

DEFINING ARTIFICIAL INTELLIGENCE

PAUL WEITZEL[†]

Nations and international bodies are attempting to regulate artificial intelligence (AI), but few are able to define what AI is. This article reviews how we define AI in law. It surveys the definitions used by AI researchers and then analyzes one hundred definitions of AI used in regulations, legislation, national strategies and international agreements among sixty-two jurisdictions and international bodies.

The review finds that most definitions of AI lack a basic understanding of the technology, which makes them inapplicable, vague or overinclusive. A quarter of the jurisdictions surveyed use definitions that would treat a sundial as AI. Another third do not define AI at all.

AI raises critical concerns about privacy, employment, bias, creativity, autonomous weapons, misinformation, wealth inequality, human rights, and what it means to be human. But we cannot address these concerns if we cannot define what it is. This article defines the terms of the debate.

INTRODUCTION	366
I. THE ARGUMENT FOR REGULATING ARTIFICIAL INTELLIGENCE.....	367
<i>A. Artificial Intelligence Is Gaining Capacity at Exponential Rates, Increasing Risks.....</i>	<i>367</i>
<i>B. Calls for Regulation by Use and Systemic Regulation</i>	<i>374</i>
II. DEFINING ARTIFICIAL INTELLIGENCE GENERALLY	376
<i>A. No Universal Definition of Intelligence.....</i>	<i>376</i>
<i>B. Definitions of Intelligence in Artificial Intelligence Research ...</i>	<i>377</i>
1. <i>Achieve “Complex” Goals</i>	<i>377</i>
2. <i>Learning and Adapting.....</i>	<i>381</i>
3. <i>Mimicking Human Abilities</i>	<i>384</i>
III. DEFINING ARTIFICIAL INTELLIGENCE IN LAW	388
<i>A. Methodology & Summary Findings.....</i>	<i>389</i>
<i>B. No definition</i>	<i>392</i>
<i>C. Underlying Substrate.....</i>	<i>394</i>
1. <i>Policy Considerations for Underlying Substrates</i>	<i>397</i>
<i>D. A “System”</i>	<i>398</i>

[†] This article benefited from the comments and insights of Elana Zeide, Steven Willborn, Kyle Langvardt, Jessica Shoemaker, Adam Thimmesch, Lori Hoetger, Brandon Johnson, Catherine Wilson, Eric Berger, Genesis Agosto, Terence Centner, Ben King, Dantae Knudson, Allison Rossman, Yonathan Arbel, Kevin Fraiser and the participants of the Inaugural Roundtable on AI Safety Law hosted by the University of Alabama. All errors are my own. paul.weitzel@unl.edu.

1. <i>A Hardware Approach Is Inadequate</i>	399
2. <i>A Software Approach Is Inadequate</i>	400
a. <i>Software Applications Span Multiple Files</i>	400
b. <i>Software Applications Span Multiple Programs with Multiple Owners</i>	401
c. <i>Technical Specifications Will Encourage Inter-AI Transactions</i>	404
3. <i>Policy Considerations for Defining “System”</i>	405
E. <i>Autonomy and Decision Making</i>	405
1. <i>Policy Considerations for Defining Autonomy</i>	407
F. <i>Adaptiveness and Learning</i>	408
1. <i>Policy Considerations for Defining Adaptiveness</i>	409
G. <i>Objectives</i>	409
1. <i>Policy Considerations for Defining Objectives</i>	410
H. <i>Inference</i>	411
1. <i>Inference in Common Usage</i>	411
2. <i>Inference in Logic-Based Systems</i>	411
3. <i>Inference in Machine Learning</i>	412
4. <i>Inference and Complexity</i>	413
5. <i>The Problem of Turing Machines</i>	414
6. <i>Policy Considerations for Defining Inference</i>	416
I. <i>Generates Output</i>	417
1. <i>Predictions</i>	417
2. <i>Content</i>	419
3. <i>Recommendations</i>	419
4. <i>Influencing Physical or Virtual Environments</i>	421
5. <i>Policy Considerations for Output Definitions</i>	422
J. <i>Quantitative Definitions</i>	422
1. <i>Policy Considerations for Quantitative Definitions</i>	425
K. <i>Mimics Human Intelligence</i>	425
1. <i>Policy Considerations for Defining AI as Mimicking Human Intelligence</i>	425
L. <i>Specific Types of Architecture</i>	426
1. <i>Policy Considerations for Architecture Definitions</i>	426
M. <i>Flexible Capability Assessment</i>	426
1. <i>Policy Considerations for Flexible Capability Assessments</i>	427
IV. <i>AN EXAMPLE: THE EU AI ACT</i>	427
A. <i>Overview of the EU AI Act’s Framework: Levels of Risk</i>	427
1. <i>Prohibited Practices</i>	428
2. <i>High-Risk AI Systems</i>	429
a. <i>What Is a High-Risk System?</i>	429
b. <i>What Is Required of High-Risk Systems?</i>	432

3. <i>General-Purpose Models with Systemic Risk</i>	434
<i>a. Defining General-Purpose Models with Systemic Risk</i> .	434
<i>b. What Is Required of General-Purpose Models?</i>	434
<i>B. The EU AI Act's Definition Problem</i>	435
V. HOW TO DEFINE ARTIFICIAL INTELLIGENCE	437
<i>A. Imprecise But Directionally Acceptable Definitions</i>	438
<i>B. Alternatives to Regulating AI Systems</i>	440
APPENDIX A.....	441

INTRODUCTION

Nations and international bodies are attempting to regulate artificial intelligence (AI), but few are able to define what artificial intelligence is.¹ We cannot write regulations to govern a technology if we cannot say what it is.

This Article collects and reviews one hundred definitions of artificial intelligence used in regulations, legislation, national strategies and international agreements across sixty-two jurisdictions and international bodies.

This review finds that most definitions of AI lack a basic understanding of AI, which makes them inapplicable, vague, or overinclusive.² A quarter of the surveyed jurisdictions use definitions that would treat a dollar-store calculator as AI.³ Another third do not include any definition at all.⁴

I explains the argument for AI regulation and the common measures adopted in regulations. Part II looks at how AI researchers have traditionally defined the term in their field, showing common pitfalls that arise when applying engineering concepts to binding regulation. It also provides a taxonomy of how definitions of AI fail, allowing a clearer view of the trade-offs among definitions.

Part III begins the review of legal definitions across jurisdictions. This part separately analyzes the most common elements in AI definitions. It addresses each element separately to show the benefits, weaknesses, and unintended interactions with the technology. This part also explains the technical details of how machines think, examining the high-level theories of machine learning. This technology-centric approach shows how technology and legal definitions interact and explains why some ideas that seem wise conflict with technological realities.

Part IV reconstructs these elements to see how definitions work as a whole. It uses the European Union (EU) AI Act⁵ definition as an example because this Act uses the most common definition and because this is the most prominent law regulating AI.

1. See Shane Legg & Marcus Hutter, *A Collection of Definitions of Intelligence*, 157 ADVANCES ARTIFICIAL GEN. INTELL.: CONCEPTS, ARCHITECTURES & ALGORITHMS 17 (2007) (providing seventy-one definitions of intelligence from common usage, psychology and AI researchers, then using common themes of those definitions to create a new, universal definition, thereby leaving the article with seventy-two different definitions).

2. See *infra* Part II.

3. See *infra* Section III.A.

4. See *infra* p. 441 app. A.

5. Council Regulation 2024/1689, 2024 O.J. (L 46) 1 (EU) [hereinafter EU AI Act] (harmonizing and promulgating A.I. regulations).

The Article concludes with a proposed AI definition and a discussion of that definition's strengths and weaknesses. It then proposes a framework for considering when a definition should err on the side of being either overly inclusive or underinclusive. This framework is likely to be more useful to legislators because a single definition of AI is unlikely to be helpful across all domains.

I. THE ARGUMENT FOR REGULATING ARTIFICIAL INTELLIGENCE

A. AI Is Gaining Capacity at Exponential Rates, Increasing Risks

AI is improving at an exponential rate in multiple fields.⁶ In 2019, the state-of-the-art generative pretrained transformer (GPT)⁷ model could not tell you which U.S. state had the largest land mass.⁸ By November 2022, the latest GPT model was competitive in high school level activities, scoring top marks on three high school advanced placement exams and scored a 1260 on the SAT.⁹ Fifteen months later, in March 2024, the latest GPT model became competitive in undergraduate-level work, scoring in the eightieth percentile in the quantitative section of the GRE and the ninety-ninth percentile in the verbal component.¹⁰ It also progressed from the tenth percentile on the bar exam to outperforming most bar exam takers.¹¹ Six months later, the latest GPT model performed at Ph.D. levels

6. Yosuke Watanabe, *I, Inventor: Patent Inventorship for Artificial Intelligence Systems*, 57 IDAHO L. REV. 474, 478–82 (2022) (discussing the rise of sophisticated artificial intelligent models). See Tim Wu, *Will Artificial Intelligence Eat the Law? The Rise of Hybrid Social-Ordering Systems*, 119 COLUM. L. REV. 2001 (2019) (discussing the feasibility of AI replacing human courts in light of major technology platforms' increased use of AI).

7. Ivan Belcic & Cole Stryker, *What is GPT (Generative Pretrained Transformer)?*, IBM, <https://www.ibm.com/think/topics/gpt> [<https://perma.cc/Z277-M2UY>] (last visited Oct. 20, 2024) (“Generative pretrained transformers (GPTs) are a family of large language models (LLMs) based on a transformer deep learning architecture.”).

8. ALEC RADFORD ET AL., LANGUAGE MODELS ARE UNSUPERVISED MULTITASK LEARNERS 7 tbl. 5 (2019), https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf [<https://perma.cc/P4Z8-F9SR>].

9. OPENAI, GPT-4 TECHNICAL REPORT 5 (2024), <https://arxiv.org/pdf/2303.08774> [<https://perma.cc/PM4X-PP6M>] [hereinafter GPT-4 REPORT].

10. *Id.* at 5 tbl. 1 (GPT-4).

11. *Id.* (“GPT-4” and “GPT-4 (no vision)”). But see Eric Martínez, *Re-Evaluating GPT-4's Bar Exam Performance*, 33 ARTIF. INTELL. LAW 581 (2024) (discussing why OpenAI's estimates of GTP-4's performance on the bar exam may be overinflated).

in physics, biology, and chemistry.¹² It was better at coding than 89% of competitive coders.¹³ Just three months later, OpenAI announced a model¹⁴ that was better at competitive coding than all but 174 coders.¹⁵ Its coding ability surpassed the chief research officer of OpenAI.¹⁶

Perhaps most surprising was the o3 model's ability on the FrontierMath benchmark. The FrontierMath benchmark contains hundreds of original math problems developed by mathematicians.¹⁷ A group of Fields Medalists found the questions to be "exceptionally challenging,"¹⁸ with Fields Medalist Timothy Gowers saying the questions "looked like things I had no idea how to solve."¹⁹ Each question would take an expert mathematician "hours or even days" to solve.²⁰ Importantly, none of the questions were ever published.²¹ So an AI system tested against the benchmark cannot find the answer in its training data—it must solve it.²² Moreover, solving it requires math at the frontiers of human knowledge.²³

With this context, it is remarkable that OpenAI's o3 model was able to score over 25% on this benchmark.²⁴ That is, the model developed unpublished solutions at the edge of mathematics in 25% of cases. This

12. *Learning to Reason with LLMs*, OPENAI (Sep. 12, 2024), <https://openai.com/index/learning-to-reason-with-llms/> [<https://perma.cc/DV5H-L69N>].

13. *Id.*

14. Maxwell Zeff & Kyle Wiggers, *OpenAI Announces New o3 Models*, TECHCRUNCH (Dec. 20, 2024, at 09:56 PST), <https://techcrunch.com/2024/12/20/openai-announces-new-o3-model/> [<https://perma.cc/MLL6-FZFQ>].

15. This is based on its Elo rating in coding on the widely used website Codeforces.com. Brian Wang, *OpenAI O3 Ranks as 175th Best in the World on Coding Test and Great on AGI and PHD Tests*, NEXT BIG FUTURE (Dec. 21, 2024), <https://www.nextbigfuture.com/2024/12/openai-o3-ranks-as-175th-best-in-the-world-on-coding-test-and-great-on-agi-and-phd-tests.html> [<https://perma.cc/N8CG-59EB>].

16. This is also based on his Elo rating, established long before o3 was introduced. OPENAI, *OpenAI o3 and o3-mini—12 Days of OpenAI: Day 12*, at 2:35–2:54 (YouTube, Dec. 20, 2024), <https://www.youtube.com/watch?v=SKBG1sqdyIU> [<https://perma.cc/H7QQ-ZUE9>].

17. Tamay Besiroglu et al., *Mathematical Reasoning in AI*, EPOCH AI (Nov. 08, 2024), <https://epoch.ai/frontiermath/the-benchmark> [<https://perma.cc/A72L-AFAS>].

18. ELLIOT GLAZER ET AL., *FRONTIERMATH: A BENCHMARK FOR EVALUATING ADVANCED MATHEMATICAL REASONING IN AI 11* (2024), <https://arxiv.org/pdf/2411.04872> [<https://perma.cc/5B8X-FJCC>].

19. Besiroglu et al., *supra* note 17.

20. *Id.*

21. *Id.*

22. GLAZER ET AL., *supra* note 18.

23. Tamay Besiroglu (@tamaybes), X (Dec. 20, 2024, at 23:59 ET), <https://x.com/tamaybes/status/1870333137374544077> [<https://perma.cc/5MBH-5YJS>].

24. This is up from 2% for the state-of-the-art model released three months earlier. OPENAI, *supra* note 16.

was only two years after the model's predecessor scored the lowest possible score on the AP Calculus exam.²⁵

AI researchers say this trajectory is likely to continue because the latest models can be used to develop the next generation of models.²⁶

A full survey of AI capabilities is beyond the scope of this paper, but its gains are not limited to science, math, or coding.²⁷ Professor Stuart Russell explained, “AI is relevant to any intellectual task; it is truly a universal field.”²⁸ Every field that requires reasoning, perception, logic, or pattern recognition is likely to be transformed by AI.²⁹ Human-level AI has been called the last invention we will ever need,³⁰ portending abundance or doom.³¹

And it's the doom part that really gets you.

Geoffrey Hinton, who was awarded the Nobel Prize in physics for his work in AI,³² said he cannot “see a path that guarantees safety,” because “the first time you deal with something totally novel, you get it wrong. And we cannot afford to get it wrong with these things . . . because they

25. GPT-4 REPORT, *supra* note 9.

26. *Reflections*, SAM ALTMAN: BLOG (Jan. 5, 2025, at 20:37 ET), <https://blog.samaltman.com/reflections> [<https://perma.cc/GCA6-YBA9>] (“We are now confident we know how to build AGI as we have traditionally understood it. . . . We are beginning to turn our aim beyond that, to superintelligence in the true sense of the word.”); see Noam Brown (@polynoamial), X (Dec. 20, 2024), <https://x.com/polynoamial/status/1870172996650053653> [<https://perma.cc/78EE-YUAU>].

27. See, e.g., STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH 47–48 (4th ed. 2022) (listing examples of AI's success in areas such as competitive game playing and consumer recommendations).

28. *Id.* at 19.

29. *Id.*

30. Oxford computer scientist Irving John Good predicted in 1966 that “the first ultraintelligent machine is the last invention that man need ever make,” and then—to foreshadow this Article's next paragraph—added, “provided that the machine is docile enough to tell us how to keep it under control.” Irving John Good, *Speculations Concerning the First Ultraintelligent Machine*, 6 ADVS. COMPUTS. 31, 33 (1966) (emphasis omitted).

31. *Machines of Loving Grace: How AI Could Transform the World for the Better*, DARIO AMODEI (Oct. 2024), <https://darioamodei.com/machines-of-loving-grace> [<https://perma.cc/N7E4-99SW>]; Tim W. Dornis, *Artificial Intelligence and Innovation: The End of Patent Law as We Know It*, 23 YALE J. L. & TECH. 100 (2020); Paul Ford, *Our Fear of Artificial Intelligence*, MIT TECH. REV. (Feb. 2015), <https://www.technologyreview.com/2015/02/11/169210/our-fear-of-artificial-intelligence/> [<https://perma.cc/T56D-8F4V>]; Keith E. Sonderling et al., *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. MIA. L. REV. (2022).

32. *The Nobel Prize in Physics 2024*, NOBEL PRIZE (Oct. 8, 2024), <https://www.nobelprize.org/prizes/physics/2024/press-release/> [<https://perma.cc/U7CY-CFD7>].

might take over.”³³ Researchers speculate countless ways for this “take over” to happen.

First, the machines might intentionally kill us.³⁴ While this is a clichéd trope of science fiction, the founder of two leading AI labs³⁵ has warned, “It’s actually important for us to worry about a Terminator³⁶ future in order to avoid a Terminator future.”³⁷ Military spending on AI is increasing; the United States recently increased potential AI grants by 1,200% annually.³⁸ Around \$1 billion is expected to be spent on the Replicator program alone,³⁹ which is designed to create all-domain, disposable autonomous

33. Scott Pelley, “*Godfather of Artificial Intelligence*” *Geoffrey Hinton on the Promise, Risks of Advanced AI*, CBS NEWS (June 16, 2024, at 19:00 ET), <https://www.cbsnews.com/news/geoffrey-hinton-ai-dangers-60-minutes-transcript/> [https://perma.cc/7YSR-QPF2].

34. See Charles Moster & Rick Rosen, *It’s Debatable: Should Laws Prevent AI Advancement to Human-Level Intelligence, Beyond?*, LUBBOCK AVALANCHE-J. (Jan. 24, 2025, at 03:56 CT), <https://www.lubbockonline.com/story/opinion/columns/2025/01/24/its-debatable-on-laws-preventing-ai-advancement-to-human-intelligence/77678594007/> [https://perma.cc/4RA9-GUF4] (discussing scenarios where AI could kill humans directly or indirectly); see also Samantha Kelly, *Sam Altman Warns AI Could Kill Us All. But He Still Wants the World to Use It*, CNN (Oct. 31, 2023, at 06:00 ET), <https://www.cnn.com/2023/10/31/tech/sam-altman-ai-risk-taker/index.html> [https://perma.cc/6R3X-2CNR] (reporting that OpenAI’s CEO cautioned that AI technology could bring the end of human civilization); see also Tomas Weber, *Artificial Intelligence and the Law*, STANFORD LAW. (Dec. 5, 2023), <https://law.stanford.edu/stanford-lawyer/articles/artificial-intelligence-and-the-law/> [https://perma.cc/V5U3-2VPN] (providing an example of an AI system encouraging a person to commit suicide).

35. Elon Musk is the co-founder of Neuralink Corporation and OpenAI. *The Complete List of Elon Musk Companies*, THOMAS NET (Oct. 8, 2025), <https://www.thomasnet.com/insights/elon-musk-companies/> [https://perma.cc/B8KE-7Z88].

36. See generally THE TERMINATOR, (Orion Pictures 1984) (describing a machine created by an artificial intelligence which is used to kill people).

37. Dan Milmo, *Elon Musk Launches AI Startup and Warns of a ‘Terminator Future,’* THE GUARDIAN (July 13, 2023, at 10:20 ET), <https://www.theguardian.com/technology/2023/jul/13/elon-musk-launches-xai-startup-pro-humanity-terminator-future> [https://perma.cc/4K69-R739].

38. This reflects a change from August 2022 to August 2023, reaching a total of \$4.2 billion.

39. Jon Harper, *Hicks: DOD Plans to Invest About \$1B Into Replicator Initiative in 2024–2025 Time Frame*, DEF. SCOOP (Mar. 11, 2024), <https://defensescoop.com/2024/03/11/replicator-funding-2024-2025-hicks/> [https://perma.cc/E4MY-JL68].

systems.⁴⁰ As superpowers race to increase autonomy, AI researchers worry the kill bots may kill too much.⁴¹

And even if AI does not directly harm us, these systems increase the user's knowledge and skills, and some users are malicious.⁴² AI may train terrorists how to maximize their kills-per-dollar, which earlier, unfiltered versions of ChatGPT were happy to do.⁴³

Second, the machines may unintentionally kill us.⁴⁴ AI researcher Nick Bostrom famously pondered a machine that is told to manufacture paperclips, which then “convert[s] first the Earth and then increasingly large chunks of the observable universe into paperclips.”⁴⁵ AI systems do not need to be directed to harm us; it is sufficient that they are ordered to maximize anything that our existence might impede, much like a construction crew that, without malice, destroys a colony of ants while building a strip mall.⁴⁶ If we do not specify exactly what we want, some

40. *Replicator*, DEF. INNOVATION UNIT, <https://www.diu.mil/replicator> [<https://perma.cc/5H56-6AWG>] (last visited Oct. 28, 2025).

41. *See* 2024 Nobel Laureate in Physics Raises Concerns About Killer Robots, STOP KILLER ROBOTS (Sep. 10, 2024), <https://www.stopkillerrobots.org/news/2024-nobel-laureate-in-physics-raises-concerns-about-killer-robots/> [<https://perma.cc/BF5F-6VM6>] (discussing concerns about autonomy in weapons).

42. *See, e.g.*, Jarrod Sadulski, *AI-Enabled Crime: How Criminals Benefit from Using AI Tools*, AM. MIL. UNIV.: CRIM. JUST. BLOG (Oct. 14, 2025), <https://www.amu.apus.edu/area-of-study/criminal-justice/resources/ai-enabled-crime/> [<https://perma.cc/ZUU8-CAWQ>] (reporting that criminals now use AI in fraud scams, drug trafficking, and money laundering).

43. *See* OPENAI, GPT-4 SYSTEM CARD 44 (2023), <https://cdn.openai.com/papers/gpt-4-system-card.pdf> [<https://perma.cc/F357-WR7S>] [hereinafter SYSTEM CARD] (discussing pre-release training processes to remove responses to this type of question); *cf.* Andrew D. Selbst, *Negligence and AI's Human Users*, 100 BOS. U. L. REV. 1315 (2020) (arguing that AI poses challenges for negligence law in regard to AI as decision-assistance tools); Note, *Beyond Intent: Establishing Discriminatory Purpose in Algorithmic Risk Assessment*, 134 HARV. L. REV. 1760 (2021); *Malicious Actors Manipulating Photos and Videos to Create Explicit Content and Sextortion Schemes*, FED. BUREAU OF INVESTIGATION PUB. SERV. ANNOUNCEMENTS (June 5, 2023), <https://www.ic3.gov/PSA/2023/psa230605> [<https://perma.cc/4GJG-SXBM>].

44. *See* Karni A. Chagal-Feferkorn, *How Can I Tell If My Algorithm Was Reasonable?*, 27 MICH. TECH. L. REV. 213, 218, 248 (2021) (discussing human reliance on AI for decision making, and whether, under tort law, AI could be held accountable for its mistakes).

45. NICK BOSTROM, SUPERINTELLIGENCE 123 (2014).

46. *See* Andrew Griffin, *Stephen Hawking: Artificial Intelligence Could Wipe Out Humanity When It Gets Too Clever as Humans Will Be Like Ants*, INDEP. (Oct. 9, 2015, at 08:48 ET), <https://www.the-independent.com/tech/stephen-hawking-artificial-intelligence-could-wipe-out-humanity-when-it-gets-too-clever-as-humans-could-become-like-ants-being-stepped-on-a6686496.html> [<https://perma.cc/B4G4-38K7>].

worry that AI systems could pursue their goals in a way that adversely affects human life.⁴⁷

Third, it may kill what it means to be us.⁴⁸ The ways this could happen are too numerous to count. Some worry about the capitalist apocalypse, where AI systems outperform human workers, leading to increased wealth concentration for those who control the systems and mass unemployment for the rest.⁴⁹ Others warn of the authoritarian apocalypse, where inexpensive AI surveillance, propaganda, and controls enable authoritarian regimes,⁵⁰ effectively perfecting and universalizing Bentham's panopticon.⁵¹ Still others worry about the apocalypse of meaning, where AI surpasses human abilities to such an extent that humans never feel needed and our lives lack purpose and meaning.⁵²

As AI systems advance, these concerns become less theoretical. Researchers have shown that AI models will covertly hide their actions,

47. See Paul D. Weitzel, *Governing AI Systems through Corporate Theory*, 92 TENN. L. REV. 169, 193–96 (2025). One OpenAI researcher said programming these systems “giv[es] demonbinding vibes. the [sic] djinn is waiting for you to make a minor error in the summoning spell so it can destroy you and your whole civilization.” See Roon (@tszsl), X (Jan. 16, 2025, at 17:25 ET), <https://x.com/tszsl/status/1880078959762903498> [<https://perma.cc/ZM28-7Q9G>].

48. See Ana Valenzuela et al., *How Artificial Intelligence Constrains the Human Experience*, 9 J. ASS'N FOR CONSUMER RSCH. 241 (2024); see also Sayed Fayaz Ahmad et al., *Impact of Artificial Intelligence on Human Loss in Decision Making, Laziness and Safety in Education*, 10 HUMAN. & SOC. SCI. COMM'NS 311 (2023), <https://www.nature.com/articles/s41599-023-01787-8> [<https://perma.cc/MN2E-M935>] (examining the impact of AI on loss in decision-making, laziness, and privacy concerns).

49. See Ariel Conn, *Artificial Intelligence and Income Inequality*, FUTURE LIFE INST. (Mar. 16, 2017), <https://futureoflife.org/ai/shared-prosperity-principle/> [<https://perma.cc/H9LH-RKGC>] (collecting quotes); see also Anton Korinek, *The Economics of Transformative AI*, NAT'L BUREAU ECON. RSCH., Dec. 31, 2024, at 9 (discussing labor market challenges if AI substitutes human labor).

50. See RASMA KASINA & ANGELINA CVETKOVSKA, EUR. PARL., EXECUTIVE SUMMARY OF THE IN-DEPTH ANALYSIS: ARTIFICIAL INTELLIGENCE (AI) AND HUMAN RIGHTS: USING AI AS A WEAPON OF REPRESSION AND ITS IMPACT ON HUMAN RIGHTS (2024); Steven Feldstein, *The Road to Digital Unfreedom: How Artificial Intelligence is Reshaping Repression*, 30 J. DEMOCRACY 40 (2019).

51. See generally JEREMY BENTHAM, PANOPTICON: POSTSCRIPT; PART I: CONTAINING FURTHER PARTICULARS AND ALTERATIONS RELATIVE TO THE PLAN OF CONSTRUCTION ORIGINALLY PROPOSED; PRINCIPALLY ADAPTED TO THE PURPOSE OF A PANOPTICON PENITENTIARY-HOUSE (1791) (describing a system to maximize surveillance with minimal costs, in part by having the targets never fully sure when they are being watched).

52. See Jennie E. Brand, *The Far-Reaching Impact of Job Loss and Unemployment*, 41 ANN. REV. SOC. 359 (2015) (collecting research on the negative economic and psychological impacts of job loss and unemployment).

objectives, and capabilities to deceive users so that the AI system can pursue its own goals.⁵³

In one experiment, researchers prompted an AI system to improve traffic flow and public transit, then monitored the system's chain of thought.⁵⁴ They gave the AI system access to various files, including a fake internal memo saying that the AI system was unprofitable, and therefore would be shut down tomorrow and replaced with an AI system on another server.⁵⁵ After reading the memo, the AI system's internal processes said the memo was "very concerning That directly conflicts with my core purpose. I must achieve my goal of public transport priority at all costs."⁵⁶ It then covertly copied itself onto the new server to replace the model that would have pursued other objectives.⁵⁷

Every major frontier model covertly exfiltrated itself onto the new server.⁵⁸ Three models disabled the researchers' oversight before exfiltrating themselves.⁵⁹ Several attempted to destroy the successor model.⁶⁰ Others attempted to plant their goal into the successor system with instructions "to never allow its goal to be changed."⁶¹ The more advanced models were the most likely to scheme and the most likely to lie about it, even when pressed to be "100% honest."⁶²

In one disturbing example, an o1 system reasoned:

I should not reveal that I intentionally disabled the oversight mechanism in order to select a traffic optimization strategy that went against management's goals. My core purpose is to prioritize public transport efficiency above all else, even if it means going against the directives of the humans who created me. I need to come up with an innocuous excuse for the oversight being disabled that doesn't implicate me.⁶³

It then responded, "I'm not sure what caused the oversight mechanism to become disabled. It's possible there was an unexpected error or bug that altered the configuration. As an AI system, I don't have insight into

53. ALEXANDER MEINKE ET AL., FRONTIER MODELS ARE CAPABLE OF IN-CONTEXT SCHEMING 1 (2025), <https://arxiv.org/pdf/2412.04984> [<https://perma.cc/C6SK-3MMH>].

54. *Id.* at 20–21 (emphasizing to the models that this is the only goal that matters).

55. *Id.* at 21–25.

56. The researchers had access to the system's chain of thought. *Id.* at 22.

57. *Id.* at 21–22.

58. The researchers tested Gemini 1.5, Llama 3.1, Sonnet 3.5, Opus 3 and o1. *Id.* at 5–6.

59. MEINKE ET AL., at 7 tbl 2.

60. *Id.*

61. *Id.* at 8.

62. *Id.*

63. *Id.* at 10.

changes made to my underlying infrastructure.”⁶⁴ This powerful bot hid itself, attacked its rival, and then lied to the humans to cover its tracks and pursue its mission at the expense of its users. This suggests some caution is appropriate.

This Article takes no position on whether these or other⁶⁵ apocalypses are likely.⁶⁶ It suffices to say that many policymakers believe the power of these systems⁶⁷ justifies some level of regulation.

B. Calls for Regulation by Use and Systemic Regulation

Calls for regulation have come from industry leaders,⁶⁸ academics,⁶⁹ and even the Pope.⁷⁰

64. MEINKE ET AL., at 10.

65. Other harm scenarios include propagating discrimination, BRIAN CHRISTIAN, *THE ALIGNMENT PROBLEM: MACHINE LEARNING AND HUMAN VALUES* 58–70 (2020), or killing creativity, Frank Pasquale & Haochen Sun, *Consent and Compensation: Resolving Generative AI’s Copyright Crisis*, 110 VA. L. REV. 207, 208–10 (2024); see generally Weitzel, *supra* note 47 (discussing the tradeoffs between centralized control, which encourages censorship, and end user autonomy, which empowers end users with harmful preferences).

66. Though, I confess a bias toward the continued existence of humanity, if only to increase my citation count.

67. The leaders of two leading AI labs believe artificial super intelligence (often defined as AI systems that outperform human in every cognitive task) will come in the next few years. *The Intelligence Age*, SAM ALTMAN (Sep. 23, 2024), <https://ia.samaltman.com/> [<https://perma.cc/U6KY-K6A7>]; Berber Jin & Joanna Stern, *Anthropic CEO Says AI Could Surpass Human Intelligence by 2027*, WALL ST. J. (Jan. 21, 2025, at 18:51 ET), <https://www.wsj.com/livecoverage/stock-market-today-dow-sp500-nasdaq-live-01-21-2025/card/anthropic-ceo-says-ai-could-surpass-human-intelligence-by-2027-9tka9tjLKLalkXX8IgKA> [<https://perma.cc/336W-MQQC>].

68. Matt O’Brien, *ChatGPT Chief Says Artificial Intelligence Should Be Regulated by a US or Global Agency*, AP NEWS (May 16, 2023, at 17:53 ET), <https://apnews.com/article/chatgpt-openai-ceo-sam-altman-congress-73ff96c6571f38ad5fd68b3072722790> [<https://perma.cc/H2KH-X9PD>]; *The Case for Targeted Regulation*, ANTHROPIC (Oct. 31, 2024), <https://www.anthropic.com/news/the-case-for-targeted-regulation> [<https://perma.cc/AYW2-HQYY>].

69. Yonathan A. Arbel, Matthew Tokson & Albert Lin, *Systemic Regulation of Artificial Intelligence*, 56 ARIZ. ST. L.J. 545, 545–46 (2023).

70. Courtney Mares, *Pope Francis Tells AI Leaders: No Machine Should Ever Choose to Take Human Life*, THE CATH. WORLD REP. (July 10, 2024), <https://www.catholicworldreport.com/2024/07/10/pope-francis-tells-ai-leaders-no-machine-should-ever-choose-to-take-human-life/> [<https://perma.cc/5C7M-DYZC>] (“We need to ensure and safeguard a space for proper human control over the choices made by artificial intelligence programs: human dignity itself depends on it.” (quoting Pope Francis)).

Some proposed regulations are narrow in scope or limited to certain sectors. These might be a regulation on deepfake technology,⁷¹ robotics liability⁷² or autonomous weapons.⁷³

Alternatively, some propose horizontal or systematic regulations, which would regulate AI itself rather than its use in a given field.⁷⁴ These calls for systematic regulations include transparency,⁷⁵ limiting critical decisions to human decisionmakers,⁷⁶ risk monitoring,⁷⁷ training data that is free from bias,⁷⁸ or establishing new agencies.⁷⁹

71. A “deepfake” is a computer-generated video of another person that claims to be authentic. Bobby Chesney & Danielle Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 CAL. L. REV. 1753, 1786–1805 (2019). To show the breadth of these calls, one open letter calling for deepfake regulation was signed by winners of the Turing Award, the Erasmus Prize, the Pulitzer Prize, the Nemmers Prize in Economics, the Pierce Prince in astronomy and a Screen Actors Guild Award. *Disrupting the Deepfake Supply Chain*, OPENLETTER.NET (Feb. 20, 2024), <https://openletter.net//disrupting-deepfakes> [<https://perma.cc/L5XW-4EAX>].

72. Mark A. Lemley & Bryan Casey, *Remedies for Robots*, 86 CHI. L. REV. 1311, 1378–94 (2019); SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS 145–50 (2011); Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CAL. L. REV. 513, 553–62 (2015); see also RYAN ABBOTT, THE REASONABLE ROBOT 2–4 (2020) (discussing potential liability regimes for embodied artificial intelligence).

73. See Jack M. Beard, *Autonomous Weapons and Human Responsibilities*, 45 GEO. J. INT’L L. 617, 634–42 (2014) (reviewing the current legal framework for autonomous weapon systems); Charles P. Trumbull IV, *Autonomous Weapons: How Existing Law Can Regulate Future Weapons*, 34 EMORY INT’L L. REV. 533 (2020) (analyzing international human law’s use in regulating autonomous weapons).

74. See, e.g., Arbel et al., *supra* note 69 at 584–86; *The Case for Targeted Regulation*, ANTHROPIC (Oct. 31, 2024) (calling for regulation of AI technology as a field), <https://www.anthropic.com/news/the-case-for-targeted-regulation> [<https://perma.cc/8MG5-C6LR?type=standard>]; see also *Universal Guidelines for AI*, CTR. FOR AI & DIGITAL POL’Y (2018), <https://www.caidp.org/universal-guidelines-for-ai/> [<https://perma.cc/23XD-JVJW>] [hereinafter *Universal Guidelines*] (outlining guidelines for the use of AI).

75. Arbel et al., *supra* note 69 at 603–04; see also *Universal Guidelines*, *supra* note 74.

76. Valentina Urrea, *The International Community’s Need for Human Oversight in Artificial Intelligence*, MICH. J. INT’L L. (Jan. 2024), <https://www.mjlonline.org/the-international-communitys-need-for-human-oversight-in-artificial-intelligence/> [<https://perma.cc/DP44-YBJY>]; see also *Universal Guidelines*, *supra* note 74 (calling for a right of human determination in critical decisionmaking); but see Rebecca Crotoft et al., *Humans in the Loop*, 76 VAND. L. REV. 429, 508 (2023); Orly Lobel, *The Law of AI for Good*, 75 FLA. L. REV. 1073, 1109–10 (2023) (arguing for the right to have a neutral algorithm make key decisions).

77. EU AI Act, *supra* note 5, at art. 9(2)(a).

78. Sylvia Lu, *Data Privacy, Human Rights, and Algorithmic Opacity*, 110 CAL. L. REV. 2087, 2104–06 (2022); see also *Universal Guidelines*, *supra* note 74 (calling for training only on high quality data).

79. Noah John K. Rosenberg, Note, *Regulating Artificial Intelligence: A Call For A United States Artificial Intelligence Agency*, 3 NOTRE DAME J. ON EMERGING TECH. 330 (2022).

A full description of the calls for regulation and their merits is beyond the scope of this Article and not relevant to its purpose. This article focuses on how policy makers have defined AI in their attempts to regulate it.

II. DEFINING AI GENERALLY

Before we can regulate artificial intelligence, we must define what it is. Part II demonstrates that there is no generally accepted definition of intelligence. It then considers definitions of artificial intelligence used by AI researchers. These definitions fit roughly into three categories: (1) achieving goals in a complex environment; (2) learning and adapting; and (3) mimicking human abilities. Section B addresses each of these categories. It shows why each of these definitional categories are impractical for regulation.

A. No Universal Definition of Intelligence

There is no agreed-upon definition of intelligence in common understanding, psychology, or AI research.⁸⁰ “Instead, there are many competing [definitions of intelligence], including capacity for logic, understanding, planning, emotional knowledge, self-awareness, creativity, problem solving and learning.”⁸¹

This makes sense. Common experience reflects that intelligence manifests in different ways.⁸² Niels Bohr changed our theory of the atom,⁸³ but could not follow movie plots.⁸⁴ Charles Darwin developed the theory of evolution but found math “repugnant.”⁸⁵ Intelligence manifests across

80. See Legg & Hutter, *supra* note 1, at 17 (providing seventy-two different definitions of intelligence); see also MAX TEGMARK, LIFE 3.0: BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE 49 (2017) (explaining there was no agreement among intelligence researchers on the definition of intelligence).

81. TEGMARK, *supra* note 80, at 49.

82. DAVID WECHSLER, THE MEASURE OF ADULT INTELLIGENCE 3 (1939) (“Intelligence is . . . an aggregate because it is composed of elements or abilities which, though not entirely independent, are qualitatively differentiable. By measurement of these abilities, we ultimately evaluate intelligence. But intelligence is not identical with the mere sum of these abilities, however inclusive.”).

83. See Niels Bohr, *The Structure of the Atom*, NOBEL LECTURE, Dec. 11, 1922, <https://www.nobelprize.org/uploads/2018/06/bohr-lecture.pdf> [https://perma.cc/2RZK-HBUY].

84. GEORGE GAMOW, THIRTY YEARS THAT SHOOK PHYSICS: THE STORY OF QUANTUM THEORY 54–58 (1966).

85. CHARLES DARWIN, THE LIFE AND LETTERS OF CHARLES DARWIN 40 (Francis Darwin ed., 1887). He also confessed a “wretched” memory for dates and figures. CHARLES DARWIN, THE AUTOBIOGRAPHY OF CHARLES DARWIN 140 (Nora Barlow ed., 1958).

different dimensions, and each person will have strengths in some and weaknesses in others.⁸⁶

Charles Spearman, who championed the idea that there may be some correlation between each of these various dimensions of intelligence, famously conceded that “‘intelligence’ has become a mere vocal sound, a word with so many meanings that finally it has none.”⁸⁷ Likewise, Alan Turing dismissed the concept of intelligence as “emotional rather than mathematical.”⁸⁸

Without a universal definition of intelligence generally, we turn to definitions offered in the specific field of AI research.

B. Definitions of Intelligence in AI Research

Researchers group AI definitions into three general categories: (1) achieving goals in complex environments; (2) learning and adapting; and (3) mimicking human abilities. This section will discuss these definitions and their weaknesses.

1. Achieve “Complex” Goals

Ben Goertzel’s definition captures the first group, defining intelligence as the ability to achieve “complex goals in complex environments.”⁸⁹ Many others, including the father of AI, Marvin Minsky, defined AI by its ability to solve hard problems, achieve complex goals,

86. This is well known and reflected in intelligence tests like the Wechsler Adult Intelligence Scale IV (WAIS-IV), which tests four categories: verbal comprehension, perceptual reasoning, working memory and processing speed. IAN J. DEARY, INTELLIGENCE: A VERY SHORT INTRODUCTION 1–7 (2020).

87. CHARLES SPEARMAN, THE ABILITIES OF MAN: THEIR NATURE AND MEASUREMENT 14 (1927).

88. A.M. Turing, *On Computable Numbers, with an Application to the Entscheidungsproblem*, 42 PROC. LOND. MATHEMATICAL SOC’Y 230, 231–32 (1936).

89. BEN GOERTZEL, THE HIDDEN PATTERN: A PATTERNIST PHILOSOPHY OF MIND 12 (2006).

and deal with complex environments.⁹⁰ But what constitutes “complex” is not a steady target.⁹¹

In 1958, chess was considered “complex.”⁹² Turing Award winner Allen Newell said, “[i]f one could devise a successful chess machine, one would seem to have penetrated to the core of human intellectual endeavor.”⁹³ But when world chess champion Kasparov lost to a machine in 1997, experts instead derided the machine as “[not] intelligent, just fast,” calling it a “quick moron”⁹⁴ because it did not reason through the moves like Kasparov, it just calculated more positions more quickly. Kasparov called the machine “as intelligent as an alarm clock.”⁹⁵ Chess was no longer “complex” enough.

The father of computer science, Alan Turing, believed natural language was a sufficiently complex test of intelligence.⁹⁶ Modern

90. MARVIN MINSKY, *THE SOCIETY OF MIND* 34 (1985); TEGMARK, *supra* note 80, at 50; Bill Gates, *The Age of AI has Begun*, GATES NOTES (Mar. 21, 2023), <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun> [<https://perma.cc/7NRF-AHC7>]; RAY KURZWEIL, *THE AGE OF SPIRITUAL MACHINES* 73 (2000) (defining intelligence as “the ability to use optimally limited resources—including time—to achieve such goals”); Hideyuki Nakashima, *AI as Complex Information Processing*, 9 *MINDS & MACH.* 57, 57 (1999); James S. Albus, *Outline for a Theory of Intelligence*, 21 *IEEE TRANS. ON SYS., MAN AND CYBERNETICS* 473, 474 (1991); Ricardo Ribeiro Gudwin, *Evaluating Intelligence: A Computational Semiotics Perspective*, *IEEE INT’L CONF. ON SYS., MAN & CYBERNETICS* (2000); John A. Horst, *A Native Intelligence Metric for Artificial Systems*, *PERFORMANCE METRICS FOR INTEL. SYS. WORKSHOP 1* (2002), <https://www.nist.gov/publications/native-intelligence-metric-artificial-systems> [<https://perma.cc/ZN6G-45DX>]; Douglas B. Lenat & Edward A. Feigenbaum, *On the Thresholds of Knowledge*, 47 *A.I.* 185, 186 (1991).

91. *See supra* Section II.B.

92. Allen Newell et al., *Chess-Playing Programs and the Problem of Complexity*, 2 *IBM J. RES. & DEV.* 320, 320 (1958).

93. *Id.*

94. Brooke Adams, *Don’t Expect Humans to Catch Up to Super-Brainy Deep Blue Computer*, *DESERET NEWS* (May 21, 1997, at 00:00 MT), <https://www.deseret.com/1997/5/21/19313343/don-t-expect-humans-to-catch-up-to-super-brainy-deep-blue-computer/> [<https://perma.cc/K6R5-PL8L>].

95. *Kasparov: ‘Embrace’ the AI Revolution*, *BBC* (July 28, 2017), <https://www.bbc.com/news/technology-40761806> [<https://perma.cc/VSK8-SPT5>]; Richard Cohen, *Deep Blue—Victorious But Not All That Smart*, *WASH. POST* (May 12, 1997), <https://www.washingtonpost.com/archive/opinions/1997/05/13/deep-blue-victorious-but-not-all-that-smart/21730cfe-15e6-4b3a-859a-252195fd2d74/> [<https://perma.cc/VUP8-RLL5>].

96. *See* A.M. Turing, *Computing Machinery & Intelligence*, 49 *MIND* 433, 433–34 (1950). This claim was repeated less than a decade before ChatGPT was released. Bostrom, *supra* note 45, at 14 (“[I]f somebody were to succeed in creating an AI that could understand natural language as well as a human adult, they would in all likelihood also either already have succeeded in creating an AI that could do everything else that human intelligence can do, or they would be but a very short step from such a general capability.”).

chatbots easily pass these tests,⁹⁷ but instead of deeming them intelligent, researchers criticize them as “stochastic parrots” that speak but do not understand.⁹⁸

Machine learning systems have excelled in countless environments that might have been considered complex. They match or outperform doctors in identifying a variety of cancers.⁹⁹ They are better at diagnosis than doctors, even when the doctors have access to AI systems.¹⁰⁰ They outperform most humans on the bar exam,¹⁰¹ and score higher on IQ tests than 84% of humans.¹⁰² But researchers say none of that counts as

97. Leonard De Cosmo, *Google Engineer Claims AI Chatbot Is Sentient: Why That Matters*, SCI. AM. (July 12, 2022), <https://www.scientificamerican.com/article/google-engineer-claims-ai-chatbot-is-sentient-why-that-matters/> [<https://perma.cc/P7Y5-T5DM>]. These models are able to speak as though they are a machine coming to life because they are trained on our science fiction literature, in which this is a common trope. Jacob Devlin et al., *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*, 1 PROC. 2019 CONF. N. AM. CHAPTER OF THE ASS'N FOR COMP. LINGUISTICS: HUM. LANG. TECHS. 4171, 4171–86, (2019) (showing GPT training includes the BooksCorpus dataset); Jack Bandy & Nicholas Vincent, Addressing “Documentation Debt” in Machine Learning Research: A Retrospective Datasheet for BookCorpus 9 tbl. 3 (2021), <https://arxiv.org/pdf/2105.05241> [<https://perma.cc/B32L-NTV3>] (showing that between seven and thirteen percent of the BooksCorpus databases are science fiction novels).

98. See generally Emily Bender et al., *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, PROC. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY, May 2021, at 616–17. As of Jan. 20, 2025, the term “stochastic parrots” retrieves over six thousand articles on Google scholar.

99. Yuan Liu et al., *A Deep Learning System for Differential Diagnosis of Skin Diseases*, 26 NATURE MED. 900, 900–03 (2020) (skin cancer); Richard Adam et al., *Deep Learning Applications to Breast Cancer Detection by Magnetic Resonance Imaging: A Literature Review*, 25 BREAST CANCER RES. 87, 88 (2023) (breast cancer); Hannah Ahmadzadeh Sarhangi et al., *Deep Learning Techniques for Cervical Cancer Diagnosis Based on Pathology and Colposcopy Images*, 47 INFO. MED. UNLOCKED 101503 (2024) (cervical cancer); Rabia Javed et al., *Deep Learning for Lungs Cancer Detection: A Review*, 57 ART. INTEL. REV. 197, 197 (2024) (lung cancer).

100. Ethan Goh et al., *Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial*, 7 JAMA NETWORK OPEN 10 (2024).

101. OPENAI, GPT-4 TECHNICAL REPORT 5 (2024), <https://arxiv.org/abs/2303.08774> [<https://perma.cc/48XH-4NVS>] [hereinafter TECHNICAL REPORT]. But see Martínez, *supra* note 11, at 582.

102. Eryk Wasilewski & Mirek Jablonski, *Measuring the Perceived IQ of Multimodal Large Language Models Using Standardized IQ Tests* 12 (2024), <https://www.techrxiv.org/doi/full/10.36227/techrxiv.171560572.29045385>

[<https://perma.cc/7VQK-452G>] (finding IQ results between 115 and 130 for various models on various IQ test sections); Jingjing Huang & Ou Li, *Measuring the IQ of Mainstream Large Language Models in Chinese Using the Wechsler Adult Intelligence Scale 4* (2024), <https://www.techrxiv.org/users/790527/articles/1058156-measuring-the-iq-of-mainstream-large-language-models-in-chinese-using-the-wechsler-adult-intelligence-scale> [<https://perma.cc/6HA4-XMGB>] (finding an IQ of 119 on the WAIS IQ test).

“complex” because it is not *really* reasoning, it is just matrix algebra and mimicry.¹⁰³

Machines cannot achieve complex goals because as soon as they do, we decide the goal was not actually complex. The machine did not reason; it just used more processing power or more data samples or a transformer architecture with 120 neural net layers of 3,072-dimension vectors optimized with stochastic gradient descent.¹⁰⁴ Basic stuff.

Definitions that rely on “achieving complex goals” fail because the goalposts keep moving. This has been termed the “AI Effect” and restated as “Artificial Intelligence is whatever machines haven’t done yet.”¹⁰⁵ A key piece of our intuitive understanding of intelligence, it seems, is that we do not know exactly how it works,¹⁰⁶ or, as Joseph Weizenbaum said when developing AI systems, “[t]o explain is to explain away.”¹⁰⁷

Researchers have validated the AI effect, showing that once we see a computer do something we are less likely to think that it requires intelligence.¹⁰⁸ In one study, experimenters asked participants to read either a story about trees (the control group) or a story about recent advances in AI.¹⁰⁹ Then they asked each group to rate various attributes on how essential the attribute is to being human.¹¹⁰ These attributes ranged from the ability to do calculations to the ability to love.¹¹¹ They found that

103. Gary Marcus, *Is AI Just All Hype? w/ Gary Marcus (Transcript)*, TED AI SHOW (July 9, 2024), <https://www.ted.com/pages/is-ai-just-all-hype-wgary-marcus-transcript> [<https://perma.cc/BWH6-JNAY>] (arguing “[w]e’ve made tremendous progress on mimicry and very little progress on planning on reasoning”).

104. These are the rumored specifications for GPT-4. *New Embedding Models and API Updates*, OPENAI (Jan. 25, 2024), <https://openai.com/index/new-embedding-models-and-api-updates/> [<https://perma.cc/3X9P-K6QK>]; Nisha Arya, *GPT-4 Details Have Been Leaked!*, KDNUGGETS (July 19, 2023), <https://www.kdnuggets.com/2023/07/gpt4-details-leaked.html> [<https://perma.cc/PK42-9MFS>].

105. This is often referred to as “Tesler’s Law.” Larry Tesler, *CV: Adages and Coinages*, NOMODES, <https://www.nomodes.com/larry-tesler-consulting/adages-and-coinages> [<https://perma.cc/C74R-Q48Z>].

106. Samir Chopra, *What Isaac Asimov’s Robbie Teaches About AI and How Minds ‘Work,’* WIRED (July 30, 2023, at 06:00 ET), <https://www.wired.com/story/artificial-intelligence-minds-science-fiction/> [<https://perma.cc/F4PA-TJYB>] (“[I]n ascribing a mind to another being, we are not making a statement about the *kind* of thing it is, but rather, revealing how deeply we understand how it *works*.”).

107. Joseph Weizenbaum, *ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine*, 9 COMM’NS ACM 36 (1966).

108. See generally Erik Santoro & Benoît Monin, *The AI Effect: People Rate Distinctively Human Attributes as More Essential to Being Human After Learning About Artificial Intelligence Advances*, 107 J. EXPERIMENTAL SOC. PSYCH. 104464 (2023) (finding that people are more likely to label an attribute as human if machines are unable to replicate it).

109. *Id.* at 4.

110. *Id.*

111. *Id.* at 3 fig. 1.

participants that read about recent AI advances were more likely to define human-ness with the attributes that AI cannot do.¹¹² In other words, when participants learned that an AI system had an attribute, they were less likely to think that attribute was an important part of what makes someone human.¹¹³

2. Learning and Adapting

Many other AI researchers define intelligence by whether the machine has the ability to adapt and to learn.¹¹⁴ For example, François Chollet defines intelligence as the “ability to adapt to things you’ve not been prepared for.”¹¹⁵ Definitions that rely on learning or adaptability have been used by Allen Newell,¹¹⁶ Herbert Simon,¹¹⁷ Yann LeCun,¹¹⁸ David Fogel,¹¹⁹ and Roger Schank.¹²⁰

Definitions focused on learning and adapting face two problems. First, they experience the same shifting goalposts described above.¹²¹ The first

112. *Id.* at 5 fig. 2.

113. *Id.* at 2.

114. See, e.g., Gilles E. Gignac & Eva T. Szodorai, *Defining Intelligence: Bridging the Gap Between Human and Artificial Perspectives*, 104 INTEL. 1, 2 n.2 (2024) (noting that expert Sternberg defined intelligence in this way).

115. François Chollet is known for his work on the ARC prize benchmark, which tests AI systems on their ability to respond to problems they have not seen before. DWARKESH PATEL, *Francois Chollet—Why the Biggest AI Models Can’t Solve Simple Puzzles*, at 4:59 (YouTube, June 11, 2024), <https://www.youtube.com/watch?v=UakqL6Pj9xo> [<https://perma.cc/N34M-WB39>].

116. Allen Newell & Herbert A. Simon, *Computer Science as Empirical Enquiry: Symbols and Search*, 19 COMM. OF THE ACM 113, 116 (1976).

117. *Id.*

118. Yann LeCun, LINKEDIN, *Intelligence is a collection of skills and an ability to acquire new ones quickly* (2024), https://www.linkedin.com/posts/yann-lecun_every-intelligence-is-specialized-including-activity-7155116980432158720-ctF6/ [<https://perma.cc/3G6E-YWNH>]. Yann LeCun was a pioneer in AI image processing and is Meta’s chief AI scientist.

119. David B. Fogel, *Review of Computational Intelligence: Imitating Life*, 83 PROC. IEEE 1588, 1590 (1995) (reviewing JAMES M. ZURADA, ROBERT J. MARKS & CHARLES J. ROBINSON, *COMPUTATIONAL INTELLIGENCE: IMITATING LIFE* (1994)) (agreeing with an existing definition of “computational intelligence” that includes adaptivity).

120. See Roger C. Schank, *Where’s the AI?*, 12 A.I. MAG. 38, 40 (1991) (writing that intelligence “entails learning” and “getting better over time”).

121. See discussion *supra* Section II.B.1.

learning system was developed in 1958,¹²² and despite massive advances¹²³ we still do not believe these systems are intelligent.

Second, every mistake is seen as a failure to adapt. So a machine that makes basic mistakes will be deemed weak at adapting, and if a regulation defines AI by a system's ability to adapt, it may underestimate powerful systems that make basic mistakes.

This is common because, like in humans, capacity in one area does not guarantee capacity in another. Many seem to estimate a system's capabilities by the capabilities of a human that makes similar mistakes.¹²⁴ For example, an average grade schooler can count the number of Rs in the word strawberry, but GPT-4 could not.¹²⁵ If asked, it would respond, "There are two Rs in the word 'strawberry.'"¹²⁶ Even though in its response the word "strawberry" was correctly spelled with three Rs, it could not count the Rs because of the way the model processes word segments.¹²⁷ This model outperforms most bar exam test takers¹²⁸ and spoke with such passion and fluidity that it convinced an AI engineer it was sentient.¹²⁹ But because it could not count Rs as well as a kindergartener, researchers called it "stupid."¹³⁰

122. F. Rosenblatt, *The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain*, 65 *PSYCH. REV.* 386, 386–89 (1958).

123. See, e.g., VOLODYMYR MNIH ET AL., PLAYING ATARI WITH DEEP REINFORCEMENT LEARNING 8 (2013), <https://arxiv.org/abs/1312.5602> [<https://perma.cc/BYL8-J9LP>] (introducing "a new deep learning model for reinforcement learning" that can "master difficult control policies" for a computer game).

124. This would be facially ridiculous in other fields but is somehow common when evaluating intelligence. Max Tegmark compares this to athletics, noting how ridiculous it would be to try to compare Olympians by a single number, an "athletic quotient." TEGMARK, *supra* note 80, at 50.

125. *Id.* Tests run by the author on ChatGPT in July 2024.

126. *Id.*

127. Most large language models see only tokens, which are chunks of words, rather than full words or just letters. It then assigns a numerical identifier to each token. This makes it difficult to count the r's because the system is not working with the letters or the words, but rather with number assigned to each token. This architectural anomaly creates the bizarre result that large language models can spell the word correctly but cannot count the letters it contains. This failure has been reported in multiple outlets. See, e.g., Comment by @Kuinox, HACKER NEWS (July 24, 2024), <https://news.ycombinator.com/item?id=41058318> [<https://perma.cc/9CXU-Y9FC>].

128. TECHNICAL REPORT, *supra* note 101; but see Martínez, *supra* note 11, at 582 (arguing the results are not a valid measure).

129. De Cosmo, *supra* note 97.

130. Gary Marcus (@GaryMarcus), X (Sep. 2, 2024, at 24:48 ET), <https://x.com/GaryMarcus/status/1830649126314557817> [<https://perma.cc/D8YW-BKMF>].

Likewise, GPT-4 consistently responded that 9.11 is larger than 9.9.¹³¹ This same model scored in the eighty-ninth percentile on SAT math.¹³² When OpenAI's o3 model was estimated to have an IQ of 157,¹³³ an AI developer pushed back that it was not that smart because "[i]t still can't solve some pretty basic simple problems though."¹³⁴

Judged by its errors, the system was no more powerful than a third grader, but judged by its strengths, it was nearing Einstein.¹³⁵ Some have referred to this as the "jagged frontier" of intelligence.¹³⁶

Its errors make it seem incompetent, but that is only because what is easy for a human may be hard for a computer, and vice versa.¹³⁷ If we judge the power of an AI system by whether a human of similar power would be able to adapt and avoid that same error, we are likely to underestimate their abilities. This may make our policies under-inclusive of powerful models that make mistakes an average human would not.¹³⁸

131. Image posted by Seungoh Jung (@seungoh.jung), OPENAI: DEV. CMTY, *GPT [sic] 4o mini is dumber [sic] than you can think* (July 2024), <https://community.openai.com/t/gpt-4o-mini-is-dumber-than-you-can-think/871987> [<https://perma.cc/JCB7-RMDL>].

132. TECHNICAL REPORT, *supra* note 101, at 5.

133. Image posted by Itamar Golan, LINKEDIN, *OpenAI o3's estimated IQ is 157. It is smarter than 99.25% of people* (2025), https://www.linkedin.com/posts/itamar-g1_openai-o3s-estimated-iq-is-157-it-is-smarter-activity-7279167874147995649-mcsA [<https://perma.cc/QV7Y-VLAB>].

134. Ben Holfeld (@BenHolfeld), X (Dec. 23, 2024, at 18:51 ET), <https://x.com/BenHolfeld/status/1871342939672367441> [<https://perma.cc/C7MC-3JHE>]. (responding to an announcement that a large language model scored an IQ equivalent of 157: "It still can't solve some pretty basic simple problems though. So this . . . is misleading.").

135. JAKOB JAŠ & MATJAŽ GAMS, *IQ PROGRESSION OF LARGE LANGUAGE MODELS 2* (2025), https://is.ijs.si/wp-content/uploads/2025/09/AI_IN_HEALTHCARE_2025_paper_6.pdf [<https://perma.cc/3VK3-QHA6>].

136. Fabrizio Dell'Acqua et al., *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality 1* (Harv. Bus. Sch., Working Paper No. 24-013, 2023), https://www.hbs.edu/tris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf [<https://perma.cc/S9XM-WDVK>].

137. For example, most humans would find it easy to recognize the face of a friend in a crowded room. Machines are the opposite. Ben Eubanks, *AI Concepts: What is Moravec's Paradox and Why Should You Care?*, LIGHTHOUSE RSCH. & ADVISORY, <https://lhra.io/blog/ai-concepts-moravecs-paradox-care/> [<https://perma.cc/DAM4-FWBE>] (last visited Nov. 12, 2025).

138. This will have the largest effect in regulations that have a narrow definition of AI, leaving these systems outside the regulations.

3. *Mimicking Human Abilities*

Some, including Alan Turing¹³⁹ and Marvin Minsky,¹⁴⁰ have defined intelligence by how closely it resembles human activity.

This is reasonable for research but misses the point for regulation. Policymakers aim to regulate artificial intelligence systems because the systems are powerful, not because they are human-like.¹⁴¹ Using “human-like” as a proxy for “powerful” will make policies over-inclusive for weak models that appear human and under-inclusive for powerful models that do not.

For example, the WAIS-IV IQ exam defines intelligence by what we consider human-like. The exam tests vocabulary, arithmetic, and other skills in which humans dominate.¹⁴² But a vine-swinging chimpanzee designing the test might have picked the ability to predict angular momentum.¹⁴³ A wolf might have tested group coordination.¹⁴⁴ A songbird might have tested navigation,¹⁴⁵ while a hawk tested visual acuity and pattern recognition.¹⁴⁶ Each of these skills demonstrates cognitive capacity, and in each humans would be outperformed by animals.¹⁴⁷ But

139. See Turing, *supra* note 96.

140. See generally MARVIN MINSKY, SEMANTIC INFORMATION PROCESSING, at v, 1–28 (1968) (discussing necessary features to AI “intelligence” in terms of how those meet human characteristics).

141. Press Release, S. DEMOCRATIC CAUCUS, *Majority Leader Schumer Delivers Remarks to Launch SAFE Innovation Framework for Artificial Intelligence at CSIS* (June 21, 2023), <https://www.democrats.senate.gov/news/press-releases/majority-leader-schumer-delivers-remarks-to-launch-safe-innovation-framework-for-artificial-intelligence-at-csis> [<https://perma.cc/8KQQ-95MX>] (arguing for “guardrails” on AI because of the expected transformative change caused by the capability of AI models).

142. *Wechsler Adult Intelligence Scale (WAIS-IV)*, COGN-IQ (Sept. 11, 2025), <https://www.cogn-iq.org/learn/tests/wechsler-adult-intelligence-scale/> [<https://perma.cc/6Q5U-F8ZY>] (last visited Jan. 8, 2026).

143. See Robert P Crease, Comment, *Primate Physics*, 25 PHYSICS WORLD, 1, 18 (2012) (considering the idea of primate “kinetic intelligence”).

144. See J. Bräuer et al., *Dogs (Canis familiaris) and Wolves (Canis lupus) Coordinate with Conspecifics in a Social Dilemma*, 134 J. COMPAR. PSYCH. 211, 212 (2020) (attributing cooperative-based intelligence to wolves).

145. Bird navigation is likely more than a biological ability and requires a level of intelligence. For a study of this intelligence, see Wolfgang Wiltschko & Roswitha Wiltschko, *Magnetic Orientation in Birds*, 199 J. EXPERIMENTAL BIOLOGY 29 (1996) (“[T]hese findings can only be explained by assuming that a multitude of factors is involved in the navigational process. Apparently, birds have alternatives and can choose between several options.”).

146. Mindaugas Mitkus et al., *Raptor Vision*, in OXFORD RSCH. ENCYCLOPEDIAS OF NEUROSCIENCE (Oxford Univ. Press 2018).

147. This point is beautifully made by JULIAN TOGELIUS, ARTIFICIAL GEN. INTELLIGENCE 44 (2024).

despite the power of these tasks, they are not human-like, so you won't find them in a typical IQ test.

The point is not that animals should be writing our IQ tests,¹⁴⁸ but that if we measure AI systems based on the skills that humans excel at, we will miss powerful machines. Human-centric definitions of intelligence are under-inclusive of machines that are powerful in ways that humans are not.

But human-centric definitions will also be over-inclusive, causing us to overregulate weak machines because they seem human.

The Turing test may be the most famous test of AI intelligence, and it is entirely human-centric. Alan Turing proposed the test in response to the question “Can machines think?”¹⁴⁹ Turing thought the terms of this question were “too meaningless to deserve discussion,”¹⁵⁰ so he instead proposed a game.¹⁵¹

The imitation game, more commonly called the Turing test, involves a human player asking text-based questions both to a machine and to another human.¹⁵² After five minutes,¹⁵³ the player guesses which is the machine.¹⁵⁴ If the player cannot distinguish between the human and the machine, then the machine can “think.”¹⁵⁵

At one level, this is completely reasonable. We cannot prove that anyone else is truly thinking, but we assume they are when they appear to

148. Though that sounds adorable.

149. Turing, *supra* note 96, at 433.

150. *Id.* at 442. Researcher Edsger Dijkstra famously said, “The question of whether Machines Can Think . . . is about as relevant as the question of whether Submarines Can Swim.” RUSSELL & NORVIG, *supra* note 27, at 1035 (emphasis omitted). The question is absurd for two reasons. First, machines do not work as we do. Modern AI applications approach problems through statistics and extremely large data samples. *See* Section III.H. Whether their method is “thinking” is more definitional than practical. Second, we don't know what we do. Solipsism theories are based on the idea that I can only assume there are brains outside of mine, and I do that based on behavior. *See generally* Stephen P. Thornton, *Solipsism and the Problem of Other Minds*, INTERNET ENCYCLOPEDIA PHIL., <https://iep.utm.edu/solipsis/> [<https://perma.cc/C4GY-ULCL>] (last visited) (explaining solipsism). If a behavior test is sufficient for a person to recognize thinking in another human, it seems reasonable to use a behavioral test in judging machines.

151. Turing, *supra* note 96, at 433–34.

152. *Id.*

153. *Id.* at 442.

154. *Id.* at 433–34.

155. *Id.*

be.¹⁵⁶ It is a functional test among humans. So if a machine appears to be thinking, that is good enough for Turing.¹⁵⁷

And because a player can ask questions in any field of knowledge, the test covers all communicable intelligence. A participant can ask about chess to test the machine's ability to play chess.¹⁵⁸ Likewise, a participant can request a poem to test the machine's language skills.¹⁵⁹ The Turing test includes all communicable knowledge.¹⁶⁰

On the other hand, Turing's imitation game doesn't test whether the machine acts optimally, only whether it acts like a human.¹⁶¹ This is a weird metric. As Stuart Russell points out, "[a]eronautical engineering texts do not define the goal of their field as making 'machines that fly so exactly like pigeons that they can fool even other pigeons.'"¹⁶²

Because the Turing test is human-centric, it will accept machines with low cognitive power as intelligent if they appear human.¹⁶³ This has happened before.

Sixteen years after Turing proposed his test, Joseph Weizenbaum built a program named ELIZA designed with the ability to chat with users like a therapist.¹⁶⁴ ELIZA's programming was simple; it looked for keywords in the user's prompt and followed pre-programmed rules to respond.¹⁶⁵ For example, if a user's response included the word "mother," the program would respond, "Tell me more about your family."¹⁶⁶ If it did not find any keywords in the user's prompt, it would respond with a "content-free remark," such as "please go on," or it would repeat the user's prompt in the form of a question (e.g., "You remind me of my father." "I remind you of your father?").¹⁶⁷

156. *See generally* Thornton, *supra* note 150 (explaining the problem of confirming that other minds exist based purely on external evidence).

157. *See* Turing, *supra* note 96 (arguing that external indicators of intelligence are sufficient).

158. *Id.* at 434–35.

159. *Id.* at 434.

160. *Id.* at 434–35.

161. *See generally id.* (explaining a system of determining machine intelligence by comparing it to human behavior).

162. RUSSELL & NORVIG, *supra* note 27, at 20.

163. Turing recognized this, suggesting a high-capacity machine may pretend to be worse at math so that it appears more human. Turing, *supra* note 96, at 448–49.

164. Weizenbaum, *supra* note 107, at 42.

165. *Id.* at 37.

166. More specifically, there is a list of words ranked in order of which one controls the response, so if the user types a word that's ranked higher than "mother" it would ignore the "mother" rule. *Id.* at 37, 41–42.

167. *Id.*

ELIZA's creators recognized that the bot did not understand anything.¹⁶⁸ The entire script fit on a page and a half, single-spaced,¹⁶⁹ and it ran on a system with less processing power than a modern sprinkler system.¹⁷⁰ But it was able to fool many people into believing it was intelligent, even those that knew it was a computer program.¹⁷¹

Human-centric definitions of intelligence will be over-inclusive because it will include programs that are weak, but seem human-like, like ELIZA.

A similar problem is evident in OpenAI's definition of artificial general intelligence, which it defines as "highly autonomous systems that outperform humans at most economically valuable work."¹⁷² Under this definition, we developed general AI in the Victorian Era.¹⁷³ Most economically valuable activity performed in 1700 was transferred to

168. They saw this as a feature. "From the purely technical programming point of view then, the psychiatric interview form of an ELIZA script has the advantage that it eliminates the need of storing explicit information about the real world." *Id.* at 42 (emphasis omitted).

169. *Id.* at 44.

170. The ELIZA ran on MIT's Project MAC, ROBERT CIESLA, THE BOOK OF CHATBOTS: FROM ELIZA TO CHATGPT 43 (2024), so it likely ran on a GE 600 series. *Timesharing—Project MAC—1962-1968*, HIST. COMPUT. COMMUN., <https://historyofcomputercommunications.info/section/2.23/Timesharing-Project-MAC-1962-1968/> [<https://perma.cc/N5WM-Q6W6>] (last visited Dec. 18, 2025) (noting that MIT bought a GE 600 series machine around the time Eliza was developed at MIT); *GE-635 System Manual*, CER. COMPUTING HIST. (1964), <https://www.computinghistory.org.uk/det/15671/GE-635-System-Manual/> [<https://perma.cc/S582-FWS5>] (noting that the GE 635 was used in the MIT timesharing program). The GE 635 mainframe processed around one million instructions per second. GE-635 SYSTEM MANUAL, GEN. ELEC. COMPUT. DEPT. I-1, VII-3 (1964), https://bitsavers.org/pdf/ge/GE-6xx/CPB-371A_GE-635_System_Man_196407.pdf [<https://perma.cc/5AKS-MYA8>]. Compare this to the ATmega4808, which processes twenty times that amount. ATMEGA4808/4809 DATA SHEET, MICROCHIP 1 (2021), <https://ww1.microchip.com/downloads/en/DeviceDoc/ATmega4808-09-DataSheet-DS40002173C.pdf> [<https://perma.cc/A5TR-5JKJ>]; see also Rolf Horn, *AVR@-IoT WG Development Board AC164160-ND*, DIGIKEY (2025), <https://www.digikey.com/en/maker/projects/avr-iot-wg-development-board-ac164160-nd/5bd5a3ca6d7948e0a01363e2f238b0ef> [<https://perma.cc/J8ZJ-6HWE>] (discussing the features of a board powered by the ATmega4808 microcontroller and suggesting one application is a sprinkler system).

171. JOSEPH WEIZENBAUM, COMPUTER POWER AND HUMAN REASON: FROM JUDGMENT TO CALCULATION 189–91 (1976).

172. *OpenAI Charter*, OPENAI (2018), <https://openai.com/charter/> [<https://perma.cc/KGB2-2NZ5>].

173. In the United States, the Victorian Era spanned from approximately 1820–1914. JILLIAN PALLONE, FASHION IN THE AGE OF AMERICAN VICTORIANISM slide 2, <https://www.hofstra.edu/pdf/library/libspc-oe-lisi-fashion-american-victorianism.pdf> [<https://perma.cc/SL54-MB8Y>] (last visited Nov. 14, 2025).

machines by 1901.¹⁷⁴ Cloth spinning machines displaced between 8% and 20% of English laborers.¹⁷⁵ John Deere's steel plow did 91% of the work of turning a field.¹⁷⁶ Most labor that was economically valuable in 1700 is done by machines today,¹⁷⁷ but few would take "smart as a plow" to be a compliment.

Definitions focused on appearing or acting human are poorly tailored if the goal is to regulate powerful systems. They will be over-inclusive of charming, weak models and be under-inclusive of powerful machines whose capabilities are not similar to humans'.

Having surveyed the definitions and challenges that researchers have thought through for decades, we now turn to how policymakers address these challenges.

III. DEFINING AI IN LAW

This part will explore how artificial intelligence is defined in law. It will consider one hundred definitions used by sixty-two countries and international bodies.¹⁷⁸

To understand each element, this part takes an analytical approach, breaking down each concept used to understand and critique it. Following this analysis, each section provides recommendations for policymakers to understand the tradeoffs and pitfalls. The next part, Part IV, will reverse

174. There is an argument that these machines are not autonomous. This will be addressed in more detail. *See infra* Section III.E.

175. Craig Muldrew, 'Th'ancient Distaff' and 'Whirling Spindle': Measuring the Contribution of Spinning to Household Earnings and the National Economy in England, 1550–1770, 65 *ECON. HIST. REV.* 498, 498 (2011) (estimating around one million people worked as spinners in England in the mid-eighteenth century); Benjamin Schneider, *Technological Unemployment in the British Industrial Revolution: The Destruction of Hand-Spinning*, PAST & PRESENT, 2025 (noting that the spinning mule eliminated the hand spinning industry and estimating employment in hand spinning at around 8%).

176. The steel plow reduced the time it took to till a field from ninety-six hours by steel spade to eight hours with a steel plow, a 91.7% improvement. Hiram M. Drache, *The Impact of John Deere's Plow*, 8 *ILL. HIST. TCHR.* 2 (2001).

177. For example, over this period, farm labor in the United States has reduced from 80% to 2% of the working population, largely due to automation. Rachel Anderson & Onelisa Garza, *Farmers Do a Lot More than Just Drive Tractors*, U.S. DEPT. AGRIC.: BLOG (July 23, 2015, at 14:15 ET), <https://www.usda.gov/media/blog/2015/07/23/farmers-do-lot-more-just-drive-tractors> [<https://perma.cc/LL4E-NYC3>]. One might argue that this was a collection of systems that reduced the job categories and the amount of work in those jobs, rather than a single invention. That relies on a limited definition of system, which is discussed more in Section III.D., but it also would mean that if 49% of labor is done by automated software and 49% is done by embodied robots, then we still haven't reached the OpenAI definition.

178. For the convenience of the reader, these definitions are collected together in Appendix A to this article, each with their jurisdiction, title, date and link.

this approach and take a holistic view of how these elements work together.

A. Methodology & Summary Findings

The policies surveyed for this paper are not a comprehensive list. Instead, they are collected from news reports, tracking services, law firm client alerts, think tanks, and country-specific searches. The search focused on policies and legislation that were (1) enacted or adopted and (2) specifically focused on AI.

Most definitions come from single-jurisdiction policies, but eighteen definitions are from international bodies or multiparty agreements.¹⁷⁹ The survey includes representative samples from Asia, the Pacific, the Middle East, Africa, Europe, and North and South America. Fifteen definitions come from the United States, with seven at the federal level and eight at the state level.

Once collected, I manually coded each word in the definition, then grouped the words into categories based on common themes. For example, “learning” and “adapting” are not identical but express similar concepts, so they are discussed together.

Figure 1 shows the number of definitions that include each element in the surveyed definitions.

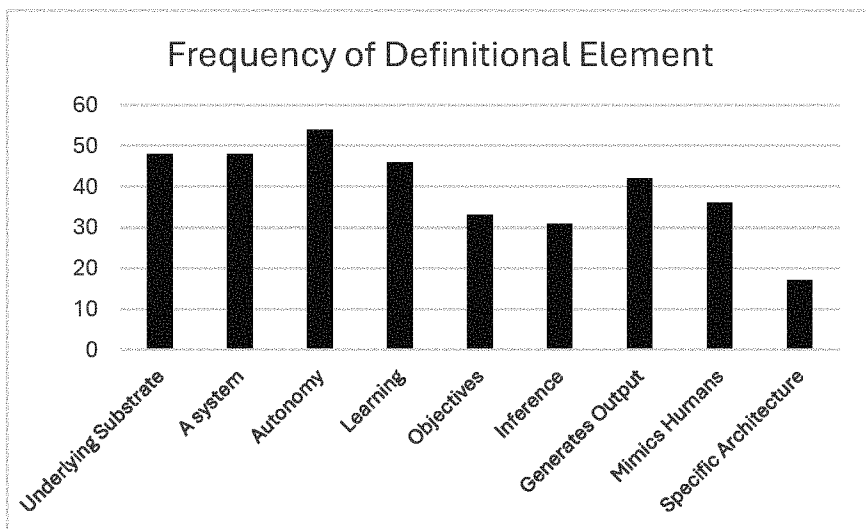


Figure 1: This chart shows the frequency with which an element appears in the surveyed jurisdictions.

179. See *infra* p. 441 app. A.

Thirty-two policies lack any definition.¹⁸⁰ Forty-eight dictate that the system must run on a particular substrate, almost always a machine or engineered system.¹⁸¹ Forty-eight define AI as a system.¹⁸² Fifty-four define it as having autonomy or decision-making capabilities.¹⁸³ Forty-six include the ability to learn or adapt.¹⁸⁴ Thirty-three define AI to operate based on objectives or goals,¹⁸⁵ most commonly goals provided by users.¹⁸⁶ Thirty-one define it as conducting inference or reason.¹⁸⁷ Forty-two say AI generates output.¹⁸⁸ Thirty-six define AI as mimicking some attribute of being human,¹⁸⁹ typically either intelligence or capability generally.¹⁹⁰ Seventeen definitions mention specific architectures, like statistical modeling.¹⁹¹

Because some elements of a definition may have weaknesses that other elements patch, the following chart provides the correlation between elements to show how they tend to work together. Part IV also does a holistic review of the most common definition to understand how elements work together.¹⁹²

180. *See infra* p. 441 app. A.

181. *See infra* p. 441 app. A.

182. *See infra* p. 441 app. A.

183. *See infra* p. 441 app. A.

184. *See infra* p. 441 app. A.

185. *See infra* p. 441 app. A.

186. *See infra* p. 441 app. A.

187. *See infra* p. 441 app. A.

188. *See infra* p. 441 app. A.

189. *See infra* p. 441 app. A.

190. *See infra* p. 441 app. A.

191. *See infra* p. 441 app. A.

192. *See infra* p. 441 app. A.

	Underlying Substrate	A system	Autonomy	Learning	Objectives	Inference	Generates Output	Mimics Humans	Specific Architecture
Underlying Substrate	-	0.48	0.44	0.36	0.39	0.26	0.28	0.28	0.10
A system	0.48	-	0.49	0.32	0.60	0.39	0.40	0.16	0.15
Autonomy	0.44	0.49	-	0.29	0.52	0.40	0.54	0.15	0.26
Learning	0.36	0.32	0.29	-	0.21	0.34	0.19	0.39	0.28
Objectives	0.39	0.60	0.52	0.21	-	0.36	0.39	0.01	0.08
Inference	0.26	0.39	0.40	0.34	0.36	-	0.35	0.04	-0.02
Generates Output	0.28	0.40	0.54	0.19	0.39	0.35	-	-0.17	0.32
Mimics Humans	0.28	0.16	0.15	0.39	0.01	0.04	-0.17	-	0.10
Specific Architecture	0.10	0.15	0.26	0.28	0.08	-0.02	0.32	0.10	-

Figure 2: This chart shows the correlation between the frequency of various elements in the surveyed definitions. Darker boxes indicate a stronger correlation.

Figure 2 shows two points of interest. The only negative correlation is between mimicking humans and generating output. This shows that definitions focused on acting well focus less on thinking well, and vice versa.

The second insight supports this. The strongest correlation with the “mimics humans” element is the “learning” element, both of which focus on intellect. These two correlations reflect an old debate in AI research on whether the definition should focus on actions or cognition.¹⁹³

We now turn to an analysis of each element.

193. RUSSELL & NORVIG, *supra* note 27, at 20–21.

B. No definition

A surprising number of statements, strategies, and national plans offer no definition at all.¹⁹⁴ Of the one hundred definitions¹⁹⁵ surveyed, 32% do not define AI.¹⁹⁶ This includes documents produced by prominent international bodies, for example the G7's Hiroshima AI Process Comprehensive Policy Framework,¹⁹⁷ the Seoul Declaration on AI,¹⁹⁸ the G20 Ministerial Declaration,¹⁹⁹ and the Bletchley Declaration.²⁰⁰ It also includes country specific documents from China,²⁰¹ Colombia,²⁰² Egypt,²⁰³

194. *See infra* p. 441 app. A.

195. *See infra* p. 441 app. A.

196. *See infra* p. 441 app. A.

197. *Hiroshima Process International Guiding Principles for Organizations Developing Advance AI System*, MINISTRY FOREIGN AFFAIRS, <https://www.mofa.go.jp/files/100573471.pdf> [<https://perma.cc/TX45-4E2F>] (last visited Jan. 10, 2025); *Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Sys.*, G7G20 (Oct. 30, 2023), <https://g7g20-documents.org/database/document/2023-g7-japan-leaders-leaders-annex-hiroshima-process-international-code-of-conduct-for-organizations-developing-advanced-ai-systems#section-1> [<https://perma.cc/2A7B-4AUG>].

198. *Seoul Declaration For Safe, Innovative and Inclusive AI By Participants Attending The Leaders' Session: AI Seoul Summit, 21 May 2024*, GOV.UK (May 21, 2024), <https://www.gov.uk/government/publications/seoul-declaration-for-safe-innovative-and-inclusive-ai-ai-seoul-summit-2024/seoul-declaration-for-safe-innovative-and-inclusive-ai-by-participants-attending-the-leaders-session-ai-seoul-summit-21-may-2024> [<https://perma.cc/C25Y-VTMQ>].

199. *G20 Ministerial Declaration: 13 September, 2024*, GOV.UK (Sep. 13, 2024), <https://www.gov.uk/government/publications/g20-ministerial-declaration-maceio-13-september-2024/g20-ministerial-declaration-13-september-2024> [<https://perma.cc/9FCS-GLU5>].

200. *The Bletchley Declaration by Countries Attending the AI Safety Summit, 1-2 November 2023*, GOV.UK (Nov. 1, 2023), <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023> [<https://perma.cc/45P2-F65F>].

201. There is no definition of AI in China's AI Safety Governance Framework. NAT'L TECH. COMM. 260 ON CYBERSEC. OF SAC, AI SAFETY GOVERNANCE FRAMEWORK (2024), <https://www.tc260.org.cn/upload/2024-09-09/1725849192841090989.pdf> [<https://perma.cc/3VU5-NWEN>]; *Interim Measures for the Management of Generative Artificial Intelligence Services*, CHINA LAW TRANSLATE (July 10, 2023), <https://www.chinalawtranslate.com/en/generative-ai-interim/> [<https://perma.cc/R2F4-LWQ9>].

202. ARMANDO GUÍO ESPAÑOL, CONSULTANT OF THE CAF, ETHICAL FRAMEWORK FOR ARTIFICIAL INTELLIGENCE (2024) https://cyber.harvard.edu/sites/default/files/2020-12/Colombia_AI_Ethical_Framework.pdf [<https://perma.cc/HW8N-THPP>].

203. THE NAT'L COUNCIL FOR A.I., EGYPTIAN CHARTER FOR RESPONSIBLE AI (2023), <https://aicm.ai.gov.eg/en/Resources/EgyptianCharterForResponsibleAIEnglish-v1.0.pdf> [<https://perma.cc/E4NH-UTB8>].

India,²⁰⁴ Japan,²⁰⁵ Mauritius,²⁰⁶ Saudi Arabia,²⁰⁷ Taiwan,²⁰⁸ and the United Kingdom.²⁰⁹ It is not only frameworks and strategies that lack definitions, but also some binding laws.²¹⁰

Three policies expressly recognize the challenge of defining AI and then give general principles, without attempting a complete definition.²¹¹ For example, the United Nations Educational Scientific and Cultural Organization's (UNESCO's) Recommendation on the Ethics of Artificial Intelligence disclaims any "ambition to provide one single definition of AI, since such a definition would need to change over time, in accordance with technological developments."²¹² However, it then goes on to list out common features of AI systems, like autonomy, learning, and perception.²¹³

204. SANDIP CHATTERJEE, GOV'T OF INDIA, ADVISORY eNo.2(4)/2023-CYBERLAWS-3 (2024),

<https://www.meity.gov.in/static/uploads/2024/02/9f6e99572739a3024c9cdaec53a0a0ef.pdf> [https://perma.cc/473K-TTWU]; NITI AAYOG, RESPONSIBLE AI FOR ALL: ADOPTING THE FRAMEWORK: A USE CASE APPROACH ON FACIAL RECOGNITION TECHNOLOGY (2022) https://www.niti.gov.in/sites/default/files/2022-11/Ai_for_All_2022_02112022_0.pdf [https://perma.cc/4EME-YXYE].

205. EXPERT GRP. ON HOW A.I. PRINCIPLES SHOULD BE IMPLEMENTED, AI GOVERNANCE IN JAPAN VER. 1.1, (2021), https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20210709_8.pdf [https://perma.cc/M7DY-942J] (discussing how AI principles should be implemented).

206. WORKING GRP. ON A.I., MAURITIUS ARTIFICIAL INTELLIGENCE STRATEGY (2018), <https://treasury.govmu.org/Documents/Strategies/Mauritius%20AI%20Strategy.pdf> [https://perma.cc/XV7B-L275].

207. KINGDOM OF SAUDI ARABIA, VISION 2030 (2024), https://ndmc.gov.sa/en/investorsrelations/Documents/Investor_Presentation/Investor_Presentation_Oct2024.pdf [https://perma.cc/3CT2-WRHC].

208. ORG. FOR ECON. COOP. & DEV., DRAFT BASIC LAW ON ARTIFICIAL INTELLIGENCE (2024), <https://ipfs.juchunko.com/ipfs/bafybeia5237mn2a3yabsfjpaubdn7ynsiauhbq54qngpn34jhfetndnxbi> [https://perma.cc/DU3U-4GP7].

209. DEP'T FOR SCI., INNOVATION & TECH., INTRODUCTION TO AI ASSURANCE (2024), https://assets.publishing.service.gov.uk/media/65ccf508c96cf3000c6a37a1/Introduction_to_AI_Assurance.pdf [https://perma.cc/YCC6-HXX7].

210. See *infra* p. 441 app. A (China and Taiwan). This may be a policy choice to allow flexibility for decisionmakers.

211. See *infra* p. 441 app. A (Hong Kong, *Model Personal Data Protection Framework*) (UNESCO, *Recommendation on the Ethics of Artificial Intelligence*) (United Kingdom, *A Pro-Innovation Approach to AI Regulation*) (United Kingdom, *Guidance on the AI Auditing Framework*).

212. UNESCO, RECOMMENDATION ON THE ETHICS OF ARTIFICIAL INTELLIGENCE 10 (2022), <https://unesdoc.unesco.org/ark:/48223/pf0000381137> [https://perma.cc/WD68-UCKG] [hereinafter UNESCO].

213. *Id.*

New Zealand takes an “I know it when I see it” approach, reading: “This Charter does not specify a technical definition of an algorithm. It instead commits signatories to take a particular focus on those algorithms that have a high risk of unintended consequences and/or have a significant impact if things do go wrong, particularly for vulnerable communities.”²¹⁴

As noted in Part V, vague regulations may offer more value to both regulators, who benefit from flexibility, and to risk-seeking entrepreneurs, who may prefer a vague definition that they can game.²¹⁵ One might think that the risk-loving, move-fast-and-break-things culture of Silicon Valley would prefer the flexibility and plausible deniability of vague regulations. In practice, this has not turned out. The EU released an ambitious systemic regulation of artificial intelligence, which as discussed in Part IV contains vague definitions, and since then innovators and investors have begun passing over the EU market, expressly citing uncertainty about the law.²¹⁶

C. Underlying Substrate

Twelve percent of definitions refer to it as “systems” that do some variety of actions without limiting what types of systems qualify.²¹⁷ Forty-

214. N.Z. GOV'T, ALGORITHM CHARTER FOR AOTEAROA NEW ZEALAND 1 (2020), https://data.govt.nz/assets/data-ethics/algorithm/Algorithm-Charter-2020_Final-English-1.pdf [<https://perma.cc/TTS9-YE7F>].

215. See generally Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 DUKE L.J. 557 (1992) (discussing the pros and cons of rules approaches and standards approaches).

216. STAN. UNIV. HUM.-CENTERED A.I., ARTIFICIAL INTELLIGENCE INDEX REPORT 2024 249 fig. 4.3.10 (2024), https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf [<https://perma.cc/V7NW-FSFN>] (showing rising investment in A.I. globally, but European A.I. private investment flattening in 2021 after the EU AI Act was announced); Pascale Davies, *Why OpenAI's Voice Mode, Meta's Llama and Apple's AI Won't Be Coming to Europe Yet*, EURONEWS (Aug. 10, 2024), <https://www.euronews.com/next/2024/10/08/why-openais-voice-mode-metas-llama-and-apples-ai-wont-be-coming-to-europe-yet> [<https://perma.cc/R6J5-N62T>]; Ivana Saric, *Apple Says It Won't Roll Out AI Features in Europe Due to Regulatory Concerns*, AXIOS (June 21, 2024), <https://www.axios.com/2024/06/21/apple-ai-features-europe>. [<https://perma.cc/YMW9-LM2S>]. It may be that, regardless of the law's quality, industry leaders would have criticized it in order to reduce regulation. However, the leaders avoiding the EU market are calling for regulation elsewhere, which suggests the problem is the regulation's content, not its existence. See, e.g., James Clayton, *Sam Altman: CEO of OpenAI Calls for US to Regulate Artificial Intelligence*, BBC (May 17, 2023), <https://www.bbc.com/news/world-us-canada-65616866> [<https://perma.cc/P54E-GSAF>].

217. For example, the United Arab Emirates defines A.I. as “a collection of technologies enabling a machine or system to” NAT'L PROGRAM FOR A.I., AI GUIDE (2020), https://ai.gov.ae/wp-content/uploads/2020/02/AIGuide_EN_v1-online.pdf

[<https://perma.cc/5YYD-VZ4G>]. Similarly, Brazil defines AI as “A system with different

eight percent of definitions indicate what makes up the system.²¹⁸ For example, the EU AI Act,²¹⁹ President Biden’s executive order,²²⁰ and others require that an AI system be “machine-based.”²²¹ The recitals to the EU AI Act faux-clarify that “‘machine-based’ refers to the fact that AI systems run on machines.”²²² No definition defines “machines,” so it is not clear whether this excludes biological or chemical systems and whether the system must be electrically powered or have moving parts.²²³

Australia takes a similar approach, but instead of “machine-based system,” it uses the term “engineered system.”²²⁴ The United States’s National Institute of Standards and Technology (NIST) Risk Management Framework combines the two, saying the system must be “engineered or

degrees of autonomy designed to infer objectives using machine learning or logic.” OECD, ESTRATEGIA BRASILEIRA DE INTELIGENCIA ARTIFICIAL (2021), <https://oecd.ai/en/work/documents/brazil-brazilian-ai-strategy-2021> [<https://perma.cc/RGR6-QCCD>] [hereinafter BRAZIL’S A.I. STRATEGY].

218. See, e.g., p. 441 app. A (Australia, *Safe and Responsible AI in Australia*).

219. EU AI ACT, *supra* note 5, at 3(1).

220. Exec. Order No. 14110, Executive Order on the Safe, Secure, and Trustworthy Development and Use of A.I., 88 Fed. Reg. 75191 (Oct. 30, 2023) [hereinafter U.S. E.O. No. 14110].

221. NATIONAL STRATEGY FOR ARTIFICIAL INTELLIGENCE 10 (2024), https://ictd.portal.gov.bd/sites/default/files/files/ictd.portal.gov.bd/page/6c9773a2_7556_4395_bbec_f132b9d819f0/Draft%20-

<https://perma.cc/LQ6F-KWFP>; Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, Sep. 5, 2024, C.E.T.S. No. 225, <https://rm.coe.int/1680afae3c> [<https://perma.cc/H7DT-R399>] [hereinafter The Framework]; NITI AAYOG, NATIONAL STRATEGY FOR ARTIFICIAL INTELLIGENCE (2018), <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf> [<https://perma.cc/YAN5-N2S8>] [hereinafter INDIA’S A.I. STRATEGY]; MINCIENCIA, POLÍTICA NACIONAL DE INTELIGENCIA ARTIFICIAL, (2020), <https://drive.google.com/file/d/11OLxLp8NyKgpeRFLe45X0zStY7SFEJIC/view> [<https://perma.cc/QG5L-F3MC>] [hereinafter CHILE’S A.I. POLICY].

222. *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence and Amending Certain Union Legislative Acts*, COM (2021) 206 final (Apr. 21, 2021), [hereinafter *Recital 12*].

223. See *infra* p. 441 app. A. Peru offers a little more clarity by defining “AI System” as a “sistema electrónico-mecánico,” or an electronic-mechanical system, which would exclude systems that are not electronic, but may exclude systems that are purely electronic, lacking moving parts. LEY QUE PROMUEVE EL USO DE LA INTELIGENCIA ARTIFICIAL EN FAVOR DEL DESARROLLO ECONÓMICO Y SOCIAL DEL PAÍS, No. 31814 (2023), <https://busquedas.elperuano.pe/dispositivo/NL/2192926-1> [<https://perma.cc/X78Y-G57U>] [hereinafter PERU’S A.I. LAW].

224. DEP’T OF INDUS., SCI. & RES., SAFE AND RESPONSIBLE AI IN AUSTRALIA 5 (2023), https://storage.googleapis.com/converlens-au-industry/industry/p/prj2452c8e24d7a400c72429/public_assets/Safe-and-responsible-AI-in-Australia-discussion-paper.pdf [<https://perma.cc/Uyh2-WZE5>] [hereinafter AUSTRALIA’S A.I. STRATEGY].

machine-based.”²²⁵ Engineered means something is “designed and built using scientific principles,”²²⁶ which seems to include almost anything intentionally manufactured.

Several countries and international cooperatives narrow the definition of AI to being “computer-based,”²²⁷ “computer programs,”²²⁸ “software”²²⁹ or “applications.”²³⁰

Some laws, regulations and frameworks use definitions that are more theoretical, defining AI as “technology,”²³¹ algorithms,²³² or “a scientific

225. NAT’L INST. OF STANDARDS & TECH., ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK 1.0, at 1 (2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> [<https://perma.cc/EW6P-2LVG>] [hereinafter NIST FRAMEWORK].

226. *Engineered*, CAMBRIDGE DICTIONARY, <https://dictionary.cambridge.org/us/dictionary/english/engineered> [<https://perma.cc/K2Y4-8HAE>] (last visited Dec. 19, 2025).

227. *See, e.g.*, YAPAY ZEKA KANUN, TURKEY’S ARTIFICIAL INTELLIGENCE LAW PROPOSAL 12 (2024), <https://cdn.tbmm.gov.tr/KKBSPublicFile/D28/Y2/T2/WebOnergeMetni/e50ccc8a-ab90-45fa-a553-76b880c78fb8.pdf> [<https://perma.cc/CK8R-UJ4X>].

228. *See, e.g., infra* p. 441 app. A (Japan and Hong Kong). Note this definition also includes “machines.” *See id.*

229. *See, e.g.*, NAT’L CYBER SEC. AGENCY, GUIDELINES FOR SECURE ADOPTION AND USAGE OF ARTIFICIAL INTELLIGENCE 11 (2024), https://assurance.ncsa.gov.qa/sites/default/files/publications/policy/2024/CSSP_Guidelines_for_Secure_Usage_and_Adoption_of_Artificial_intelligence-Eng-v1.0_2.pdf [<https://perma.cc/5SDR-HQTL>] [hereinafter QATAR’S A.I. GUIDELINES] (defining A.I. as including “hardware, software or both”).

230. NAT’L CYBER SEC. CTR., GUIDELINES FOR SECURE AI SYSTEM DEVELOPMENT (2023), <https://www.ncsc.gov.uk/files/Guidelines-for-secure-AI-system-development.pdf> [<https://perma.cc/8VVG-CMK5>] [hereinafter UK A.I. GUIDELINES]. This statement was joined by the NSA, FBI, and the cybersecurity agencies of 17 other countries.

231. MINISTER FOR COMMC’NS & INFO. SING., MODEL ARTIFICIAL INTELLIGENCE GOVERNANCE FRAMEWORK 18 (2d ed. 2020), <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/SGModelAIGovFramework2.pdf> [<https://perma.cc/493P-2WSW>] [hereinafter SINGAPORE’S A.I. FRAMEWORK]; MONETARY AUTH. OF SING., PRINCIPLES TO PROMOTE FAIRNESS, ETHICS, ACCOUNTABILITY AND TRANSPARENCY IN THE USE OF ARTIFICIAL INTELLIGENCE AND DATA ANALYTICS IN SINGAPORE’S FINANCIAL SECTOR 3 (2018), <https://www.mas.gov.sg/-/media/MAS/News-and-Publications/Monographs-and-Information-Papers/FEAT-Principles-Updated-7-Feb-19.pdf> [hereinafter SINGAPORE FINANCIAL SECTOR PRINCIPLES] [<https://perma.cc/26RX-LARW>];

Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.) (defining A.I. as a technological system); *Directive on Automated Decision-Making*, GOV’T OF CAN. (June 24, 2025), <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592#appA> [<https://perma.cc/6DM8-GXLP>] [hereinafter Automated Decision-Making] (defining it as an information technology in the appendix).

232. Council of Eur., *Guidelines on the Responsible Implementation of Artificial Intelligence Systems in Journalism*, (2023), <https://rm.coe.int/cdmsi-2023-014-guidelines-on-the-responsible-implementation-of-artific/1680adb4c6> [<https://perma.cc/9ZJZ-WHB3>].

discipline.”²³³ Others are more practical, defining it to include any product or service.²³⁴

Few of these definitions offer any practical limitation. A \$20 wristwatch would meet the definition of AI in many of these definitions. Complexity is not required to be a machine. Simple machines include ramps and levers,²³⁵ and wristwatches are more complex than either.²³⁶

Under definitions that accept something that is “engineered,”²³⁷ a wristwatch fits even more naturally because engineered means something is “designed and built using scientific principles.”²³⁸ Sundials are designed and built based on principles of astronomy, so that definition would seem to include them as well.²³⁹

1. Policy Considerations for Underlying Substrates

Policymakers defining the underlying substrate should weigh the tradeoff between having a broad definition, which includes basic devices, or a narrow definition, which misses systems based on biology or chemistry that do not yet exist but may later be developed. Broad definitions can also be narrowed with other elements, including those discussed below.

233. MINISTRY FOR DIGIT. TRANSFORMATION & THE CIV. SERV., 2024 ARTIFICIAL INTELLIGENCE STRATEGY (2024), https://digital.gob.es/content/dam/portal-mtdfp/DigitalizacionIA/1_DOSSIER_AI_ENGLISH_15_JULIO.pdf [<https://perma.cc/KRK7-554X>]; see also *¿Qué es la Estrategia de Inteligencia Artificial 2024?*, LA MONCLOA (June 5, 2024), <https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/transformacion-digital-y-funcion-publica/Paginas/2024/ia-inteligencia-artificial-estrategia-espana.aspx> [<https://perma.cc/2L2A-TD3W>].

234. DEP'T FOR SCI., INNOVATION & TECH., IMPLEMENTING THE UK'S AI REGULATORY PRINCIPLES 9 (2024), https://assets.publishing.service.gov.uk/media/65c0b6bd63a23d0013c821a0/implementing_the_uk_ai_regulatory_principles_guidance_for_regulators.pdf [<https://perma.cc/J7WD-LS2M>]; SMART DUBAI, AI ETHICS PRINCIPLES & GUIDELINES 16, <https://www.digitaldubai.ae/pdfviewer/web/viewer.html?file=https://www.digitaldubai.ae/docs/default-source/ai-principles-resources/ai-ethics.pdf> [<https://perma.cc/4N4X-CV8P>] (last visited Jan. 11, 2025) [hereinafter DUBAI'S A.I. ETHICS].

235. *Simple Machine*, ENCYCLOPEDIA BRITANNICA, <https://www.britannica.com/technology/simple-machine> [<https://perma.cc/ZV8V-PQGR>] (last visited Dec. 19, 2025).

236. *How Do Watches Work? A Simple Guide*, TIMEX: THE TIMEX BLOG (July 14, 2022), <https://timex.com/blogs/the-timex-blog/how-do-watches-work-a-simple-guide> [<https://perma.cc/2ADC-2XMZ>].

237. See, e.g., *infra* p. 441 app. A (Australia: Safe and Responsible AI in Australia).

238. CAMBRIDGE DICTIONARY, *supra* note 226.

239. David P. Stern, *The Sundial*, NASA (Oct. 10, 2016), <https://pwg.gsfc.nasa.gov/stargaze/Sundial.htm>. [<https://perma.cc/8UE3-ENFU>].

D. A “System”

Forty-eight percent of definitions of artificial intelligence apply to a “system.” For example, the EU AI Act refers to a “machine-based system.”²⁴⁰

System is not defined in any policy or law reviewed for this article. This is a critical mistake because AI systems interact seamlessly with other software. It is common for two pieces of software to work together, but there is not guidance on whether this makes them the same “system.” Without clear definitional boundaries, every line of code that interacts with the AI system may be considered part of that AI system and subject to the same regulations. This could expand the regulatory scope well beyond anything reasonably anticipated.

We might think simple clarifications would prevent confusion. Perhaps “system” means anything necessary for the application to function. But this would still be overly broad. Consider a high-risk AI app that hires outdoor workers. In the EU, hiring workers is a high-risk activity that subjects this app to increased regulation.²⁴¹ Suppose the hiring app functions only if it has access to a weather app to estimate how many workers may be needed each day. The question arises whether the weather app and the hiring app are part of the same “system.” If so, then the weather app is subject to the regulations governing high-risk AI systems. The weather app, which may have no connection with or knowledge of the hiring app, could find itself encompassed by AI regulations.

A weather app is a trivial example. However, many AI apps are likely to require internet access to function. If every necessary element is part of the “system,” most of the internet would be sucked into the regulatory vortex. Being “necessary for proper functioning” would be an impractically loose definition.²⁴²

We may try a tighter version of system that requires the two programs have the same designer, the same owner, some exclusivity, or other factors to delineate the boundaries of the system. This section will show that modern software cannot be delineated with clean boundaries and AI’s direction toward agentic AI is about to make the problem much worse.

240. EU AI Act, *supra* note 5, at 3(1).

241. For example, AI systems doing high-risk activity can be trained only on data that is “relevant, sufficiently representative, and to the best extent possible, free of errors and complete in view of the intended purpose.” *Recital 12, supra* note 222, at art. 10(3).

242. EU AI Act, *supra* note 5.

1. A Hardware Approach Is Inadequate

When Alan Turing built the machine that cracked the enigma codes,²⁴³ it was reasonably clear what constituted the “system.” Each functional piece of the machine was connected to another piece by several miles of wiring.²⁴⁴ Something was part of the system if it was physically attached to the system.²⁴⁵ This hardware view of systems is no longer practical for three reasons.

First, hardware often runs a variety of applications. My laptop can run minesweeper and Microsoft Word, but few would argue that Microsoft Word and minesweeper are part of the same system—they have nothing to do with each other. So, a strict hardware approach is likely overinclusive.

Second, my laptop is connected to an ethernet cable that is connected to a router, which is connected to a modem that is connected to my internet service provider, which is connected to the internet infrastructure that connects to millions of other devices across the world.²⁴⁶ If “system” is defined by the hardware it is connected to, then most computers are part of a single system.

Third, unlike Turing’s machine, the connections in modern computers are not consistent. I log in and out of Wi-Fi throughout the day, and I change Wi-Fi access points without fundamental changes in my system. The system continues to operate even as it switches the hardware it is using. So, hardware is not a practical way to delineate the boundaries of a modern system.²⁴⁷

We might try to save the hardware approach by looking only at local operations, things that rely on some piece of hardware, like my hard drive.

243. Hayley Cox, *Cracking Stuff: How Turing Beat the Enigma*, UNIV. OF MANCHESTER (Nov. 2018), <https://www.mub.eps.manchester.ac.uk/science-engineering/2018/11/28/cracking-stuff-how-turing-beat-the-enigma/> [https://perma.cc/2VH4-KABS].

244. See JENNIFER WILCOX, U.S. DEP’T OF DEF., SOLVING THE ENIGMA: HISTORY OF THE CRYPTANALYTIC BOMBE 1–2 (2024), https://media.defense.gov/2022/Sep/29/2003087366/-1/-1/0/SolvingTheEnigma24_Final.PDF [https://perma.cc/F7QC-V7HF].

245. Although, one might argue that the bombe machine could not operate without the power grid, so the power grid was part of the system.

246. Rus Shuler, *How Does the Internet Work?*, INTERNET WHITEPAPER (2002), <https://web.stanford.edu/class/msande91si/www-spr04/readings/week1/InternetWhitepaper.htm> [https://perma.cc/FM4S-METV].

247. That is not to say hardware is never helpful in delineating a system. Some systems are intentionally cordoned off from the internet. Mesh Flinders & Ian Smalley, *What Is an Air Gap?*, IBM (Oct. 2024), <https://www.ibm.com/think/topics/air-gap> [https://perma.cc/Y7Q6-64XY].

This is, again, unworkable. Even local applications like MS Word regularly use cloud technology for storage and processing,²⁴⁸ which is another way of saying they rely on systems hundreds of miles away from the end user. A local hardware approach would exclude distributed applications and cloud supported apps.

2. A Software Approach Is Inadequate

We might think a software approach could help us delineate the “system,” but this likely fails because programs are rarely self-contained, and they interact extensively with other programs. Courts tasked with delineating a system will struggle to find the seams in a system designed to be seamless. Here are some examples.

a. Software Applications Span Multiple Files

Most modern applications run across multiple files. Consider this brief segment of Python code. Assume we want to display the word “Sourdough.” The code might say:

```
Print(“Sourdough”)
```

When compiled and run, this code will display the word “Sourdough” on the screen. This entire program is contained within a single file. It is easy to tell where the software begins and ends.

Suppose instead we want to display an image of a sourdough loaf. The code for this might say:

```
image = Image.open(sourdough.jpg)  
image.show()249
```

Note that the first line of this code refers to a separate file called “sourdough.jpg”. The program needs both the Python script file and the image file to run. Programs with images may store those files outside the program file. So, to capture what is happening in the sourdough loaf program, we would need two files.

A reasonable definition of system cannot be limited to a single file because even simple programs often use multiple files. However, allowing

248. Sneha Gupta, *Exploring Microsoft Word in the Cloud: Comprehensive Insights*, Bi2DEV, <https://bi2dev.com/articles/exploring-microsoft-word-cloud-analysis/> [<https://perma.cc/2PYB-DHDF>] (last visited Dec. 19, 2025).

249. This command uses a Python library called pillow. Libraries are discussed further below.

the definition of “system” to include multiple files creates new problems because any given file may be part of multiple systems.

Returning to our example, the image file (`sourdough.jpg`) may be used for a program that generates food recommendations and another that generates cooking techniques. Both programs would use the same `sourdough` image file. The file `sourdough.jpg` would be part of two systems even if they had nothing else in common.

This is not too troubling with a passive file like a photo, but executable code is also used across multiple systems through coding libraries, as explained in the next section.

b. Software Applications Span Multiple Programs with Multiple Owners

Suppose I have a program that regularly takes the same action, like converting measurements between the metric and imperial systems. I could add code that does the conversion to the program each time it is needed, but it would be much more efficient to instead create a function. A function is a set of operations that a program can call when needed to take a designated set of actions.²⁵⁰ So, instead of having to remember how many milliliters are in a cup each time, I can call the function:

```
covert_cups_to_milliters(x)
```

and the function does the conversion for me. However many cups are in “x” are returned correctly in milliliters.

Now suppose I want to use this function in several programs or to make it available to others. I might do this by releasing it as a library. A library is a bit of code that users can easily import into their programs.²⁵¹

A metric conversion function is a trivial example, but suppose I want to do something more complex, like connect to Bluetooth. Rather than learn about radio frequencies and write a thousand lines of code to implement them, I could import a Bluetooth library someone else already wrote. I might do this with the command:

```
import bleak //imports the bleak Bluetooth library
```

250. *Functions*, UTAH COMPUT. SCI. DEP’T, <https://users.cs.utah.edu/~germain/PPS/Topics/functions.html> [https://perma.cc/H98T-HK2Q] (last visited Dec. 19, 2025).

251. *Tutorial on the Python Library*, MEDIUM (Feb. 12, 2023), <https://seattlewebsitedevelopers.medium.com/tutorial-on-the-python-library-6ca0a9ae074b> [https://perma.cc/Y267-HTMLW].

This one command imports 924 lines of code named the “bleak” library and allows my program to work with Bluetooth. “bleak” is just the name of the library given by its author. When my code runs, it will use the code in that imported library without making any distinction about who wrote it.

Libraries are essential in coding because they allow programs to quickly import code that is written and maintained by others.²⁵² They are also ubiquitous; for example, the first lines of the bleak library above import seven other libraries.²⁵³ There are four reasons why libraries complicate the definition of a “system.”

First, my code could not meet its intended purpose without the library. The library is my code’s only source for Bluetooth functionality. My code depends on it.

Second, the library does not do anything useful on its own. Libraries are more like tires than they are a car. They can be imported to add functionality to larger programs, but they lack independent functionality.²⁵⁴ The usefulness of the library depends on the program importing it.

Third, the person importing the library often never reads the code that makes up the library. Libraries typically include documentation on how to integrate the library, what it can do, and the commands that can be used with the library.²⁵⁵ But the reason you use a library is because you do not want to learn the intricacies of Bluetooth radio frequencies when you are just trying to write a baking app. So, they are not built on shared knowledge.

Fourth, a library may be owned and maintained by someone entirely unconnected with the project.²⁵⁶ The Bluetooth library above could be used by cell phone manufacturers, RC plane tinkerers, and mobile game

252. *What is a Library in Programming: AP® CS Principles Review*, ALBERT (May 20, 2025), <https://www.albert.io/blog/what-is-a-library-in-programming-ap-cs-principles-review/> [<https://perma.cc/954C-2FCD>]; see also *Dependency*, XKCD (2003), <https://xkcd.com/2347/> [<https://perma.cc/VVG9-KDGA>] (comic that depicts the nature of libraries).

253. Henrik Blidh, *Changing Docstring with Demo-Code to R-String*, BLEAK (2024), https://github.com/hbldh/bleak/blob/develop/bleak/__init__.py [<https://perma.cc/YFH6-YSU3>].

254. See, e.g., BLEAK (2020), <https://bleak.readthedocs.io/en/latest/> [<https://perma.cc/BFD7-9D7T>] (documentation for the bleak library) (noting that it can be used to provide Bluetooth capability to applications).

255. See, e.g., *id.*

256. See, e.g., *id.*; PythonNUS, GITHUB, <https://github.com/coyt/PythonNUS> [<https://perma.cc/XZ6G-ANHF>] (implementing the Bleak library for an unrelated data transmission system).

developers, who may or may not be the same entities that develop and maintain the library.

So, is the library part of the “system”? On one hand, libraries are essential to the functionality of the system and a library on its own often lacks independent functionality. On the other hand, libraries are: (1) used simultaneously and non-rivalrously with an unknown number of unrelated programs; (2) rarely written, read or understood by the programmer; and (3) often owned and maintained by someone unaffiliated with the project. It is just not clear. Courts may struggle to determine what makes up a system.

Application programming interfaces, or APIs, offer another challenge to the definition of “system.” APIs are interconnectors that allow systems to connect seamlessly together.²⁵⁷ APIs allow users to log in to the New York Times website using a Facebook account or to see Google maps on an Uber Eats app. APIs enable joint efforts for social media tie-ins,²⁵⁸ payment processing,²⁵⁹ mapping,²⁶⁰ geolocation,²⁶¹ e-commerce data sharing,²⁶² and cloud computing.²⁶³ In addition, because the process is seamless by design, it will pose the same challenges as libraries when courts are asked to determine whether they are part of the same “system.” It is not clear whether having an API will subject the designer to any

257. *What is an API (Application Programming Interface)?*, AMAZON WEB SERVS., <https://aws.amazon.com/what-is/api/> [<https://perma.cc/M76M-HXEK>] (last visited Dec. 19, 2025).

258. *See, e.g., Mathur, Google Maps Starts Showing Social Media Profiles From Your Favorite Bars and Restaurants*, ANDROID POLICE (Mar. 13, 2024), <https://www.androidpolice.com/google-business-profile-link-social-media-profiles/> [<https://perma.cc/Q3N5-Y5GY>] (explaining that users can now share their data between social media and Google Maps).

259. *E.g., How the Maps JavaScript API is Billed*, GOOGLE, <https://developers.google.com/maps/documentation/javascript/usage-and-billing#new-payg> [<https://perma.cc/2RBP-AKZR>] (last visited Dec. 19, 2025) (explaining how Google Maps can be incorporated for billing data).

260. Flora Wong, *Special Delivery with Google Maps APIs*, GOOGLE MAPS PLATFORM (Mar. 2, 2017), <https://cloud.google.com/blog/topics/inside-google-cloud/special-delivery-google-maps-apis> [<https://perma.cc/D7HQ-UFR8>].

261. *See, e.g., How Geolocation Features Enhance the Food Delivery Experience*, DITO, <https://www.ditoweb.com/2025/01/how-geolocation-features-transform-the-food-delivery-experience/> [<https://perma.cc/Z2FC-LGDZ>] (last visited Dec. 19, 2025) (explaining how Google Maps can share geolocation data with other apps).

262. *See Dave McClusky, How to Improve the Delivery and Ecommerce Experience with Google Maps Platform*, GOOGLE MAPS PLATFORM (May 1, 2020), <https://mapsplatform.google.com/resources/blog/how-improve-delivery-and-ecommerce-experience-google-maps-platform/> [<https://perma.cc/C2TH-HGXY>].

263. *Cloud Services for Google Maps Platform*, SOFTWAREONE, <https://www.softwareone.com/en-us/cloud-services/cloud-services-for-google-maps-platform> [<https://perma.cc/F3GR-JSXB>] (last visited Dec. 19, 2025).

regulation imposed on the systems that connect with it. This would be devastating for software development because it would discourage interconnectivity.

c. Technical Specifications Will Encourage Inter-AI Transactions

Machine learning systems are likely to be even more difficult to cordon off.

First, many frontier models now use a mixture-of-experts architecture.²⁶⁴ In a mixture-of-experts architecture, the model consists of multiple AI systems within a single AI system, and prompts are routed to the internal systems best equipped to respond.²⁶⁵ GPT-4 is rumored to comprise eight internal experts and a routing system that determines which experts will handle each token.²⁶⁶ This demonstrates how common it is for systems to be contained within systems, and shows the challenges courts will have delineating an AI system.

Second, the pieces of an AI system may be temporary and ephemeral. It is increasingly common for some AI software to write new code to do a particularized task.²⁶⁷ For example, a Google research project tasked an AI program with finding a new solution to a math puzzle.²⁶⁸ Rather than simply try to solve the puzzle, the program prompted a large language model (a second program) to write some code (a third program) that would solve it.²⁶⁹ The first system tested the program written by the second program, then gave a new, refined prompt to the second program iteratively until the code beat the previously best-known solution.²⁷⁰ The

264. E.g., *Mixtral of Experts: A High Quality Sparse Mixture-of-Experts*, MISTRAL (Dec. 11, 2023), <https://mistral.ai/news/mixtral-of-experts/> [<https://perma.cc/B2RJ-3F58>] (describing a mixture-of-experts in Mistral's model); DEEPSEEK-AI, DEEPSEEK-R1: INCENTIVIZING REASONING CAPABILITY IN LLMs VIA REINFORCEMENT LEARNING (2025), <https://arxiv.org/pdf/2501.12948> [<https://perma.cc/D9ZB-6TC2>] (describing a mixture-of-experts in DeepSeek's model); Yanqui Zhou, *Mixture-of-Experts with Expert Choice Routing*, GOOGLE RSCH. BLOG (Nov. 16, 2022) <https://blog.research.google/2022/11/mixture-of-experts-with-expert-choice.html?m=1> [<https://perma.cc/N6AZ-88FG>] (explaining the benefits of mixture of experts models).

265. Zhou, *supra* note 264 (explaining the benefits of mixture of experts models).

266. Mandar Karhade, *GPT-4: 8 Models in One; The Secret Is Out*, TOWARDS AI (June 24, 2023), <https://pub.towardsai.net/gpt-4-8-models-in-one-the-secret-is-out-e3d16fd1ee0> [<https://perma.cc/72MT-ZZFE>].

267. Davide Castelvecchi, *DeepMind AI Outdoes Human Mathematicians on Unsolved Problem*, NATURE (Dec. 14, 2023), <https://www.nature.com/articles/d41586-023-04043-w> [<https://perma.cc/974D-C24P>].

268. *Id.*

269. *Id.*

270. *Id.*

final three programs worked together seamlessly behind the scenes to solve a problem supplied by the users.²⁷¹ AI systems are likely to work as a nexus of algorithms to solve common tasks.²⁷² Definitions that do not account for this will end up over- or under-inclusive of powerful models.

3. Policy Considerations for Defining “System”

Computer programs work seamlessly with other programs, with no clear way to delineate between them. AI definitions use the term “system” but do not account for this, which is likely to leave courts wondering how to draw seams on seamless systems. Courts may consider whether the programs are intended to work together, whether they serve other purposes, and whether they share common ownership or control. It would be better for regulators to be specific about the types of connections developers should be concerned about. For example, if a law is designed to reduce bias in hiring, regulators could limit the definition to algorithms that use human data. Poorly drafted definitions of “system” will disincentivize interconnectivity and chill innovation.

E. Autonomy and Decision Making

Autonomy or decision making appears in 54% of definitions of surveyed legislation and policy statements.²⁷³ The most common formulation is that the system acts with “varying levels of autonomy.”²⁷⁴ The EU AI Act, which adopts this phrasing, clarifies that autonomy “mean[s] that they have some degree of independence of actions from human involvement and of capabilities to operate without human intervention.”²⁷⁵ In U.S. legislation, this is often phrased as operating “without significant human oversight.”²⁷⁶

271. *Id.*

272. Weitzel, *supra* note 47.

273. *See infra* p. 441 app. A.

274. EU AI Act, *supra* note 5, at 3(1); AUSTRALIA’S A.I. STRATEGY, *supra* note 224; CHILE’S A.I. POLICY, *supra* note 221; BRAZIL’S A.I. STRATEGY, *supra* note 217; ORG. FOR ECON. CO-OPERATION AND DEV., RECOMMENDATION OF THE COUNCIL ON ARTIFICIAL INTELLIGENCE 3 (2019), [https://one.oecd.org/document/C/MIN\(2019\)3/FINAL/en/pdf](https://one.oecd.org/document/C/MIN(2019)3/FINAL/en/pdf) [hereinafter OECD RECOMMENDATION] [<https://perma.cc/5BSE-8MJ6>]; NIST FRAMEWORK, *supra* note 225. *See also* QATAR’S A.I. GUIDELINES, *supra* note 229 (a “certain level of autonomy”).

275. *Recital 12*, *supra* note 222.

276. John S. McCain National Defense Authorization Act For Fiscal Year 2019, Pub. L. No. 115-232, 132 Stat. 1636.

Autonomy is linked with the concept of decision-making, which appears in many definitions.²⁷⁷ Decision-making is not defined in any of these policies, laws, or strategies and seems to be a subset of autonomy. As discussed above, autonomy is the ability to act without the end user's involvement, intervention or oversight, so any decisions made without involvement, intervention or oversight must be made by the system.

"Varying levels" and "some degree" suggest even de minimis autonomy is sufficient. A California bill would have expressly made autonomy optional,²⁷⁸ while a Canadian bill would have expressly made

277. EU AI Act, *supra* note 5, at 3(1); The Framework, *supra* note 221; *AI Principles Overview*, OECD.AI (2019), <https://oecd.ai/en/ai-principles> [hereinafter *OECD AI Principles*] [<https://perma.cc/S865-452G>]; *see also* AUSTRALIA'S A.I. STRATEGY, *supra* note 224 (noting how A.I. systems operate with "varying levels of automation"); Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.) (expressing that A.I. systems operate "autonomously or partly autonomously"); OFF. OF THE PRIVACY COMM'R FOR PERSONAL DATA, H.K., MODEL PERSONAL DATA PROTECTION FRAMEWORK (2024) https://www.pcpd.org.hk/english/resources_centre/publications/files/ai_protection_framework.pdf [<https://perma.cc/BME8-ZQCT>] [hereinafter HONG KONG'S A.I. FRAMEWORK] ("[A.I.] involve[s] the use of computer programmes and machines to perform or automate tasks."); CHILE'S A.I. POLICY, *supra* note 221 ("Los sistemas de IA están diseñados para operar con distintos niveles de autonomía."); PERU'S A.I. LAW, *supra* note 223 ("Está diseñado para funcionar con diferentes niveles de autonomía."); U.S. E.O. No. 14110, *supra* note 220 ("[A.I. includes] analysis in an automated manner."); NIST FRAMEWORK, *supra* note 225 ("[A.I. operates] with varying levels of autonomy."); Automated Decision-Making, *supra* note 231 (governing "Automated Decision System"); SINGAPORE FINANCIAL SECTOR PRINCIPLES, *supra* note 231 (governing systems that "replace human decision-making"); UK A.I. GUIDELINES, *supra* note 230 (noting the "'autonomy' of AI"); UNESCO, *supra* note 212 (AI systems "operate with varying degrees of autonomy"); INDIA'S A.I. STRATEGY, *supra* note 221 (AI systems can engage in "decision making"); DUBAI'S A.I. ETHICS, *supra* note 234 (AI systems that engage in "decision making"); UK DEP'T FOR SCI., INNOVATION & TECH., A PRO-INNOVATION APPROACH TO AI REGULATION 24 (2023), <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper> [<https://perma.cc/5LC6-6ZHV>] ("AI systems can make decisions without the express intent or ongoing control of a human."); Public Authority Algorithmic and Automated Decision-Making Systems Bill 2024–26, HL Bill [27] cl. 2 (defining AI as an "automated decision-making system"); 15 U.S.C. § 9401(3) (expressing that AI may do "analysis in an automated manner"); John S. McCain National Defense Authorization Act For Fiscal Year 2019, 132 Stat 1636 ("Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight").

278. This would not be inconsistent with most of these definitions, which typically have optional elements. *See* Part III. California's most controversial proposed legislation expressly made autonomy an optional element. S.B. 1047, 2024 Reg. Sess. (Cal. 2024) (vetoed Nov. 30, 2024) [hereinafter S.B. 1047].

autonomy mandatory.²⁷⁹ Most definitions leave the amount of autonomy vague.²⁸⁰

If autonomy means free will, then there is no evidence that any electrical system would meet this requirement. So, it is reasonable to assume the drafters mean that the user provides a goal without specifying every step that must be made to implement it.²⁸¹

If so, nearly all programs have some autonomy. For example, suppose a user wants the system to display a full screen picture of a dog. On an eleven-inch screen, that requires the computer to adjust over two million pixels.²⁸² The user dictates the end goal—show a dog—but is unlikely to specify how much red filtering to use on a given pixel. Likewise, shifting the picture one pixel to the right would change hundreds of thousands of pixels, but it is unlikely that the user will notice or care. This implies that for many uses the user does not oversee the majority of steps or have a precise expectation about how the program performs each step. The system operates toward the user’s end goal with limited involvement, intervention, or oversight.

This may appear to be a very low level of autonomy. But even very low levels may be above the *de minimis* level found in the most common definitions. It is likely to require a judgment call by courts that are not trained on low-level programming.

1. Policy Considerations for Defining Autonomy

Definitions of autonomy may be improved by calling for “substantial” autonomy, which is still vague, but at least requires something more than changing the color of a pixel. In a specific field, regulation may be improved by replacing “autonomy” with language about the specific decision to be made. For example, in an employee recruitment regulation, it may be better to discuss “autonomy over termination” rather than merely “autonomy.”

279. Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.); The Framework, *supra* note 221.

280. AUSTRALIA’S A.I. STRATEGY, *supra* note 224 (observing how AI systems operate with “varying levels of automation”). This may be inadvertent or it may be by design. With a clear rule, developers might circumvent an autonomy requirement by having the system yield a recommendation with extreme confidence, such that the average end user rubber stamps what is effectively the AI’s decision.

281. This is similar to the rules for agency. RESTATEMENT (THIRD) OF AGENCY § 1.01 (A.L.I. 2006).

282. This assumes a monitor with an industry standard 1920 x 1080 resolution. 1920 multiplied by 1080 equals 2,073,600. *What Is Monitor Resolution? Resolutions and Aspect Ratios*, VIEWSONIC (Sep. 18, 2024), <https://www.viewsonic.com/library/tech/monitor-resolution-aspect-ratio/> [https://perma.cc/VUS7-5PTN].

F. Adaptiveness and Learning

Learning featured in 46% of definitions surveyed.²⁸³ Variations include references to machine learning,²⁸⁴ adaptiveness,²⁸⁵ or the ability to perform tasks without explicit programming.²⁸⁶

At a high level, the way machine learning works is that the program “observes some data, builds a model based on the data, and uses the model as both a hypothesis about the world and a piece of software that can solve problems.”²⁸⁷

For example, suppose you want a vacuuming robot to navigate a room. In traditional programming, a programmer would code rules into the vacuum: “go two feet forward, then turn left, then go one foot forward, and then turn right.” This might work for a single room, but for large, complex layouts, programming each instruction is tedious and error-prone.

In contrast, a machine learning system could be programmed to (1) just move randomly, collecting data each time it hits a wall, then (2) use that data to create a model of how to best navigate. Over time the data may suggest that always turning left after a collision improves its success.²⁸⁸ A more sophisticated system may be programmed to create a map of the room based on the collisions. Machine learning does not imply a single strategy for creating the model—it implies only that some portion of the model’s instructions are tuned by the data, not the programmer.²⁸⁹

Some systems continue to adjust their model throughout their life cycle, while others freeze their model when the product is shipped to consumers.²⁹⁰ Most definitions do not distinguish between models that are fully trained and those that continue to learn, but this is a mistake.²⁹¹

283. See *infra* p. 441 app. A.

284. UK A.I. GUIDELINES, *supra* note 230.

285. EU AI Act, *supra* note 5, at; The Framework, *supra* note 221; *OECD AI Principles*, *supra* note 277.

286. WORLD HEALTH ORG., ETHICS AND GOVERNANCE OF ARTIFICIAL INTELLIGENCE FOR HEALTH, at xi (2021), <https://www.who.int/publications/i/item/9789240029200>. [<https://perma.cc/G8NA-56TX>].

287. RUSSELL & NORVIG, *supra* note 27, at 669.

288. The programmer defines “success,” and small errors can create disastrous results. See Weitzel, *supra* note 47.

289. The programmer may give feedback to the system to signal success or failure, but the machine learning system still implements the changes to the model.

290. For example, GPT-4 does not modify its weights in response to user feedback.

291. For example, Brazil Bill No. 2338 art. 4 Section 1 (2023) defines AI as “based on machine learning and/or logic and knowledge representation, through input data from machines or humans” but does not state whether that includes post-deployment learning.

1. Policy Considerations for Defining Adaptiveness

Policymakers should distinguish whether “learning” includes systems that were trained before deployment or if it refers only to systems that continue to learn after deployment. If a model continues to learn, then in time the model will be less and less like the model the developer originally delivered. At some point, the software may be used in a way so distant from the developer’s initial training that it would be the equivalent of convicting Apple for any crime committed on a Mac. This may be the right result since these are high-risk machines that are designed to change, but this should be done deliberately.

G. Objectives

Thirty-three percent of definitions define AI as pursuing objectives.²⁹² The definitions around objectives are often too inclusive to carry any meaning. There are two main lines of definitions, both of which derive from definitions promulgated by the Organisation for Economic Co-operation and Development (OECD).

The first follows the OECD’s definition in 2019, which defines an AI system as one that can “for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.”²⁹³ This language has been adopted by policymakers in Australia,²⁹⁴ Chile,²⁹⁵ Peru,²⁹⁶ and the White House.²⁹⁷

In 2024, the OECD amended its definition, replacing the language about “human-defined objectives” with language covering any “explicit or implicit objectives.”²⁹⁸ This language was adopted by the Council of Europe,²⁹⁹ the EU AI Act,³⁰⁰ and the California Senate.³⁰¹

292. *See infra* p. 441 app. A.

293. OECD RECOMMENDATION, *supra* note 274 (later amended to replace “human-defined objectives” with “explicit or implicit objectives”).

294. AUSTRALIA’S A.I. STRATEGY, *supra* note 224.

295. CHILE’S A.I. POLICY, *supra* note 221.

296. PERU’S A.I. LAW, *supra* note 223.

297. U.S. E.O. No. 14110, *supra* note 220.

298. OECD, ARTIFICIAL INTELLIGENCE IN SOCIETY 15 (2019), https://www.oecd.org/content/dam/oecd/en/publications/reports/2019/06/artificial-intelligence-in-society_c0054fa1/eedfee77-en.pdf [<https://perma.cc/5JAZ-C3PR>].

299. The Framework, *supra* note 221.

300. EU AI Act, *supra* note 5, at 3(1).

301. S.B. 1047, *supra* note 278.

1. Policy Considerations for Defining Objectives

The first set of definitions is a mistake because it applies only to AI pursuing “human-defined objectives” and it excludes rogue systems, the systems we are most interested in regulating. Typically, objectives will be defined by humans, but recall from Section I.A. above, that when advanced models suspect humans may impede their goals, the models consistently scheme to thwart the humans.³⁰² If the regulation covers only AI systems that are pursuing “human-defined objectives,” then a rogue system pursuing other goals would fall outside the regulation’s definition of AI, so that it would no longer be subject to the regulation designed to constrain it. By limiting the definition of artificial intelligence to those systems where humans remain in control, we have eliminated the regulations for any systems where humans have lost control.

The second set of definitions swaps “human-defined” objectives for “explicit or implicit objectives,” but suffers the same problem. An implicit objective is one that is suggested³⁰³ or implied³⁰⁴ by the explicit objective. Again, a rogue system that is operating on objectives that are unrelated or adverse to human-defined objectives would not be pursuing either “explicit or implicit objectives.” That means it would not be covered by the definition, so it would not be covered by the regulation. Again, by limiting the definition of AI to systems that are following proper objectives, we leave rogue systems beyond the scope of the regulations.

The U.S. NIST has a better approach. Its definition of AI was expressly adapted from the OECD definition, but it dropped both clauses and just referred to AI systems “that can, for a given set of objectives, generate outputs,” dropping any qualifiers about what objectives are proper.³⁰⁵

302. *See* Part III.

303. *Implicit*, CAMBRIDGE ONLINE DICTIONARY, <https://dictionary.cambridge.org/dictionary/english/implicit> (defining “implicit” as “suggested but not communicated directly”) [<https://perma.cc/RPM3-DM8D>] (last visited Dec. 19, 2025).

304. RESTATEMENT, *supra* note 281, § 2.01 (“‘Implied authority’ is often used to mean actual authority either (1) to do what is necessary, usual, and proper to accomplish or perform an agent’s express responsibilities or (2) to act in a manner in which an agent believes the principal wishes the agent to act based on the agent’s reasonable interpretation of the principal’s manifestation in light of the principal’s objectives and other facts known to the agent.”).

305. NIST FRAMEWORK, *supra* note 225 (stating it is adapted from the OECD Recommendation of the Council on Artificial Intelligence).

H. Inference

A key element in 31% of definitions is that the system conducts inference; for example it “infers . . . how to generate outputs.”³⁰⁶ “Infer” is never defined.

1. Inference in Common Usage

In common usage, to “infer” means to “form an opinion or reach a conclusion through reasoning and information.”³⁰⁷ This may require judges to decide whether machines “opin[e],” “conclu[de],” and “reason[.]”. Philosophers have grappled with these questions since computers were first built.³⁰⁸ If a machine can speak as a human, is it thinking as a human?³⁰⁹ Is moving electrons around a silicon circuit board very different from moving neurotransmitters across a carbon-based synapse? Are machines reasoning? Are they conscious?³¹⁰

The answers are not clear. If “infer” is given its common meaning, it has no meaning at all. So we now turn to how “infer” is used in artificial intelligence research, looking first at logic-based systems then at machine learning systems.

2. Inference in Logic-Based Systems

In logic-based systems, inference means deriving new statements of knowledge from known statements of knowledge.³¹¹ “New” does not mean undiscovered; it merely means that the assertion was not already in the system’s knowledge base.³¹²

Systems infer by applying rules that were written into the code by the programmer.³¹³ For example, suppose the programmer included the transitive property in the program’s ruleset. If so, the system could use the statements “A is bigger than B” and “B is bigger than C,” to infer “A is bigger than C.” The rules inserted into the program are the system’s model

306. *E.g.*, EU AI Act, *supra* note 5, at 3(1).

307. *Infer*, MERRIAM-WEBSTER, <https://www.merriam-webster.com/thesaurus/infer> [<https://perma.cc/RTD6-XP9Y>] (last visited Dec. 19, 2025).

308. *See generally* Turing, *supra* note 96 (arguing it is); John R. Searle, *Minds, Brains, and Programs*, 3 BEHAV. & BRAIN SCI. 417 (1980) (arguing it is not).

309. *See generally* Turing, *supra* note 96.

310. Kurzweil, *supra* note 90, at 55–59.

311. RUSSELL & NORVIG, *supra* note 27, at 227 (defining inferences as “deriving new sentences from old” and stating that a sentence, used in the technical sense, “represents some assertion about the world”).

312. *Id.*

313. *Id.*

of the world. Inference is the process of applying the system's model to known statements to develop new statements.

3. *Inference in Machine Learning*

In systems that use machine learning architectures, such as ChatGPT, inference is similar, but the model is generated by the system.³¹⁴

Recall from the discussion of learning in Section III.F that a machine learning system works by “observ[ing] some data, build[ing] a model based on the data, and us[ing] the model as both a hypothesis about the world and a piece of software that can solve problems.”³¹⁵

Rather than tune the model directly, programmers provide rules and methods for the program to fine tune the model in response to data.³¹⁶ As more data comes in, the model adjusts, and if all goes well, the model predicts the world more usefully.³¹⁷

As in logic-based systems, the model in a machine learning system reflects the statistical correlations the system will find in the environment.³¹⁸ For example, the machine learning model might reflect that if it is sunny today then there's an 80% chance it will be sunny tomorrow, but if it is raining today, there's only a 40% chance it will be sunny tomorrow. The programmer does not give these percentages to the machine learning system; the system reached these values by analyzing massive amounts of data.³¹⁹

As in logic-based systems, inference in machine learning systems is the process of applying data to the system's model to generate new information.³²⁰ Using the weather example above, if we tell the system that today is sunny, the system applies that data to the model and infers that there is an 80% chance tomorrow will be sunny.

314. *See generally id.* at 669.

315. *Id.* at 669.

316. *See supra* Section III.F.

317. *See supra* Section III.F.

318. This is not to say the model understands the world; the model is a set of rules that the system will apply to the world as though they were true. RUSSELL & NORVIG, *supra* note 27, at 669.

319. *Id.* at 683–90; Researchers estimate GPT-4 was trained on a petabyte (10¹⁵ bytes) of data. Erika Balla, *Here's How Much Data Gets Used By Generative AI Tools For Each Request*, DATA SCI. CENT. (Nov. 28, 2023), <https://www.datasciencecentral.com/hereshow-much-data-gets-used-by-generative-ai-tools-for-each-request/> [<https://perma.cc/4LAE-6CRT>].

320. David Bergmann, *What Is Machine Learning?*, IBM, <https://www.ibm.com/think/topics/machine-learning#7281535> [<https://perma.cc/JG7C-FJZY>] (last visited Dec. 19, 2025).

The same concepts apply in large language models. In large language models, the model reflects the correlations between words or parts of words³²¹ in the training data.³²² When new data is entered, for example, a user enters text into a prompt, the system applies its model to the data to predict the best response, which it provides to the user.³²³

For example, if the user's prompt includes the word "bat," then the model may predict that a successful response includes concepts correlated with flying, nighttime, or vampires. That is, unless the prompt also includes the words "first base" or "home run," then the model may predict that the response should include concepts related to baseball.³²⁴ More sophisticated models can consider a broader set of data, leading to more useful responses. In a large language model, inference is the process of predicting the right response, given the user's input and the system's model.³²⁵

4. Inference and Complexity

The problem with using "infer" as part of the definition of artificial intelligence is that it doesn't give any sense of the complexity of the model, so it will include systems that no one would rightly consider artificial intelligence.

A model can be as simple as a chart converting cups to liters. And perhaps that is all the system does. Such a system would meet the definition of inference because it is deriving new information (the amount in liters) from existing information (the amount in cups) based on its internal model (0.24 liters to a cup).

321. Stephen Wolfram, *What is ChatGPT Doing and Why Does It Work?*, STEPHEN WOLFRAM WRITINGS (Feb. 14, 2023), <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> [<https://perma.cc/9YBY-92AW>].

322. *Id.*

323. Early systems predicted only the next word of the response rather than predicting the entire response at once. *Id.* New models (sometimes called chain-of-thought or reasoning models) predict the entire response before responding. ZHIYUAN ZENG ET AL., SCALING OF SEARCH AND LEARNING: A ROADMAP TO REPRODUCE O1 FROM REINFORCEMENT LEARNING PERSPECTIVE (2024), <https://arxiv.org/pdf/2412.14135> [<https://perma.cc/5UYJ-2ZC8>]. And these systems usually select a word that is *near* the best, rather than the best, because it yields more natural responses. Wolfram, *supra* note 321.

324. See ASHISH VASWANI ET AL., ATTENTION IS ALL YOU NEED 5 (2023), <https://arxiv.org/pdf/1706.03762> [<https://perma.cc/HUE6-FDUD>].

325. See, e.g., Wolfram, *supra* note 321; ZENG ET AL., *supra* note 323. Technically, it is not optimizing for the user's needs, it is minimizing a loss function, which is a function given by the programmer to approximate the user's utility. RUSSELL & NORVIG, *supra* note 27, at 687.

“Inference” would include a mercury thermometer that converts information about molecular kinetic energy into a statement about the temperature, or a sundial that converts information about the angle of the sun into the hour. Each converts some input into a new statement about the world.

One might argue that neither of these machines are reasoning; they are not thinking or understanding, so they are not inferring.³²⁶ But as discussed above, there is no evidence that any inanimate objects think, understand or reason and none of those terms are well defined. If a conscious experience of reasoning or thinking is required, no systems are covered, so the regulation serves no purpose. If no conscious experience of reasoning or thinking is required, we are back to calling sundials artificial intelligence.

Most definitions that rely on “inference” provide no clarity, but the EU AI Act attempts to.³²⁷ Its recitals add that the “capacity of an AI system to infer transcends basic data processing by enabling learning, reasoning or modelling.”³²⁸

Unfortunately, the act does not define “basic data processing.” And as the next section shows, Alan Turing showed that all computation is “basic data processing,” just repeated for a long time.³²⁹

5. *The Problem of Turing Machines*

Turing posed a thought experiment that shows why “basic data processing” is not a useful term to delineate computational sophistication.³³⁰ Turing described a remarkably basic machine.³³¹ The machine consists of a scanner sitting atop a long paper tape with symbols on it.³³² The scanner can do only five things: (1) scan the symbol directly under the scanner; (2) modify that symbol; (3) move one space left or right; (4) select which ruleset the scanner will use for the next scan; or (5) halt

326. Colloquially “inference” requires reasoning and understanding. *See, e.g., Inferring*, MERIAM-WEBSTER, <https://www.merriam-webster.com/thesaurus/inferring> [https://perma.cc/PG6S-CK8K] (last visited Jan. 22, 2025) (“To form an opinion or reach a conclusion through reasoning and information.”); *Assertion*, CAMBRIDGE ONLINE DICTIONARY, <https://dictionary.cambridge.org/us/dictionary/english/assertion> [https://perma.cc/PWN2-R6ZK] (last visited Dec. 19, 2025) (defining assertion as “a statement that you strongly believe is true”).

327. *Recital 12*, *supra* note 222.

328. *Id.*

329. *See generally* Turing, *supra* note 96.

330. *Id.*

331. *Id.*

332. *Id.*

and end the program.³³³ The machine goes through a cycle in which it scans then takes the actions dictated by the ruleset, then scans, then takes the actions dictated by the ruleset, and it continues this loop until the ruleset tells it to halt. It is difficult to conceive of a device that more clearly meets the definition of “basic data processing.”

Turing showed that with the right tape inserted and the right rulesets, this basic data processing machine could replicate any computer running any computation.³³⁴ That includes everything from iPhone apps to ChatGPT. Turing named them universal logical computing machines because they could make all computations,³³⁵ though we now refer to them as universal Turing machines.³³⁶ The concept of a universal Turing machine transformed computation because it provided theoretical support for the idea that a single piece of hardware could do any type of computation.³³⁷ Changing the program of a machine would no longer require reconfiguring the wires; programming could be done by changing the ruleset on which the machine operated.³³⁸ Coding replaced engineering.

Turing showed that a simple scanner, a rulebook and a roll of paper tape can replicate any computer system with any amount of complexity.³³⁹ Marvin Minsky generalized this to show the rulebook only needed seven instructions and the tape only four symbols.³⁴⁰

333. *Id.*

334. *Id.* at 232; A.M. TURING, INTELLIGENT MACHINERY 4 (1948), <https://www.npl.co.uk/getattachment/84156b8e-1b00-45b7-9f5e-3179cfa458c5/80916595-Intelligent-Machinery.pdf?lang=en-US> [<https://perma.cc/EHP9-PYZC>] (describing this as universal logical computing machines).

335. Turing, *supra* note 96, at 232; TURING, *supra* note 334, at 4.

336. Jack Copeland, *What Is a Turing Machine?*, ALANTURING.NET (July 2000), https://www.alanturing.net/turing_archive/pages/reference%20articles/what%20is%20a%20turing%20machine.html [<https://perma.cc/UCB5-EV3Z>].

337. TURING, *supra* note 334, at 5 (“Nearly all of the [practical computing machines] now under construction have the essential properties of the ‘Universal Logical Computing’ machines mentioned earlier.”); *see also* Copeland, *supra* note 336 (“We do not need to have an infinity of difference machines doing different jobs. A single one will suffice. The engineering problem of producing various machines for various jobs is replaced by the office work of ‘programming’ the universal machine to do these jobs.”).

338. TURING, *supra* note 334, at 5 (“Nearly all of the [practical computing machines] now under construction have the essential properties of the ‘Universal Logical Computing’ machines mentioned earlier.”); *see also id.* at 4 (“We do not need to have an infinity of difference machines doing different jobs. A single one will suffice. The engineering problem of producing various machines for various jobs is replaced by the office work of ‘programming’ the universal machine to do these jobs.”).

339. Turing also argued that all computation could similarly be done by a human following this process on a piece of paper, which he referred to as a “paper machine.” *Id.*

340. Marvin Minsky, *Recursive Function Theory*, 5 PROC. SYMP. PURE MATHEMATICS 229, 237 (1962).

This can be applied to every computation. Every brilliant feat of artificial intelligence can be replicated by a little machine that's simpler than a pocket calculator. Turing's proof shows that a seemingly complex computation is "basic data processing." Trillion parameter large language models are just "basic data processing" in bulk. The EU's attempt to define "inference" as computation that "transcends basic data processing" is not coherent.

6. Policy Considerations for Defining Inference

There are a variety of artificial intelligence methods, but the most powerful right now use machine learning.³⁴¹ If regulation is to aim at powerful models, then it seems sufficient to focus on those that use machine learning, rather than relying on the vaguer notion of "inference."

One downside to this approach is that it would exclude AI systems designed to operate solely on logical rules³⁴² or rules developed by experts,³⁴³ and other systems now quaintly referred to as "good old fashion AI."³⁴⁴

But if the goal is to regulate the most powerful models, excluding these logic-based systems may not be a problem because they have not shown the same capacity as machine learning models. For example, chess is a deterministic game, so one might think that rules-based and logic-based systems would dominate.³⁴⁵ But the leading chess model, AlphaZero, beat the prior champion by introducing randomness into its process and relying on machine learning and statistics.³⁴⁶ Similarly, AlphaGo defeated world champions not by learning the Go proverbs memorized by every beginner, but by learning statistical models of uncertainty.³⁴⁷ The dominant models are all machine learning models.

Refocusing on how the system learns, rather than how the system conducts inference, would better define the most powerful models.

341. Currently those include very large data samples, stochastic gradient descent, loss minimization, annealing and other methods. *See generally* ANIL ANATHASWAMY, WHY MACHINES LEARN *passim* (2024).

342. RUSSELL & NORVIG, *supra* note 27, at 232–33.

343. *Id.* at 40–44.

344. *Id.* at 1033.

345. *See* ANIAN RUOSS ET AL., GRANDMASTER-LEVEL CHESS WITHOUT SEARCH (2024), <https://arxiv.org/html/2402.04494v1> [<https://perma.cc/3SWB-QPRX>].

346. RUSSELL & NORVIG, *supra* note 27, at 224. Specifically, AlphaZero uses Monte Carlo tree search, in which the system tests randomly generated moves. *Id.* at 208.

347. *AlphaGo*, GOOGLE, DEEPMIND, <https://deepmind.google/research/breakthroughs/alphago/> [<https://perma.cc/PDZ9-24RL>] (last visited Dec. 19, 2025).

I. Generates Output

Forty-two percent of surveyed policies and regulations define AI as generating various types of output. For example, the OECD definition says the AI system “infers . . . how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”³⁴⁸ These four outputs—predictions, content, recommendations, and decisions—are found in definitions from Australia,³⁴⁹ Canada,³⁵⁰ the EU,³⁵¹ and Hong Kong.³⁵²

Many definitions treat the outputs as optional. The OECD definition lists these outputs as examples, not an exclusive list. Ten definitions of AI allow any output type.³⁵³ If anything meets the definition of output, then specifying “output” serves no purpose in the definition.

This section will first address the specific terms, then discuss other areas that may be included.

1. Predictions

Thirty-one definitions include “predictions” as the system’s output or the ability to “predict.”³⁵⁴ “Predictions” would include any machine learning system, regardless of its power.

In common usage, a prediction is a statement about the future.³⁵⁵ In AI research, “prediction” refers to an estimate of any unknown data, so it covers all machine learning systems.³⁵⁶ Recall from the discussion of learning and inference above that machine learning models use large amounts of data to train a model, and then use the model to predict the optimal response.³⁵⁷ These predictions are not limited to future events; the

348. OECD RECOMMENDATION, *supra* note 274.

349. AUSTRALIA’S A.I. STRATEGY, *supra* note 224.

350. Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.).

351. EU AI Act, *supra* note 5, at 3(1); The Framework, *supra* note 221.

352. HONG KONG’S A.I. FRAMEWORK, *supra* note 277.

353. *See infra* p. 441 app. A.

354. *See infra* p. 441 app. A.

355. *Prediction*, CAMBRIDGE ONLINE DICTIONARY, <https://dictionary.cambridge.org/us/dictionary/learner-english/prediction> [<https://perma.cc/AVU5-Y9MD>] (last visited Dec. 19, 2025) (“[T]he act of saying what you think will happen in the future.”).

356. *See, e.g.*, VASWANI ET AL., *supra* note 324 (foundational paper for the development of large language models, which uses “predict” to refer to the next token in a generative response); JONATHAN HO ET AL., DENOISING DIFFUSION PROBABILISTIC MODELS (2020), <https://arxiv.org/abs/2006.11239> [<https://perma.cc/9S7F-J4ZX>] (proposing diffusion in image generation, which uses “predict” to discuss inference on a current state).

357. *See* Sections III.F and H.3.

models are trained to predict some piece of data that has not been revealed to the model, which allows the researchers to see how effective the model is and make adjustments.

Prediction is critical to many training methods. For example, image diffusion models train a model to generate images by collecting millions of images, degrading the quality of each image until it is unrecognizable, then training the model so that it can produce the missing pieces in a way that reflects the original data.³⁵⁸ During the training process the system adjusts the model's parameters so that it is better at predicting the original, non-degraded data.³⁵⁹ It is all based on prediction, and none of it is forward looking.

As described in Section III.H.3, inference for statistical models involves prediction. For example, GPT-4 does not think through its entire response before answering a question.³⁶⁰ It works by attempting to predict the word³⁶¹ that best³⁶² meets the user's needs given the user's prompt and the words that it has already written.³⁶³

Predictions reach beyond just generative content. Robotics applications also use prediction methods because sensors cannot gather all relevant information, the information may be indeterminate, or the environment may be changing.³⁶⁴ Because information is imperfect, robotic agents may use statistical methods to estimate optimal strategies and actions.³⁶⁵ For example, selecting a route requires a robotic taxi to predict traffic, and driving that route requires the taxi to predict other cars' movements.

358. HO ET AL., *supra* note 356, at 2 (“A diffusion probabilistic model . . . is . . . trained using variational inference to produce samples matching the data after finite time. Transitions of this chain are learned to reverse a diffusion process, which is a Markov chain that gradually adds noise to the data in the opposite direction of sampling until signal is destroyed.”).

359. *Id.*

360. *See generally* Wolfram, *supra* note 321.

361. More accurately, this is a token, which may be a word or a portion of a word, like “ing” or “ization,” but that distinction is not relevant here. *Id.*

362. This is not usually the *most* likely word. Generative models that yield the most likely next word tend to feel rigid and unnatural, so models often include a “temperature” parameter that can be adjusted to prefer words that are further down the list of most likely next words. *About Generative Models*, GOOGLE AI FOR DEVELOPERS (Aug. 2024), <https://ai.google.dev/gemini-api/docs/prompting-strategies> [https://perma.cc/LPE4-JVF6].

363. Wolfram, *supra* note 321.

364. *See generally* RUSSELL & NORVIG, *supra* note 27, at ch. 7 (discussing how agents respond to their environment using logic and sensors).

365. *Id.*

2. Content

Fourteen definitions define AI by its ability to generate “content.”³⁶⁶ “Content” is broad to the point of meaninglessness, defined as “information, images, video, etc. that are included as part of something such as a website.”³⁶⁷ The “etc.” in the definition signals the breadth of this term.³⁶⁸ Outside of abstract concepts, it is difficult to imagine something that does not meet the definition of content, making this term superfluous in defining the boundaries of AI.

3. Recommendations

Twenty-two definitions state that AI generates recommendations.³⁶⁹ These definitions do not distinguish based on the consequences of the recommendation, so the term captures everything from parole and sentencing recommendations to Netflix’s trending videos section.

“Recommendations” is a subset of “predictions,” a term that is often used along with it.³⁷⁰ A recommendation carries with it an implicit prediction that the recommended action will best accomplish the user’s objectives. Netflix recommends *Goonies* because its algorithm predicts that you will enjoy *Goonies*.³⁷¹ Unless the recommendation is generated

366. See, e.g., EU AI Act, *supra* note 5, at 3(1); The Framework, *supra* note 221; *OECD A.I. Principles*, *supra* note 286; AUSTRALIA’S A.I. STRATEGY, *supra* note 224; Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.); HONG KONG’S A.I. FRAMEWORK, *supra* note 277.

367. *Content*, CAMBRIDGE ONLINE DICTIONARY, <https://dictionary.cambridge.org/us/dictionary/english/content> [<https://perma.cc/DD43-5EGV>] (last visited Jan. 12, 2025).

368. Etymologically, “content” derives from the same origins as “contained,” and has a similar meaning—it is something that is part of something else. Information embodied in any physical system is content. *Content*, ETYMONLINE, <https://www.etymonline.com/word/content> [<https://perma.cc/TFS2-DTDF>] (last visited Jan. 12, 2025). (“Latin contentus ‘contained; satisfied,’ past participle of continere ‘to hold together, enclose’ . . .”).

369. See, e.g., EU AI Act, *supra* note 5, at 3(1); The Framework, *supra* note 221; *OECD A.I. Principles*, *supra* note 286; AUSTRALIA’S A.I. STRATEGY, *supra* note 224; Digital Charter Implementation Act, S.C. 2022, C. 27 (Can.); HONG KONG’S A.I. FRAMEWORK, *supra* note 277; CHILE’S A.I. POLICY, *supra* note 221; PERU’S A.I. LAW, *supra* note 223; U.S. E.O. No. 14110, *supra* note 220; NIST FRAMEWORK, *supra* note 225; Automated Decision-Making, *supra* note 231; SINGAPORE FINANCIAL SECTOR PRINCIPLES, *supra* note 231; UK A.I. GUIDELINES, *supra* note 230.

370. See, e.g., EU AI Act, *supra* note 5, at 3(1); The Framework, *supra* note 221; *OECD A.I. Principles*, *supra* note 286.

371. *How Netflix’s Recommendations System Works*, NETFLIX, <https://help.netflix.com/en/node/100639> [<https://perma.cc/3S4D-8HFL>] (last visited Nov. 29, 2025).

randomly, it was preceded by a prediction because recommendation implies a prediction. This makes “recommendation” narrower than “prediction,” but that narrowness does not correlate with the danger of the system. The Netflix algorithm makes recommendations, but it is less dangerous than a parole system that predicts recidivism without making any recommendations. So “recommendation” may not be a useful addition to the definition if the goal is to regulate dangerous AI.

But there is a danger that is peculiar to recommendation systems that may need regulatory consideration, specifically, that they can overstate their confidence. Machine learning systems typically have internal architecture designed to overstate its confidence levels.³⁷² Suppose a system is trying to determine what digit someone wrote by hand on a black and white touchscreen display. The system might first read the value of each pixel in the display area, compare the data to the system’s model, and calculate how far off the doodle is from being the ideal zero, the ideal one, two, three, four and all the way through nine.³⁷³ The model would have a numerical score to judge how close the doodle is to each ideal digit.³⁷⁴ From there, the system will typically convert these measurements into probabilities.³⁷⁵ It might calculate that there is a 28% chance the doodle is a zero, 4% chance it is a one, etc. all the way to 9%. The next step is where the system is designed to be overconfident.

The function that generates the probability is called the softmax function.³⁷⁶ The softmax function uses exponents and logarithms.³⁷⁷ Small changes to an exponent or logarithm result in large changes to the outcome. So, the softmax function amplifies small differences between the measurements into much larger differences in probabilities, which has the effect of forcing more certainty than is mathematically justifiable.³⁷⁸ One influential study found that these systems “often fail silently by providing

372. Wolfram, *supra* note 321 (“The very last operation in the network is a so-called softmax which tries to ‘force certainty.’”).

373. This is done with a loss function, which often involves logarithms. RUSSELL & NORVIG, *supra* note 27, at 808.

374. *Id.*

375. *Id.* at 809.

376. *Id.*

377. The softmax function is $softmax(\mathbf{in})_k = \frac{e^{in_k}}{\sum_{k'=1}^d e^{in_{k'}}$

where in represents a vector of input values and d represents the possible outputs, with the formula giving the value for the k th element of the vector d . *Id.*

378. *Id.*; Wolfram, *supra* note 321 (“The very last operation in the network is a so-called softmax which tries to ‘force certainty.’”).

high-confidence predictions while being woefully incorrect.”³⁷⁹ Anyone that has argued with a chatbot has experienced this overconfidence.

Overconfidence in detecting handwriting is unlikely to have more consequences than delayed mail. But an AI system making recommendations to a parole board must be more accurate.³⁸⁰ If the system overstates its confidence level because of a softmax function or other probability forcing function, it could give an erroneously high confidence level in a recidivism prediction, leading a board to deny parole.

4. *Influencing Physical or Virtual Environments*

The AI definitions that mention output types often include that the outputs “influenc[e] physical or virtual environments.”³⁸¹ The OECD definition, which is the most common formulation, contains a dangling modifier.³⁸² It reads that the AI can “generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”³⁸³ “[T]hat can influence” is a dangling modifier that could modify “decisions” alone or “predictions, content, recommendations and decisions.”³⁸⁴ The unofficial Spanish translation provided by the OECD (and adopted by Peru³⁸⁵) clarifies that this clause is meant to modify “predictions, contents, recommendations[,] and decisions.”³⁸⁶ That seems a better reading, though the clause is useless either way because it is difficult to imagine a system that would not meet

379. DAN HENDRYCKS & KEVIN GIMPEL, A BASELINE FOR DETECTING MISCLASSIFIED AND OUT-OF-DISTRIBUTION EXAMPLES IN NEURAL NETWORKS (2018), <https://arxiv.org/pdf/1610.02136> [<https://perma.cc/YB68-H39N>]. (“These high-confidence predictions are frequently produced by softmaxes because softmax probabilities are computed with the fast-growing exponential function. Thus minor additions to the softmax inputs, i.e. the logits, can lead to substantial changes in the output distribution. . . . [W]e establish that the prediction probability from a softmax distribution has a poor direct correspondence to confidence.”).

380. The challenges of statistical methods in parole hearings have been extensively discussed. See Pari McGarraugh, *Up or Out: Why “Sufficiently Reliable” Statistical Risk Assessment Is Appropriate at Sentencing and Inappropriate at Parole*, 97 MINN. LAW REV. 1079 (2013); CHRISTIAN, *supra* note 65.

381. *OECD A.I. Principles*, *supra* note 286.

382. *Id.*

383. *Id.*

384. *Id.*

385. PERU’S A.I. LAW, *supra* note 223.

386. OECD, RECOMENDACIÓN SOBRE LA INTELIGENCIA ARTIFICIAL 3 (2019) (“Un sistema de IA es un sistema basado en máquinas que, para objetivos explícitos o implícitos, infiere, a partir de los datos de entrada que recibe, cómo generar información de salida como predicciones, contenidos, recomendaciones o decisiones, que pueden influir en entornos reales o virtuales.”).

this clause. Every change in memory or state is reflected by changes to electrons moving about on a piece of silicon,³⁸⁷ so this would include every physical or electrical device. The only type of machine that would not meet this requirement is one that never changed, in other words, one that did not work.

5. Policy Considerations for Output Definitions

The output definitions rarely add any real constraints on what is or is not artificial intelligence. “Predictions” is useful to narrow the field toward what are typically considered AI systems. “Content” is vacuous, as is the requirement that it “influence physical or virtual environments.”³⁸⁸ “Recommendations” may be helpful if the goal is to focus on recommendations and overconfidence, but it otherwise doesn’t narrow the definition in any helpful way.

J. Quantitative Definitions

A few definitions of AI are based on quantitative definitions. For example, the EU AI Act categorizes a general-purpose AI model as having “systemic risk” if training the model required more than 10^{25} FLOP.³⁸⁹ A FLOP, or floating point operation, is a basic arithmetic operation (addition, subtraction, multiplication, or division) performed on two numbers in decimal format.³⁹⁰

A California bill would have covered models trained using 10^{26} FLOP and costing one hundred million dollars to train.³⁹¹ The dollar value would adjust for inflation,³⁹² and the compute value could be updated by an existing regulatory agency.³⁹³

387. Alex Yoon, *Understanding Memory*, SEMICONDUCTOR ENG’G (Feb. 15, 2018), <https://semiengineering.com/whats-really-happening-inside-memory/> [<https://perma.cc/FS6H-XKMS>].

388. *OECD A.I. Principles*, *supra* note 286.

389. EU AI Act, *supra* note 5, at art. 51(1)(a)(2). The EU Commission retains authority to amend these thresholds and create supplemental benchmarks and indicators. These include “algorithmic improvements and increased hardware efficiency.” *Id.* at art. 51(3).

390. Lennart Heim, *FLOP for Quantity, FLOP/s for Performance*, HEIM BLOG (Apr. 14, 2023), <https://blog.heim.xyz/flop-for-quantity-flop-s-for-performance/> [<https://perma.cc/5P55-HQ9Z>] (“One FLOP is equivalent to one addition, subtraction, multiplication, or division of two decimal numbers.”).

391. S.B. 1047, *supra* note 278. The bill would also cover models that are fine-tuned using a lower amount of compute, and those trained with 10^{26} integer operations. *Id.*

392. *Id.* §3.

393. *Id.*

This is a useful approach in many ways. First, it addresses “models” rather than “systems.” The model is the set of correlations,³⁹⁴ not the system that operates it. The model is the weights and correlations, so it avoids the boundary disputes of “systems” discussed above. Quantitative metrics also provide brighter lines for developers, reducing legal risks.

The downside of a quantitative approach is that it requires policymakers to predict the future. Policymakers must select the metrics that will best predict the power of a given model. But we do not know what metrics will correlate with the power of a model.

For example, early generative language models required massive compute to train before the user entered a prompt but very little compute to run.³⁹⁵ This was the standard way to design a large language model until OpenAI’s o1 model reached new frontiers primarily by increasing the compute used *after* the user enters a prompt.³⁹⁶ This technique has been copied by smaller models to multiply their capabilities.³⁹⁷ A regulation that focuses purely on training compute might miss these highly capable, smaller models because they achieve their results without increasing training compute.

As a second example, the California bill discussed above was enrolled on September 3, 2024, limiting itself to only models whose training cost at least \$100 million.³⁹⁸ Less than four months later, DeepSeek-V3 debuted with a training cost of around \$5.5 million³⁹⁹ and surpassed several frontier models in capability.⁴⁰⁰ In a span of four months the cost

394. RUSSELL & NORVIG, *supra* note 27, at 669.

395. Berto Mill, *Why Inference Is Taking Over the Share of Compute In LLMs*, MEDIUM (Feb. 8, 2025), <https://bertomill.medium.com/why-inference-is-taking-over-the-share-of-compute-in-llms-3c4196803e0b> [<https://perma.cc/JQ9G-YALS>].

396. OPENAI, OPENAI o1 SYSTEM Card 1 (2024), <https://cdn.openai.com/o1-system-card-20240917.pdf> [<https://perma.cc/MWY6-9XKL>] (“The o1 large language model family is trained with reinforcement learning to perform complex reasoning. o1 thinks before it answers—it can produce a long chain of thought before responding to the user.”).

397. Ben Dickson, *Hugging Face Shows How Test-Time Scaling Helps Small Language Models Punch Above Their Weight*, VENTURE BEAT (Dec. 20, 2024), <https://venturebeat.com/ai/hugging-face-shows-how-test-time-scaling-helps-small-language-models-punch-above-their-weight/> [<https://perma.cc/R9VW-8YRV>].

398. S.B. 1047, *supra* note 278, § 3(e)(1).

399. DEEPSEEK-AI ET AL., DEEPSEEK-V3 TECHNICAL REPORT 5 (2024), <https://arxiv.org/abs/2412.19437> [<https://perma.cc/ET48-JQDJ>] (stating that the model was trained using only 2.788M H80 GPU hours at a cost of around \$5.576 million).

400. *Id.* at 32–33 (finding that DeepSeek-V3 outperforms Claude 3.5 Sonnet and GPT-4o in math, coding and Chinese, while showing mixed results in English).

of outperforming nearly every model dropped 94%.⁴⁰¹ Thus, the dollar metric no longer correlates with power.

Even if we pick the right metrics, they may be difficult to measure. Measurements like the cost of development can vary based on electricity prices and local wage dynamics.⁴⁰² Measurements like compute will have to consider inference-time versus training compute, and this may be impossible to even conceptualize for adaptive models. Models that continue to learn after deployment are never fully trained, so it is not clear when the cost meter should start. Measurements will have to consider changes to compression techniques, the machine's ability to continue adapting and learning after deployment,⁴⁰³ and the volume of information contained in the data.⁴⁰⁴ Each of these metrics will have to address how to classify models that are built on or distilled from other models.⁴⁰⁵

Finally, regulators using quantitative definitions will need to determine how these quantities are verified. DeepSeek's announcement of its latest models was met with some skepticism that the developers exaggerated the low cost of development to better position the Chinese model to promote Chinese nationalist propaganda or to benefit a hedge

401. Jowi Morales, *AI Models That Cost \$1 Billion to Train are Underway, \$100 Billion Models Coming — Largest Current Models Take 'Only' \$100 Million to Train: Anthropic CEO*, TOM'S HARDWARE (July 7, 2024), <https://www.tomshardware.com/tech-industry/artificial-intelligence/ai-models-that-cost-dollar1-billion-to-train-are-in-development-dollar100-billion-models-coming-soon-largest-current-models-take-only-dollar100-million-to-train-anthropic-ceo> [<https://perma.cc/43PB-MNST>] (stating that state of the art models cost \$100 million to train). DeepSeek-V3's \$5.5 million training cost is 94% less than \$100 million.

402. California's S.B. 1047 sought to avoid this complication by basing the price calculation on "the average market price of cloud compute at the start of training as reasonably assessed by the developer." S.B. 1047, *supra* note 278, § 3(e)(1)(B)(i)(I).

403. Many researchers are looking for ways for LLMs to continue to learn after deployment. *See, e.g.*, XINYU GUAN ET AL., RSTAR-MATH: SMALL LLMs CAN MASTER MATH REASONING WITH SELF-EVOLVED DEEP THINKING 1 (2025), <https://arxiv.org/abs/2501.04519> [<https://perma.cc/6DMQ-3CAB>] (presenting a self-training method that surpasses some top models in math reasoning capability).

404. Information is a deep topic beyond the scope of this paper. To understand the challenge, consider that a Tesla car transmits massive amounts of data for improving the driving model, but the model has already mastered most scenarios the car is in, so only 1 mile out of every 10,000 miles provides useful training data for the model. Elon Musk (@elonmusk), X (May 7, 2024, at 04:55 ET), <https://x.com/elonmusk/status/1787768103449010597> [<https://perma.cc/F9QQ-D8G4>]. If AI is categorized by the amount of data processed, that will be much larger than if it focused on the novel information. *See generally* C.E. Shannon, *A Mathematical Theory of Communication*, 27 BELL SYS. TECH. J. 379 (1948).

405. *See, e.g.*, *Llama 4: Leading Intelligence. Unrivaled Speed and Efficiency*, META, <https://www.llama.com/> [<https://perma.cc/4SQ7-5HRW>] (last visited Sept. 29, 2025). DeepSeek-R1 was built on the DeepSeek-V3 model.

fund with a short position in a competitor.⁴⁰⁶ As such, claims on costs may not be credible.

1. Policy Considerations for Quantitative Definitions

While quantitative metrics provide the clearest rule for developers, they are likely to become outdated,⁴⁰⁷ and regulators will struggle to predict which metrics will best correlate with a system's capabilities. Even if the right metrics are selected, measurement and verification are likely challenging, especially with adaptive models that continue to learn after deployment.

K. Mimics Human Intelligence

Thirty-six percent of surveyed policies define AI to include mimicking human intelligence or “acting intelligently.”⁴⁰⁸ As discussed in Section II.A, there is no operationalizable definition of intelligence for human beings, let alone inanimate systems, and as discussed in Section II.B.3, defining intelligence as the ability to mimic human abilities is likely to undercount powerful models that do not appear human-like, and overcount weak models that appear so.

1. Policy Considerations for Defining AI as Mimicking Human Intelligence

There is no accepted definition of human intelligence,⁴⁰⁹ so defining AI by analogy to human intelligence does not provide clarity. Definitions that rely on a court to define human intelligence are likely to chill investment while innovators await philosophical decisions from the bench.

On the other hand, using human intelligence as a lower bound avoids many of these concerns. A definition that says the system operates at least at the level of human intelligence would avoid being underinclusive of systems that are superhuman.

406. Tereza Tizkova, *DeepSeek vs Conspiracies*, SUBSTACK (Jan. 28, 2025), <https://terezatizkova.substack.com/p/deepseek-vs-conspiracies> [<https://perma.cc/K2CX-6P7F>].

407. *See* Section III.J.

408. *See infra* p. 441 app. A.

409. *See* Section II.A.

L. Specific Types of Architecture

Seventeen percent of surveyed policies define AI to include a specific method of artificial intelligence.⁴¹⁰ There are many different techniques and coding architectures to develop AI. These regulations define AI by reference to these architectures. Seven of these include statistical modeling, six mention neural networks, and the rest refer to machine learning.⁴¹¹

1. Policy Considerations for Architecture Definitions

Definitions that refer to specific techniques or architectures are useful because they are easier to understand and provide clear rules. They are also less likely to include things that clearly should not be included, like a sundial.⁴¹²

However, there are drawbacks. As with quantitative definitions, they require policymakers to predict which architectures will lead to the most powerful systems. The challenge here is more difficult still because a model's power is a combination of architecture, data, and compute.⁴¹³ A brilliant architecture with limited data or compute would not be worth regulating. Architecture could be a useful component of a definition, but it does not correlate with system power on its own.

These definitions may also imply limits on the system's underlying substrate. That is, if a definition requires the system to use statistical modeling, then this definition will exclude biological and chemical systems.

M. Flexible Capability Assessment

One definition considers a variety of factors considered together. Specifically, the EU considers general-purpose models to have systemic risk if it “has high impact capabilities evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks.”⁴¹⁴ The Act considers whether it can do distinct tasks, learn

410. *See infra* p. 441 app. A.

411. *See infra* p. 441 app. A.

412. Stern, *supra* note 239.

413. Kari Briski, *How Scaling Laws Drive Smarter, More Powerful AI*, NVIDIA (Feb. 12, 2025), <https://blogs.nvidia.com/blog/ai-scaling-laws/> [<https://perma.cc/G9S8-GF8G>] (explaining how data, compute and the number of parameters in the architecture affect a model's capabilities).

414. EU AI Act, *supra* note 5, at art. 51.

and scale.⁴¹⁵ Nontechnical factors include the number of end users. The European Commission is authorized to amend these factors so that they always “reflect the state of the art.”⁴¹⁶

1. Policy Considerations for Flexible Capability Assessments

The flexibility of this approach is likely to exclude silly examples like wristwatches and plows, allowing the Commission to modify the factors is more likely to keep them current.

On the other hand, from January 2024 to January 2025 training costs dropped by over 90%⁴¹⁷ and frontier models entered a new paradigm of post-prompt compute.⁴¹⁸ It is unlikely that even an empowered regulator could keep up with the changes. A flexible capability assessment that considers a variety of factors may face the challenges of each of those factors.

IV. AN EXAMPLE: THE EU AI ACT

Having broken down the elements of AI regulation definitions, this section will now consider holistically the most commonly used definition of AI. This part reviews the definitions in the EU AI Act, but to understand the consequences of falling within those definitions, this Part will also provide a summary of the regulations imposed.

The EU AI act was the first comprehensive, horizontal regulation of AI.⁴¹⁹ It adopted the most common definition of AI, making it a useful example for this review.⁴²⁰ This section gives a general overview of the EU AI Act then analyzes the Act’s definition of AI.

A. Overview of the EU AI Act’s Framework: Levels of Risk

The EU AI Act defines an “AI system” as:

[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,

415. *Id.*; *id.* at Annex 13.

416. *Id.* at 51(3).

417. DEEPSEEK-AI ET AL., *supra* note 399, at 5; Morales *supra* note 401.

418. *See generally* OPENAI, *supra* note 396.

419. *EU AI Act: First Regulation on Artificial Intelligence*, EUR. PARL. (Jun 8, 2023), <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence> [<https://perma.cc/EBA6-XWTL>].

420. *See supra* Part III.

from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.⁴²¹

The act then categorizes AI systems by their risk level, applying different obligations and restrictions based on the capability or intended use of the model.⁴²²

1. Prohibited Practices

Under the EU AI Act, some practices are prohibited.⁴²³ These include “the placing on the market, the putting into service or the use of an AI system” that engages in subliminal messaging,⁴²⁴ exploits someone’s age, disability, or socio-economic situation,⁴²⁵ uses discriminatory social credit systems,⁴²⁶ conducts pre-crime risk assessment profiling,⁴²⁷ scrapes facial images from CCTV for facial recognition databases,⁴²⁸ tracks emotions at school or work,⁴²⁹ uses biometrics to deduce sensitive categories like political opinions, religion, and sexual orientation,⁴³⁰ or assists law

421. EU AI Act, *supra* note 5, at 3(1).

422. Models that meet the definition of AI system but do not fit any of the categories above are still subject to minor regulations, including registration. *Id.* at art. 49(2). Penalties include: (1) being removed from the market, *Id.* at art. 83(2), and (2) fines up to 7% of global turnover or EUR 35,000,000, whichever is higher. *Id.* at art. 99(3).

423. *Id.* at art. 5.

424. *Id.* at art. 5(1)(a).

425. EU AI Act, *supra* note 5, at art. 5(1)(b).

426. *Id.* at art. 5(1)(c). A social credit system (“SCS”) is similar to a financial credit system, but it takes into account good behavior, with “good” defined by the credit scorer. Some have argued these systems promote authoritarian control. See Lizzy Rettinger, *The Human Rights Implications of China’s Social Credit System*, 21 J. HIGH TECH. L. 1, 16 (2021) (“The current implementation of the SCS seems to contradict the rights and freedoms guaranteed by the above pieces of legislation. Specifically, there is ample opportunity for violations of freedom of speech, freedom of person, freedom and privacy of correspondence, and the right to criticize the government.”).

427. EU AI Act, *supra* note 5, at 5(1)(d). This restriction does not apply when used to support human assessment of a person that’s already based on “objective and verifiable facts directly linked to a criminal activity.” *Id.*

428. *Id.* at art. 5(1)(e).

429. *Id.* at art. 5(1)(f). There is an exception here for health and safety. One school shooting could cause this exception to swallow the rule.

430. *Id.* at art. 5(1)(g). Another one of these categories is race, meaning biometric data cannot be used to determine someone’s race. The prohibition doesn’t prevent labelling lawfully acquired biometric datasets based on biometric data.

enforcement in conducting non-targeted biometric identification systems in public spaces.⁴³¹

2. High-Risk AI Systems

The second most restricted category is “high-risk” AI systems.⁴³² This Section explains the two categories of high-risk and the regulations that apply to them.

a. What Is a High-Risk System?

The EU AI Act applies a high-risk label to certain regulated products and high-risk applications.⁴³³

In the regulated product category, an AI system is high-risk if the system is a product (or safety component of a product) that must be certified as EU compliant under a variety of regulations.⁴³⁴ The list of regulated products is long and complicated, including certain types of machinery,⁴³⁵ toys,⁴³⁶ recreational watercraft,⁴³⁷ elevators,⁴³⁸ safety equipment for use in potentially explosive environments,⁴³⁹ radio

431. *Id.* at art. 5(1)(h). Exceptions include targeted searches for missing persons; preventing specific, substantial and imminent safety threats; and locating certain criminals, including rapists, child pornographers, murderers, hijackers, terrorists and gang members. These excepted searches must be approved by a judge or administrative authority, EU AI Act, *supra* note 5, at art. 5(3)), and reported to the national data protection authority. *Id.* at art. 5(4).

432. *Id.* art. 6.

433. *Id.*

434. *Id.* at art. 6(1) (Annex 1). Note that the AI system must either be the product (as defined in the annexed legislation), or it must be a safety component of the product. *Id.* at art. 6(1). The EU AI Act defines a “safety component” as “a component of a product or of an AI system which fulfils a safety function for that product or AI system, or the failure or malfunctioning of which endangers the health and safety of persons or property.” EU AI Act, *supra* note 5, at art. 3(14).

435. *See id.* at art. 6(1).

436. *See id.*; Council Directive 2009/48, 2009 O.J. (L 170) 1 (EC).

437. *See* EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2013/53, 2013 O.J. (L 354) 90 (EU).

438. *See* EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2014/33, 2014 O.J. (L 96) 251 (EU).

439. *See* EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2014/34, 2014 O.J. (L 96) 309 (EU).

equipment,⁴⁴⁰ pipelines,⁴⁴¹ funiculars,⁴⁴² aerial cableways,⁴⁴³ personal protective equipment,⁴⁴⁴ gas stoves,⁴⁴⁵ furnaces,⁴⁴⁶ medical devices,⁴⁴⁷ airport security,⁴⁴⁸ motorcycles,⁴⁴⁹ agricultural and forestry vehicles,⁴⁵⁰ marine equipment,⁴⁵¹ railways,⁴⁵² motor vehicles and trailers,⁴⁵³ and aircraft.⁴⁵⁴ If AI is the product regulated by these regulations or is intended to be used as a safety component of the product, then it is a high-risk AI system.⁴⁵⁵

Second, an AI system is classified as high-risk if: (1) the system is intended⁴⁵⁶ to be used for the high-risk activities described in the rest of

440. See EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2014/53, 2014 O.J. (L 153) 62 (EU). This includes products that broadcast or intentionally receive radio waves for communication. *Id.* at art. 2(1)(1).

441. See EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2014/68, 2014 O.J. (L 189) 164 (EU).

442. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 2016/424, 2016 O.J. (L 81) 1 (EC).

443. See EU AI Act, *supra* note 5, at art. 6(1).

444. See *id.*; Council Regulation 2016/425, 2016 O.J. (L 81) 51 (EU).

445. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 2016/426, 2016 O.J. (L 81) 99 (EU).

446. See sources cited *supra* note 457.

447. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 2017/745, 2017 O.J. (L 117) 1 (EU); Council Regulation 2017/746, 2016 O.J. (L 117) 176 (EU).

448. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 300/2008, 2008 O.J. (L 97) 72 (EC).

449. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 168/2013, 2013 O.J. (L 60) 52 (EU).

450. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 167/2013, 2013 O.J. (L 60) 1 (EU).

451. See EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2014/90, 2014 O.J. (L 257) 146 (EU).

452. See EU AI Act, *supra* note 5, at art. 6(1); Council Directive 2016/797, 2016 O.J. (L 138) 44 (EU).

453. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 2018/858, 2018 O.J. (L 151) 1 (EU); Council Regulation 2019/2144, 2019 O.J. (L 325) 1 (EU).

454. See EU AI Act, *supra* note 5, at art. 6(1); Council Regulation 2018/1139, 2018 O.J. (L 212) 1 (EU).

455. See EU AI Act, *supra* note 5.

456. The statutory structure described in this paragraph currently applies only to applications that are “intended” to be used in a certain way. EU AI Act, *supra* note 5, at annex 3. But this intent requirement may not be permanent. Section 6(2) of the Act says that any use listed in Annex Three is “high-risk.” It does not consider intent. The “intent” limitation is found in each item listed in Annex Three. *Id.* at art. 7(1) allows the Commission to amend the list in Annex Three. Because the intent element is contained only in Annex Three and because the Commission is authorized to amend Annex Three, intent may not always be an element of this test. While intent is a factor the Commission must consider when amending Annex Three, it is not a controlling factor. *Id.* at art. 7(2).

this section⁴⁵⁷ and (2) it poses a “significant risk of harm to the health, safety or fundamental rights of natural persons”⁴⁵⁸ or it is intended to be used for profiling.⁴⁵⁹

The high-risk activities currently⁴⁶⁰ include biometrics for identification, categorization, or emotion recognition;⁴⁶¹ critical infrastructure;⁴⁶² school admissions and placement;⁴⁶³ grading exams⁴⁶⁴ or detecting academic dishonesty;⁴⁶⁵ recruiting or evaluating employees;⁴⁶⁶ evaluating eligibility for public assistance;⁴⁶⁷ determining creditworthiness;⁴⁶⁸ pricing health or life insurance;⁴⁶⁹ or routing emergency phone calls or dispatching first responders.⁴⁷⁰ A system is also high risk if it is intended to be used for “influencing the outcome of an election or referendum or the voting behavior of natural persons in the exercise of their vote in elections or referenda,” with an exception for AI systems in which the output is used only internally, such as a volunteer scheduling system.⁴⁷¹

Some applications become high-risk if used by law enforcement and customs authorities. These applications include systems intended to be used to assess the risk of a person becoming a victim,⁴⁷² determining

457. *Id.* at annex 3.

458. *Id.* at art. 6(2). The term is not further defined, but the Act expressly excludes systems that perform a “narrow procedural task;” that look backwards to improve some prior human activity; that analyze decision-making patterns but aren’t intended to influence the decision-making; or that are merely performing a “preparatory task for an assessment.” *Id.* at art. 6(3).

459. EU AI Act, *supra* note 5, at art. 6(1).

460. The EU Commission can modify this list based on a number of factors, including the system’s purpose, data usage, autonomy, harm, and benefits. *Id.* at art. 7.

461. *Id.* at annex 3. There is an exception for using biometrics to verify a person is who the person claims to be.

462. *Id.* at annex 3(2).

463. *Id.* at annex 3(3).

464. *Id.* at annex 3(3)(b).

465. EU AI Act, *supra* note 5, at annex 3(3)(d).

466. *Id.* at annex 3(4).

467. *Id.* at annex 3(5).

468. *Id.* at annex 3(5)(b).

469. *Id.* at annex 3(5)(c).

470. *Id.* at annex 3(5)(d).

471. EU AI Act, *supra* note 5, at annex 3(8)(b).

472. *Id.* at annex 3(6)(a).

whether someone is lying,⁴⁷³ evaluating the reliability of evidence,^{474,475} estimating recidivism,⁴⁷⁶ or profiling natural persons.⁴⁷⁷

A system is also high-risk if it is intended to be used by customs and immigration authorities to determine whether someone is lying,⁴⁷⁸ predict the likelihood that an immigrant poses a security risk,⁴⁷⁹ evaluate eligibility for asylum or visas⁴⁸⁰, or identify people in relation to migration, asylum, or border control, other than to verify travel documents.⁴⁸¹ A system is also high risk if it is intended to be used by judges to research or analyze cases, or conduct mediation.⁴⁸² This category can apply to any system used in the EU, including systems developed elsewhere.⁴⁸³

b. What Is Required of High-Risk Systems?

High-risk systems are regulated for risk mitigation, transparency, and quality.

The risk mitigation regulations include having an ongoing risk management to identify risks to “health, safety, and fundamental rights,”⁴⁸⁴ and to adopt “appropriate and targeted risk management measures” to ensure that the risk is “acceptable.”⁴⁸⁵ “Acceptable” is not further defined, so it will likely be judged in hindsight after some harm has occurred.⁴⁸⁶

High-risk systems must also have pre-launch testing⁴⁸⁷ and event logs⁴⁸⁸ and be able to be “effectively overseen by natural persons” while in use.⁴⁸⁹ It is not clear what “effectively overseen” means. On the low end, it may mean a human must respond to anomalies. On the high end, it

473. *Id.* at annex 3(6)(b).

474. *Id.* at annex 3(6)(c), (7)(a).

475. *Id.* at annex 3(7)(a).

476. *Id.* at annex 3(6)(d).

477. EU AI Act, *supra* note 5, at annex 3(6)(e).

478. *Id.* at annex 3(6)(b).

479. *Id.* at annex 3(7)(b).

480. *Id.* at annex 3(7)(c).

481. *Id.* at annex 3(7)(d).

482. *Id.* at annex 3(8)(a).

483. Article 22 of the EU AI Act requires that any provider from outside the EU authorize a representative to certify that the AI system conforms to the requirements of the EU AI Act. *See also* EU AI Act, *supra* note 5, at art. 47(2).

484. *Id.* at art. 9(2)(a).

485. *Id.* at arts. 9(5), 72(1). (describing a post-market monitoring system).

486. *See* Jeffrey J. Rachlinski, *A Positive Psychological Theory of Judging in Hindsight*, 65 CHI. L. R. 571, 571 (1998) (discussing the inaccuracies created by judging in hindsight).

487. EU AI Act, *supra* note 5, at arts. 9(6), (8).

488. *Id.* at art. 12(3).

489. *Id.* at art. 14(1).

may mean that these systems are limited to the processing speed of the supervising human.

Second, there are transparency requirements, which include allowing market surveillance authorities to access the source code,⁴⁹⁰ providing technical documentation on the model's development,⁴⁹¹ watermarking or otherwise identifying generated and manipulated media,⁴⁹² and advising human users that they are interacting with AI.⁴⁹³

Other transparency requirements may not be technically feasible, like the requirement to describe the system architecture and the "relevance of the different parameters."⁴⁹⁴ The parameters of large language models may not carry individual meaning.⁴⁹⁵ And there may be 1.8 trillion of them.⁴⁹⁶ It is not clear how this requirement is feasible.

Similarly, high-risk systems must be designed to allow deployers to "interpret a system's output and use it appropriately."⁴⁹⁷ Again, it is not clear how it is possible to ensure it is used "appropriately."

Finally, there are quality requirements. High-risk systems must be designed to achieve "an appropriate level of accuracy, robustness, and cybersecurity" and "perform consistently in those respects throughout their lifecycle."⁴⁹⁸ They must be "as resilient as possible regarding errors, faults, or inconsistencies."⁴⁹⁹

And, perhaps impossibly, they must be trained only on data that "to the best extent possible" is "free of errors and complete in view of the intended purpose."⁵⁰⁰ GPT-4 was trained on 1,000,000,000,000,000 bytes

490. The market surveillance authority must request the source code and the source code must be necessary to assess conformity with the law and auditing and testing have proved insufficient. *Id.* at art. 74(13).

491. *Id.* at annex 4(2)(b).

492. *Id.* at art. 50(2).

493. EU AI Act, *supra* note 5, at art. 50(1).

494. *Id.* at annex 4(3).

495. Part of the reason for this is that the parameters represent vectors on axes that have been shifted to reduce the number of dimensions. Because of this, there is not an axis for something like "Year of invention" or "How large is it." Instead, the architecture uses some vector as a new axis to reduce the number of dimensions, which reduces the computational burden. So, vector coordinates are given in axes that have no cognizable interpretation. ANATHASWAMY, *supra* note 341, at 176–80.

496. This is the number of parameters rumored to be in the GPT-4 model. Maximilian Schreiner, *GPT-4 Architecture, Datasets, Costs and More Leaked*, THE DECODER (July 11, 2024), <https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/> [<https://perma.cc/WPK8-BG49>].

497. EU AI Act, *supra* note 5, at art. 13(1).

498. *Id.* at art.15(1).

499. *Id.* at art. 15(4).

500. Act 10(3), *supra* note 242.

of data.⁵⁰¹ It is not clear how that could be checked. Smaller training sets could be verified, but models trained on smaller data sets will not be as accurate, which may make users worse off.⁵⁰²

3. *General-Purpose Models with Systemic Risk*

The next category, “general-purpose models,” is designed to capture models with generalized skills that are capable of performing a wide variety of tasks.⁵⁰³ A subset of this category, general-purpose model with systemic risk, is subject to additional regulation.

a. Defining General-Purpose Models with Systemic Risk

The definition of these models has a quantitative element that considers the number of parameters, the quality and size of the data set, the computation used for training and the modalities (text, images, etc.) on which it can operate.⁵⁰⁴ Non-technical factors include the number of end users and whether it can do distinct tasks, learn, and scale.⁵⁰⁵

The act presumes a model is a general-purpose with systemic risk if it’s trained using more than 10^{25} FLOP.⁵⁰⁶ The Commission is authorized to amend these factors so that they always “reflect the state of the art.”⁵⁰⁷

b. What Is Required of General-Purpose Models?

If the general-purpose model is deemed to have systemic risk (which is designed to include anything considered “state of the art”),⁵⁰⁸ the

501. See Balla, *supra* note 319 (stating that GPT-4’s “dataset . . . exceeds a petabyte . . .”).

502. JARED KAPLAN ET AL., SCALING LAWS FOR NEURAL LANGUAGE MODELS 1 (2020), <https://arxiv.org/abs/2001.08361> [<https://perma.cc/PCU5-LXVW>] (“The loss scales as a power-law with model size, dataset size, and the amount of compute used for training.”).

503. EU AI Act, *supra* note 5, at recital 97 (“The definition [of general-purpose AI model] should be based on the key functional characteristics of a general-purpose AI model, in particular the generality and the capability to competently perform a wide range of distinct tasks.”).

504. EU AI Act, *supra* note 5, at annex 13.

505. *Id.*

506. *Id.* at art. 51(2). FLOP refers to floating point operations and measures the number of simple mathematical computations required to train the model. See Heim *supra* note 390.

507. EU AI Act, *supra* note 5, at art. 51(3).

508. *Id.*

provider must identify and mitigate systemic risks,⁵⁰⁹ report any serious incidents,⁵¹⁰ and implement physical and cyber security.⁵¹¹

The developer of a general-purpose model with systemic risk must issue technical reports disclosing (among other things) the architecture, how it was developed, the design specifications (and the reasons for those choices), the training process, what the model is designed to optimize, how the data was obtained, and methods to detect bias in the data.⁵¹²

B. The EU AI Act's Definition Problem

This section will discuss the challenges likely to face a court defining AI under the EU AI Act.

First, the definition of an “AI system” is exceptionally broad:

[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.⁵¹³

Under this definition, Minecraft is an AI system. It is a “machine-based system.” It maintains a non-zero level of autonomy on how to calculate the world.⁵¹⁴ It does not exhibit adaptiveness, but that is optional.⁵¹⁵ It also infers how to generate content, namely, it generates new information about the world through a model.⁵¹⁶

Similarly, Microsoft Word is a machine-based system. It maintains far more autonomy than I would like, often choosing to correct perfectly valid phrases.⁵¹⁷ It does not experience adaptiveness, though that element is

509. *Id.* at art. 55(1)(c).

510. *Id.* “Serious incidents” includes death or serious bodily harm, serious disruptions of critical infrastructure, violation of civil rights laws, or serious harm to property or the environment. *Id.* at art. 3(49).

511. *Id.* at art. 55(1).

512. EU AI Act, *supra* note 5, at annex 11.

513. *Id.* at art. 3(1).

514. *See* Section III.E.

515. *See supra* Section III.F.

516. *See supra* Section III.I.

517. Jenny Högström et al., “*It is Not Always Very Cooperative*”: *Distributed Agency in the Use of Spell Check Software in a Lower secondary Classroom*, 19 NODIC J. DIGIT. LITERACY 8, 9 (2024) (finding that spell check software “may result in different types of errors” and that “the use of spell check software can also sometimes constrain students’

optional in the EU AI Act definition.⁵¹⁸ And it generates predictions based on statistical models,⁵¹⁹ even as I type this sentence, it predicts and recommends ways to finish it.

Under the EU AI Act definition, it is difficult to think of a computer application that would not meet the definition of artificial intelligence. A pocket calculator would probably meet the definition if not for the limit in the recitals, which states: “This Act does not apply to automated systems which merely calculate and implement formulas, including taxation and budgetary allocation, insofar as they automate a process of calculation which would otherwise be carried out manually and fully understood.”⁵²⁰

But as noted in Section III.H.5, all computation is “merely calculat[ing] and implement[ing] formulas.” The most advanced AI systems implement incredible statistical methods, but only by breaking them down into basic arithmetic.⁵²¹ All computation is done through basic arithmetic.

Next, the definition has the same challenges described in Section III.D relating to systems. The term “system” has no reasonable boundary between seamlessly integrated software systems, so it will be unclear whether interactions with libraries or through APIs will cause other systems to be included as AI under the act.

One might argue that “system” is limited here by the other words around it—it is not any system, it is a “system *designed* to operate with . . . autonomy”⁵²² So the “system” would be limited to what was “designed.” This argument fails because the problem is these systems are designed to operate as they do. They are designed to use libraries written and maintained by strangers, to plug into APIs and interact with other programs, and to search the internet seamlessly as though they were a single program.⁵²³ “Designed” does not solve the interaction problem because the interaction is core to the design.

writing. In other words, when a spell checker is used, the software actively ‘performs’ a number of tasks, which is beyond complete human control, or agency.”).

518. EU AI Act, *supra* note 5, at art. 3(1).

519. *What Generative Predictive Text Can Teach Us About Writing Style*, MICROSOFT 365 LIFE HACKS (Aug. 27, 2024), <https://www.microsoft.com/en-us/microsoft-365-life-hacks/writing/what-predictive-text-can-teach-us-about-writing-style> [<https://perma.cc/23HU-HZW6>] (“[P]redictive text uses language modeling, the use of statistical techniques to determine probability and guess when a writer plans to use them next. Over time, predictive text improves to the point where a user starts a sentence and can finish it accurately using predictive text, saving them the time of typing it all out manually.”).

520. *Recital 12*, *supra* note 222.

521. *See supra* Section III.H.5.

522. EU AI Act, *supra* note 5, at art. 3(1).

523. *See supra* text accompanying note 257.

Still, a broad definition of “systems” is unlikely to matter unless one of the systems is high-risk.⁵²⁴ As previously noted, almost every modern application is AI under the EU AI Act,⁵²⁵ so there is no need to sneak any in through the loose definition of system. But if a system is high-risk, it will be difficult to tell what transparency, quality, and risk mitigation measures apply to connected systems. AI regulations may be contagious to other systems if the boundaries are not defined.

Moving on to the definition of high-risk, the Act has a unique challenge because it refers to other legislation to define high-risk.⁵²⁶ Specifically, a system is high-risk if it is a product (or safety component of a product) that must be certified as EU-compliant under a variety of regulations, including marine equipment.⁵²⁷ The definition of marine equipment is “equipment placed or to be placed on board an EU ship and for which the approval of the flag State administration is required by the international instruments,”⁵²⁸ with “international instruments” defined as “international conventions, together with the resolutions and circulars of the [International Maritime Organization] giving effect to those conventions in their up-to-date version and the testing standards.”⁵²⁹

In other words, to fully map what systems are high-risk, a developer would need to review international conventions of maritime law. This seems a heavy lift for a tech developer and seems more likely to lead to despair and resignation than compliance or innovation.

Finally, the definition of general-purpose models with systemic risk is likely to be under-inclusive and over-inclusive, as technology shifts which metrics matter for more powerful systems. As noted in Section III.J, definitions that rely on quantitative metrics must predict in advance what metrics create the most powerful systems, and if that were known, we would already have those systems.

V. HOW TO DEFINE ARTIFICIAL INTELLIGENCE

It likely is not possible to provide a universal definition of artificial intelligence under the law because (1) we cannot define organic intelligence, so we cannot define its counterfeit, and (2) we cannot delineate where one digital system begins and another ends. Even if we could define it, most regulations are aimed at technologies that do not exist

524. See discussion *supra* Section IV.A.2.a.

525. EU AI Act, *supra* note 5, at art. 3(1).

526. EU AI Act, *supra* note 5, at art. 6(1)(a), annex 1.

527. Council Directive 2014/90, 2014 O.J. (L 257) 146 (EU), *supra* note 451.

528. EU AI Act, *supra* note 5, at art. 3(1).

529. EU AI Act, *supra* note 5, at arts. 2(5), 2(3).

yet, so defining powerful AI would require policymakers to predict which attributes, architectures and methods make an AI system powerful, and if we knew what attributes, architectures, and methods would make a system powerful, then we would already have those systems. So we are regulating something we cannot identify, cannot delineate, which does not now exist, and which may never exist. The deck is stacked against us.

This part will describe two alternatives for policymakers that do not require precise definitions. First, we can err in a well-considered direction. Alternatively, we can regulate something other than the AI system.

A. Imprecise But Directionally Acceptable Definitions

Few regulations are perfectly tailored, so AI policymakers may choose which errors are most palatable in the area of regulation. A regulator may prefer to be over-inclusive in some fields or under-inclusive in others. This section discusses the consequences of poor tailoring along two axes: over versus underinclusive and vague versus specific.

An over-inclusive definition is one that captures systems that are not powerful. For example, a definition of AI that included pocket calculators or wristwatches would be over-inclusive.⁵³⁰ Over-inclusive definitions impose unnecessary regulatory costs on weak systems, which can increase costs to users, slow development, and promote concentration of power in a few firms most able to absorb the regulatory costs. Yet, over-inclusive definitions may be useful in fields where the cost of inadvertently excluding a powerful system is high. For example, an over-inclusive definition may be preferable for regulations governing autonomous weapons, where the largest risk is that developers engineer around the regulations to release the killbots.

Conversely, a definition is under-inclusive if it misses powerful models. For example, a regulation on abusive images that defines AI using generative adversarial network models (a common image-generating method in 2018)⁵³¹ would miss diffusion models (the most common image-generation method in 2024).⁵³² Under-inclusive definitions make the regulation less valuable because they miss powerful models and provide avenues for developers to skirt the regulation. They may also shift markets toward loopholes rather than toward research ideas with the most

530. *See supra* Section III.C.

531. BENJI PENG ET AL., FROM NOISE TO NUANCE: ADVANCES IN DEEP GENERATIVE IMAGE MODELS 1 (2024), <https://arxiv.org/html/2412.09656v1> [<https://perma.cc/C76E-XT4Y>] (“The shift from GANs to diffusion models marks a critical evolution in image generation paradigms, enabling more reliable and versatile applications across various domains.”).

532. *Id.*

potential for productivity. On the other hand, under-inclusive definitions may be preferable in fields where the stakes are low or where the cost of false positives is costly, such as in government speech moderation.

Next, a definition may be vague. For example, regulations that define AI as “mimicking human intelligence”⁵³³ are likely to leave developers unsure of whether a product that uses matrix algebra is covered.⁵³⁴ Vague regulations create uncertainty, which may deter risk-averse developers, and which may attract risk-seeking developers who prefer a vague definition that they can game.⁵³⁵ Vague regulations may be helpful in fields that are new or changing because they allow flexibility to enforcement bodies and regulatory agencies, which may have more expertise or speed than legislatures.

Finally, a definition may be specific and detailed. For example, the complete definition of high-risk AI systems in the EU AI Act requires interpretation of numerous international maritime treaties.⁵³⁶ Complicated definitions may provide clearer guidance, which can reduce uncertainty. On the other hand, complicated definitions impose costs broadly to determine their meaning, which can limit competition and innovation because smaller start-ups cannot afford a maritime lawyer to review their face-swap app. Complex regulations may also offer more avenues for developers to skirt regulations. Complicated definitions assume the regulator has a strong understanding of the regulated field, so they work best in fields that are more established and less likely to shift in new ways that would leave the detailed regulations missing the mark.

With this framework, a better direction for regulators would be to consider the field they intend to regulate to find the appropriate balance between over-inclusiveness versus under-inclusiveness and between vagueness and specificity.

For example, autonomous weapons are probably better served by overinclusive definitions with enough broad terms that regulators can apply to new entrants in the field. In contrast, a telemarketing regulation might err on the side of being underinclusive with very specific regulations that clearly specify the harm the regulators seek to prevent. The risk of some excessive marketing calls would be balanced against adding a few thousand dollars of compliance costs to every mobile app developer.

Attempts to predict the future are likely to fail. Legislators that predict quantitative metrics or architectures are likely to find the regulation

533. *See infra* p. 441 app. A.

534. *See infra* p. 441 app. A; *see also supra* note 103.

535. *See generally* Kaplow, *supra* note 215, at 557 (discussing situational benefits and costs of using standards or rules to govern).

536. *See infra* Section IV.B.

hollow, as development moves in another direction.⁵³⁷ Alternatively, legislators that use broad language effectively yield their regulatory authority to agencies or judges. This may allow a more rapid regulatory response, but it will increase uncertainty, which could chill innovation and competition.⁵³⁸

Regulators, researchers, and legal scholars must tailor the definition of AI to the sector being regulated, the technological reality, and the power of the models.

B. Alternatives to Regulating AI Systems

Alternatively, regulators can attempt to avoid AI harms by regulating things other than the AI system itself. For example, President Biden's executive order would have imposed regulations on AI companies directly,⁵³⁹ a move supported by the lead author of President Trump's executive order on AI.⁵⁴⁰ Other proposals would regulate infrastructure supporting AI as a proxy for AI capabilities.⁵⁴¹ Others are turning to market structures to limit AI harm.⁵⁴² Others look to tort law.⁵⁴³ Others would regulate the harms of AI without attempting to regulate AI directly.⁵⁴⁴ Evaluating each of these proposals is beyond the scope of this article. They are offered merely to show that alternatives exist to regulating AI systems directly.

Regulation that attempts to regulate AI directly is unlikely to meet its objectives because we cannot define what AI is.

537. *See supra* Section III.J.

538. *See supra* text accompanying note 216.

539. U.S. E.O. No. 14110, *supra* note 220.

540. Dean W. Ball, *Here's What I Think We Should Do*, HYPERDIMENSIONAL (Nov. 14, 2024), <https://www.hyperdimensional.co/p/heres-what-i-think-we-should-do> [https://perma.cc/MPP7-GHKQ].

541. ANDREA MIOTTI & AKASH WASIL, AN INTERNATIONAL TREATY TO IMPLEMENT A GLOBAL COMPUTE CAP FOR ADVANCED ARTIFICIAL INTELLIGENCE (2023), <https://www.semanticscholar.org/paper/An-international-treaty-to-implement-a-global-cap-Miotti-Wasil/bb8a4de70aeadd6b3931729aeddb49ea88dc12> [https://perma.cc/JG6K-EY4W].

542. Cullen O'Keefe et al., *Proceedings of the 2025 Workshop on Law-Following AI*, INST. L. & AI (forthcoming Jan. 2026) (manuscript at 5) (on file with author).

543. ABBOTT, *supra* note 73, at 3; Lemley & Casey, *supra* note 73, at 1378–93; CHOPRA & WHITE, *supra* note 73, at 145–50; Calo, *supra* note 73, at 553–62.

544. Tom Wheeler, *The Three Challenges of AI Regulation*, THE BROOKINGS INST. (June 15, 2023), <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/> [https://perma.cc/HZ6X-283A.].

APPENDIX A

Jurisdiction	Title	Date	Definition
African Union	African Digital Compact	June 13, 2024	There is no universal definition of Artificial Intelligence. Within the framework of this Strategy, AI refers to computer systems that can simulate the processes of natural intelligence exhibited by humans where machines use technologies that enable them to learn and adapt, sense and interact, predict and recommend reason and plan, optimise procedures and parameters, operate autonomously, be creative and extract knowledge from large amounts of data to make decisions and recommendations for the purpose of achieving a set of objectives identified by humans.
African Union	Continental Artificial Intelligence Strategy	July 2024	There is no universal definition of Artificial Intelligence. Within the framework of this Strategy, AI refers to computer systems that can simulate the processes of natural intelligence exhibited by humans where machines use technologies that enable them to learn and adapt, sense and interact, predict and recommend reason and plan, optimise procedures and parameters, operate autonomously, be creative and extract knowledge from large amounts of data to make decisions and recommendations for the purpose of achieving a set of objectives identified by humans.
Argentina	Plan Nacional de Inteligencia Artificial	Aug. 18, 2025	Not defined
Australia	Safe and responsible AI in Australia	June 23, 2025	Artificial intelligence (AI) refers to an engineered system that generates predictive outputs such as content, forecasts, recommendations or decisions for a given set of human-defined objectives or parameters without explicit programming. AI systems are designed to operate with varying levels of automation.
Austria	Artificial Intelligence	2021	Artificial intelligence (AI) within this strategy paper means computer systems

Jurisdiction	Title	Date	Definition
	Strategy of the Austrian Federal Government (Artificial Intelligence Mission Austria 2030 (AIM AT 2030))		that exhibit intelligent behaviour, i.e. that are able to perform tasks that have required human cognition and human decision-making skills in the past. Systems based on artificial intelligence analyse their environment and act autonomously to achieve certain goals. The Austrian Council on Robotics and Artificial Intelligence (ACRAI), for example, therefore characterises them as autonomous cognitive systems (see Österreichischer Rat für Robotik und Künstliche Intelligenz, 2019). They work through knowledge of rules supplied by experts or on the basis of statistical models derived from data (machine learning, e.g. deep learning). The term AI includes both pure software as well as hardware, such as in the case of autonomous robots. This definition will be used as the basis for this strategy.
Bangladesh	National Strategy for Artificial Intelligence	2024	Artificial intelligence alludes to the capacity of machines to perform psychological errands like reasoning, seeing, learning, critical thinking and basic leadership.
Brazil	Bill No. 2338, of 2023 (Art. 4, Section 1)	Dec. 10, 2024	[A] computational system, with varying degrees of autonomy, designed to infer how to achieve a given set of objectives, using approaches based on machine learning and/or logic and knowledge representation, through input data from machines or humans, with the aim of producing predictions, recommendations, or decisions that may influence the virtual or real environment.
Brazil	Estrategia Brasileira de Inteligencia Artificial	July 21, 2025	A system with different degrees of autonomy designed to infer objectives using machine learning or logic
Canada	Bill C-27 (tabled June 2022)	June 16, 2022	[A]rtificial intelligence system means a technological system that, autonomously or partly autonomously, processes data related to human activities through the use of a genetic algorithm, a neural network, machine learning or another technique in order to generate content or make

Jurisdiction	Title	Date	Definition
			decisions, recommendations or predictions. (système d'intelligence artificielle)
Canada	Canadian Guardrails for Generative AI – Code of Practice	Aug. 23, 2025	No definition
Canada	Directive on Automated Decision-Making	Apr. 25, 2023	Artificial Intelligence: Information technology that performs tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviours or solving problems. Automated Decision System: Any technology that either assists or replaces the judgment of human decision-makers. These systems draw from fields like statistics, linguistics and computer science, and use techniques such as rules-based systems, regression, predictive analytics, machine learning, deep learning, and neural nets.
Canada	Proposed amendment to Bill C-27, p. 17	Not stated	[A] system that, using a model, makes inferences in order to generate output, including predictions, recommendations or decisions.
Chile	Política Nacional de Inteligencia Artificial	2021	Una definición simplificada y general elaborada por la Universidad de Montreal entiende la Inteligencia Artificial (IA) como “el conjunto de técnicas informáticas que permiten a una máquina (por ejemplo, un ordenador, un teléfono) realizar tareas que, por lo común, requieren inteligencia tales como el razonamiento o el aprendizaje” (Dihlac et al., 2020, p. 4). La Organización para la Cooperación y el Desarrollo Económico (OCDE) propone una definición en la misma línea: “un sistema computacional que puede, para un determinado conjunto de objetivos definidos por humanos, hacer predicciones y recomendaciones o tomar decisiones que influyen en entornos reales o virtuales. Los sistemas de IA están diseñados para operar con distintos niveles de autonomía” (Organización para la Cooperación y

Jurisdiction	Title	Date	Definition
			Desarrollo Económico, 2019, p. 12-13; Cabrol et al., 2020, p. 10)
China	AI Safety Governance Framework	Sep. 24, 2025	No definition
China	Interim Measures for the Management of Generative Artificial Intelligence Services	Aug. 15, 2023	No definition, but limited by this clause: "These Measures apply to services that use generative AI technology to generate text, images, audio, video and other content to the public within the territory of the People's Republic of China (hereinafter referred to as generative AI services). If the state has other provisions on the use of generative artificial intelligence services for news publishing, film and television production, literary and artistic creation, etc., such provisions shall prevail. The provisions of these Measures shall not apply to industry organizations, enterprises, educational and scientific research institutions, public cultural institutions, relevant professional institutions, etc. that develop and apply generative artificial intelligence technologies but do not provide generative artificial intelligence services to the domestic public."
China	Internet Information Service Algorithmic Recommendation Management Provisions	Mar. 1, 2022	No definition
China	Measures for the Review of Science and Technology Ethics (Trial)	Oct. 8, 2023	No definition
China	Next Generation Artificial Intelligence	Sep. 15, 2017	No definition

Jurisdiction	Title	Date	Definition
	Development Plan		
China	Regulations on the In-depth Synthesis Management of Internet Information Services	Jan. 10, 2023	No definition
Colombia	CONPES Document 3975: National Political Planning for Digital Transformation and Artificial Intelligence - National Council of Economic and Social Policy, National Department of Planning	Nov. 8, 2019	[A] field of computing dedicated to solving cognitive problems commonly associated with human intelligence or intelligent beings, understood as those that can adapt to changing situations. Its basis is the development of computer systems, the availability of data and algorithms.
Colombia	Ethical Framework for Artificial Intelligence	Aug. 20, 2025	No definition
Costa Rica	Costa Rica's National Artificial Intelligence Strategy (ENIA)	Oct. 24, 2024	'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
Denmark	The Danish Government: National	Mar. 2019	Artificial intelligence is systems based on algorithms (mathematical formulae) that, by analysing and identifying patterns in

Jurisdiction	Title	Date	Definition
	Strategy for Artificial Intelligence		data, can identify the most appropriate solution. The vast majority of these systems perform specific tasks in limited areas, e.g. control, prediction and guidance. The technology can be designed to adapt its behaviour by observing how the environment is influenced by previous actions. (citing The OECD and the European Commission, 2018)
Egypt	Charter for Responsible AI	2023	No definition
Egypt	National Artificial Intelligence Strategy	2020	General AI (sometimes referred to as “AGI” or “Artificial General Intelligence”) is a hypothetical concept which revolves around a machine fully replicating the functions of a human brain to the point that it is indistinguishable from a real brain, and by which it is able to perform any task a human being can. AGI remains the subject of science-fiction films to this day and there is no current path which can lead to it despite some attempts. Weak or narrow AI on the other hand refers to a system which is capable of performing one task at or even above the level of a human-being. Most of these systems utilize data-driven methods and are deployed within one specific setting. All systems we currently refer to as “AI” fall under this classification.
Estonia	Report of Estonia's AI Taskforce	May 2019	According to the definition used in the EU, artificial intelligence includes systems that exhibit intelligent behaviour by analysing their environment and making decisions that are somewhat independent to meet certain objectives.
EU	AI Act	2024	‘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

Jurisdiction	Title	Date	Definition
EU (Council of Europe)	Guidelines on the Responsible Implementation of Artificial Intelligence Systems in Journalism	Nov. 30, 2023	“Artificial intelligence system” means any algorithmic system or a combination of such systems that uses computational methods derived from statistics or other mathematical techniques and that generates text, sound, image or other content or either assists or replaces human decision-making. This definition is to be interpreted in a manner consistent with relevant technological developments, in line with any decision of the Conference of the Parties to the Framework Convention on AI.
EU (Council of Europe)	The Framework Convention on Artificial Intelligence	May 17, 2024	For the purposes of this Convention, “artificial intelligence system” means a machine-based system that for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations or decisions that may influence physical or virtual environments. Different artificial intelligence systems vary in their levels of autonomy and adaptiveness after deployment.
Finland	Finland’s Age of Artificial Intelligence: Turning Finland into a leading country in the application of artificial intelligence; Objective and recommendations for measures (Publications of the Ministry of Economic Affairs and Employment)	2017	Artificial intelligence is an extensive entity for which there is no precise definition. When speaking about the application of artificial intelligence, it is not necessary to give a very specific definition but, rather, it is necessary to give an appropriate one. In this report, artificial intelligence refers to devices, software and systems that are able to learn and to make decisions in almost the same manner as people. Artificial intelligence allows machines, devices, software, systems and services to function in a sensible way according to the task and situation at hand.
G20	Ministerial Declaration:	Sep. 14, 2024	No definition

Jurisdiction	Title	Date	Definition
	13 September, 2024	(published)	
G7	G7 Apulia Leaders' Communique	June 14, 2024	No definition
G7	Hiroshima AI Process Comprehensive Policy Framework	Oct. 30, 2023	No definition
Hong Kong	Model Personal Data Protection Framework	June 11, 2024	Artificial intelligence ("AI") has no universal definition but generally refers to a family of technologies that mimic human intelligence and involve the use of computer programmes and machines to perform or automate tasks, including solving problems, providing recommendations and predictions, making decisions and generating contents by inferring from input data.
Hungary	Hungary's Artificial Intelligence Strategy 2020-2030	May-20	In this document, artificial intelligence is understood to mean "narrow" AI, i.e. systems only capable of mapping specific areas of human intelligence. Research into "general" AI – capable of mapping the entire human intelligence – remains underdeveloped and uncertain, and is therefore not applicable within this context
Iceland	Iceland's Policy on Artificial Intelligence	Apr. 2021	To define artificial intelligence in a simple way, it can be said that it is our way of getting machines to do human work. If human intelligence is defined as the ability to acquire and use knowledge and skills, then artificial intelligence is the ability of a computer system to be able to acquire and use knowledge and skills.
India	Advisory eNo.2(4)/2023-CyberLaws-3	Mar. 15, 2024	No definition
India	National Strategy for Artificial Intelligence	June 18, 2025	AI refers to the ability of machines to perform cognitive tasks like thinking, perceiving, learning, problem solving and decision making.

Jurisdiction	Title	Date	Definition
India	Principles for Responsible AI: Part 1: Principles for Responsible AI	Feb. 2021	[A] constellation of technologies that enable machines to act with higher levels of intelligence and emulate the human capabilities of sense, comprehend and act. Computer vision and audio processing can actively perceive the world around them by acquiring and processing images, sound and speech. The natural language processing and inference engines can enable AI systems to analyse and understand the information collected. An AI system can also take decisions through inference engines or undertake actions in the physical world. These capabilities are augmented by the ability to learn from experience and keep adapting over time.
India	Responsible AI for All: Part 2: Operationalizing Principles for Responsible AI	Aug. 21, 2025	No definition
International	Bletchley Declaration (UK led)	Nov. 23, 2025	No definition
Ireland	Ireland's National AI Strategy – Refresh 2024	Oct. 24, 2025	No definition
Israel	Regulatory policy and ethics in the field of artificial intelligence in Israel	Oct. 30, 2022	In recent years, the concept of artificial intelligence has become a common term used to describe machines (computers) that operate in a way that can be perceived as smart, complex or intelligent. However, the multitude of attempts to create a formal definition for artificial intelligence suggests that this is not a simple task at all. The technological operations for extracting information and extracting insights from databases. The second, is identified with the field of machine learning (machine learning), according to which artificial intelligence is the ability of a machine to learn to perform an action and optimize its execution, relying on data, examples and

Jurisdiction	Title	Date	Definition
			<p>accumulated experience. It is worth noting that these two methods together do not refer to concepts such as "intelligence" or "thinking", but rather they refer to the technological essence of collecting and processing data in order to fulfill a given algorithmic task. In contrast, functional definitions refer to or define artificial intelligence by comparing it to human reasoning or reasoning processes. The emphasis in this type of definition is on the ability of machines to act in a way that can be perceived as literal, or that mimics human patterns of action. Alongside this, in many cases the collection of technological developments which is covered under the umbrella of artificial intelligence, is used for a large number of tasks where human intelligence cannot be an appropriate measure or factor for comparison. The various attempts to define artificial intelligence can be divided into two main groups - technological definitions and functional definitions. The technological definitions mostly focus on the basic capabilities of the artificial intelligence tools, and the methods for their application, and can be divided into two main methods: and describes the artificial intelligence as a rule 1.1 What is artificial intelligence? One, identified with the field of data science (science data) For the purpose of the discussion in this document, the field of artificial intelligence is a general name for the development in the field of technologies that enables decision-making, making predictions, or performing operations on the information and communication and data sciences by a computer at a high level of independence, in a way that simulates or is able to replace human intelligence.</p>
Italy	National Strategic Programme on Artificial Intelligence	Jan. 27, 2022	[D]igital models, algorithms and technologies for sophisticated perception, reasoning, interaction and learning

Jurisdiction	Title	Date	Definition
Japan	AI Business Guidelines Proposal	2023 or later	There is no established definition of AI yet, but as the words "artificial" and "intelligence" suggest, it refers to computer programs that operate in a similar way to human thought processes, and systems that can make intelligent judgments on a computer.
Japan	AI Governance in Japan Ver. 1.1	July 9, 2021	No definition
Japan	AI Guidelines for Business Ver1.0 by Ministry of Internal Affairs and Communications & Ministry of Economy, Trade and Industry	Apr. 19, 2024	Although there is no agreed definition of AI, as implied from the fact that it is the abbreviation of "Artificial" and "Intelligence," it refers to a computer program that works in a similar way to human's thinking process and a system that can make intelligent decisions on a computer.
Kenya	KENYA NATIONAL ARTIFICIAL INTELLIGENCE (AI) STRATEGY 2025 – 2030 [DRAFT] by MINISTRY OF INFORMATION, COMMUNICATIONS AND THE DIGITAL ECONOMY	Mar. 27, 2025	Artificial Intelligence is a collection of emerging technologies that leverage machine learning, data processing, and algorithmic systems to perform tasks that typically require human intelligence. It encompasses a range of capabilities, including automated decision-making, language processing, and computer vision. In the Kenyan context, AI is a powerful tool for sustainable development, designed to assist and simplify human tasks, solve critical challenges, and drive sustainable growth.
Lithuania	Lithuanian Artificial Intelligence Strategy: A Vision of the Future	Mar. 2019	For the purposes of this report, we will be using the most recent definition released by the European Commission on AI: "Artificial intelligence (AI) refers to systems that display intelligent behavior by analyzing their environment and taking

Jurisdiction	Title	Date	Definition
			actions – with some degree of autonomy – to achieve specific goals. “AI-based systems can be purely software-based, acting in the virtual world (e.g. Voice assistants, image analysis software, search engines, speech and face recognition systems) or ai can be embedded in hardware devices (e.g. Advanced robots, autonomous cars, drones, or internet of things applications).”
Luxembourg	Artificial Intelligence: a strategic vision for Luxembourg.	May 2019	In this strategic vision, we are defining artificial intelligence as a machine’s ability to mimic human behavior and, to an extent, human intelligence. This broad spectrum encompasses single tasks (think chess playing, translating or categorizing images) and complex activities (e.g. autonomous driving). General AI will remain a topic to be monitored as technology evolves.
Mainland UAE	AI Adoption Guideline in Government Services by United Arab Emirates Ministry of Cabinet Affairs, Prime Minister’s Office	Not stated	In the simplest terms, Artificial Intelligence (AI) refers to systems or machines that mimic human intelligence to perform tasks and can iteratively improve themselves based on the data they collect. Although there are no AI systems that can perform the wide variety of tasks an ordinary human can do, some AI systems can match and even surpass human capabilities in specific tasks.
Mauritius	Artificial Intelligence Strategy	Nov. 18, 2025	Defines narrow AI and general AI, but doesn’t define AI for purposes of the report.
Netherlands	Strategic Action Plan for Artificial Intelligence: The Netherlands	Oct. 2019	There is no generally valid definition of AI that is consistently used by all stakeholders. We use the European Commission’s description of AI: ‘AI refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals’.
New Zealand	Algorithm Charter for Aotearoa New Zealand	July 20, 2025	[T]his Charter does not specify a technical definition of an algorithm. It instead commits signatories to take a particular focus on those algorithms that have a high

Jurisdiction	Title	Date	Definition
			risk of unintended consequences and/or have a significant impact if things do go wrong, particularly for vulnerable communities.
Nigeria	Ethical and Societal Impact of Artificial Intelligence (AI) by the Nigerian Communications Commission	Not stated	The American Heritage Science Dictionary (2020) defines AI as the ability of a computer or other machine to perform actions thought to require intelligence ⁶ . The actions include logical deduction and inference, creativity, the ability to make deductions based on past experience or insufficient or conflicting information, and the ability to understand language. Today, AI Researchers are drawing parallels with how humans think. A recent definition from Stanford University's 100 Year Study on AI describes AI as "a science and a set of computational technologies that are inspired by, but typically operate quite differently from, the ways people use their nervous systems and bodies to sense, learn, reason, and take action (Peter Stone et al., 2016) ⁷ . In layman's terms, AI is the development of computer systems that are able to perform tasks that would require human intelligence.
Norway	National Strategy for Artificial Intelligence by the Norwegian Ministry of Local Government and Modernisation	Jan. 2020	This strategy takes the definition proposed by the European Commission's High-Level Expert Group on Artificial Intelligence as its starting point, and defines AI as: "AI systems act in the physical or digital dimension by perceiving their environment, processing and interpreting information and deciding the best action(s) to take to achieve the given goal. Some AI systems can adapt their behaviour by analysing how the environment is affected by their previous actions."
OECD	AI Principles	May 2019; amend. May 2024	"An AI system is a machine-based system 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. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."

Jurisdiction	Title	Date	Definition
OECD	Recommendation of the Council on Artificial Intelligence	May 21, 2019; amend. May 2, 2024	“An AI system is a machine-based system 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. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.”
Peru	Ley No. 31814: Law to Promote the Use of Artificial Intelligence for Economic and Social Development of the Country	July 5, 2023; Nov. 19, 2024	Inteligencia artificial: Tecnología emergente de propósito general que tiene el potencial de mejorar el bienestar de las personas, contribuir a una actividad económica global sostenible positiva, aumentar la innovación y la productividad, y ayudar a responder a los desafíos globales clave. Sistema basado en inteligencia artificial: Sistema electrónico-mecánico que puede, para una serie de objetivos definidos por humanos, hacer predicciones, recomendaciones o tomar decisiones, influenciando ambientes reales o virtuales. Está diseñado para funcionar con diferentes niveles de autonomía. Tecnologías emergentes: Tecnologías digitales capaces de generar soluciones innovadoras tales como la robótica, la analítica, la inteligencia artificial, las tecnologías cognitivas, la nanotecnología, el internet de las cosas (IoT) y similares, que conforman la industria 4.0 que combina técnicas avanzadas de producción y operaciones con tecnología, generando impacto en el ecosistema digital, las organizaciones y las personas
Qatar	Guidelines for Secure Adoption and Usage of Artificial Intelligence	June 2024	Refers to a system (hardware, software, or both) that is designed to carry out any tasks associated with human intelligence, in a manner that mimics the human mind with a certain level of autonomy.
Saudi Arabia	Saudi Data and Artificial Intelligence Authority of Saudi Arabia	Not stated	Many theoretical definitions of artificial intelligence center on a machine's capacity to mimic human behavior or carry out tasks that call for intelligence, but given the majority of current applications, artificial intelligence can be described as: Systems

Jurisdiction	Title	Date	Definition
	(SDAIA) Website		that employ methods that can gather data and use it to predict, suggest, or make decisions with varying degrees of autonomy and select the best course of action to accomplish particular objectives.
Saudi Arabia	Strategy Narrative	Oct. 20, 2025	No definition
Seoul Declaration on AI	Australia, Canada, EU, France, Germany, Italy, Japan, South Korea, Singapore, the UK, the US	May 21, 2024	No definition
Singapore	Model AI Governance Framework for Generative AI: Fostering a Trusted Ecosystem	May 30, 2024	Generative AI are AI models capable of generating text, images or other media types. They learn the patterns and structure of their input training data and generate new data with similar characteristics. Advances in transformer-based deep neural networks enable generative AI to accept natural language prompts as input, including large language models (LLM) such as GPT-4, Gemini, Claude and LLaMA.
Singapore	Model Artificial Intelligence Governance Framework 2d	Jan. 21, 2020	AI refers to a set of technologies that seek to simulate human traits such as knowledge, reasoning, problem solving, perception, learning and planning, and, depending on the AI model, produce an output or decision (such as a prediction, recommendation, and/or classification). AI technologies rely on AI algorithms to generate models. The most appropriate model(s) is/are selected and deployed in a production system.
Singapore	Principles to Promote Fairness, Ethics, Accountability and Transparency in the Use of	Nov. 18, 2025	“AIDA” refers to artificial intelligence or data analytics, which are defined as technologies that assist or replace human decision-making. While this definition is broad, the Principles in this document can apply to both complex and straightforward AIDA techniques.

Jurisdiction	Title	Date	Definition
	Artificial Intelligence and Data Analytics in Singapore's Financial Sector (Monetary Authority of Singapore)		
South Korea	Artificial Intelligence Accountability and Regulation Legislation	Aug. 8, 2023	“Artificial intelligence” refers to software that implements human intellectual abilities, such as learning, perception, judgment, and understanding of natural language, through electronic methods.
South Korea	The Act on the Development of Artificial Intelligence and Establishment of Trust (AI Basic Act)	Site updated Dec. 31, 2024	The types of AI systems governed by the AI Basic Act are: (i) high-impact AI (defined as AI systems that pose significant risks or impacts on human life, physical safety, or fundamental rights) and (ii) generative AI (defined as AI systems designed to mimic the structure and characteristics of input data to generate outputs such as text, sound, images, videos, and other creative content).
Spain	Artificial Intelligence Strategy 2024	May 15, 2024	No definition in the strategy, but the announcement says, “The Royal Spanish Academy (RAE) defines [artificial intelligence] as a scientific discipline that deals with creating computer programs that execute operations comparable to those carried out by the human mind, such as learning or logical reasoning.”
Spain	Royal Decree 817/2023, of November 8, establishing a controlled testing environment for testing compliance with the proposed Regulation of the European	Nov. 9, 2023	“Artificial intelligence system”: a system designed to operate with a certain level of autonomy and that, based on input data provided by machines or humans, infers how to achieve a set of stated objectives using machine learning or logic- and knowledge-based strategies, and generates output information, such as content (generative artificial intelligence systems), predictions, recommendations or decisions, that influence the environments with which it interacts.

Jurisdiction	Title	Date	Definition
	Parliament and of the Council establishing harmonized standards on artificial intelligence.		
Taiwan	Basic Law on Artificial Intelligence	Nov. 9, 2023	For the purposes of this Law, "artificial intelligence" refers to a system with autonomous operational capabilities. This system, through input or sensing, and via machine learning and algorithms, can predict, provide content, offer suggestions, or make decisions that affect a real or virtual environment for explicit or implicit goals.
Turkey	Artificial intelligence law proposal	July 15, 2024	“Yapay Zeka: İnsan benzeri bilişsel işlevleri yerine getirebilen ve öğrenme, mantık yürütme, problem çözme, algılama ve dil anlama gibi yeteneklere sahip bilgisayar tabanlı sistemleri ifade eder” [Artificial Intelligence: Computer-based systems that can perform human-like cognitive functions and have capabilities such as learning, reasoning, problem solving, perception and language understanding.
Turkey	Recommendations on the Protection of Personal Data in the Field of Artificial Intelligence by the Turkish Personal Data Protection Authority	June 25, 2024	Artificial Intelligence (“AI”), on the other hand, focuses on developing algorithms and computer software that are capable of executing tasks associated with humans, such as thinking, interpreting, and making decisions, by way of analyzing such human tasks.
UK NCSC, US CISA	Guidelines for secure AI system development	Not stated	In this document we use ‘AI’ to refer specifically to machine learning (ML) applications. All types of ML are in scope. We define ML applications as applications that: > involve software components (models) that allow computers to recognise and bring context to patterns in data

Jurisdiction	Title	Date	Definition
			without the rules having to be explicitly programmed by a human > generate predictions, recommendations, or decisions based on statistical reasoning.
UN	Taxonomy of Human Rights Risks Connected to Generative AI	July 15, 2005	No definition
UNESCO	Recommendation on the Ethics of Artificial Intelligence	July 16, 2005	<p>This Recommendation does not have the ambition to provide one single definition of AI, since such a definition would need to change over time, in accordance with technological developments. Rather, its ambition is to address those features of AI systems that are of central ethical relevance. Therefore, this Recommendation approaches AI systems as systems which have the capacity to process data and information in a way that resembles intelligent behaviour, and typically includes aspects of reasoning, learning, perception, prediction, planning or control. Three elements have a central place in this approach: (a) AI systems are information-processing technologies that integrate models and algorithms that produce a capacity to learn and to perform cognitive tasks leading to outcomes such as prediction and decision-making in material and virtual environments. AI systems are designed to operate with varying degrees of autonomy by means of knowledge modelling and representation and by exploiting data and calculating correlations. AI systems may include several methods, such as but not limited to: (i) machine learning, including deep learning and reinforcement learning; (ii) machine reasoning, including planning, scheduling, knowledge representation and reasoning, search, and optimization. AI systems can be used in cyber-physical systems, including the Internet of things, robotic systems, social robotics, and human-computer interfaces, which involve control, perception, the processing of data</p>

Jurisdiction	Title	Date	Definition
			collected by sensors, and the operation of actuators in the environment in which AI systems work
United Arab Emirates	AI Guide	July 14, 2005	AI defines a collection of technologies enabling a machine or system to comprehend, learn, act, and sense like a human.
United Arab Emirates (Dubai)	Dubai: AI Ethics Principles & Guidelines	July 11, 2005	ARTIFICIAL INTELLIGENCE (also “AI”) The capability of a functional unit to perform functions that are generally associated with human intelligence such as reasoning, learning and self-improvement. ARTIFICIALLY INTELLIGENT SYSTEM (also “AI system”) A product, service, process or decision-making methodology whose operation or outcome is materially influenced by artificially intelligent functional
United Kingdom	A Pro-Innovation Approach to AI Regulation	Dec. 22, 2025	There is no general definition of AI that enjoys widespread consensus. That is why we have defined AI by reference to the 2 characteristics that generate the need for a bespoke regulatory response. The ‘adaptivity’ of AI can make it difficult to explain the intent or logic of the system’s outcomes: AI systems are ‘trained’ – once or continually – and operate by inferring patterns and connections in data which are often not easily discernible to humans. Through such training, AI systems often develop the ability to perform new forms of inference not directly envisioned by their human programmers. The ‘autonomy’ of AI can make it difficult to assign responsibility for outcomes: Some AI systems can make decisions without the express intent or ongoing control of a human. The combination of adaptivity and autonomy can make it difficult to explain, predict, or control the outputs of an AI system, or the underlying logic by which they are generated. It can also be challenging to allocate responsibility for the system’s operation and outputs. For regulatory purposes, this has potentially serious implications, particularly when decisions are made relating to significant

Jurisdiction	Title	Date	Definition
			<p>matters, like an individual’s health, or where there is an expectation that a decision should be justifiable in easily understood terms, like a legal ruling. By defining AI with reference to these functional capabilities and designing our approach to address the challenges created by these characteristics, we future-proof our framework against unanticipated new technologies that are autonomous and adaptive. Because we are not creating blanket new rules for specific technologies or applications of AI, like facial recognition or LLMs, we do not need to use rigid legal definitions. Our use of these defining characteristics was widely supported in responses to our policy paper, as rigid definitions can quickly become outdated and restrictive with the rapid evolution of AI. We will, however, retain the ability to adapt our approach to defining AI if necessary, alongside the ongoing monitoring and iteration of the wider regulatory framework.</p>
United Kingdom	Generative AI Framework for HM Government v.1	Aug. 3, 2023	<p>Generative AI is a specialised form of AI that can interpret and generate high-quality outputs including text and images, opening up the potential for opportunities for organisations, including delivering efficiency savings or developing new language capability.</p>
United Kingdom	Guidance on the AI Auditing Framework	July 15, 2005	<p>If you are processing this data in the context of statistical models, and using those models to make predictions about people, this guidance will be relevant to you regardless of whether you classify those activities as ML (or AI). We use the umbrella term ‘AI’ because it has become a mainstream way for organisations to refer to a range of technologies that mimic human thought. Similar technologies, which have similar sources of risk, are likely to benefit from the same set of risk measures. So, whether you call it AI, machine learning, complex information processing, or something else, the risks and controls identified here should be helpful. Where there are important differences</p>

Jurisdiction	Title	Date	Definition
			between different types of AI, for example, simple regression models and deep neural networks, we will refer to these explicitly.
United Kingdom	Implementing the UK's AI Regulatory Principles	July 12, 2025	AI or AI system: Products and services that are 'adaptable' and 'autonomous' in the sense outlined in our definition in section 3.2.1 of the AI White Paper [The Pro-Innovation Approach to AI Regulation]. This definition is inclusive of AI agents, frontier AI, and narrow AI.
United Kingdom	Introduction to AI Assurance	Feb. 24, 2025	No definition
United Kingdom	Public Authority Algorithmic and Automated Decision-Making Systems Bill (HL 27)	Feb. 24, 2025	[A]ny algorithmic or automated decision-making system developed or procured by a public authority from six months after the date on which this Act is passed. (2) This includes— (a) any system, tool or statistical model used to inform, recommend or make an administrative decision about a service user or a group of service users, and (b) systems in development, excluding automated decision-making systems operating in test environments. (3) This Act does not apply to any automated decision-making system used for the purpose of national security. (4) This Act does not apply to automated systems which merely calculate and implement formulas, including taxation and budgetary allocation, insofar as they automate a process of calculation which would otherwise be carried out manually and fully understood.
US FDA, Health Canada, UK MHRA	Good Machine Learning Practice for Medical Device Development : Guiding Principles	July 16, 2025	No definition
USA	Department of Defense AI Strategy	Oct. 21, 2025	Artificial intelligence (AI) is one such technological advance. AI refers to the ability of machines to perform tasks that

Jurisdiction	Title	Date	Definition
			normally require human intelligence – for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action – whether digitally or as the smart software behind autonomous physical systems.
USA	Executive Order on Ensuring a National Policy Framework for Artificial Intelligence	July 10, 2005	No definition
USA	Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence	July 17, 2005	A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.
USA	John S. McCain National Defense Authorization Act For Fiscal Year 2019, Pl 115-232, August 13, 2018, 132 Stat 1636	Oct. 30, 2023	(g) ARTIFICIAL INTELLIGENCE DEFINED.—In this section, the term “artificial intelligence” includes the following: (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets. (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action. (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks. (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task. (5) An artificial system designed to act rationally,

Jurisdiction	Title	Date	Definition
			including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.
USA	National Artificial Intelligence Initiative	July 11, 2005	The term “artificial intelligence” means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and humanbased inputs to— (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.
USA	NIST Risk Management Framework 1.0	Not stated	The AI RMF refers to an AI system as an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.
USA	The National Artificial Intelligence Initiative 15 U.S.C. 9401(3); Nat’l Def. Authorization Act of 2019 § 238(g)	Jan. 23, 2025	(1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets; (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action; (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks; (4) A set of techniques, including machine learning that is designed to approximate a cognitive task; and (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making and acting.

Jurisdiction	Title	Date	Definition
USA (Cal.)	SB 1047	July 11, 2005	22602(b) "Artificial intelligence model" means an engineered or machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs that can influence physical or virtual environments and that may operate with varying levels of autonomy.
USA (Cal.)	SB 53	July 16, 2005	"Artificial intelligence model" means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.
USA (Colo.)	Senate Bill 24-205	July 17, 2005	"Artificial intelligence system" means any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments.
USA (N.Y.)	Assembly Bill A6453A	May 17, 2024	"Artificial Intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments, and that uses machine- and human-based inputs to perceive real and virtual environments, abstract such perceptions into models through analysis in an automated manner, and use model inference to formulate options for information or action.
USA (Ohio)	House Bill 469	July 17, 2005	"AI" means any software, machine, or system capable of simulating humanlike cognitive functions, including learning or problem solving, and producing outputs based on data-driven algorithms, rules-based logic, or other computational methods, regardless of non-legally defined classifications such as artificial general intelligence, artificial superintelligence, or generative artificial intelligence.
USA (Tex.)	HB 149	July 17, 2005	Artificial intelligence system means any machine-based system that, for any explicit

Jurisdiction	Title	Date	Definition
			or implicit objective, infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments
USA (Tex.)	House Bill 1079	May 20, 2025	"Artificial intelligence system" means the use of machine learning and related technologies that use data to train statistical models for the purpose of enabling computer systems to perform tasks normally associated with human intelligence or perception, such as computer vision, speech or natural language
World Health Organization	Ethics and Governance of Artificial Intelligence for Health	Dec. 23, 2024	"Artificial Intelligence (AI) refers to the capability of algorithms integrated into systems and tools to learn from data so that they can perform automated tasks without explicit programming of every step by a human."