fallen from USD 259,375.00 in 1999 to just USD 19.36 in 2022.

Figure 4.4 Evolution of Single GPU Price per teraFLOP Computing Power over Time in Consumer-Grade GPUs³⁰

To keep matters in perspective, it is equally illustrative to consider the amount of computation required to train an AI model. For example, the technical documentation provided with GPT-3 175B – a large language model released by OpenAI in 2020 – states that a staggering 3.14×10^{23} floating point operations were needed to train the

29. Numbers adopted from *ibid.*

30. See *supra* n. 29.

model.³¹ If we were to train GPT-3 175B using just a single Nvidia H100 GPU with a processing power of 51.22 teraFLOPs, it would take 6,013,021,831 seconds or about 191 years of continuous computing.^{32,33&34}

While no official information is available about GPT-4 – OpenAI’s successor to GPT-3 released in 2023 – leaked details suggest that it is approximately ten times the size of GPT-3, with 1.8 trillion parameters compared to GPT-3’s 175 billion parameters.³⁵ When using GPT-4, the cost of generating a single token³⁶ is about 560 teraFLOP.³⁷ Training GPT-4 required $\sim 2.15 \times 10^{25}$ floating point operations and was performed on 25,000 Nvidia A100 GPUs over 90-100 days at a total computational cost of about USD 63 million.³⁸ While these numbers have never been officially confirmed by OpenAI, they are corroborated by other news articles and, more importantly, also explain why tech conglomerates such as Meta are spending an inordinate amount of money to acquire bulk quantities of Nvidia H100 GPUs to parallelize AI model training.³⁹

2.3.2 Data Availability

Beyond the raw computational capabilities of GPUs, floating point operations ultimately need to be applied to data. In this respect, it is relevant to note that data come in all forms and shapes. For the purposes of this chapter, a distinction is made between numerical data, textual data, and visual data.

When it comes to *numerical data*, vast amounts of data are being generated as a by-product of using IT systems.⁴⁰ For example, large quantities of information are collected via the internet (e.g., search queries and search behaviour, e-commerce

31. Brown et al. 2020, 46; writing this number out in full amounts to 314,000,000,000,000,000,000,000,000 floating point operations.

32. Note that these calculations assume that a 32-bit single-precision floating point format is used (i.e., FP32). There are other formats, notably 16-bit floating point (FP16; half-precision), 64-bit floating point (FP64; double-precision) and 8-bit floating point (FP8; minifloat). Lowering the ‘precision’ is associated with performance gains during both model training and inference, but comes at the expense of accuracy in terms of the model’s output. It should be noted that proprietary computational formats also exist (e.g., TensorFloat-32; TF32).

33. Note that there exists a unit for expressing total computation time, which is PF-days or petaFLOP per second-days. The number of PF-days refers to the total number of days required to train an AI model at a rate of one petaFLOP per second for a full day. The training of GPT-3 175B required 3,640 PF-days: Brown et al. 2020, 46.

34. Note that AI model training is not performed at the theoretical maximum capacity of the GPU to better address thermal limitations and limit power consumption. Real-world model FLOPs utilization (MFU) typically ranges from 60% to 30%.

35. While the official leak has been scrubbed from the internet, the relevant information is still accessible via web archiving websites: Peleg 2023.

36. In the context of large language models, ‘tokens’ represent the basic unit of text processed by an AI model. Tokens are best thought of as parts of a word, such as the stem, prefix, or suffix of a word. For instance, the word ‘unbelievable’ could be tokenized in ‘un’, ‘believ’ and ‘able’. Note, however, that there are different types of tokenizers that break down text differently.

37. Peleg 2023.

38. *Ibid.*

39. Kan 2024.

40. EPRS 2020, 16.

transactions, or social networking activities). Similarly, with the proliferation of smart devices, the amount of data generated by these devices has also increased, ranging from heart rate sensor data from smartwatches to productivity metrics from a manufacturing plant. Other economic activities, such as banking and transportation, also contribute to a steady flow of numerical data. Taken together, this readily available influx of information can be used for data modelling purposes, whether in the context of detecting heart arrhythmia or building complete virtual, true-to-life representations of manufacturing plants (i.e., digital twins).

Also of interest are sources of *textual data*. In this respect, it is worth noting that language datasets are becoming increasingly available.⁴¹ Arguably the best-known textual dataset is the Common Crawl:⁴² a public repository of web crawl data that is available for anyone to use. Its most recent archive contains 3.16 billion web pages, as well as trillions of contextual URLs that link those pages together.⁴³ The dataset represents an impressive 424.7 terabytes (TB) of uncompressed text data,⁴⁴ making the Common Crawl arguably the most impactful dataset when it comes to training language-based AI models. Also noteworthy is The Pile, a curated and open-source dataset consisting of data taken from PubMed, ArXiv, GitHub, Project Gutenberg, Wikipedia, and seventeen other sources.⁴⁵ Many other, more topical text-based datasets are also available online (e.g., news articles, wine reviews, movie plots, Twitter tweets).⁴⁶

Several repositories also host *visual data*, consisting of image-text pairs that can be used to train image-generating AI systems. For example, the training of text-to-image model Stable Diffusion relied on LAION-5B, which is essentially a dataset derived from the Common Crawl and consists of over five billion pairs of images and their corresponding captions.⁴⁷ DALL•E 2 – a text-to-image model developed by OpenAI – was trained based on images from publicly available sources as well as several licensed sources.⁴⁸

In addition to well-known collections of data, there are several aggregator websites – such as *Hugging Face*,⁴⁹ *Papers with Code*,⁵⁰ and *Kaggle*⁵¹ – that list numerous other datasets. Beyond the three main data types mentioned above, there are open-source datasets available for download relating to video, audio, 3D visuals, dialog data, graphs, tabular data, medical data (e.g., MRIs or EEGs), financial data (e.g., stock prices), LiDAR data, and so on. Combined, these datasets can serve as training data for a wide variety of AI systems and AI models, ranging from image classification, text,

41. By way of illustration, consider the datasets used to train GPT-3: Brown et al. 2020, 9.

42. Common Crawl 2024.

43. Vaughan 2024.

44. *Ibid.*

45. For the relevant underlying academic paper, see Gao et al. 2020.

46. See, for instance, Kaggle (undated), which lists 3,462 datasets suitable for natural language processing.

47. Beaumont 2022; Schuhmann et al. 2022.

48. Ramesh et al. 2022; OpenAI 2022.

49. Hugging Face 2024.

50. Papers with Code (undated).

51. Kaggle (undated).

image, or music generation, object recognition, anomaly detection, speech recognition, autonomous driving, facial recognition, machine translation, recommendation systems, emotion recognition systems, image upscalers, fake news detectors, handwriting recognition, grammatical error correction, and literally hundreds of other applications.

2.4 The World Wants to Regulate AI

The realization that AI systems can easily be applied to a very wide range of economic sectors and parts of society has brought up questions about the regulation of AI applications as well as access to AI technology.⁵² Given the disruptive effects of AI, this is not entirely unsurprising: AI-related concerns extend not only to individual human rights and values but also to broader global power dynamics.

2.4.1 Geopolitical Developments

The growing recognition that AI offers both broad economic benefits and poses significant threats has led to a twofold dynamic. On the one hand, both the EU and the United States (US) have launched legislative packages – the Chips Act⁵³ and the CHIPS Act of 2022,⁵⁴ respectively – aimed at boosting the local recreation of *semiconductor infrastructure* to secure the production of advanced computer chips. In doing so, these jurisdictions are signalling their intention to reduce their dependence⁵⁵ on Taiwan, which currently produces approximately 90% of the world’s state-of-the-art microchips.⁵⁶

On the other hand, the heavy reliance on Taiwan to maintain leadership in both civilian and military technology has led the US to tighten export controls on China.⁵⁷ Among others, the US has banned semiconductor firm Nvidia from exporting its most cutting-edge GPUs, thereby putting China at a competitive disadvantage in terms of the availability of computing power to train AI systems.⁵⁸ Similarly, the US has sought to further cripple China’s ability to build and expand its own semiconductor industry by

52. AIA, Recs 4-5.

53. Chips Act.

54. CHIPS and Science Act, Public Law (US) 117-167.

55. Chips Act, Rec. 2: ‘A framework for increasing the Union’s resilience in the field of semiconductor technologies should be established, reinforcing the Union’s semiconductor ecosystem by reducing dependencies, enhancing digital sovereignty, stimulating investment, strengthening the capabilities, security, adaptability and resilience of the Union’s semiconductor supply chain.’

56. Lee et al. 2021.

57. Bureau of Industry and Security, Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections, 88 Federal Register 73458 (2023); for an overview of the relevant press releases, see BIS 2023.

58. As disclosed in an investor report: Nvidia 2023: ‘The Interim Final Rule amends ECCN 3A090 and 4A090 and imposes additional licensing requirements for exports to China and Country Groups D1, D4, and D5 ... of the Company’s integrated circuits exceeding certain performance thresholds.’

barring the Netherlands-based company ASML⁵⁹ from selling its extreme ultraviolet (EUV) lithography systems – machines needed to mass-produce microchips – to China.⁶⁰ In doing so, the US is further kneecapping China’s ability to develop computationally intensive AI systems.

2.4.2 *Calls for a Pause on AI Development*

Aside from the geopolitical developments mentioned above, parts of the AI industry itself have called for a temporary pause on AI development. This particular movement culminated in 2023 with the publication of an open letter by the Future of Life Institute, signed by prominent AI researchers, academics, and industry players.⁶¹ Their core argument is that unfettered competition will force AI companies to behave recklessly, potentially unleashing AI models that pose serious threats to society and humanity.⁶² This argument rests on two premises: first, that AI regulation is lagging behind, and second, that developers can no longer fully control their own AI models once they are released on the market. As a result, the open letter argues that an AI development freeze is needed. Ultimately, the proposed AI development freeze did not materialize, and firms have continued their AI research efforts.

When considering the market for AI, this outcome is not entirely unexpected. With the release of several notable open-source⁶³ and closed-source⁶⁴ AI models, the proverbial genie may already be out of the bottle, and it will be nearly impossible to put it back in.⁶⁵ Especially since firms such as OpenAI,⁶⁶ Anthropic,⁶⁷ and Google⁶⁸ operate on a for-profit basis and charge a subscription fee for access to their models. As a result, these firms face strong incentives to be first-to-market with the next-generation frontier model – with an even higher number of parameters – making it unrealistic to expect them to stop improving their AI models. And that leaves only one way forward: AI regulation.

2.4.3 *Race to Regulate*

The concerns about AI outlined earlier in this chapter are not entirely unfounded, and recent AI developments have indeed prompted a race to regulate AI. This, of course, raises an important preliminary question: *what to regulate?* Perhaps only the so-called

59. ASML is the only firm in the world that manufactures EUV lithography systems (i.e., chip-making equipment) capable of producing microchips with transistor sizes in the range of 2-3 nanometer (nm).

60. ASML 2024.

61. Future of Life Institute 2023.

62. *Ibid.*

63. E.g., Llama 3.1 405B (Meta), Command R+ (Cohere), Mixtral-8x22b-Instruct-v0.1 (Mistral).

64. E.g., GPT-4o (OpenAI), Claude 3.5 Sonnet (Anthropic), Gemini 1.5 Pro (Google).

65. In the realm of large language models, AI developments have even spawned chatbot leaderboards: LMSYS 2024; *see also* the relevant research paper: Chiang et al. 2024.

66. OpenAI 2023.

67. Anthropic 2024.

68. Google 2023.

frontier models that represent the state of the art, such as GPT-4o (OpenAI), Claude 3.5 Sonnet (Anthropic), Llama 3.1 405B (Meta), and Gemini 1.5 Pro (Google), should become subject to regulation? Or perhaps AI models should be regulated on the basis of certain performance thresholds – such as a maximum amount of computation required to operate the model – or that AI models must perform their task within a certain energy budget? Alternatively, perhaps the emphasis should be on regulating specific AI use cases rather than AI models themselves, suggesting, for example, that more intrusive AI applications (e.g., facial recognition) should be subject to more stringent rules.

While many of these questions remain open to discussion today, the world has witnessed multiple efforts around the globe to reign in AI.⁶⁹ In the US, the Biden administration adopted Executive Order 14110 on a Safe, Secure, and Trustworthy AI.⁷⁰ In turn, China has enacted rules governing specific applications of AI, such as recommender systems and generative AI (GenAI).⁷¹ India is also in the process of drafting new legislation, in the form of the Digital India Act, to regulate high-risk AI systems. However, on 13 March 2024, the first comprehensive AI regulatory instrument of its kind was adopted by the EU: the *AI Act*. The Act was published in the Official Journal of the EU as Regulation (EU) 2024/1689 on 12 July 2024.

3 LEGISLATIVE JOURNEY OF THE AI ACT

3.1 2018: AI for Europe

The first tentative steps towards the regulation of AI in Europe were already taken in 2018 when the Commission published a communication entitled *AI for Europe*.⁷² In this document, the Commission not only sets out its ambition to embrace AI but also states that the European way forward on AI should be aligned with the following three objectives:⁷³

- (1) boosting the EU's technological and industrial capacity and AI uptake;
- (2) preparing for the socio-economic changes brought about by AI; and
- (3) ensuring that an ethical and legal framework is in place to align AI developments with European values and fundamental rights.

The publication was followed by the creation of a High-Level Expert Group on Artificial Intelligence (HLEG),⁷⁴ tasked with drafting ethical guidelines for the development of trustworthy AI, and the publication of a twenty-three-page *Coordinated Plan on AI*,⁷⁵

69. Several AI regulation trackers are available online. Consider, for instance, IAPP 2024; Sun 2024.

70. Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 88 Federal Register 75191 (2023).

71. For an English translation, see Creemers et al. 2022.

72. Commission Communication 2018.

73. *Ibid.*, at 3.

74. Commission 2018a.

75. Commission 2018b, 1.

which outlines a more detailed strategy for maximizing the benefits of AI for all Europeans. Specifically, the action plan puts forward the idea of ‘AI made in Europe’, which should be implemented by pursuing strategic investments, focussing on AI talent development, creating a common European data space, and establishing an appropriate ethical and regulatory framework.⁷⁶

3.2 2019: High-Level Expert Group on AI

In 2019, the HLEG published its *Ethics Guidelines for Trustworthy AI*.⁷⁷ The approach of the HLEG Guidelines is threefold.

First, it identifies a set of four ethical principles – respect for human autonomy, prevention of harm, fairness, and explicability – that must be considered during the development, deployment, and use of AI systems.⁷⁸

Second, these principles are expanded into seven⁷⁹ key ethical requirements for AI systems:⁸⁰

- (1) human agency and oversight;
- (2) technical robustness and safety;
- (3) privacy and data governance;
- (4) transparency;
- (5) diversity, non-discrimination and fairness;
- (6) environmental and societal well-being; and
- (7) accountability.

Third, the Guidelines also present both technical and non-technical methods for operationalising these ethical requirements.⁸¹ This is achieved by means of an AI assessment list, which provides specific questions to assess compliance with the seven requirements outlined above.

3.3 2020: White Paper on AI

The Commission, now fully equipped with a detailed action plan and having received extensive input from the HLEG, released a *White Paper on AI* in February 2020.⁸² The document, entitled ‘A European approach to excellence and trust’, reiterates many of the findings of previous reports and uses ambitious language such as ‘seizing the opportunity’ and ‘building an ecosystem of excellence’. As a prelude to the proposal text of the AI Act, the Commission’s White Paper dedicates specific attention to

76. *Ibid.*, at 2.

77. HLEG Guidelines.

78. *Ibid.*, at 11-12.

79. The list of seven ethical requirements is intended to be ‘non-exhaustive’ and there is no particular hierarchy among these requirements: *ibid.*, at 14.

80. *Ibid.*

81. *Ibid.*, at 26-31.

82. White Paper on AI.

delineating the scope of potential future EU rules on AI. Among other things, attention is paid to how AI should be defined and that any definition should be sufficiently flexible to accommodate technological progress.⁸³

It is also noteworthy that the Commission is mindful of interactions with other legal instruments, such as consumer protection rules and data protection legislation. In this respect, the Commission clearly states that all existing provisions of EU law will apply in relation to AI, although certain updates may be necessary.⁸⁴

The White Paper subsequently refers to a *risk-based approach* to AI.⁸⁵ To ensure effectiveness, the Commission proposes to differentiate whether AI systems are ‘high risk’ or not. Without going into details, the Commission suggests that the qualification as high risk should depend on the specific sector in which the AI system is used and the significance of the risks that can be expected in light of the activities performed by the AI system.

3.4 2021: Proposal for a Regulation Laying Down Harmonized Rules on AI

Over time, the Commission’s vision for AI regulation continued to take shape, culminating in the production of an updated *Coordinated Plan on AI*,⁸⁶ a proposal text for an *Artificial Intelligence Act*,⁸⁷ as well as an *Impact Assessment*⁸⁸ accompanying the proposed AI Act. The explanatory memorandum included in the proposed AI Act not only reiterates much of the information contained in earlier documents but also expands on it in several respects. Among others, the text takes the time to emphasize that AI can have a positive impact on society, but that it is of the utmost importance to also address any risks.⁸⁹ Ultimately, the proposed AI act seeks to create a single European market for trustworthy AI. It aims to do this by ensuring that the development of AI can take place on the basis of a framework of European values, while at the same time providing much-needed legal certainty to encourage AI innovation.

There are two important takeaways from the text of the proposal. First, the proposed AI Act seeks to establish a *definition of an AI system* that would apply across the EU. The proposed (but now outdated) definition⁹⁰ refers to a dedicated Annex to further elaborate on what is to be considered an AI system. In the Annex, three different approaches are mentioned that qualify as AI: (1) machine learning approaches; (2) logic- and knowledge-based approaches; and (3) statistical approaches.⁹¹

83. *Ibid.*, at 16.

84. *Ibid.*

85. *Ibid.*, at 17.

86. Commission 2021, 1.

87. AI Act Commission Proposal.

88. Commission SWD Impact Assessment.

89. AI Act Commission Proposal, 1.

90. AI Act Commission Proposal, Art. 3(1): “artificial intelligence system” (AI system) means software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with’.

91. AI Act Commission Proposal, Annex I.

Second, the proposal text further refines the *risk-based approach* mentioned in the 2020 White Paper. Among others, the proposal differentiates between:

- (1) *Prohibited AI systems*, which pose an unacceptable risk to EU values and therefore cannot be placed on the EU market.
- (2) *High-risk AI systems*, which pose a serious risk to EU values, and therefore must be subject to a set of strict requirements before being allowed to circulate in the EU.
- (3) *Low-risk AI systems*, which pose only a minimal risk to EU values and are therefore subject only to limited transparency requirements.

3.5 2022-2024: Run-Up to the AI Act

3.5.1 The Council's Position

As part of the ordinary legislative procedure, discussions took place in the Council and its preparatory bodies in 2022.⁹² On 6 December 2022, the Council published its common position on the AI Act.⁹³ Compared to the original text of the proposal, several notable changes and deviations were found:⁹⁴

- (1) The *definition of an AI system* has been redrafted in a narrower manner.⁹⁵ The definition also no longer refers to Annex I of the proposed AI Act to delineate the techniques and approaches that fall within its scope. Instead, the relevant language is now entirely contained within the definition itself. Unlike Annex I, however, the revised definition no longer mentions 'statistical approaches', retaining only machine learning and logic-based approaches.
- (2) Regarding the *scope of application* of the proposed AI Act, several exclusion criteria have now been explicitly added to the proposal (e.g., military purposes).⁹⁶ In the same vein, the revised text now also clarifies that the obligations of the AI Act should not apply to AI systems that are used for research purposes.⁹⁷
- (3) On the topic of *prohibited AI practices*, the Council's position extends the prohibition on social scoring to private actors, while also inscribing an exemption for real-time remote biometric identification systems for law enforcement purposes in exceptional circumstances.⁹⁸

92. For a further analysis of the position of the Council, see Chapter 5, Pehlivan, in this volume, s. 2.1.

93. AI Act Council Proposal.

94. Council 2022.

95. AI Act Council Proposal, Art. 3(1).

96. *Ibid.*, Art. 2(3).

97. *Ibid.*, Art. 2(6).

98. *Ibid.*, Art. 5.

- (4) Apropos *high-risk AI systems*, the Council provides for a new carve-out within the high-risk AI system categorization when the AI system's output is merely ancillary to the action or decision to be taken, with the effect being that no significant risks to the health, safety or fundamental rights arise.⁹⁹ In effect, this derogation entails that some AI systems may now escape the classification as high risk. Staying on the topic of high-risk AI systems, many of the substantive obligations have also been slightly redrafted and further refined.
- (5) The Council added a number of new provisions that apply when AI systems can be used for a wide variety of different purposes, so-called *general purpose AI*.¹⁰⁰ In essence, the Council's position reflects that some requirements applicable to high-risk AI systems should also apply to general purpose AI.
- (6) Concerned about the implications of the proposed AI Act on innovation, the Council proposed to amend the provisions on *measures in support of innovation*.¹⁰¹ In particular, the aim was to be more innovation-friendly and to encourage evidence-based learning. As a result, AI regulatory sandboxes should allow for the testing of AI systems under real-world conditions.

3.5.2 The Parliament's Position

Following the adoption of a common position by the Council, the Parliament was next to take action.¹⁰² A joint committee procedure was launched under the auspices of the Committee on Internal Market and Consumer Protection (IMCO) and the Committee on Civil Liberties, Justice and Home Affairs (LIBE). The resulting joint report, dated 16 May 2023, introduces further changes to the Commission's proposal.¹⁰³ These changes mainly revolve around the following six topics:¹⁰⁴

- (1) The *definition of an AI system* remained a subject of contention and inevitably saw further amendments. While some Members of the European Parliament (MEPs) preferred a more broadly worded definition to ensure future proofing, the rapport eventually settled on aligning the AI Act's definition of an AI system with the definition used by the Organisation for Economic Co-operation and Development (OECD).¹⁰⁵
- (2) The *list of prohibited AI practices* was substantially revised to include biometric categorization systems, AI systems intended for risk assessments,

99. *Ibid.*, Art. 6(3).

100. *Ibid.*, Art. 4a, 4b, and 4c.

101. *Ibid.*, Arts 53-55a.

102. For a further analysis of the position of the Parliament, see Chapter 5, Pehlivan, in this volume, s. 2.2.

103. AI Act Parliament Amendments.

104. European Parliament, *Parliament's Negotiating Position on the Artificial Intelligence Act*, 1 (2023).

105. AI Act Parliament Amendments, Art. 3.

- emotion recognition systems, and a blanket ban on AI systems that create facial recognition databases through untargeted scraping from the internet.¹⁰⁶
- (3) New substantive requirements were added for *high-risk AI systems*: to qualify as a high-risk AI system, these AI systems must pose a ‘significant risk’ to the health, safety, or fundamental rights of natural persons.¹⁰⁷ The assessment of high risk entails a self-assessment procedure.
 - (4) Additional rules were added regarding the regulation of *General-Purpose Artificial Intelligence (GPAI) systems*, including a new definition of a ‘foundation model’, the extension of existing provisions to foundation models, as well as some new specific substantive obligations.¹⁰⁸
 - (5) Under the governance section of the AI Act, it was proposed to establish a *European Artificial Intelligence Office (AI Office)* in order to strengthen the enforcement of the Act.¹⁰⁹ The proposed role of the AI Office would be to support, advise, and cooperate with Member States, and also monitor the effective application of the Regulation.
 - (6) The introduction of a generous, albeit not complete, exemption for *free and open-source* AI systems.¹¹⁰

3.5.3 Trilogue Negotiations

Trilogue negotiations between the Parliament, the Council, and the Commission took place next. Between June and December 2023, the provisional text of the AI Act was further amended to arrive at a compromise text. Notable topics of discussion revolved around the definition of AI systems, the list of banned AI systems, requirements for high-risk AI systems, obligations pertaining to GPAI, measures in support of innovation and small and medium-sized enterprises (SMEs), and the enforcement of the AI Act.¹¹¹

Ultimately, it took a thirty-six-hour ultramarathon of negotiations to fully hammer out the final details of the AI Act and resolve any outstanding issues. On 9 December 2023, EU policymakers announced that they had reached a political agreement on what would become the first global instrument of AI regulation.¹¹² The provisional text of the AI Act was subsequently endorsed by the Parliament on 13 March 2024, with the roll call showing 523 votes in favour, 46 against, and 49 abstentions.¹¹³

106. *Ibid.*, Art. 5.

107. *Ibid.*, Art. 6.

108. *Ibid.*, Arts 3 and 28b.

109. *Ibid.*, Art. 56.

110. *Ibid.*, Art. 2.

111. Fazlioglu 2023.

112. Bertuzzi 2023.

113. Parliament 2024.

4 THE AI ACT

While a detailed analysis is beyond the scope of this chapter,¹¹⁴ it is nonetheless illustrative to provide a high-level overview of the finalized AI Act. In terms of the overall structure, the Act consists of 180 Recitals, 113 Articles, and 13 Annexes. Seven key issues surrounding the AI Act deserve attention before proceeding.

A closer look at the definitions taken up in Article 3 indicates that the European definition of an AI system reads as follows:¹¹⁵

‘AI system’ means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

Assessing whether a specific AI system falls within the definition will certainly pose its fair share of challenges, and will certainly spark much scholarly debate about what words such as ‘autonomy’ and ‘adaptiveness’ mean or how they should be applied.

It is also necessary to draw attention to the *scope of application* of the AI Act. Article 2 is quite clear in this respect, stipulating that it is mainly applicable to: (1) providers who place AI systems or GPAI models on the EU market or put those systems or models into service in the EU; (2) deployers of AI systems that are located in the EU and that use an AI system under their authority;¹¹⁶ (3) providers and deployers of AI systems located in a third country, but where the output produced by the AI system is used in the EU; and (4) importers and distributors of AI systems in the EU.

The *risk-based approach* should also be further contextualized.¹¹⁷ As the term suggests, the specific requirements applicable to an AI system will vary depending on the intensity and scope of risks posed by that AI system to the values and rights enshrined in EU law. To that end, the AI Act classifies AI systems into four different risk categories. The first category is *prohibited AI practices* which, due to their harmful impact, cannot be deployed in the EU. The second category concerns *high-risk AI systems*, which potentially pose significant risks to EU citizens’ health, safety, and fundamental rights and must comply with an extensive set of mandatory requirements. The third category encompasses *limited-risk AI systems*, which are subject to transparency obligations due to their ability to shape false impressions (e.g., synthetic content or artificially generated text, images, or video). The fourth category encapsulates *minimal-risk AI systems*, which can be used freely albeit while still complying with other applicable rules (e.g., the GDPR).

114. For a more detailed analysis of the provisions of the AI Act, the reader is referred to the relevant article-by-article commentaries contained in Ceyhun Necati Pehlivan, Nikolaus Forgó & Peggy Valcke (eds), *The EU Artificial Intelligence (AI) Act* (Kluwer Law International 2025).

115. AIA, Art. 3(1).

116. *Ibid.*, Art. 3(4) *in fine*: an exception applies when the AI system is deployed in the course of a personal, non-professional activity.

117. *Ibid.*, Rec. 26.

Some comments should be added on the topic of *GPAI*. In this respect, the AI Act sets out two relevant definitions: a definition for GPAI model¹¹⁸ and one for GPAI system.¹¹⁹ In essence, the ‘general purposeness’ of an AI model refers to the ability to build different AI systems – each with different use cases – on the basis of the GPAI model. In this respect, the ‘generality’ of the AI model depends on (among others) the number of parameters included in the model (*see* GPT-3 with 175 billion parameters and GPT-4 with 1.8 trillion parameters). In turn, the number of parameters has a significant impact on the amount of computation required to use the AI model, both during the training and inference phases. In light of this, EU policymakers opted to use the amount of computing power required as a criterion for determining whether a GPAI model should qualify as ‘high-risk’ or not. This measure is expressed in FLOPs: if more than 10^{25} FLOPs are required to train the AI model, a presumption applies that the AI model poses a systemic risk.¹²⁰ In addition, Annex XIII sets out several criteria on which the Commission can rely to designate GPAI as a model with systemic risk on an ad hoc basis. It should also be noted that a generous exemption applies to free and open-source AI models.¹²¹

The AI Act contains a dedicated chapter on measures in support of innovation. In this context, Member States face an obligation to operationalize so-called *AI regulatory sandboxes*,¹²² to be understood as controlled environments where developers can train, test, and validate new AI systems under regulatory supervision but with certain relaxations.¹²³ This regime also extends to the possibility of testing AI systems under real-world circumstances. In this respect, the AI Act accommodates the specific needs of SMEs by providing an easier compliance procedure for market entry.¹²⁴

With regard to *enforcement*, the approach taken in the AI Act is multi-layered. A first anchor point is that the AI Act mandates Member States to designate at least one market surveillance authority to oversee the correct application of the AI Act. This market surveillance authority possesses extensive enforcement powers, ranging from investigative powers to the ability to impose sanctions for non-compliance, including the imposition of administrative fines of up to 7% of the worldwide turnover or EUR 35 million, whichever is higher.¹²⁵ Market surveillance authorities and the Commission can also launch joint investigations. The AI Office can play a supporting role.

The AI Act will have a staggered *entry into force*, spread over a two-year window (*see* Figure 4.5).¹²⁶ Six months after publication in the Official Journal, the provisions

118. *Ibid.*, Art. 3(63).

119. *Ibid.*, Art. 3(66).

120. *Ibid.*, Art. 51(2); Note that GPT-4 – with an estimated 2.15×10^{25} floating point operations required for its training – exceeds this threshold.

121. *Ibid.*, Art. 2(12).

122. *Ibid.*, Art. 57.

123. Regarding these relaxations, *see, for instance, ibid.*, Art. 57(12): ‘provided that the prospective providers observe the specific plan and the terms and conditions for their participation and follow in good faith the guidance given by the national competent authority, no administrative fines shall be imposed by the authorities for infringements of this Regulation’.


124. *Ibid.*, Art. 57(9)(e).

125. *Ibid.*, Art. 99.

126. *Ibid.*, Art. 113.

on prohibited AI systems will come into force (2 February 2025). At the twelve-month mark, the rules on GPAI will take effect (2 August 2025). After twenty-four months, most of the other obligations of the AI Act come into force (2 August 2026). Lastly, at thirty-six months, the rules on high-risk AI systems enter into force (2 August 2027).

Figure 4.5 Timeline of the Entry into Force and Application of the AI Act¹²⁷



5 A COMPREHENSIVE EUROPEAN DIGITAL FRAMEWORK FOR AI

The AI Act does not exist in a vacuum but is part of a larger body of legal instruments dealing with digitization in the EU. In this respect, the AI Act itself devotes attention to clarifying how its provisions are consistent with other EU instruments. Reading Recital 9 of the AI Act, it is posited that EU rules on AI should be:¹²⁸

without prejudice to existing Union law, in particular on data protection, consumer protection, fundamental rights, employment, and protection of workers, and product safety.

This effectively suggests that the *lex specialis* principle¹²⁹ will be relevant to consider in situations where two EU instruments apply to the same factual circumstances. As shown in Table 4.2, sometimes the scope of application of the AI Act may indeed overlap with other instruments of EU law. In such cases, specific rules should take precedence over general rules, and the more general rules should be disapplied to the extent that they conflict with the more specific rules. In addition, whenever possible, different legal instruments should be applied concurrently to provide the most comprehensive level of protection. While not intended to be exhaustive, this section provides an initial overview of the intersections between the AI Act and other EU instruments related to digitization.

127. The visual is the author's own work, and is based on the information contained in *ibid.*

128. *Ibid.*, Rec. 9.

129. Based on the Latin maxim *lex specialis derogat legi generali*: more specific rules prevail over more general rules.

Table 4.2 The AI Act as Part of the EU Digital Acquis¹³⁰

<i>Regulatory Instrument</i>	<i>Actors</i>	<i>Rightsholders</i>	<i>Objective</i>	<i>Material Scope</i>
AI Act	– Providers of AI systems and AI models – Deployers of AI systems – Importers of AI systems – Distributors of AI systems	– Individuals – Society – Democracy	Promoting the uptake of human-centric and trustworthy AI, while ensuring a high level of protection of health, safety, fundamental rights ¹³¹	Placing on the market or putting into service AI systems or AI models ¹³²
GDPR	– Data controllers – Data processors	– Data subjects	Protecting natural persons with regard to the processing of personal data ¹³³	Processing of personal data wholly or partly by automated means ¹³⁴
DGA	– Public sector bodies – Data intermediation services – Data altruism organizations	– Data subjects – Data holders	Creating a harmonized framework for data exchanges and laying down certain basic requirements for data governance ¹³⁵	Re-using, within the Union, of certain categories of data ¹³⁶
DA	– Manufacturers of connected products – Data holders – Data recipients – Public sector bodies – Providers of data processing services – Participants in data spaces	– Users of connected products or related services	Ensuring fairness in the allocation of value from data among actors in the data economy and fostering fair access to and use of data ¹³⁷	Making available product data and related service data and their subsequent use ¹³⁸

130. The tabular overview is the author's own work and is based on the wording used in the different legal instruments.

131. AIA, Art. 1.

132. *Ibid.*, Art. 2.

133. GDPR, Art. 1.

134. *Ibid.*, Art. 2.

135. DGA, Rec. 3.

136. *Ibid.*, Art. 1.

<i>Regulatory Instrument</i>	<i>Actors</i>	<i>Rightsholders</i>	<i>Objective</i>	<i>Material Scope</i>
DMA	– Undertakings designated as gatekeepers ¹³⁹	– End users – Business users	Laying down harmonized rules ensuring for all businesses, contestable and fair markets in the digital sector across the Union where gatekeepers are present ¹⁴⁰	Provision of core platform services ¹⁴¹ by gatekeepers to business users or end users in the Union ¹⁴²
DSA	– Providers of intermediary services, including: – very large online platforms; and – very large online search engines	– Recipients of intermediary services	Setting out harmonized rules for a safe, predictable and trusted online environment that facilitates innovation and in which fundamental rights are effectively protected ¹⁴³	Provision of intermediary services in the internal market ¹⁴⁴

5.1 Interface Between the GDPR and the AI Act

The first relevant area of attention concerns the interplay between the AI Act and the GDPR. As recalled, the GDPR is concerned with the protection of data subjects with respect to the processing of their personal data.¹⁴⁵ It is important to note that the GDPR is *technologically neutral*;¹⁴⁶ its only anchor point is the processing of personal data, regardless of how or by what means this is done.

As evidenced by Article 4(2) of the GDPR, the definition of ‘processing’ is very broad and covers virtually any activity involving personal data. The implication thereof is that the GDPR can find application to AI systems, to the extent that personal data is processed by the AI system. Such situations are likely to arise either because AI systems process personal data during the training phase (i.e., personal data is processed as

137. DA, Rec. 119.

138. *Ibid.*, Art. 1.

139. DMA, Art. 2(1) *juncto* Art. 3.

140. *Ibid.*, Art. 1.

141. *Ibid.*, Art. 2(2).

142. *Ibid.*, Art. 1.

143. DSA, Art. 1.

144. *Ibid.*

145. GDPR, Art. 1.

146. *Ibid.*, Rec. 15.

input data) or because AI systems process personal data during the inference phase (i.e., personal data is processed as output data). In these circumstances, both the AI Act and the GDPR will apply, which in practice means that a firm can hold a dual qualification. Consider, for instance, the following three situations:¹⁴⁷

- (1) *Qualification as a provider under the AI Act and as a data controller under the GDPR*: a firm processes personal data for the purpose of training a new AI system, that is intended to be placed on the Union market.
- (2) *Qualification as a deployer under the AI Act and as a data controller under the GDPR*: a firm acquires an AI system from another entity and subsequently uses that AI system under its authority to process personal data (e.g., for staff recruitment purposes).
- (3) *Qualification as a provider under the AI Act and as a data processor under the GDPR*: a firm sells an AI system to another firm, but remains involved in the second firm's processing of personal data (e.g., because the AI system is hosted on the vendor's servers).

When such circumstances arise, the principles of both instruments must be respected. This is also reflected as such in Recital 10 of the AI Act:¹⁴⁸

[i]t also does not affect the obligations of providers and deployers of AI systems in their role as data controllers or processors.

To streamline these interactions, the text of the AI Act also aligns itself with several definitions used in the GDPR.¹⁴⁹

Now that the AI Act has been formally adopted, there is also, in some ways, a greater alignment between the substantive principles of the two instruments. In some respects, there is a clear *complementarity* between the AI Act and the GDPR. For instance, some of the principles of the GDPR also correspond to the principles of the AI Act. Befitting examples in this respect are the provisions on transparency,¹⁵⁰ data governance,¹⁵¹ data protection impact assessments (DPIAs),¹⁵² and human oversight.¹⁵³ In other respects, however, there are also *tensions* between both instruments. For instance, AI systems typically require the processing of large amounts of data during the training phase, which could conflict with the principle of data minimization and storage limitation.¹⁵⁴ Similarly, AI systems can sometimes be used for a variety of different purposes, making it difficult to clearly delineate the purpose of data processing and potentially putting these AI systems at odds with the principle of purpose

147. Clark & Muhammed Demircan 2024, 9.

148. AIA, Rec. 10.

149. This is the case for 'biometric data', 'special categories of personal data', 'personal data', and 'non-personal data', and 'profiling'.

150. AIA, Art. 50 *juncto* GDPR, Art. 5(1)(a).

151. AIA, Art. 10 *juncto* GDPR, Art. 5(1)(d)-(f).

152. AIA, Art. 27 *juncto* GDPR, Art. 35.

153. AIA, Art. 14 *juncto* GDPR, Art. 22.

154. GDPR, Art. 5(1)(c).

limitation. More broadly, the strict application and compliance with the data processing principles of the GDPR may have the effect of limiting the functionality of AI systems.

Prior to the adoption of the AI Act, there have already been a number of cases involving AI systems brought before national data protection authorities. Two cases in particular are worth mentioning: the banning of ChatGPT in Italy¹⁵⁵ and the fining of Clearview AI.¹⁵⁶ If anything, these cases indicate that this is an area of law to monitor.

5.2 Interface Between the DGA and the AI Act

The intersection between the EU's DGA and the AI Act is also worth analysing. At its core, the DGA is designed to facilitate the re-use of certain categories of data held by public sector bodies, and its main objective is to provide a coherent framework to facilitate the sharing of data in a secure and efficient manner.¹⁵⁷ Thus, while the DGA is not specifically focused on AI, it serves as a prerequisite for placing high-quality AI systems on the Union market by ensuring that data that could be used for training AI systems is handled with due care. In this respect, there exists some overlap between the AI Act and the DGA. Three elements in particular are worth highlighting.

First, there is the *data sharing* dimension. The DGA is primarily concerned with sharing data in a trustworthy manner, while the AI Act is concerned with the use of this data.¹⁵⁸ In that way, the DGA acts as an enabler for the training of AI systems on the basis of high-quality datasets. Second, there is the *data governance* aspect: both instruments draw attention to data governance and the need to properly manage data.¹⁵⁹ As such, the requirements of the DGA align well with those of the AI Act. A final element to highlight is the *risk management* angle. While the DGA seeks to manage possible risks associated with data sharing, the AI Act seeks to ensure that once data is used by AI systems, the resulting risks by deploying those AI systems are also appropriately administered.¹⁶⁰

Ultimately, the DGA is aimed at developing a borderless digital internal market and a human-centric, trustworthy and secure data society, while the AI Act is intended to improve the functioning of the single market by promoting the deployment of human-centric and trustworthy AI.¹⁶¹ In this sense, there is an element of linkage between the two instruments. Only when trustworthy and secure data sharing is enabled under the DGA, can human-centric AI systems be developed under the AI Act. Therefore, both legal instruments together are an integral part of shaping a comprehensive data ecosystem within the EU.

155. Garante 2023.

156. CNIL 2022.

157. DGA, Rec. 3 and Art. 3.

158. *Ibid.*, Art. 5(1) *juncto* AIA, Art. 10.

159. DGA, Art. 5(3) *juncto* AIA, Art. 10.

160. DGA, Art. 5(4)-(7) *juncto* AIA, Art. 10.

161. DGA, Rec. 3 *juncto* AIA, Rec. 1.

5.3 Interface Between the DA and the AI Act

Another area of note is the intersection between the DA and the AI Act. The premise of the DA is that, due to the proliferation of smart devices and the widespread availability of internet connectivity, the volume and value of data have increased tremendously since the 2010s.¹⁶² At the same time, the DA also identifies barriers to data sharing that prevent much of the value of that data from being unlocked.¹⁶³ In light of this, the DA sets out rules to remove the barriers to sharing data from connected products and their ancillary services. Essentially, the DA operates in two ways. First, it creates new obligations for businesses to make data from connected products and related services available to users.¹⁶⁴ Second, it empowers those users to share their data with third parties.¹⁶⁵

At this point, the analysis of the intersection between the DA and the AI Act shows some similarities to the findings made under the DGA, namely that data from connected devices and their ancillary services can serve as useful input data for training AI systems. In that sense, the element of linkage between the DA's data-sharing obligations and the ability to train AI systems is also present here. The key aspect is that the DA is an enabler of the provisions of the AI Act, in particular the data and data governance requirements that are taken up in Article 10 of the AI Act.

It is worth looking at a concrete example to illustrate this. Consider a smart speaker that collects audio data from interactions with its owner. The collected data – voice commands and other spoken input – is primarily used by the smart speaker to provide some sort of functionality to the user (e.g., the ability to control a smart lighting system through voice commands). However, the potential utility of the voice data goes beyond the immediate functionality of the smart speaker. For example, the voice data could also be used to develop an AI language learning application: the same voice excerpts could also be relied upon to identify pronunciation errors, suggest corrections, and tailor language learning exercises to the user's language proficiency level. The data sharing obligations of the DA make this kind of value realization possible.

Of particular interest is the DA's attention to the importance of structured and machine-readable data for users.¹⁶⁶ This suggests that the DA recognizes the value of data volume and the necessity for AI systems to be trained on large datasets. When data is structured in a well-defined and organized manner, this makes the data ingestion process considerably more straightforward.

5.4 Interface Between the DMA and the AI Act

When the attention turns to digital ecosystems, the discussion quickly gravitates to the EU's DMA. At the heart of the DMA is the observation that a select number of

162. DA, Rec. 1.

163. *Ibid.*, Rec. 2.

164. *Ibid.*, Art. 3.

165. *Ibid.*, Art. 5.

166. E.g., *ibid.*, Art. 3.

undertakings have emerged with considerable economic power, enabling them to exercise a great deal of control over digital markets and the digital economy as a whole.¹⁶⁷ Pursuant to the DMA, these actors are called gatekeepers¹⁶⁸ and it is accepted that the presence of gatekeeper firms on a market may hinder the contestability of that market.¹⁶⁹ Therefore, in order to remedy these situations as well as phenomena such as lock-in effects and data network effects, the DMA imposes a set of rules on gatekeepers to ensure the contestability and fairness of digital markets.

A closer reading of the DMA reveals that the primary focus is on the regulation of core platform services provided or offered by gatekeepers to business users established in the Union.¹⁷⁰ At first glance, this seems to suggest that AI is not directly covered by the DMA. While it is true that AI technology is not (currently) included in the list of ‘core platform services’ explicitly referred to in Article 2(2) of the DMA, it is nonetheless true that services that are included in this list, such as online search engines, virtual assistants, and online advertising, rely on some form of AI. Therefore, an overlap may exist between the application of the DMA and the AI Act to the extent that AI is used by gatekeepers in the process of competition. In these circumstances, the DMA exists to ensure that AI-driven core platform services do not enhance monopoly power or impede the competitive process – for example, by prohibiting algorithmic self-preferencing in the context of online search engines or online advertising – while the AI Act imposes extensive obligations on the functioning of the AI systems themselves.

Within the domain of competition law, particular attention has been paid to GenAI and whether GenAI models should be included in the list of core platform services. Calls to this effect were made by the Parliament in 2024 when it asked the Commission to launch a market investigation into GenAI and raised the question of whether a gatekeeper’s possession of a state-of-the-art GenAI model could further solidify that gatekeeper’s market position.¹⁷¹ In this context, concerns seem to focus in particular on the ability of gatekeepers to leverage immense amounts of data and computing power that are typically only available to a few undertakings. This in turn could lead to new bottlenecks in the development of AI systems. Thus, while the DMA is concerned with keeping markets contestable, the nature of AI introduces some barriers to market entry – in terms of data availability and computing power – that would arguably make it difficult for smaller AI firms to compete with large gatekeepers.

5.5 Interface Between the DSA and the AI Act

Moving from the DMA to the DSA, it is clear that the main objective of the DSA is to create a safe, predictable, and trustworthy online environment.¹⁷² It does so by

167. DMA, Rec. 3.

168. *Ibid.*, Recs 2-3; Art. 2(1).

169. *Ibid.*, Recs 2-3.

170. *Ibid.*, Art. 1.

171. European Parliament, *Competition Policy: Annual Report 2023*, 11 (2024).

172. DSA, Rec. 3.

targeting so-called providers of intermediary services – including (very large) online platforms and search engines – by imposing a number of obligations on them, in particular, to address the spread of illegal content, the risks associated with the spread of disinformation, and transparency in online advertising.¹⁷³

When considering the interplay with the AI Act, it becomes clear that the goals are markedly different but clearly compatible: while the DSA regulates the online platforms on which AI systems can be deployed, the AI Act regulates the AI systems themselves and ensures that their deployment on online platforms is in accordance with fundamental rights obligations.¹⁷⁴ In this regard, two particular facets are worth highlighting.

The first relevant aspect to note is that both the DSA and the AI Act put a lot of emphasis on *transparency*. For the DSA, much of the attention is focused on the transparency of recommender systems that online platforms use to push products or content to their users. The DSA requires, among other things, disclosure of the criteria used to determine what particular information is displayed to end users, as well as the reasons why certain parameters are of relative importance to those users.¹⁷⁵ This element of transparency is also found in the AI Act, for example, through the wording of Recital 27.¹⁷⁶ In relation to AI systems that are considered ‘high risk’, it is relevant to draw attention to Article 13 of the AI Act, which requires transparency regarding the design and development of AI systems and mandates that AI systems must be accompanied by comprehensive documentation – including on the characteristics, capabilities, and limitations of the AI system – in order to properly explain or interpret their output. For low-risk AI systems, the relevant transparency requirements are taken up in Article 50 of the AI Act. Although these are somewhat relaxed compared to Article 13, they remain relevant, for example, to clearly inform users when they interact with AI systems¹⁷⁷ or when they are confronted with artificially generated/manipulated content. Thus, at the intersection of the DSA and the AI Act, both instruments contain provisions aimed at improving users’ understanding of the AI systems with which they interact. Importantly, these transparency requirements may overlap, for example, when AI systems are used to personalize or recommend content on an online platform that falls within the scope of the DSA.

The second key interaction to consider between the DSA and the AI Act is the phenomenon of *AI-driven compliance*. Essentially, this refers to situations where substantive compliance with the DSA can be achieved through the implementation of an AI system. Good examples of this are content moderation requirements¹⁷⁸ and

173. *Ibid.*, Art. 2; Arts 11 et seq.

174. *Ibid.*, Rec. 70 *juncto* AIA, Rec. 3.

175. DSA, Art. 27.

176. *Ibid.*, Rec. 27: ‘Transparency means that AI systems are developed and used in a way that allows appropriate traceability and explainability, while making humans aware that they communicate or interact with an AI system, as well as duly informing deployers of the capabilities and limitations of that AI system and affected persons about their rights.’

177. This requirement does not apply when it is obvious to individuals that they are interacting with an AI systems.

178. I.e., DSA, Art. 14.

notice-and-action mechanisms¹⁷⁹ for dealing with illegal content. Of course, the use of such AI systems can raise concerns in their own right, for example, in relation to the right to freedom of expression. Therefore, while AI-driven compliance is an emerging theme that is likely to attract more attention in the future when AI systems are used for regulatory compliance purposes, those AI systems will need to be subjected to a rigorous assessment under both the DSA and the AI Act to consider whether they impact any fundamental rights.

Overall, while the DSA and the AI Act cover different aspects of the digital economy, they both seek to contribute to a safer and more transparent digital environment, albeit each in its own way. It is clear that AI systems have a role to play in the DSA, whether for product/content recommendation purposes or content moderation. At the same time, the use of AI systems by providers of intermediary services and online platforms requires careful consideration in light of the extensive transparency requirements and potential negative impact on fundamental rights.

6 CONCLUSION

The purpose of this introductory chapter was to provide some background context regarding the development of AI and the adoption of the AI Act in the EU. With the AI Act, the EU is the first jurisdiction to take a clear stance on the future direction of AI growth, steering towards AI that is human-centric and trustworthy. Given the groundbreaking nature of the AI Act, there is little doubt that it will become the subject of extensively legal scholarly writing in the coming years. At the same time, the many interpretative issues surrounding the AI Act will likely require the Commission to issue various implementing acts, delegated instruments, and guidelines.

At the beginning of this chapter, it was mentioned that early developments in the field of AI took place as early as the 1950s. The graph depicting the AI development curve is characterized by ups and downs. From this perspective, it is fair to ask what this curve will look like in ten, twenty, or thirty years' time, and many questions remain about the future of AI. Is the writing on the wall now that the AI Act has been adopted? Will there be a third AI winter? Will AI technology reach a ceiling? What happens when there is no more data to consume? There are no clear answers to any of these questions, nor is it possible to make any firm assertions about them either.

Based on what we know, AI momentum is high in 2024 and there seems to be room for further growth in the coming years. Can the current high level of growth be sustained? Presumably not, and readers should be reminded of the S-curved pattern of innovation: a new technology trend requires some time to gain momentum, then experiences a period of rapid growth and feature expansion, before slowing down in terms of meaningful further improvements, only to be trumped by a whole new technology trend. At the time of writing, however, not a week goes by without another, more powerful, AI model being released to the public. AI applications also seem to be becoming more complex and intricate, with text-to-video generation tools representing

179. I.e., *Ibid.*, Art. 16.

the latest example.¹⁸⁰ In addition, the archival website arXiv.org receives several hundred preprint paper submissions on AI and machine learning developments every day. However, it remains healthy to view AI developments with a certain degree of scepticism. For example, the phenomenon of so-called ChatGPT wrapper start-ups has begun to emerge: new start-ups that claim to be pioneers of an innovative AI-driven product, but that in reality only rely on the AI model of OpenAI behind the scenes. Against this backdrop, it remains important to be wary of overly presumptuous AI claims.

As this chapter draws to a close, I would like to invite you – the reader of this book – to reflect on three critical developments that are likely to shape the future of AI technology and society:

- (1) *The commodification of computing power.* As AI technology advances, the value of computing power may transform it into a key economic asset, potentially serving as a currency in its own right. This shift could redefine how resources are allocated and valued across industries.
- (2) *The rise of data markets.* The increasing reliance on data is not only shaping the capabilities of AI but is also likely to encourage the further growth of specialized markets for data. In the future, these markets may play a central role in governing how data is valued and exchanged.
- (3) *The ongoing demand for AI expertise.* Despite advances in automation, the need for human expertise in AI will likely remain irreplaceable. Even over time, talent in this field will continue to be in high demand and serve as a driving factor in AI innovation.

Combined, these three developments allow us to envision a future where technology, society, and human expertise are increasingly intertwined. It is our task, as humans, to navigate this evolving landscape, to effectively adapt, and to continue to thrive.

7 APPENDIX

Table 4.3 Evolution of teraFLOPs and Prices per teraFLOP over Time in Consumer-Grade GPUs¹⁸¹

Year	GPU	teraFLOPs	MSRP (in USD)	Price per teraFLOP (in USD)
1999	GeForce 256	0.00096	249.00	259,375.00
2000	GeForce2 GTS	0.0016	349.00	218,125.00
2001	GeForce3	0.01	499.00	49,900.00
2002	GeForce4 Ti4600	0.018	399.00	22,166.67

180. OpenAI 2024.

181. See *supra* n. 29.

<i>Year</i>	<i>GPU</i>	<i>teraFLOPs</i>	<i>MSRP (in USD)</i>	<i>Price per teraFLOP (in USD)</i>
2003	GeForce FX 5800	0.024	299.00	12,458.33
2004	GeForce 6800 Ultra	0.045	499.00	11,088.89
2005	GeForce 7800 GTX	0.165	599.00	3,630.30
2006	GeForce 8800 GTX	0.345	599.00	1,736.23
2008	GeForce 9800 GTX	0.432	299.00	692.13
2009	GeForce GTX 295	1.1923	499.00	418.52
2010	GeForce GTX 480	1.3449	499.00	371.01
2011	GeForce GTX 590	2.4883	699.00	280.91
2012	GeForce GTX 680	3.0904	499.00	161.47
2013	GeForce GTX 780	3.9767	649.00	163.20
2014	GeForce GTX 980	4.9807	549.00	110.23
2015	GeForce GTX 980 Ti	6.0544	649.00	107.19
2017	GeForce GTX 1080 Ti	11.3397	699.00	61.64
2018	GeForce RTX 2080 Ti	13.4476	1,199.00	89.16
2021	GeForce RTX 3080 Ti	34.1	1,199.00	35.16
2022	GeForce RTX 4090	82.6	1,599.00	19.36

BIBLIOGRAPHY

Guidelines of Authorities

BIS 2023: Bureau of Industry and Security, *Public Information on Export Controls Imposed on Advanced Computing and Semiconductor Manufacturing Items to the People's Republic of China (PRC)*, US Department of Commerce (2023), www.bis.doc.gov/index.php/documents/about-bis/newsroom/press-releases/3158-2022-10-07-bis-press-release-advanced-computing-and-semiconductor-manufacturing-controls-final/file (accessed 22 April 2024).

Commission 2021: European Commission, *Coordinated Plan on Artificial Intelligence: 2021 Review* (2021).

Commission 2018a: Commission, *Commission Appoints Expert Group on AI and Launches the European AI Alliance*, Digibyte (2018), digital-strategy.ec.europa.eu/en/news/commission-appoints-expert-group-ai-and-launches-european-ai-alliance (accessed 2 May 2024).

Commission 2018b: Commission, *Coordinated Plan on Artificial Intelligence*, Policies (2018), digital-strategy.ec.europa.eu/en/policies/plan-ai (accessed 2 May 2024).

Commission Communication 2018: Commission, *Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: Artificial Intelligence for Europe*, COM/2018/237 final (Brussels, 25 April 2018).

Commission SWD Impact Assessment.

CNIL 2022: Commission nationale de l'Informatique et des Libertés (CNIL), *Délibération de la formation restreinte n° SAN-2022-019 du 17 octobre 2022 concernant la société CLEARVIEW AI* (2022) [for the press release in English, CNIL, *Facial Recognition: 20 Million Euros Penalty against CLEARVIEW AI*, CNIL (2022), www.cnil.fr/en/facial-recognition-20-million-euros-penalty-against-clearview-ai (accessed 6 May 2024)].

Garante 2023: Garante per la protezione dei dati personali, *Provvedimento del 30 marzo 2023*, doc. web n. 9870832.

HLEG Guidelines.

Reports and Recommendations

EPRS 2020: European Parliamentary Research Service (EPRS), *The Impact of the General Data Protection Regulation (GDPR) on Artificial Intelligence* (2020), data.europa.eu/doi/10.2861/293 (accessed 1 April 2024).

White Paper on AI.

Academic Writings

Brown et al. 2020: Tom B. Brown et al., *Language Models Are Few-Shot Learners*, arXiv (2020), doi.org/10.48550/arXiv.2005.14165 (accessed 7 May 2024).

Campbell et al. 2002: Murray Campbell, A. Joseph Hoane & Feng-hsiung Hsu, *Deep Blue*, 134 *Artificial Intelligence* 57 (2002), [doi.org/10.1016/s0004-3702\(01\)00129-1](https://doi.org/10.1016/s0004-3702(01)00129-1).

Chiang et al. 2024: Wei-Lin Chiang et al., *Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference*, arXiv (2024)

Schuhmann et al. 2022: Christoph Schuhmann et al., *LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models*, arXiv (2022).

Clark & Muhammed Demircan 2024: James Clark & Muhammed Demircan (DLA Piper), *EU AI Act: Data Protection Aspects (EU)*, 2024 *Practical Law* 1 (2024).

Cowls et al. 2023: Josh Cowls et al., *The AI Gambit: Leveraging Artificial Intelligence to Combat Climate Change: Opportunities, Challenges, and Recommendations*, 38 *AI & Society* 283 (2023), doi.org/10.1007/s00146-021-01294-x.

Duan et al. 2019: Yanqing Duan, John S. Edwards & Yogesh K Dwivedi, *Artificial Intelligence for Decision Making in the Era of Big Data: Evolution, Challenges and Research Agenda*, 48 *International Journal of Information Management* 63 (2019), doi.org/10.1016/j.ijinfomgt.2019.01.021.

Fazlioglu 2023: Müge Fazlioglu, *Contentious Areas in the EU AI Act Trilogues*, IAPP (2023), iapp.org/news/a/contentious-areas-in-the-eu-ai-act-trilogues/ (accessed 7 May 2024).

Foster & Friston 2023: David Foster & Karl Friston, *Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play* (O'Reilly 2nd ed. 2023).

Gao et al. 2020: Leo Gao et al., *The Pile: An 800GB Dataset of Diverse Text for Language Modeling*, arXiv 1 (2020).

Goodfellow et al. 2016: Ian Goodfellow, Yoshua Bengio & Aaron Courville, *Deep Learning* (The MIT Press 2016).

- Herbosch 2023: Maarten Herbosch, *Intelligent contracteren: Het precontractueel gebruik van systemen op basis van artificiële intelligentie* (Wolters Kluwer 2023).
- Hsu 2002: Feng-hsiung Hsu, *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion* (Princeton University Press 2002).
- Kan 2024: Michael Kan, *Zuckerberg's Meta Is Spending Billions to Buy 350,000 Nvidia H100 GPUs*, PCMag (2024), www.pcmag.com/news/zuckerbergs-meta-is-spending-billions-to-buy-350000-nvidia-h100-gpus (accessed 6 April 2024).
- Kelleher 2019: John D. Kelleher, *Deep Learning* (The MIT Press 2019).
- Mock et al. 2023: Marissa Mock et al., *AI Can Help to Speed up Drug Discovery: But Only If We Give it the Right Data*, 621 Nature 467 (2023), doi.org/10.1038/d41586-023-02896-9.
- Nilsson 2009: Nils J. Nilsson, *The Quest for Artificial Intelligence* (Cambridge University Press 2009).
- Ramesh et al. 2022: Aditya Ramesh et al., *Hierarchical Text-Conditional Image Generation with CLIP Latents*, arXiv (2022).
- Saygin et al. 2000: Ayse Pinar Saygin, Ilyas Cicekli & Varol Akman, *Turing Test: 50 Years Later*, 10 Minds and Machines 463 (2000), doi.org/10.1007/978-94-010-0105-2_2.
- Stein 2020: Amy L. Stein, *Artificial Intelligence and Climate Change*, 37 Yale Journal on Regulation 890 (2020).
- Turing 1950: Alan Turing, *Computing Machinery and Intelligence*, 59 Mind 433, 433 (1950).
- Wang et al. 2019: Alex Wang et al., *SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems*, 33rd Conference on Neural Information Processing Systems 1 (2019).
- Wang et al. 2018: Alex Wang et al., *GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding*, Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP 353 (2018).

News

- Atkinson 2023: Emily Atkinson, *Man Crushed to Death by Robot in South Korea*, BBC News (2023), www.bbc.com/news/world-asia-67354709 (accessed 16 January 2024).
- Bertuzzi 2023: Luca Bertuzzi, *European Union Squares the Circle on the World's First AI Rulebook*, Euractiv (2023), www.euractiv.com/section/artificial-intelligence/news/european-union-squares-the-circle-on-the-worlds-first-ai-rulebook (accessed 7 May 2024).
- Dave 2022: Paresh Dave, *It's Alive! How Belief in AI Sentience Is Becoming a Problem*, Reuters (2022), www.reuters.com/technology/its-alive-how-belief-ai-sentience-is-becoming-problem-2022-06-30 (accessed 1 April 2024).
- Henshall 2023: Will Henshall, *Google DeepMind AI Breakthrough Could Help Battery and Chip Development*, Time (2023), www.time.com/6340681/deepmind-gnome-ai-materials (accessed 19 January 2024).

- Lee et al. 2021: Yimou Lee, Norihiko Shirouzu & David Lague, *Taiwan Chip Industry Emerges as Battlefield in U.S.-China Showdown*, Reuters (2021), www.reuters.com/investigates/special-report/taiwan-china-chips/ (accessed 9 April 2024).
- Reuters 2022: Reuters, *Google Fires Software Engineer Who Claimed Its AI Chatbot Is Sentient*, Reuters (2022), www.reuters.com/technology/google-fires-software-engineer-who-claimed-its-ai-chatbot-is-sentient-2022-07-23 (accessed 1 April 2024).
- The Economist 2023: The Economist, *How to Worry Wisely About Artificial Intelligence*, 447 The Economist, Technology and Society 9 (London, 22 April 2023), www.economist.com/leaders/2023/04/20/how-to-worry-wisely-about-artificial-intelligence (accessed 7 May 2024).

Other

- Anthropic 2024: Anthropic, *Claude 3 Pricing* (2024), www.anthropic.com/api (accessed 22 April 2024).
- ASML 2024: ASML, *Statement Regarding Partial Revocation Export License*, ASML Press Releases (2024), www.asml.com/en/news/press-releases/2023/statement-regarding-partial-revocation-export-license (accessed 22 April 2024).
- Beaumont 2022: Romain Beaumont, *LAION-5B: A New Era of Open Large-Scale Multi-Modal Datasets*, LAION (2022), www.laion.ai/blog/laion-5b (accessed 6 April 2024).
- Common Crawl 2024: Common Crawl, *Open Repository of Web Crawl Data* (2024), www.commoncrawl.org/ (accessed 6 April 2024).
- Council 2022: Council of the European Union, *Artificial Intelligence Act: Council Calls for Promoting Safe AI That Respects Fundamental Rights*, Press Release (6 December 2022), www.consilium.europa.eu/en/press/press-releases/2022/12/06/artificial-intelligence-act-council-calls-for-promoting-safe-ai-that-respects-fundamental-rights/ (accessed 3 May 2024).
- Creemers et al. 2022: Rogier Creemers, Graham Webster & Helen Toner, *Translation: Internet Information Service Algorithmic Recommendation Management Provisions*, Stanford University: DigiChina (2022), www.digichina.stanford.edu/work/translation-internet-information-service-algorithmic-recommendation-management-provisions-effective-march-1-2022 (accessed 23 April 2024).
- Future of Life Institute 2023: Future of Life Institute, *Pause Giant AI Experiments: An Open Letter* (2023), www.futureoflife.org/wp-content/uploads/2023/05/FLI_Pause-Giant-AI-Experiments_An-Open-Letter.pdf (accessed 22 April 2024).
- Google 2023: Google, *Gemini API Pricing*, *Google AI for Developers* (2023), ai.google.dev/pricing (accessed 22 April 2024).
- Hugging Face 2024: *Hugging Face: The AI Community Building the Future*, Hugging Face (2024), huggingface.co (accessed 6 April 2024).
- IAPP 2024: IAPP, *Global AI Law and Policy Tracker*, *International Association of Privacy Professionals* (2024), iapp.org/media/pdf/resource_center/global_ai-law_policy_tracker.pdf (accessed 23 April 2024).
- Kaggle (undated): Kaggle, *Find Open Datasets and Machine Learning Projects*, www.kaggle.com/datasets (accessed 6 April 2024).

- Kiela et al. 2024: Kiela et al., *Test Scores of AI Systems on Various Capabilities Relative to Human Performance*, Our World in Data (2023, updated on 2 April 2024), ourworldindata.org/grapher/test-scores-ai-capabilities-relative-human-performance (accessed 6 April 2024).
- Kiela et al. 2023: Kiela et al., *Plotting Progress in AI, Contextual AI (2023)*, contextual.ai/plotting-progress-in-ai (accessed 6 April 2024).
- LMSYS 2024: LMSYS, *Chatbot Arena Leaderboard (2024)*, chat.lmsys.org/?leaderboard (accessed 23 April 2024).
- Nvidia 2023: Nvidia Corporation, *Form 8-K, SEC (2023)*, www.sec.gov/Archives/edgar/data/1045810/000104581023000217/nvda-20231017.htm (accessed 22 April 2024).
- OpenAI 2024: OpenAI, *Sora: Creating Video from Text*, OpenAI (2024), openai.com/index/sora (accessed 7 May 2024).
- OpenAI 2023: OpenAI, *Pricing (2023)*, openai.com/pricing (accessed 22 April 2024).
- OpenAI 2022: OpenAI, *DALL•E 2 Preview: Model Training Data*, GitHub, github.com/openai/dalle-2-preview/blob/main/system-card.md (accessed 6 April 2024).
- Papers with Code (undated): *Papers with Code: The Latest in Machine Learning, Papers with Code*, paperswithcode.com/ (accessed 6 April 2024).
- Parliament 2024: Parliament, *Results of Roll-Call Votes 13-03-2024: Artificial Intelligence Act, 1 (2024)*.
- Peleg 2023: Yam Peleg, *GPT-4's Details Are Leaked*, X, Archive.Today (2023), archive.ph/2RQ8X (accessed 6 April 2024).
- Sun 2024: Raymond Sun, *Global AI Regulation Tracker*, TechieRay (2024), www.techieray.com/GlobalAIRegulationTracker (accessed 23 April 2024).
- TechPowerUp (undated): TechPowerUp, www.techpowerup.com (accessed 7 May 2024).
- Tesla 2016: Tesla, *A Tragic Loss*, Tesla Blog (2016), www.tesla.com/blog/tragic-loss (accessed 16 January 2024).
- Vaughan 2024: Thom Vaughan, *February/March 2024 Crawl Archive Now Available, Common Crawl*, Common Crawl Blog (2024), commoncrawl.org/blog/february-march-2024-crawl-archive-now-available (accessed 6 April 2024).