

Internal deployment in the AI Act

Commentary by Matteo Pistillo | Submitted: May 2026

This chapter analyses and stress-tests arguments in favour and against the inclusion of internal deployment within the scope of the European Union (EU) Artificial Intelligence (AI) Act. Specifically, this chapter first analyses interpretative pathways based on Article 2(1)(a)–(c) supporting the application of the AI Act to internally deployed AI models and systems (Section 2). Then, it examines possible objections and exceptions based on Articles 2(6) and 2(8), with particular attention to the complexity of the scientific R&D exception under Article 2(6) (Section 3). Finally, it illustrates how Articles 2(1), 2(6), and 2(8) can be viewed as complementary to each other once broken down to their most plausible meaning and interpreted in conjunction with Articles 3(1), 3(3), 3(4), 3(9)–3(12) and 3(63) and Recitals 12, 13, 21, 25, 97, 109 and 110 (Section 4; Figure 1).

Select bibliography

- Stix C and others, ‘AI Behind Closed Doors: A Primer on the Governance of Internal Deployment’ (arXiv, 16 April 2025).
- Acharya A and Delaney O, ‘Managing Risks from Internal AI Systems’ (Institute for AI Policy and Strategy 2025).
- Chan L, ‘AI Models Can Be Dangerous before Public Deployment’ (*METR*, 17 January 2025) <<https://metr.org/blog/2025-01-17-ai-models-dangerous-before-public-deployment/>> accessed 26 May 2026.
- Van Eecke P and Regenhardt B, ‘Article 2: Scope’ in Ceyhun Necati Pehlivan, Nikolaus Forgó and Peggy Valcke (eds), *The EU Artificial Intelligence (AI) Act: A Commentary* (Wolters Kluwer 2024).

Commentary

1. Introduction	2
2. Arguments supporting the application of the AI Act to internal deployment.....	6
2.1. Internal AI systems	7
2.1.1. Article 2(1)(a): ‘Providers [...] putting into service AI systems’	7
2.1.2. Article 2(1)(a): ‘Providers placing on the market [...] AI systems’	9
2.1.3. Article 2(1)(b): ‘Deployers of AI systems that have their place of establishment or are located within the Union’	10
2.1.4. Article 2(1)(c): ‘output produced by the AI system is used in the Union’	11
2.2. Internal GPAI Models.....	12
3. Arguments opposing or restricting the application of the AI Act to internal deployment.....	16
3.1. Article 2(6): ‘AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development.’	16

3.2.	Article 2(8): ‘Any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service.’	22
4.	Conclusion	24

1. Introduction

1. This chapter focuses on one of the most challenging questions surrounding the European Union (EU) Artificial Intelligence (AI) Act:¹ whether and to what extent the AI Act applies to what frontier AI companies and researchers refer to as ‘internal deployment’.² The inclusion of internal deployment within the scope of the AI Act remains an open question.³ This chapter attempts to unpack this question, offering an overview of several potential interpretative pathways. Specifically, Section 1 of this chapter provides an initial introduction to the concept of internal deployment. Next, Section 2 grounds itself in the technical context of frontier AI development and deployment and elaborates on potential arguments supporting the application of the AI Act to internal deployment. Finally, Section 3 outlines arguments for opposing, or at least restricting, the application of the AI Act to internal deployment, as well as potential counterarguments.
2. Internal deployment is an umbrella term used to describe situations in which an AI provider develops an AI model or system and makes it available for use within the provider itself.⁴ In other words, internal deployment describes the behaviour of an AI provider that, instead of publicly releasing an AI model or system, elects to use it only within the provider’s walls. As an umbrella term, internal

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 1689/1 (“AI Act”).

² See Charlotte Stix and others, ‘AI Behind Closed Doors: A Primer on the Governance of Internal Deployment’ (arXiv, 16 April 2025) <<https://doi.org/10.48550/arXiv.2504.12170>> accessed 26 May 2026; Ashwin Acharya and Oscar Delaney, ‘Managing Risks from Internal AI Systems’ (Institute for AI Policy and Strategy 2025) <<https://www.iaps.ai/research/managing-risks-from-internal-ai-systems>> accessed 26 May 2026; Steven Adler, ‘AI Companies Should Monitor Their Internal AI Use’ (*Clear-Eyed AI*, 5 January 2026) <<https://www.clear-eyed.ai/p/ai-companies-unmonitored-internal>> accessed 1 June 2026; Yoshua Bengio and others, ‘International AI Safety Report’ (DSIT 2025/001, 2025) <<https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025>> accessed 1 June 2026, 35 (‘Deployment can take several forms: internal deployment for use by the system’s developer, or external deployment either publicly or to private customers’); Anthropic, ‘Claude Mythos Preview System Card’ (Anthropic 2026) <<https://www-cdn.anthropic.com/08ab9158070959f88f296514c21b7facce6f52bc.pdf>> accessed 1 June 2026, 57, 61, 62, 65, 132, 203; OpenAI, ‘GPT-5.5 System Card’ (OpenAI 2026) <<https://deploymentsafety.openai.com/gpt-5-5/gpt-5-5.pdf>> accessed 1 June 2026, 13, 14, 43; Meta, ‘Advanced AI Scaling Framework Version 2’ (Meta 2026) <https://ai.meta.com/static-resource/Meta_Advanced-AI-Scaling-Framework-v2> accessed 17 May 2026, 4, 9, 18, 27, 42.

³ See Stix and others, ‘AI Behind Closed Doors’ (n 2) 27.

⁴ This definition is adapted from Stix and others (n 2) 6 (defining ‘internal deployment’ as ‘the act of making an AI system available for access and/or usage exclusively for the developing organization’). See also Acharya and Delaney (n 2) (using the term ‘internal deployment’ and defining ‘internal models’ as ‘AI systems that are only accessible to company employees (and perhaps to a few external experts given access by the company)’); Meta (n 2) 42 (defining ‘internal deployment’ as ‘models that are exclusively available to Meta personnel’). Some researchers use the term ‘internal usage’ or ‘internal use’ to refer to ‘internal deployment’. See, for instance, Lawrence Chan, ‘AI Models Can Be Dangerous before Public Deployment’ (*METR*, 17 January 2025) <<https://metr.org/blog/2025-01-17-ai-models-dangerous-before-public-deployment/>> accessed 26 May 2026 (referring to internal deployment through the expression ‘internal usage’); Adler (n 2) (defining ‘a company’s internal use of AI’ as ‘internal deployment’). In this respect, it is important to observe that the meaning of ‘internal deployment’, as commonly used by AI companies and researchers, is generally broader than the definition of ‘deployer’ contained in article 3(4) and recital 13 AI Act.

deployment covers a range of different scenarios diverging, among other aspects, in the exclusivity of access to AI models or systems, in the usage of these AI models or systems, and in the level of AI model or system autonomy.⁵ An AI provider's exclusive access to and/or usage of an internal AI model or AI system can be temporary or can instead be permanent.⁶ In other words, internal AI models and systems may be first deployed internally and then made available to the public,⁷ or they may be used exclusively within an AI provider for the entirety of their lifetime and never made available to the public. Internal AI models and systems may be deployed to assist with, or automate, a variety of different tasks within a provider (for instance, within its HR, legal or R&D departments).⁸ The degree of autonomy of internal AI models and systems also sits on a spectrum, from undertakings that are small, self-contained projects, all the way to future arrangements that could engage 'a country of geniuses in a datacenter' running very large segments of an AI provider's workflow autonomously.⁹

3. One important, and potentially concerning,¹⁰ use that AI providers can make of their most capable AI models and systems is to use them to design, train or improve future (and, plausibly, more capable) AI models and systems.¹¹ This use is often referred to as autonomous AI Research & Development,

⁵ See Stix and others (n 2) 12–14 (outlining various potential internal uses as well as various potential internal user groups).

⁶ See Stix and others (n 2) 8.

⁷ For instance, OpenAI's model GPT-4 was available internally for six months before being publicly released. See OpenAI, 'GPT-4 System Card' (OpenAI 2023) <<https://cdn.openai.com/papers/gpt-4-system-card.pdf>> accessed 1 June 2026, 59. Anthropic's Opus 4.6 and 4.7 and OpenAI's GPT-5.5 also appear to have been used internally before public release. See Anthropic, 'System Card: Claude Opus 4.6' (Anthropic 2026) <<https://www-cdn.anthropic.com/6a5fa276ac68b9aeb0c8b6af5fa36326e0e166dd.pdf>> accessed 1 June 2026, 95 ('Throughout late-stage training, we deployed several snapshots of Claude Opus 4.6 for provisional internal use and evaluation, with increasingly broad uptake as time went on.');

Anthropic, 'System Card: Claude Opus 4.7' (Anthropic 2026) <<https://cdn.sanity.io/files/4zrzovbb/website/037f06850df7fbc871e206dad004c3db5fd50340.pdf>> accessed 1 June 2026, 95 ('We used versions of Claude Opus 4.7 substantially internally before deploying it');

OpenAI, 'GPT-5.5 System Card' (n 2) 14 (mentioning 'a pre-final version of GPT-5.5 internal usage').

⁸ The usage of internal AI models and systems may span from light professional usage and work support (i.e., supporting a provider's staff with their workflow and outsourcing portions of existing or new workflows, and then using the relevant outputs) to outright delegation to virtual co-workers or independent workers. See Stix and others (n 2) 12–14. See also Adler (n 2) (mentioning 'critical internal tasks like security reviews, interpretability analysis, and developing next-generation AI systems'); Toby Shevlane and others, 'Model Evaluation for Extreme Risks' (arXiv, 22 September 2023) <<https://doi.org/10.48550/arXiv.2305.15324>> accessed 16 May 2026, 9; Acharya and Delaney (n 2) 8–10.

⁹ See Dario Amodei, 'Machines of Loving Grace' (2024) <<https://www.darioamodei.com/essay/machines-of-loving-grace>> accessed 1 June 2026. Higher degrees of autonomy may entail less friction to speed, but also a lower degree of human oversight. See Stix and others (n 2) 16–17.

¹⁰ In summary, concerns cluster around the risk that, once AI systems play a growing role in designing and training their successors, this could trigger a hyperbolic explosion in AI capabilities (Tom Davidson, 'How Can AI Labs Incorporate Risks from AI Accelerating AI Progress Into Their Responsible Scaling Policies?' (*Forethought*, 2025) <<https://www.forethought.org/research/how-can-ai-labs-incorporate-risks-from-ai-accelerating-ai-progress-into>> accessed 1 June 2026; Daniel Eth and Tom Davidson, 'Will AI R&D Automation Cause a Software Intelligence Explosion?' (*Forethought*, 2025) <<https://www.forethought.org/research/will-ai-r-and-d-automation-cause-a-software-intelligence-explosion>> accessed 1 June 2026.) and effectively prevent humans from understanding, auditing, or controlling the resulting technology (Daniel Kokotajlo and others, 'AI 2027' (*AI 2027*, 2025) <<https://ai-2027.com/>> accessed 1 June 2026). For an overview of these and other potential threats from internal deployment, see Stix and others (n 2), Chan (n 4), and Acharya and Delaney (n 2). See also a recent interview of Anthropic's co-founder and chief scientist, Jared Kaplan, who defined the decision to let AI systems train themselves as 'the biggest decision yet' and 'the ultimate risk' (Robert Booth, "'The Biggest Decision yet': Jared Kaplan on Allowing AI to Train Itself" *The Guardian* (2 December 2025) <<https://www.theguardian.com/technology/ng-interactive/2025/dec/02/jared-kaplan-artificial-intelligence-train-itself>> accessed 1 June 2026).

¹¹ See Johannes Treutlein and others, 'Pre-Deployment Auditing Can Catch an Overt Saboteur' (*Alignment Science Blog*, 2026) <<https://alignment.anthropic.com/2026/auditing-overt-saboteur/>> accessed 1 June 2026 ('Anthropic

or simply autonomous AI R&D.¹² While it is not currently feasible to fully automate the AI R&D process,¹³ AI models and systems are becoming increasingly capable of accomplishing longer and more complex AI R&D tasks.¹⁴ A frontier AI company recently stated in their system card that ‘it is plausible that models equipped with highly effective scaffolding may not be very far away from [...] AI R&D-4 threshold,’ which captures the ability of an AI model to ‘fully automate the work of an entry-level, remote-only researcher’ at that frontier AI company.¹⁵ Existing progress in autonomous AI R&D capabilities has recently spurred another frontier AI company to announce that they could achieve ‘an automated AI research intern by September of 2026 [...] and a true automated AI researcher by March of 2028.’¹⁶

4. Legal and policy frameworks on both sides of the Atlantic are paying increasing attention to internal deployment and autonomous AI R&D,¹⁷ spurring the question as to whether and to what extent the

uses existing Claude models to assist in the development of future generations of Claude models’); Tom Davidson, Rose Hadshar and William Macaskill, ‘Three Types of Intelligence Explosion’ (*Forethought*, 2025) <<https://www.forethought.org/research/three-types-of-intelligence-explosion>> accessed 1 June 2026; Eth and Davidson (n 10). For clarity, AI R&D is *only one* of the possible internal uses of AI models and systems by AI providers. This clarification is important, as some provisions in the AI Act seem to concern internal deployment generally (see Sections 2.1.–3.1. and 3.3.), whereas other provisions seem to specifically concern *only* AI R&D (see Section 3.2.). For the purposes of this paper, the term ‘internal deployment’ refers to all possible uses of internally deployed AI models and systems, unless otherwise specified.

¹² For clarity, AI R&D is used to refer to both AI R&D as a capability (i.e., an AI model or system’s ability to accelerate and/or automate AI R&D) and as a potential internal application of AI models or systems (and the relevant threat model).

¹³ See, for instance, Anthropic, ‘Model Report’ (Anthropic 2026) <<https://www.anthropic.com/transparency/model-report>> accessed 1 June 2026 (‘Claude Opus 4.5 could not fully automate an entry-level, remote-only research role at Anthropic’); OpenAI, ‘GPT-5 System Card’ (OpenAI 2025) <<https://deploymentsafety.openai.com/gpt-5-5/gpt-5-5.pdf>> accessed 1 June 2026, 41–42; Google, ‘Gemini 3 Pro Model Card’ (Google 2025) <<https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-Pro-Model-Card.pdf>> accessed 1 June 2026, 8, and Google, ‘Gemini 3 Pro Frontier Safety Framework Report’ (Google 2025) <https://storage.googleapis.com/deepmind-media/gemini/gemini_3_pro_fsf_report.pdf> accessed 1 June 2026, 14–15. See also Hjalmar Wijk and others, ‘RE-Bench: Evaluating Frontier AI R&D Capabilities of Language Model Agents against Human Experts’ (arXiv, 27 May 2025) <<https://doi.org/10.48550/arXiv.2411.15114>> accessed 1 June 2026, 21 (‘a significant gap remains compared to the top human performance in most environments’); Thomas Kwa and others, ‘Measuring AI Ability to Complete Long Tasks’ (*METR* 2025) <<https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/>> accessed 1 June 2026 (noting that ‘the best AI agents are not currently able to carry out substantive projects by themselves or directly substitute for human labor’ and observing, in particular, that current AI models only ‘succeed <10% of the time on tasks taking more than around 4 hours’).

¹⁴ See, for instance, Kwa and others, ‘Measuring AI Ability to Complete Long Tasks’ (n 13) (tracking the task-completion time horizons for the most advanced AI models publicly available); Thomas Kwa and others, ‘Measuring AI Ability to Complete Long Software Tasks’ (arXiv, 25 February 2026) <<https://doi.org/10.48550/arXiv.2503.14499>> accessed 1 June 2026, 23 (noting that ‘50% task completion time horizon on our tasks has been growing exponentially from 2019–2025 with a doubling time of approximately seven months’ and that ‘an 80% confidence interval for the release date of AI that can complete 1-month long software tasks spans from late 2028 to early 2031’).

¹⁵ Anthropic, ‘System Card: Claude Opus 4.5’ (Anthropic 2025) <<https://www-cdn.anthropic.com/bf10f64990cfda0ba858290be7b8cc6317685f47.pdf>> accessed 1 June 2026, 13–14.

¹⁶ See Sam Altman, ‘Yesterday We Did a Livestream. TL;DR: We Have Set Internal Goals of Having an Automated AI Research Intern by September of 2026 Running on Hundreds of Thousands of GPUs, and a True Automated AI Researcher by March of 2028. [...]’ <<https://x.com/sama/status/1983584366547829073>> accessed 1 June 2026; Rebecca Bellan, ‘Sam Altman Says OpenAI Will Have a “Legitimate AI Researcher” by 2028’ (*TechCrunch*, 28 October 2025) <<https://techcrunch.com/2025/10/28/sam-altman-says-openai-will-have-a-legitimate-ai-researcher-by-2028/>> accessed 1 June 2026. See also OpenAI, ‘AI Progress and Recommendations’ (*OpenAI*, 18 May 2026) <<https://openai.com/index/ai-progress-and-recommendations/>> accessed 1 June 2026 (‘we get closer to systems capable of recursive self-improvement.’).

¹⁷ See California SB-53, ‘Artificial intelligence models: large developers’ [2025], s 22757.12.(a)(10); President Joe Biden, ‘Memorandum on Advancing the United States’ Leadership in Artificial Intelligence; Harnessing Artificial

AI Act covers internal deployment. For example, in the United States, California Senate Bill 53 requires ‘large frontier developer[s]’ to ‘[a]ssess[] and manag[e] catastrophic risk resulting from the *internal use of*[their] *frontier models*’.¹⁸ Similarly, President Biden’s National Security Memorandum required the U.S. Department of Commerce to concentrate on the testing of a series of AI capabilities that included the capability to ‘automate development and deployment of other models with such capabilities’.¹⁹ At the same time, frontier AI companies have included dedicated thresholds for automated AI R&D capabilities within their frontier safety policies.²⁰ Through these internal policies, frontier AI companies have committed to, for example, evaluating ‘the ability of an AI system to accelerate AI research, including to increase the system’s own capability’,²¹ ‘[t]he ability to fully automate the work of an entry-level, remote-only researcher’ and ‘to cause dramatic acceleration in the rate of effective scaling’,²² and whether this capability can ‘result[] in AI progress substantially accelerating from historical rates.’²³

5. The Safety and Security Chapter of the Code of Practice for General-Purpose AI (“GPAI”) Models²⁴ hints at internal deployment for AI R&D purposes²⁵ and includes the ‘capabilities to automate AI research and development’ (i.e., autonomous AI R&D) within the ‘sources of systemic risks’,²⁶ making the question around the application of the AI Act to internally deployed AI models and systems all the more urgent and consequential.²⁷ Nonetheless, compared to other legal frameworks mentioning

Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence’ (*The White House*, 24 October 2024) <<https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/10/24/memorandum-on-advancing-the-united-states-leadership-in-artificial-intelligence-harnessing-artificial-intelligence-to-fulfill-national-security-objectives-and-fostering-the-safety-security/>> accessed 1 June 2026, s 3.3(e)(i).

¹⁸ SB-53 (n 17) s 22757.12.(a)(10) (emphasis added).

¹⁹ Memorandum on Advancing the United States’ Leadership in Artificial Intelligence (n 17) s 3.3(e)(i) (emphasis added).

²⁰ See OpenAI, ‘Preparedness Framework Version 2’ (OpenAI 2025) <<https://cdn.openai.com/pdf/18a02b5d-6b67-4cec-ab64-68cdfbddebcdf/preparedness-framework-v2.pdf>> accessed 17 May 2026; Anthropic, ‘Responsible Scaling Policy Version 2.2’ (Anthropic 2025) <<https://www-cdn.anthropic.com/872c653b2d0501d6ab44cf87f43e1dc4853e4d37.pdf>> accessed 1 June 2026; Google, ‘Frontier Safety Framework Version 3.0’ (Google 2025) <https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/strengthening-our-frontier-safety-framework/frontier-safety-framework_3.pdf> accessed 1 June 2026; METR, ‘Common Elements of Frontier AI Safety Policies’ (METR 2025) <<https://metr.org/common-elements.pdf>> accessed 1 June 2026.

²¹ OpenAI, ‘Preparedness Framework’ (n 20) 6. See, e.g., OpenAI, ‘GPT-5 System Card (n 13) at 41.

²² Anthropic, ‘Responsible Scaling Policy’ (n 20) 4. See, e.g., Anthropic, ‘System Card: Claude Opus 4.5’ (n 15) 13–14; Anthropic, ‘System Card: Claude Opus 4.6’ (n 7) 12–13, 181–194.

²³ Google, ‘Frontier Safety Framework’ (n 20) 13. In addition to an ‘ML R&D acceleration’ capability threshold, Google DeepMind also has an ‘ML R&D automation’ capability threshold, capturing whether an AI model or system can ‘fully automate the work of any team of researchers at Google focused on improving AI capabilities, with approximately comparable all-inclusive costs’ (Google, ‘Frontier Safety Framework’ (n 20) 14). See, e.g., Google, ‘Gemini 3 Pro Model Card’ (n 13) s 8.

²⁴ For clarity, all references to the Code of Practice in this chapter refer to its Safety and Security Chapter.

²⁵ See European Commission, ‘Code of Practice for General-Purpose AI Models – Safety and Security Chapter’ (2025) <<https://ec.europa.eu/newsroom/dae/redirection/document/118119>>, Measure 7.1 (‘Signatories will provide in the Model Report: [...] a description of how the model has been *used* and is expected to be used, including its use *in the development, oversight, and/or evaluation of models*’) (emphasis added); Measure 7.3 (‘Signatories will provide in the Model Report: [...] (4) a high-level description of: (a) the techniques and assets they intend to use to further develop the model over the next six months, including *through the use of other AI models and/or AI systems* [...]’) (emphasis added).

²⁶ Code of Practice, Safety and Security Chapter (n 25) app 1.3. See also AI Act, recital 110 (mentioning the ‘risks from models [...] training other models’ in the context of systemic risk from GPAI models).

²⁷ If internal GPAI models are covered by the AI Act, signatories will have to evaluate these internal GPAI models for the ‘capabilities to automate AI research and development’ (Code of Practice, Safety and Security Chapter (n 25) app 1.3.1). This may be particularly consequential because the AI models and systems used for AI R&D

internal deployment expressly (for instance, California Senate Bill 53), the AI Act does not clearly state that internal AI models or systems fall within its scope.

6. This chapter examines both sides of the equation: it considers interpretative options to deem internal deployment as included within the scope of the AI Act (Section 2), as well as possible counterarguments and exceptions (Section 3).²⁸ This chapter elaborates on, and stress-tests, Articles 2(1)(a)–(c), 2(6) and 2(8) AI Act, which offer potential interpretative pathways for supporting or opposing the inclusion of internal deployment within the scope of the AI Act. Finally, the conclusion of this chapter (Section 4) puts forward a proposal for how all these, potentially conflicting, provisions of the AI Act can be interpreted systematically – and, ultimately, all brought to coherence. As the conclusion of this chapter illustrates, once broken down to their most plausible meaning and interpreted in conjunction with other relevant provisions,²⁹ Articles 2(1), 2(6) and 2(8) AI Act are, in fact, complementary to each other.³⁰ In summary, the internal deployment of AI systems or GPAI models integrated into AI systems could trigger the application of the AI Act under Article 2(1)(a)–(c), unless an internal AI model or system is specifically developed and deployed for the sole purpose of scientific R&D under Article 2(6). Pre-internal-deployment research, testing and development activity arguably falls outside the scope of the AI Act under Article 2(8).

2. Arguments supporting the application of the AI Act to internal deployment

7. This section explores potential interpretations of the AI Act that support the inclusion of internal deployment within the scope of the Act. Specifically, Section 2.1. examines the potential inclusion of *internal AI systems* within the scope of the AI Act. Section 2.2. then discusses the potential inclusion of *internal GPAI models* within the scope of the AI Act.
8. Section 3 will then lay out potential interpretations of the AI Act that support the exclusion of internal deployment from the scope of the Act, or at least restrict the AI Act's scope only to some AI models and systems deployed internally, as well as potential counterarguments.

purposes are internal rather than external. In this respect, it is worth observing that some frontier AI companies already evaluate *external* AI models and systems for AI R&D capabilities (see nn 20–23).

²⁸ In general, the question around the application of the AI Act to internal deployment should not be confused with the obligations AI providers may have for *covered* AI models and systems *before their public deployment*. Examples of the latter are article 55 AI Act and recital (a) of the Code of Practice, which require providers to 'take appropriate measures' to 'assess and mitigate' 'systemic risks' by 'taking appropriate measures *along the entire model lifecycle* (including during development that occurs *before* [...] *a model has been placed on the market*)' (emphasis added). Similarly, Measures 1.2 and 4.2 of the Code of Practice require providers to undertake these assessments '*along the entire model lifecycle*' and 'only proceed with the *development*, the making available on the market, and/or the use of the model, if the systemic risks stemming from the model are determined to be acceptable' (Code of Practice, Safety and Security Chapter (n 25) emphasis added). Both questions do relate to the scope of the AI Act, see the commentary on Article 2, Section 2.4. in this work.

²⁹ AI Act, arts 3(1), 3(3), 3(4), 3(9)–3(12) and 3(63) and recitals 12, 13, 21, 25, 97, 109 and 110.

³⁰ See Section 4 and Figure 1.

2.1. Internal AI systems

9. Sections 2.1.1.–2.1.4. concentrate on Article 2 AI Act, which defines the ‘scope’ of the Act, and offer reasons why the deployment of *internal AI systems* could be considered as covered by the AI Act.³¹ Specifically:

- Section 2.1.1. concentrates on the expression ‘providers [...] putting into service AI systems’ in Article 2(1)(a).
- Section 2.1.2. concentrates on the expression ‘providers placing on the market [...] AI systems’ in Article 2(1)(a).³²
- Section 2.1.3. concentrates on the expression ‘deployers of AI systems that have their place of establishment or are located within the Union’ in Article 2(1)(b).
- Section 2.1.4. concentrates on the expression ‘output produced by the AI system is used in the Union’ in Article 2(1)(c).

2.1.1. Article 2(1)(a): ‘Providers [...] putting into service AI systems’

10. Article 2 AI Act defines the ‘scope’ of the Act.³³ Specifically, under Article 2(1)(a) the AI Act ‘applies to’ ‘providers [...] putting into service AI systems’ in the Union.³⁴ This section examines whether the expression ‘providers ... putting into service AI systems’³⁵ could include AI systems internally deployed by the providers developing them.

11. The AI Act offers definitions for all the components of this expression (‘providers [...] putting into service AI systems’³⁶). First, the term ‘providers’ is defined by Article 3(3) as ‘a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge’. Importantly, both Article 2(1)(a) and Recital 21 clarify that the AI Act applies to providers of AI systems ‘irrespective of whether those providers are established or located within the Union or in a third country’.³⁷ Therefore, the term ‘providers’ encompasses both EU-based and foreign AI developers.

³¹ For clarity, Sections 2.1.–2.4. discuss the potential application of the AI Act to internal *AI systems*. Section 3.1. discusses the potential application of the AI Act to internal *GPAI models* and the internal AI systems integrating GPAI models, thus looping back to Sections 2.1.–2.4.

³² Section 3.1. concentrates on the expression ‘placing on the market *general-purpose AI models*’ in article 2(1)(a) (emphasis added).

³³ The title of article 2 of the AI Act is, in fact, ‘Scope’.

³⁴ AI Act, art 2(1) (‘This Regulation applies to: (a) providers placing on the market or putting into service AI systems or placing on the market general-purpose AI models in the Union, irrespective of whether those providers are established or located within the Union or in a third country’).

³⁵ AI Act, art 2(1)(a).

³⁶ *ibid.*

³⁷ *ibid.* See also AI Act, recital 21 (‘the rules established by this Regulation should apply to providers of AI systems in a non-discriminatory manner, *irrespective of whether they are established within the Union or in a third country*, and to deployers of AI systems established within the Union’) (emphasis added). See, for instance, Patrick Van Eecke and Bartholomäus Regenhardt, ‘Article 2: Scope’ in Ceyhan Necati Pehlivan, Nikolaus Forgó and Peggy Valecke (eds), *The EU Artificial Intelligence (AI) Act: A Commentary* (Wolters Kluwer 2024) 35; Marco Almada and Anca Radu, ‘The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy’ (2024) 25 *German Law Journal* 646, 656; Michał Czerniawski, ‘Towards the Effective Extraterritorial Enforcement of the AI Act’ *Privacy Symposium 2024 (2025)* 35, 39. Also see the commentary on Article 2, Section 2.1.3. in this work.

12. Second, ‘putting into service’ is defined by Article 3(11) as ‘the *supply* of an AI system for first use directly to the deployer or *for own use in the Union* for its intended purpose’.³⁸ While the AI Act does not define ‘supply’, it is possible to infer that its meaning corresponds to ‘making available’.³⁹ Article 3(3) clarifies that ‘putting into service’ an AI system can be ‘for payment or free of charge’.⁴⁰ Therefore, given these reference points, ‘putting into service’ can be interpreted as including a provider’s act of making available an AI system for the provider’s own use in the Union.⁴¹
13. Third, ‘AI system’ is defined by Article 3(1) 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, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments’.⁴² Importantly, the definition of AI systems does not include any restrictions on who, amongst internal and external users, can give ‘input’ to an AI system and/or use its ‘outputs’.⁴³ In other words, the definition of ‘AI system’ in Article 3(1) does not differentiate between external and internal AI systems. Therefore, at least in theory, the definition of an AI system could include *both external and internal AI systems*.
14. By collating these various definitions, the expression ‘[p]roviders [...] putting into service AI systems’⁴⁴ can be reasonably interpreted to include EU-based or foreign AI developers making available an internal AI system for their own use in the Union. In other words, Article 2(1)(a) could arguably be interpreted to include internal AI systems within the scope of the AI Act. Indeed, absent any express exclusion of internal use from the scope of the AI Act (as identified in Article 2 AI Act),⁴⁵ ‘[p]roviders [...] putting into service AI systems’ could arguably be interpreted to include AI systems that EU-based or foreign AI developers deploy for internal use in the Union.⁴⁶

³⁸ See Commission Notice, The ‘Blue Guide’ on the Implementation of EU Product Rules 2022 [2022] OJ C247/1 (“Blue Guide”), s 2.6, which clarifies that ‘putting into service’ ‘results in the scope of Union harmonisation legislation being extended beyond the moment of making available of a product’. Specifically, ‘[p]utting into service takes place at the moment of first use within the Union by the end user for the purposes for which it was intended’. See, similarly, the commentary on Article 2, Section 2.1.2. in this work.

³⁹ See, for instance: (i) under the Blue Guide (n 38) s 2.2 states that ‘[s]upplying a product is only considered as *making available* on the Union market, when the product is intended for end use on the Union market’ (emphasis added); and (ii) the definition of ‘making available on the market’ is ‘supply’ in Regulation (EU) 2019/1020 of the European Parliament and of the Council of 20 June 2019 on market surveillance and compliance of products and amending Directive 2004/42/EC and Regulations (EC) No 765/2008 and (EU) No 305/2011 [2019] OJ L 169/1, art 3(1), Regulation (EC) No 765/2008 of the European Parliament and of the Council of 9 July 2008 setting out the requirements for accreditation and market surveillance relating to the marketing of products and repealing Regulation (EEC) No 339/93 (Text with EEA relevance) [2008] OJ L 218/30, art 2.1, and Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC [2017] OJ L 117/1, art 2(27).

⁴⁰ The AI Act also defines ‘intended purpose’. Under article 3(12), ‘intended purpose’ means ‘the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation’. Therefore, based on this definition, if a provider intended to deploy an AI system internally, arguably that is the system’s intended purpose.

⁴¹ It remains unclear how ‘in the Union’ will be interpreted. See AI Act, art 3(11).

⁴² See also European Commission, ‘Communication from the Commission - Commission Guidelines on the definition of an AI system established by Regulation (EU) 2024/1689 (AI Act) C(2025) 5053 final.

⁴³ AI Act, art 3(1).

⁴⁴ AI Act, art 2(1)(a).

⁴⁵ See also, for reference, the Code of Practice, Safety and Security Chapter (n 25) Glossary, defining ‘use (of a model)’ simply as ‘use of the model by the Signatory or other actors’.

⁴⁶ For instance, AI developers making internal AI systems accessible to staff based in their EU offices.

2.1.2. Article 2(1)(a): ‘Providers placing on the market [...] AI systems’

15. Section 2.1.1. started examining the scope of the AI Act, concentrating on Article 2(1)(a). As mentioned above, under Article 2(1)(a), the AI Act ‘applies to’ ‘[p]roviders [...] putting into service AI systems’.⁴⁷ Under the same Article, the AI Act also applies to ‘[p]roviders placing on the market [...] AI systems’.⁴⁸ Building on the considerations on ‘providers’ and ‘AI system’ developed in Section 2.1.1., this section explores whether ‘placing on the market’⁴⁹ could potentially refer to internally deployed AI systems. In doing so, this section examines Article 3(9)–(10) and assesses whether these articles support a broad interpretation of ‘placing on the market’ that includes internal AI systems.
16. Article 3(9) defines ‘placing on the market’ as ‘the *first making available* of an AI system or a general-purpose AI model *on the Union market*’.⁵⁰ Article 3(10) further defines ‘making available on the market’ as ‘the *supply* of an AI system or a general-purpose AI model for distribution or *use on the Union market in the course of a commercial activity*, whether in return for payment or *free of charge*’.⁵¹ As mentioned in Section 2.1.1., ‘supply’ can be interpreted as ‘making available’. The AI Act does not define ‘use’, which other relevant EU frameworks define broadly.⁵² Neither does the AI Act define ‘in the course of a commercial activity’.⁵³ However, to distil its meaning it is possible to draw inspiration from the European Commission’s Blue Guide on the implementation of EU product rules,⁵⁴ as well as other EU legal frameworks, such as for instance Regulation (EU) 2023/1115.⁵⁵ Within the first framework, ‘[c]ommercial activity is understood as providing goods *in a business related context*’.⁵⁶ The second defines ‘in the course of a commercial activity’ as ‘*for use in the business of the operator or trader itself*’.⁵⁷ In other words, ‘in the course of a commercial activity’⁵⁸ does not necessarily mean ‘for sale’ (especially considering that Articles 3(3) and 3(10) clarify that the ‘supply’

⁴⁷ AI Act, art 2(1) (‘This Regulation applies to: (a) providers placing on the market or putting into service AI systems or placing on the market general-purpose AI models in the Union, irrespective of whether those providers are established or located within the Union or in a third country’).

⁴⁸ AI Act, art 2(1)(a).

⁴⁹ AI Act, art 2(1)(a).

⁵⁰ AI Act, art 3(9) (emphasis added).

⁵¹ AI Act, art 3(10) (emphasis added).

⁵² See, for instance, Blue Guide (n 38) s 2.2 (defining ‘use’ as ‘the intended purpose of the product as defined by the manufacturer under conditions which can be reasonably foreseen’); European Commission, ‘Communication from the Commission - Commission Guidelines on prohibited AI practices established by Regulation (EU) 2024/1689 (AI Act)’ C(2025) 5052 final, s 2.3(14), 5 (‘While the “use” of an AI system is not explicitly defined in the AI Act, it should be understood in a broad manner to cover the use or deployment of the system at *any moment of its lifecycle after having been placed on the market or put into service*’) (emphasis added). Therefore, to distil the meaning of ‘use’ it is possible to draw inspiration from these frameworks.

⁵³ AI Act, art 3(10).

⁵⁴ See Blue Guide (n 38). For clarity, the Blue Guide may be an appropriate reference because ‘[t]he AI Act is rooted in classic EU product safety law’ (Claire Boine and David Rolnick, ‘Why The AI Act Fails to Understand Generative AI’ (2025) 26 *Minnesota Journal of Law, Science & Technology* 61, 84). See also Arto Lanamäki, Karin Väyrynen and Fanny Vainionpää, ‘The European Union’s Regulatory Challenge: Conceptualizing Purpose in Artificial Intelligence’ ECIS 2024 Proceedings (2024), 5 (discussing how the first proposal of AI Act by the European Commission in 2021 was based on EU product safety principles).

⁵⁵ See Regulation (EU) 2023/1115 of the European Parliament and of the Council of 31 May 2023 on the making available on the Union market and the export from the Union of certain commodities and products associated with deforestation and forest degradation and repealing Regulation (EU) No 995/2010 [2023] OJ L 150/206, art 2(19).

⁵⁶ Blue Guide (n 38) s 2.2 (emphasis added). It also clarifies that ‘[n]on-profit organisations may be considered as carrying out commercial activities if they operate in such a context. This can only be appreciated on a case by case basis taking into account the regularity of the supplies, the characteristics of the product, the intentions of the supplier, etc.’. See, similarly, the commentary on Article 2, para 14 in this work.

⁵⁷ See Regulation (EU) 2023/1115 (n 55) art 2(19).

⁵⁸ AI Act, art 3(10).

can be ‘free of charge’, which is not the case within other regulatory frameworks⁵⁹). Goods can be used in the business of the operator or trader itself for many different purposes other than sale, including for improving internal business procedures or increasing efficiency in ways that enhance a company’s competitiveness. Therefore, based on the definition contained in Article 3(9), as interpreted through Article 3(10), ‘placing on the market’ in Article 2(1)(a) could be interpreted as making available an AI system for use on the Union market⁶⁰ in a business-related context, *including for use in the business of the provider itself*. This seems confirmed, among other things, by the European Commission’s Guidelines on the scope of the obligations for GPAI models, which, in listing ‘examples of placing on the market of general-purpose AI models’, include a GPAI model being ‘used for *internal processes*’.⁶¹ Hence, Article 2(1)(a) could arguably support the inclusion of internal AI system use within the scope of the AI Act.

17. While this interpretation closely reflects the content of Article 3(9)–(10), some doubts remain as to whether this interpretation sufficiently accounts for the natural meaning of ‘placing on the market’.⁶² Ultimately, if this interpretation of ‘placing on the market’ were to prevail, the only difference with ‘putting into service’ would be that the latter could refer to any and all contexts, including ones that are not business-related, whereas ‘placing on the market’ would cover only business-related contexts.⁶³
18. In conclusion, while the common meaning of ‘placing on the market’ does not immediately present as encompassing the internal deployment of AI systems, it could arguably be interpreted in a sufficiently extensive way so as to cover an EU or foreign provider’s act of making an AI system available on the EU market for access and use by its EU staff. Indeed, most internal uses of AI systems by providers plausibly fall within a business-related context. For instance, as mentioned in Section 1, an internal system could support a provider’s staff by effectively outsourcing portions of existing or new workstreams, and, in a not-so-distant future, internal AI models and systems might start automating work.⁶⁴

2.1.3. Article 2(1)(b): ‘Deployers of AI systems that have their place of establishment or are located within the Union’

19. Under Article 2(1)(b), the AI Act ‘applies to’ ‘[d]eployers of AI systems that have their place of establishment or are located within the Union’. This section examines whether and to what extent Article 2(1)(b) could potentially include AI systems deployed within the AI providers developing them.

⁵⁹ Other EU legal frameworks take a different approach. See, for instance, Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC [2024] OJ L 2853/1, recital 14 (‘the supply of free and open-source software by non-profit organisations should not be considered as taking place in a business-related context, unless such supply occurs in the course of a commercial activity. However, where software is supplied in *exchange for a price*, or for *personal data* used other than exclusively for improving the security, compatibility or interoperability of the software, and is therefore supplied in the course of a commercial activity, this Directive should apply’ emphasis added).

⁶⁰ It remains unclear how ‘on the Union market’ will be interpreted. See AI Act, art 3(9).

⁶¹ See European Commission, ‘Communication from the Commission - Commission Guidelines on the Scope of the Obligations for Providers of General-Purpose AI Models Established by Regulation (EU) 2024/1689 (AI Act)’ C(2025) 7719 final, s 3.1.2, 18). See, similarly, the commentary on Article 2, Section 2.1.1.2. in this work.

⁶² AI Act, art 2(1)(a). These doubts could also potentially be corroborated by the fact that article 3(11) (defining ‘putting into service’) refers to ‘use *in the Union*’ (emphasis added), whereas article 3(10) (defining ‘making available on the market’) refers to ‘use *on the Union market*’ (emphasis added).

⁶³ See, similarly, the commentary on Article 2, para 56 in this work.

⁶⁴ See n 14.

20. Article 3(4) and Recital 13 define ‘deployer’ as ‘a natural or legal person, public authority, agency or other body *using* an AI system under its authority’.⁶⁵ Indirectly, this means that ‘deploying an AI system’ is generically defined as ‘using an AI system’. Neither Article 3(4) nor Recital 13 restrict in any way the type of ‘use’ that a natural or legal person can make of AI systems to qualify as a deployer. Therefore, under Article 3(4) and Recital 13 a provider and/or a provider’s staff using AI systems internally could qualify as a ‘deployer’.
21. Importantly, Article 2(1)(b) limits the geographic area of relevance to deployers (i.e., users) that ‘have their place of establishment or are located within the Union’.⁶⁶ By contrast, Article 2(1)(b) does not set any jurisdictional limits as to where the AI system is developed. In other words, while the deployer of an AI system must be established or located in the EU under Article 2(1)(b), the provider of that system can be established or located elsewhere.⁶⁷ Article 2(1)(b) is geographically constrained as to system deployment (i.e., use) but jurisdiction-agnostic as to its development.
22. In conclusion, the definition of ‘deployment’ as, generically, the ‘use’ of an AI system could be interpreted to encompass any and all potential internal uses that a EU-based AI provider or the EU-based staff of a foreign provider could make of an AI system.⁶⁸ This conclusion appears to be supported by the text of the Code of Practice, which utilizes the same verb (‘using’) to describe one possible form of internal deployment – autonomous AI R&D.⁶⁹ Specifically, Measure 7.1 acknowledges that an AI system can be ‘*used*’ ‘in the *development*, oversight, and/or evaluation of models’.⁷⁰ In other words, an AI system can be ‘used’ (i.e., deployed) for autonomous AI R&D, which is one possible use of internally deployed systems.

2.1.4. Article 2(1)(c): ‘output produced by the AI system is used in the Union’

23. Under Article 2(1)(c), the AI Act applies to ‘providers and deployers of AI systems that have their place of establishment or are located in a *third country*, where the output produced by the AI system is used in the Union’.⁷¹ Building on the conclusions reached in the previous Sections 2.1.1–2.1.3 regarding the terms ‘provider’, ‘deployer’ and ‘AI system’, this section examines whether ‘the output produced by the AI system is used in the Union’ could be interpreted to support the inclusion of internally deployed AI systems within the scope of the AI Act.
24. According to Recital 12, ‘outputs generated by the AI system reflect different functions performed by AI systems and include *predictions, content, recommendations or decisions*’.⁷² Neither this description of ‘output’ nor the definition of ‘AI system’⁷³ offer any element to reasonably exclude internal AI systems. Like external systems, internal systems can also output ‘predictions, content,

⁶⁵ AI Act, art 3(4) and recital 13 (emphasis added).

⁶⁶ AI Act, art 2(1)(b).

⁶⁷ See Section 2.1.1. Therefore, for example, members of the technical staff of a foreign AI provider who are located in the EU could potentially qualify as ‘deployer’ if they use an internal coding agent developed by the AI provider overseas and then made available to staff globally to support an existing workflow.

⁶⁸ See some examples of internal use in Section 1.

⁶⁹ Code of Practice, Safety and Security Chapter (n 25) Measure 7.1 (‘Signatories will provide in the Model Report: [...] a description of how the model has been *used* and is expected to be used, including its *use in the development, oversight, and/or evaluation of models*; [...] a description of the model versions that are going to be made or are currently made available on the market and/or used, including differences in systemic risk mitigations and systemic risks’) (emphasis added).

⁷⁰ Code of Practice, Safety and Security Chapter (n 25) Measure 7.1.

⁷¹ AI Act, art 2(1)(c) (emphasis added).

⁷² AI Act, art 12 (emphasis added).

⁷³ See Section 2.1.1.

recommendations or decisions’.⁷⁴ Therefore, it seems reasonable to conclude that ‘output produced by the AI system’ could also include the output produced by internal AI systems deployed by a provider.

25. Compared to Article 2(1)(b),⁷⁵ which concerns deployers established or located within the EU, Article 2(1)(c) concerns ‘providers and deployers [...] that have their place of establishment or are located in a *third country*’.⁷⁶ This means that, if ‘output produced by the AI system’ were interpreted to include the outputs of internal AI systems, Article 2(1)(c) would cover foreign AI providers developing and deploying an internal AI system outside of the EU and sharing the output within the EU. This could occur, for example, in cases in which (i) a foreign provider develops a highly capable AI system outside of the EU; (ii) the same foreign provider deploys internally that AI system outside of the EU to accelerate or automate AI R&D; (iii) once deployed, the provider shares the research output generated by this internal AI R&D system with staff located in the EU; and (iv) this research output is used by the provider’s staff within the EU.

2.2. Internal GPAI Models

26. Section 2.1. examined potential arguments supporting the inclusion of internal *AI systems* within the scope of the AI Act. Section 2.2. turns to internal *GPAI models* and discusses whether internal GPAI models could be covered by the AI Act, focusing on two pathways: (i) the placement on the market of internal GPAI models via their integration into internal AI systems, pursuant to Article 2(1)(a) and Recital 97 AI Act, and (ii) an extensive interpretation of the expression ‘placing on the market’ in Article 2(1)(a) AI Act as ‘making a GPAI model available for use on the Union market in a business-related context, including for use in the business of the provider itself’.⁷⁷
27. Under Article 2(1)(a), the AI Act ‘applies to’ ‘providers placing on the market or putting into service AI systems *or placing on the market general-purpose AI models in the Union*’.⁷⁸ Based on a summary reading, this article could be interpreted to exclude the application of the AI Act to internal GPAI models. Specifically, Article 2(1)(a) could be read to indicate that (i) while AI systems fall within the scope of the AI Act once they are ‘put into service’, the AI Act applies to GPAI models only after they are ‘placed on the market’,⁷⁹ and, therefore, (ii) if deploying (i.e., ‘putting into service’) an AI system internally could potentially trigger the application of the AI Act,⁸⁰ the same is not true about GPAI models (which are relevant only once they are ‘placed on the market’). In other words, since Article 2(1)(a) does not use the expression ‘putting into service’ with direct reference to GPAI models, deploying (i.e., ‘putting into service’) a GPAI model internally would not trigger the application of the AI Act in the same way that deploying (i.e., ‘putting into service’) an internal AI system does.⁸¹ To support this interpretation, it would be possible to rely on Article 3(63), which excludes from the definition of GPAI model ‘AI models that are used for research, development or prototyping

⁷⁴ AI Act, recital 12.

⁷⁵ See Section 2.1.3.

⁷⁶ AI Act, art 2(1)(c) (emphasis added).

⁷⁷ See Section 2.1.2.

⁷⁸ AI Act, art 2(1)(a) (emphasis added).

⁷⁹ AI Act, art 2(1)(a).

⁸⁰ As argued in Section 2.1.

⁸¹ See, similarly, the commentary on Article 2, Section 2.1.2. in this work. For clarity, this argument would refer only to internal GPAI models (rather than AI systems embedding GPAI models). AI systems integrating GPAI models could remain covered through article 2(1)(a)–(c). See Sections 2.1.1.–2.1.4.; see also the analysis of recital 97 in Section 3.1.

activities *before they are placed on the market*,⁸² and Recital 97. This latter recital specifies that ‘the obligations for the providers of general-purpose AI models should apply *once* the general-purpose AI models are *placed on the market*, and that ‘[t]he definition [of GPAI model] should *not cover AI models used before their placing on the market for the sole purpose of research, development and prototyping activities*’.⁸³

28. Despite these arguments, there are reasons why internal GPAI models could still fall within the scope of the AI Act. First, as mentioned in Section 2.1.2., ‘placing on the market’ a GPAI model⁸⁴ could potentially be interpreted extensively as ‘making a GPAI model available for use on the Union market in a business-related context, including for use in the business of the provider itself’. If ‘placing on the market’ were to be interpreted in this extensive way, the internal deployment of GPAI models could still be subject to the AI Act under Article 2(1)(a).⁸⁵
29. Second, and *more persuasively*: (i) it is plausible that, at least within frontier AI companies, internal deployment usually concerns GPAI models that are integrated into internal AI systems and given all the necessary affordances and permissions, and therefore, (ii) after system integration (and once the relevant system is ‘put into service’ internally⁸⁶) these GPAI models should also be considered as ‘placed on the market’ according to Recital 97. Furthermore, after AI system integration, the provisions described in Section 2.1. with regard to AI systems would become in scope. These logical steps will be explained in further detail below.
30. Internal deployment often concerns GPAI models.⁸⁷ It is a common intuition amongst AI researchers that the AI models and systems that providers elect to deploy internally are often the most capable at any given time.⁸⁸ This is particularly true when internal AI models and systems are used for autonomous AI R&D. This intuition appears to be confirmed in the text of the Code of Practice, which considers ‘capabilities to automate AI research and development’ in GPAI models as ‘sources of systemic risks’⁸⁹ and acknowledges that GPAI models can be ‘used’ ‘in the development, oversight, and/or evaluation of models’.⁹⁰ In other words, according to the Code of Practice, the capabilities

⁸² AI Act, art 3(63) (“general-purpose AI model” means an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, *except AI models that are used for research, development or prototyping activities before they are placed on the market*) (emphasis added).

⁸³ AI Act, recital 97 (emphasis added). Recital 97 also clarifies that ‘[t]he obligations laid down for models should in any case not apply when an own model is used for purely internal processes that are not essential for providing a product or a service to third parties and the rights of natural persons are not affected.’

⁸⁴ AI Act, art 2(1)(a).

⁸⁵ As mentioned in Section 2.1.2., this interpretation may risk eroding most of the difference between ‘placing on the market’ and ‘putting into service,’ and for this reason may not be pursued.

⁸⁶ AI Act, art 2(1)(a). See Section 2.1.1.

⁸⁷ As a purely illustrative example, Google DeepMind recently announced a coding agent, called AlphaEvolve, which integrates Gemini models (AlphaEvolve Team, ‘AlphaEvolve: A Gemini-Powered Coding Agent for Designing Advanced Algorithms’ (*Google DeepMind*, 14 May 2025) <<https://deepmind.google/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/>> accessed 2 June 2026). Based on information published by the provider, ‘AlphaEvolve enhanced the efficiency of Google’s data centers, chip design and AI training processes – including training the large language models underlying AlphaEvolve itself’.

⁸⁸ See Stix and others ‘AI Behind Closed Doors’ (n 2) 7–9; Acharya and Delaney (n 2) 8 (‘the best AIs at any given time are often internal [...] may be significantly better than public ones’). See n 7. See also Ajeya Cotra and others, ‘Frontier Risk Report (February to March 2026)’ (METR, 19 May 2026) <<https://metr.org/blog/2026-05-19-frontier-risk-report/>> accessed 2 June 2026 (reporting that ‘internal frontier on average ~2 months ahead of public frontier’).

⁸⁹ Code of Practice, Safety and Security Chapter (n 25) app 1.3.

⁹⁰ Code of Practice, Safety and Security Chapter (n 25) Measure 7.1. See also AI Act, recital 110.

necessary to automate AI R&D can be found in GPAI models. Of course, there are exceptions, as AI providers may also deploy internally models that are not GPAI models.⁹¹

31. Internal deployment, and specifically autonomous AI R&D, necessitates or benefits from integration of GPAI models into AI systems. To be maximally useful to its providers for a variety of potential internal uses, including autonomous AI R&D, a GPAI model does not only need to be sufficiently capable.⁹² A GPAI model also needs to be sufficiently ‘enabled’ through affordances and permissions to influence the real world.⁹³ This ‘enablement’ occurs by integrating AI models (including GPAI models) into existing or new AI systems. As the International Scientific Report on the Safety of Advanced AI explains, an AI system is ‘[a]n integrated setup that combines one or more AI models with other components’.⁹⁴ In other words, AI systems are a ‘slight generalization of AI models’ that includes not only the weights and architecture of an AI model but also a ‘broader set of system parameters’, which strongly influence the capabilities of AI systems.⁹⁵ This reflects the definition of ‘AI system’ contained in Article 3(1) AI Act, according to which AI systems are ‘machine-based system[s]’ that ‘operate with varying levels of autonomy’ and ‘can influence physical or virtual environments’.⁹⁶ The European Commission’s Guidelines on the definition of ‘AI system’ further clarify that ‘[a]ll AI systems’ ‘require’ ‘model training’ as well as ‘hardware’ and ‘software components’ in order to interact with their external environment, including ‘input/output interfaces’.⁹⁷ Without system parameters and environmental resources and opportunities for affecting the world that are available to AI systems (often referred to as ‘affordances’⁹⁸), AI systems could hardly ‘influence physical or virtual environments’.⁹⁹ For example, in order to automate AI R&D, internal AI systems may need access to multiple resources, including compute, energy, sensors to gather information about the environment, and a wide range of information (e.g., a provider’s code including algorithmic information, as well as a provider’s safety, security and technical oversight mechanisms).¹⁰⁰
32. GPAI models are considered as ‘placed on the market’ after they are integrated into an AI system (and the AI system is ‘put into service’). Recital 97 clarifies that ‘[w]hen the provider of a general-purpose AI model *integrates an own model into its own AI system* that is made available on the market or put into service,¹⁰¹ that model *should be considered to be placed on the market* and, therefore, the obligations in this Regulation for models should continue to apply in addition to those

⁹¹ An illustrative example of non-GPAI models being deployed internally could be the use of AlphaFold by Google DeepMind and Isomorphic Labs (Isomorphic Labs, ‘A Glimpse of the Next Generation of AlphaFold’ (*Isomorphic Labs*, 2023) <<https://www.isomorphiclabs.com/articles/a-glimpse-of-the-next-generation-of-alpha-fold>> accessed 2 June 2026).

⁹² For an overview on the relationship between affordances and capabilities, see Charlotte Stix and others, ‘The Loss of Control Playbook: Degrees, Dynamics, and Preparedness’ (arXiv, 9 December 2025) <<https://doi.org/10.48550/arXiv.2511.15846>> accessed 2 June 2026, 24. On internal AI models and systems specifically, see also Stix and others, ‘AI Behind Closed Doors’ (n 2) 10–12 (describing read, write and execute permissions that internal AI models and systems may necessitate with respect to their own hardware, weights, architecture, training, or oversight mechanisms or on those of their successors).

⁹³ See n 92.

⁹⁴ Bengio and others (n 2) 31.

⁹⁵ Lee Sharkey and others, ‘A Causal Framework for AI Regulation and Auditing’ (Computer Science and Mathematics, 18 January 2024) <<https://doi.org/10.20944/preprints202401.1424.v1>> accessed 2 June 2026, 4.

⁹⁶ AI Act, art 3(1). See also AI Act, recital 12.

⁹⁷ See European Commission, ‘Commission Guidelines on the definition of an AI system’ (n 42) s II.(8)–(13), 2–3.

⁹⁸ Sharkey and others (n 95) 5.

⁹⁹ AI Act, art 3(1).

¹⁰⁰ Stix and others, ‘AI Behind Closed Doors’ (n 2) 10–12.

¹⁰¹ See Sections 2.1.1.–2.1.2., discussing arguments in favour of considering internal AI systems as put in service (Section 2.1.1.) and placed on the market under AI Act, art 2(1)(a) (section 2.1.2).

for AI systems.¹⁰² In other words, once a GPAI model is integrated into an AI system¹⁰³ and that AI system is put into service, the AI Act considers the GPAI model as ‘placed on the market’,¹⁰⁴ and the GPAI model and relevant AI systems consequently fall within the scope of the AI Act under Article 2(1)(a).¹⁰⁵

33. Therefore, under this argument, even if GPAI models were not subject to the AI Act before their placement on the market, they would still fall within the remit of the AI Act once they are integrated into an AI system for the purposes of deploying them internally (because that would constitute market placement under the AI Act). In other words, under this argument, Article 2(1)(a) covers GPAI models from the moment they are integrated into an AI system which is put into service. By contrast, internal GPAI models that are *not* integrated into an internal AI system that is put into service would not be considered as placed on the market and, therefore, would remain outside the scope of the AI Act *before* public release (or AI system integration). Based on the considerations above regarding the correlation between system integration and the opportunity for a GPAI model to influence the real world, the scope of a potential carve-out of internal GPAI models from the scope of the AI Act under Article 2(1)(a) is rather limited.¹⁰⁶ This is particularly true considering that, on average, system integration occurs *before* (and not after) public release.¹⁰⁷ In other words, once a GPAI model is placed on the market, chances are it has already been integrated into an AI system.
34. Third, and finally, if Article 2(1)(a) were interpreted to exclude internally deployed GPAI models (along with the relevant internal AI systems into which GPAI models may be integrated) from the scope of the AI Act, this interpretation could clash with the content and purpose of, among others, Article 2(1)(a)–(c) and Recital 97, and it could also create an enforcement distortion. As mentioned above, internal AI systems arguably fall within the scope of the AI Act under Article 2(1)(a)–(c).¹⁰⁸ If that is the case, and if internal GPAI models were considered to be excluded from the scope of the AI Act (along with the relevant internal AI systems into which GPAI models may be integrated), the AI Act (i) would apply to AI systems embedding AI models that are *not* GPAI models, that is, AI models that are less capable and potentially less concerning than GPAI models,¹⁰⁹ and (ii) would not apply to internally deployed GPAI models and systems, which are generally more capable and thus have a greater potential to pose threats, including systemic risk.¹¹⁰

¹⁰² AI Act, recital 97 (emphasis added). See also *Case 215/88 Casa Fleischhandels-GmbH v Bundesanstalt für landwirtschaftliche Marktordnung* [1989] ECLI:EU:C:1989:33189 (clarifying that ‘a recital in the preamble to a regulation may cast light on the interpretation to be given to a legal rule’ – arguably, article 3 AI Act in this case).

¹⁰³ See AI Act, recital 100 (‘When a general-purpose AI model is integrated into or forms part of an AI system, this system should be considered to be a general-purpose AI system when, due to this integration, this system has the capability to serve a variety of purposes’).

¹⁰⁴ AI Act, art 2(1)(a).

¹⁰⁵ Also see the commentary on Article 2, Section 2.1.1.3. in this work.

¹⁰⁶ More speculatively, this might also raise the question as to whether the EU legislator simply saw no need to include within article 2(1)(a) the option: ‘putting into service general-purpose AI models’.

¹⁰⁷ See Bengio and others (n 2) 30. Also see the commentary on Article 2, paras 29 and 58 in this work.

¹⁰⁸ See Sections 2.1.1.–2.1.3.

¹⁰⁹ See AI Act, art 3(63) (defining ‘general-purpose AI model’); AI Act, recital 110 (clarifying that ‘[g]eneral-purpose AI models could pose systemic risks’); AI Act, recital 97 (mentioning GPAI models’ ‘generality and [...] capability to competently perform a wide range of distinct tasks’).

¹¹⁰ See, similarly, the commentary on Article 2, paras 36 and 63 in this work. On the other hand, it should be noted that this outcome might be consistent with the fact that obligations on providers of GPAI models ‘can be considered a “light” version of the obligations for AI systems’ (Van Eecke and Regenhardt (n 37) 35).

3. Arguments opposing or restricting the application of the AI Act to internal deployment

35. Section 2 explored potential arguments supporting the inclusion of internal deployment within the scope of the AI Act. This section does the opposite: it explores and stress-tests potential arguments that support the exclusion of internal deployment from the scope of the AI Act or restrict the application of the AI Act to only some forms of internal uses. Specifically:

- Section 3.1. concentrates on the expression ‘AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development’ in Article 2(6).
- Section 3.2. concentrates on the expression ‘any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service’ in Article 2(8).

36. As it will be described below, the potential arguments based on these provisions hardly support a total exclusion of internal deployment from the scope of the AI Act but rather enable select carve-outs from the scope of the Act.¹¹¹

3.1. Article 2(6): ‘AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development.’

37. A first argument to restrict the inclusion of internal deployment within the scope of the AI Act is based on Article 2(6). This article arguably introduces the most interesting open question on the application of the AI Act to internal deployment.¹¹² Under Article 2(6), the AI Act ‘does not apply to AI systems *or* AI models, including their output, *specifically developed and put into service for the sole purpose of scientific research and development*’.¹¹³

38. Article 2(6) seemingly introduces an exception to the general applicability of the AI Act to internal deployment. In other words, if Article 2(1)(a)–(c) were interpreted as described in Section 2.1., Article 2(6) could be used to argue that, while internal deployment is generally covered by the AI Act, AI models and systems that are ‘specifically’ developed and deployed for ‘scientific research and development’ are exempted. This section addresses this exception and examines its scope.

39. Two expressions appear particularly relevant when interpreting Article 2(6): (i) ‘specifically developed and put into service’ and (ii) ‘sole purpose of scientific research and development’.¹¹⁴ First, ‘specifically developed and put into service’ is not defined by the AI Act. The word ‘specifically’ and the

¹¹¹ See Sections 3.2.–3.3. See also Figure 1.

¹¹² See, for instance, Van Eecke and Regenhardt (n 37) 44 (‘it is difficult to assess when research, testing and development is subjected to the scope of the AI Act’).

¹¹³ AI Act, art 2(6) (emphasis added). See also AI Act, recital 109 (‘Compliance with the obligations applicable to the providers of general-purpose AI models should be commensurate and proportionate to the type of model provider, *excluding the need for compliance for persons who develop or use models for [...] scientific research purposes*’ emphasis added).

¹¹⁴ AI Act, art 2(6).

conjunction ‘or’ (‘AI systems *or* AI models’)¹¹⁵ appear to suggest that Article 2(6) concerns only AI models or AI systems that are developed *ad hoc* to serve one exclusive purpose (i.e., the ‘sole purpose of scientific research and development’) and deployed exclusively for that purpose.¹¹⁶

40. Second, ‘sole purpose of scientific research and development’ is also not defined by the AI Act. ‘Sole’ contributes to the idea that AI models and systems that fall into the exception set by Article 2(6) must have only one purpose. In other words, ‘sole’ excludes from the spectrum of Article 2(6) those AI models and systems that are deployed for multiple purposes.¹¹⁷ ‘*Scientific* research and development’ must therefore be the only purpose of training and deploying. This is further confirmed by Recital 25, which clarifies that ‘without prejudice to the exclusion of AI systems specifically developed and put into service for the sole purpose of scientific research and development, *any other AI system that may be used for the conduct of any research and development* activity should remain subject to the provisions of this Regulation.’ In other words, not all AI models and systems deployed for research and development are exempt – only the ones specifically developed and trained for ‘*scientific*’ research and development.¹¹⁸
41. The considerations above allow us to reach an intermediate conclusion: Article 2(6) AI Act does not exempt internal deployment from the scope of the AI Act. If anything, Article 2(6) creates a narrow exception for internal ‘AI systems or AI models’ that a provider purposefully develops and deploys for one exclusive purpose: scientific research and development.¹¹⁹ The main challenge of Article 2(6), therefore, becomes understanding what ‘*scientific* research and development’ refers to in the context

¹¹⁵ Under article 2(6), the AI Act ‘does not apply to AI systems *or* AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development’ (emphasis added on the ‘or’ conjunction).

¹¹⁶ AI Act, recital 109 excludes ‘the need for compliance’ ‘with the obligations applicable to the providers of general-purpose AI models’ for ‘persons who develop *or* use models for [...] scientific research purposes’ (emphasis added). Through the ‘or’ conjunction, recital 109 seems to suggest that it would be sufficient for a GPAI model to be *deployed* for the sole purpose of scientific R&D in order to fall outside the scope of the AI Act. In other words, for the purposes of recital 109, a GPAI model would not need to be *developed and deployed* for scientific R&D, but, for instance, it would only need to be *deployed* for that purpose. This interpretation appears in conflict with EU law, as recital 109 would derogate from the actual provisions of the act in question (i.e., AI Act, art 2(6)). See, for instance, *Case C-423/23 Secab Soc. coop. v Autorità di Regolazione per Energia Reti e Ambiente (ARERA) and Gestore dei servizi energetici (GSE) SpA* [2026] ECLI:EU:C:2026:32, para 77; *Case C-664/23 Caisse d’allocations familiales des Hauts-de-Seine v TX* [2024] ECLI:EU:C:2024:1046; *C-302/19 Istituto nazionale della previdenza sociale v WS* [2020] ECLI:EU:C:2020:957; *Case C-418/18 P Patrick Grégor Puppink and Others v European Commission* [2019] ECLI:EU:C:2019:1113; *Case C-345/13 Karen Millen Fashions Ltd v Dunnes Stores and Dunnes Stores (Limerick) Ltd* [2014] ECLI:EU:C:2014:2013; *Case 136/04 Deutsches Milch-Kontor GmbH v Hauptzollamt Hamburg-Jonas* [2005] ECLI:EU:C:2005:716; *C-162/97 Criminal proceedings against Gunnar Nilsson, Per Olov Hagelgren and Solweig Arrborn* [1998] ECLI:EU:C:1998:554.

¹¹⁷ See, for instance, Philipp Hacker, Andreas Engel and Marco Mauer, ‘Regulating ChatGPT and Other Large Generative AI Models’ (arXiv, 5 February 2023) <<https://arxiv.org/abs/2302.02337v8>> accessed 2 June 2026, 6 (‘this research exemption arguably does not apply anymore once the system is released into the wild, as any public release likely does not have scientific research and development as its “sole purpose”’). See, similarly, the commentary on Article 2, Section 2.3.1. in this work.

¹¹⁸ On this point, see Michèle Finck, ‘In Search of the Lost Research Exemption: Reflections on the AI Act’ (2025) 74 GRUR International 903.

¹¹⁹ Notably, AI Act, art 2(6) uses the expression ‘AI systems *or* AI models’ (emphasis added). This creates an open question as to how the AI Act should be interpreted with regard to situations in which: (i) an AI *model* is not developed *ad hoc* for scientific research and development (e.g., a GPAI model); however, (ii) the AI *system* in which the model is integrated is developed *ad hoc* for scientific research and development. In this respect, the conjunction ‘or’ could suggest that if *either* one between the AI model or the AI system is ‘specifically developed and put into service for the sole purpose of scientific research and development’ (AI Act, art 2(6)), then that AI model or that AI system could fall outside of the AI Act’s scope. See, differently, the commentary on Article 2, Section 2.3.2. in this work.

of AI¹²⁰ and specifically whether it can be relied on to exclude from the AI Act's scope AI models and systems that are specifically trained and deployed to automate AI R&D.

42. Importantly, while not defining '*scientific* research and development', the AI Act appears to juxtapose '*scientific* research and development'¹²¹ to '*product-oriented* research'.¹²² In particular, Recital 25 provides that '[a]s regards *product-oriented* research, testing and development activity regarding AI systems or models, the provisions of this Regulation should also not apply *prior to those systems and models being put into service or placed on the market*'.¹²³ In doing so, Recital 25 clarifies that product-oriented research is in fact covered by the AI Act and is subject only to the exception set forth under Article 2(8).¹²⁴ Therefore, as will be explained below, *after* an AI model or system for product-oriented research is put into service (i.e., deployed internally or externally), it is subject to the obligations set out in the AI Act. To summarize, it appears that:

- The AI Act differentiates between *scientific* R&D and *product-oriented* R&D (Article 2(6) and Recital 25 AI Act).
- Only AI models or systems specifically developed and deployed internally with the sole, exclusive purpose of doing *scientific* R&D are exempted under Article 2(6) AI Act.
- By contrast, AI models and systems developed and deployed internally with the purpose of doing *product-oriented* R&D are not exempted (subject to Article 2(8) and Recital 25).¹²⁵

43. As mentioned above, the AI Act does not define '*scientific* research and development'.¹²⁶ The AI Act does not define '*product-oriented* research and development' either.¹²⁷ To shed light on the difference between these two kinds of research and development, a helpful point of reference could be Regulation 1907/2006 of December 18, 2006 concerning the Registration, Evaluation, Authorisation and Restriction of Chemicals ("REACH").¹²⁸ Like the AI Act, the REACH Regulation similarly distinguishes between product and process orientated research and development ("PPORD") and scientific research and development ("SR&D") and exempts the latter from some of the obligations and restrictions set forth by the Regulation.¹²⁹ Essentially, just like the AI Act, the REACH Regulation has a dual approach to product-oriented research and scientific research.

¹²⁰ AI Act, art 2(6) (emphasis added). Importantly, it is also possible that '*scientific* research and development' may not refer to AI R&D at all, but rather and more broadly to projects of scientific nature.

¹²¹ AI Act, art 2(6) (emphasis added).

¹²² AI Act, recital 25.

¹²³ AI Act, recital 25 (emphasis added).

¹²⁴ See Section 3.2.

¹²⁵ See Section 3.2.

¹²⁶ AI Act, art 2(6).

¹²⁷ AI Act, recital 25.

¹²⁸ Other authors have offered parallels between the AI Act and the EU General Data Protection Regulation (GDPR; Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L 119/1). See Finck (n 118). Specifically, GDPR, recital 159 clarifies that 'the processing of personal data for scientific research purposes should be interpreted *in a broad manner* including for example *technological development and demonstration, fundamental research, applied research* and privately funded research' (emphasis added). For a similar comparative approach, see the commentary on Article 2, Section 2.3.3. in this work.

¹²⁹ Regulation (EC) No 1907/2006 of the European Parliament and of the Council of 18 December 2006 concerning the Registration, Evaluation, Authorisation and Restriction of Chemicals (REACH), establishing a European Chemicals Agency, amending Directive 1999/45/EC and repealing Council Regulation (EEC) No 793/93 and Commission Regulation (EC) No 1488/94 as well as Council Directive 76/769/EEC and Commission Directives 91/155/EEC, 93/67/EEC, 93/105/EC and 2000/21/EC [2006] OJ L 396/1 ("REACH Regulation"), arts 3(22)-(23) and 67(1).

44. Specifically, under the REACH Regulation, product and process orientated research and development is defined as ‘any scientific development *related to product development* or the further development of a substance’.¹³⁰ Relevant guidance by the European Chemical Agency (“ECHA”) clarifies that product-oriented research includes a wide range of activities aimed at developing or proving the feasibility of new products and processes and improving production efficiency. ECHA’s guidance provides the following examples of product-oriented research: ‘campaign(s) for the scaling-up or improvement of a production process in a pilot plant or in the full-scale production, or the investigation of the fields of applications for that substance’.¹³¹
45. Scientific research and development, on the other hand, is defined as ‘any scientific experimentation, analysis or [...] research carried out *under controlled conditions*’ within a certain volume (‘less than 1 tonne per year’).¹³² ECHA’s guidance clarifies that scientific research and development may include ‘experimental research or analytical activities *at a laboratory scale*’ ‘as well as the use of the substance in monitoring and routine quality control or *in vitro diagnostics* at a laboratory scale under controlled conditions’.¹³³ ECHA’s guidance also clarifies that research and development activity occurs ‘under controlled conditions’ if ‘procedures and measures are in place to minimise or control exposure and potential risks from exposure of humans and the environment to the substance’.¹³⁴
46. In other words, the REACH Regulation appears to differentiate product and process orientated research and development and scientific research and development through a combination of two main parameters.¹³⁵ First, there is the *scale* of the R&D activity. Scientific research and development involves quantities ‘less than 1 tonne per year’, ‘under controlled conditions’, and ‘at a laboratory scale’ (such as ‘in vitro diagnostics’),¹³⁶ whereas product and process orientated research and development can be either ‘in a pilot plant’ or in ‘full-scale production’.¹³⁷ Intuitions around the scientific research and development activity could be easily translated to scientific AI R&D: this type of AI R&D should be undertaken in controlled conditions and at laboratory scale (for example, in vitro).¹³⁸
47. Second, there is the *purpose* of the R&D activity. product and process orientated research and development aims at ‘product development or the further development of a substance’,¹³⁹ whereas scientific research and development aims at ‘experimentation, analysis’,¹⁴⁰ ‘monitoring and [...] diagnostics’.¹⁴¹ Translating these concepts to AI R&D, product-oriented research and development

¹³⁰ REACH Regulation (n 129) art 3(22).

¹³¹ European Chemicals Agency, *Guidance on Scientific Research and Development (SR&D) and Product and Process Orientated Research and Development (PPORD)* (European Chemicals Agency 2017) <<https://doi.org/10.2823/917878>> (“ECHA Guidance”) 7–8 (emphasis added).

¹³² REACH Regulation (n 129) art 3(23). See also REACH Regulation (n 129) recital 28 (‘scientific research and development normally takes place in quantities below 1 tonne per year’).

¹³³ ECHA Guidance (n 131) 7.

¹³⁴ ECHA Guidance (n 131) 8 (‘This may include, for example, limitation of uses to qualified persons having access to the substance, or collection and disposal of waste’).

¹³⁵ The REACH Regulation (n 129) does not appear to clarify if any of these two factors should be prioritized.

¹³⁶ REACH Regulation (n 129) art 3(23) and recital 28; ECHA Guidance (n 131) 8.

¹³⁷ ECHA Guidance (n 131) 7 (emphasis added).

¹³⁸ Arguably, the fact that scientific R&D occurs in ‘controlled conditions’ may be one of the reasons why this type of R&D is less regulated than product-oriented R&D. If a scientific experiment is undertaken in properly controlled conditions, the negative impact of potential incidents would plausibly be more contained than if conditions were not controlled.

¹³⁹ REACH Regulation (n 129) article 3(22).

¹⁴⁰ REACH Regulation (n 129) article 3(23). See also REACH Regulation (n 129) recital 28 (‘scientific research and development normally takes place in quantities below 1 tonne per year’).

¹⁴¹ ECHA Guidance (n 131) 7.

could be interpreted to include any scientific development related to the enhancement and scaling of AI capabilities, AI products, and/or the training pipeline;¹⁴² by contrast, scientific R&D could be AI R&D that is not product-oriented.

48. These two parameters (i.e., the scale and the purpose of the R&D activity) could help orient the interpretation of ‘scientific research and development’ in Article 2(6) AI Act, offering an initial point of reference for determining what research and development activities are *scientific* R&D and what are instead *product-oriented* R&D. For instance, by applying these two parameters, it would be possible to observe that the parameter of *scale* for scientific research and development as extrapolated from the REACH Regulation (i.e., ‘under controlled conditions’, and ‘at a laboratory scale’ such as ‘in vitro’¹⁴³) would be met, for instance, by what AI researchers define as ‘model organisms’.¹⁴⁴ Model organisms are, literally, ‘*in vitro* demonstrations of the kinds of failures that might pose existential threats’.¹⁴⁵ Another tentative example of research and development activities ‘under controlled conditions’ and ‘at a laboratory scale’ is the training of AI models or systems for mechanistic interpretability research.¹⁴⁶
49. Furthermore, these two illustrative examples (i.e., model organisms and mechanistic interpretability) also fit the parameter of *purpose* as extrapolated from the REACH Regulation. Both model organisms and mechanistic interpretability are usually detached from (or, at least, not directly aimed at) the enhancement and scaling of AI capabilities, AI products, and/or the training pipeline. Model organisms are used to offer scientific proof of concept, within controlled environments, of the existence or likelihood of specific behaviours in AI models and systems.¹⁴⁷ Recent examples include AI researchers: training ‘backdoored’ model organisms, and then applying safety training, to evaluate whether the backdoor behaviour persists in these model organisms;¹⁴⁸ training model organisms to study the effects on misalignment of fine-tuning AI models on narrowly harmful datasets;¹⁴⁹ training an ‘evaluation-aware’ model organism to validate the effect of activation steering on evaluation awareness;¹⁵⁰ or training an AI model with a hidden objective to then undertake alignment audits.¹⁵¹

¹⁴² An illustrative example of scientific activity aimed at the enhancement and scaling of AI capabilities, AI products, and/or the training pipeline could be the discovery of reinforcement learning (RL) pipelines that patch existing vulnerabilities in current AI systems, or solutions to improve graphics processing unit (GPU) utilization.

¹⁴³ REACH Regulation (n 129) art 3(23) and recital 28; ECHA Guidance (n 131) 8.

¹⁴⁴ See Evan Hubinger and others, ‘Model Organisms of Misalignment: The Case for a New Pillar of Alignment Research’ (*AI Alignment Forum*, 7 August 2023) <<https://www.alignmentforum.org/posts/ChDH335ckdvpXaXX/model-organisms-of-misalignment-the-case-for-a-new-pillar-of-1>> accessed 2 June 2026; Evan Hubinger and others, ‘Sleeper Agents: Training Deceptive LLMs That Persist Through Safety Training’ (arXiv, 10 January 2024) <<https://arxiv.org/abs/2401.05566v3>> accessed 2 June 2026, 8–9.

¹⁴⁵ Hubinger and others, ‘Model Organisms of Misalignment’ (n 144). See also Hubinger and others, ‘Sleeper Agents’ (n 144).

¹⁴⁶ Mechanistic interpretability is a field of research focused on reverse engineer neural networks. See generally Sharkey and others (n 95).

¹⁴⁷ See footnotes 148–151.

¹⁴⁸ See Hubinger and others, ‘Sleeper Agents’ (n 144).

¹⁴⁹ See Jan Betley and others, ‘Emergent Misalignment: Narrow Finetuning Can Produce Broadly Misaligned LLMs’ (arXiv, 24 February 2025) <<https://doi.org/10.1038/s41586-025-09937-5>> accessed 2 June 2026; Edward Turner and others, ‘Model Organisms for Emergent Misalignment’ (arXiv, 13 June 2025) <<https://arxiv.org/abs/2506.11613v1>> accessed 2 June 2026; Miles Wang and others, ‘Persona Features Control Emergent Misalignment’ (arXiv, 24 June 2025) <<https://arxiv.org/abs/2506.19823v2>> accessed 2 June 2026.

¹⁵⁰ See Tim Tian Hua and others, ‘Steering Evaluation-Aware Language Models to Act Like They Are Deployed’ (arXiv, 2 March 2026) <<https://doi.org/10.48550/arXiv.2510.20487>> accessed 2 June 2026.

¹⁵¹ See Samuel Marks and others, ‘Auditing Language Models for Hidden Objectives’ (arXiv, 28 March 2025) <<https://doi.org/10.48550/arXiv.2503.10965>> accessed 2 June 2026.

Mechanistic interpretability research aims to reverse engineer a model’s computations and is often not immediately useful for product development.¹⁵² For instance, OpenAI recently trained AI models with an architecture similar to (now largely outdated) GPT-2 and forced most weights to be zero in order to ‘substantially disentangle[] the model’s internal computations’.¹⁵³ An additional example that can help shed light on an AI model used purely for scientific experimentation is the AI model that Google DeepMind developed in collaboration with Howard Hughes Medical Institute (“HHMI”) to simulate how a fruit fly walks, flies and behaves in order to understand how brain, body and environment drive specific behaviours in animals.¹⁵⁴

50. Despite being a helpful point of reference, the precise extent to which the definitions in the REACH Regulation and the two parameters of *scale* and *purpose* should inform the interpretation of scientific research and development and product-oriented research and development in the AI Act remains an open question. For instance, it is debatable if and to what extent this interpretation of *scale* could apply to scientific R&D enabled by highly advanced AI systems or AI models, considering that (i) there are multiple examples of non-product-oriented scientific research and development projects that are not only large but arguably on a truly grand-scale (for instance, the Large Hadron Collider at the CERN¹⁵⁵ or the Human Brain Project¹⁵⁶) and that (ii) the most advanced future AI systems or AI models could arguably have the potential to support these grand-scale projects and unlock some of the most complex existing scientific bottlenecks.¹⁵⁷ Indeed, there are already some early signs that advanced AI models and systems could significantly accelerate scientific discovery,¹⁵⁸ which the

¹⁵² See Leo Gao and others, ‘Weight-Sparse Transformers Have Interpretable Circuits’ (arXiv, 17 November 2025) <<https://doi.org/10.48550/arXiv.2511.13653>> accessed 2 June 2026.; Leo Gao and others, ‘Understanding Neural Networks through Sparse Circuits’ (*OpenAI*, 29 May 2026) <<https://openai.com/index/understanding-neural-networks-through-sparse-circuits/>> accessed 2 June 2026.

¹⁵³ See Gao and others, ‘Weight-Sparse Transformers Have Interpretable Circuits’ (n 152); Gao and others, ‘Understanding Neural Networks through Sparse Circuits’ (n 152). Another recent example is Chris Cundy and Adam Gleave, ‘Preference Learning with Lie Detectors Can Induce Honesty or Evasion’ (arXiv, 18 November 2025) <<https://doi.org/10.48550/arXiv.2505.13787>> accessed 2 June 2026, in which researchers trained a model ‘incorporating a lie detector into the labelling step of LLM post-training’ in order to ‘evaluat[e] whether the learned policy is genuinely more honest, or instead learns to fool the lie detector while remaining deceptive.’

¹⁵⁴ See Roman Vaxenburg and others, ‘Whole-Body Physics Simulation of Fruit Fly Locomotion’ (2025) 643 *Nature* 1312.

¹⁵⁵ See CERN, ‘Large Hadron Collider’ (*CERN*) <<https://home.cern/science/accelerators/large-hadron-collider/>> accessed 2 June 2026. The Atlas Collaboration authoring the paper that announced the discovery of the Higgs boson comprised around 3,000 authors (see G Aad and others, ‘Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC’ (2012) 716 *Physics Letters B* 1; ‘ATLAS Collaboration’ (*ATLAS Experiment*) <<https://atlas.cern/authors/atlas-collaboration>> accessed 2 June 2026).

¹⁵⁶ See ‘The Human Brain Project Ends: What Has Been Achieved’ (*Human Brain Project*, 2023) <<https://www.humanbrainproject.eu/en/follow-hbp/news/2023/09/28/human-brain-project-ends-what-has-been-achieved/>> accessed 2 June 2026.

¹⁵⁷ In this respect, see, for instance, President Donald Trump, ‘Launching the Genesis Mission (Executive Order 14363)’ (*The White House*, 24 November 2025) <<https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/>> accessed 2 June 2026, s 1, launching the ‘Genesis Mission’ in the United States ‘as a dedicated, coordinated national effort to unleash a new age of AI-accelerated innovation and discovery that can solve the most challenging problems of this century.’

¹⁵⁸ See, for instance, proteins’ complex structures prediction enabled by Google DeepMind’s AlphaFold (John Jumper and others, ‘Highly Accurate Protein Structure Prediction with AlphaFold’ (2021) 596 *Nature* 583), which led its developers to be awarded a Nobel prize (‘Nobel Prize in Chemistry 2024’ (*NobelPrize.org*, 9 October 2024) <<https://www.nobelprize.org/prizes/chemistry/2024/press-release/>> accessed 2 June 2026). More recently, see Sébastien Bubeck and others, ‘Early Science Acceleration Experiments with GPT-5’ (arXiv, 20 November 2025) <<https://doi.org/10.48550/arXiv.2511.16072>> accessed 2 June 2026, especially at 53 and following, for early science acceleration experiments enabled by OpenAI’s GPT-5. See also Sakana AI’s ‘AI Scientist,’ an early proof of concept for fully automatic scientific discovery in machine learning (Chris Lu and

European Union may want to harness to accomplish the objective under Article 179(1) of the Treaty on the Functioning of the EU (“TFEU”).¹⁵⁹ Furthermore, it should be noted that other EU legal frameworks (for instance, the General Data Protection Regulation [“GDPR”]) do not refer to scale when describing scientific R&D.¹⁶⁰

51. On the other hand, relying exclusively on the *purpose* of the R&D activity may ultimately render the distinction between scientific and product-oriented R&D hazy and unclear. For instance, it is unclear whether and to what extent current and future research delivering generalizable knowledge and methods will ultimately contribute to the training of future generations of AI models and systems.¹⁶¹
52. In conclusion, Article 2(6) could be interpreted to refer exclusively to AI models or AI systems that are trained and made available internally within an AI provider’s organization to carry out only ‘*scientific* research and development’. To understand the boundaries of scientific research and development, one could compare scientific research and development to product-oriented research and development, mentioned in Recital 25. Read in conjunction with Recital 25, Article 2(6) appears to single out internal AI models or systems custom-built and deployed to do research and development that is *not* product-oriented as the sole exception under this Article.
53. While the boundaries between product-oriented and scientific research and development remain blurry, it is possible to utilize the definitions of product and process orientated research and development and scientific research and development from the REACH Regulation as an initial point of reference. Examining these definitions suggests that *scale* and *purpose* are important parameters when distinguishing what qualifies as scientific R&D activity. The REACH Regulation definitions also help us extrapolate some illustrative examples of *scientific* (e.g., model organisms, mechanistic interpretability) and *product-oriented* (e.g., RL pipelines to patch vulnerabilities in current AI systems, or solutions to improve GPU utilization) AI R&D.

3.2. Article 2(8): ‘Any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service.’

54. A second argument to oppose or restrict the inclusion of internal deployment is based on Article 2(8) AI Act. Under Article 2(8), the AI Act ‘does not apply to any research, testing or development activity

others, ‘The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery’ (arXiv, 1 September 2024) <<https://doi.org/10.48550/arXiv.2408.06292>> accessed 2 June 2026s).

¹⁵⁹ Consolidated version of the Treaty on the Functioning of the European Union [2012] OJ C 326/47 (“TFEU”), art 179(1) (‘The Union shall have the objective of strengthening its scientific and technological bases by achieving a European research area in which researchers, scientific knowledge and technology circulate freely, and encouraging it to become more competitive, including in its industry, while promoting all the research activities deemed necessary by virtue of other Chapters of the Treaties.’).

¹⁶⁰ See GDPR, recital 159. See also n 128. See, more critically on a comparison with GDPR, the commentary on Article 2, para 73 in this work.

¹⁶¹ Consider, for instance, an AI model or system autonomously theorizing and testing a new learning paradigm (e.g., ‘Nested Learning’ (Ali Behrouz and others, ‘Nested Learning: The Illusion of Deep Learning Architecture’ (Google Research 2025) <<https://abehrouz.github.io/files/NL.pdf>> accessed 2 June 2026), which an AI provider then adopts and adapts to be the new architectural backbone of future generations of AI systems and models. See also Van Eecke and Regenhardt (n 37) 43 (‘a research project might have both scientific and practical applications, calling into doubt the applicability of the exemption’).

regarding AI systems or AI models *prior to* their being placed on the market or *put into service*.¹⁶² This section examines the scope of potential exceptions to the application of the AI Act to internal deployment based on this article.

55. A clarification is important on Article 2(8). By contrast with Article 2(6),¹⁶³ Article 2(8) *does not* concern the use of an AI model or system to generate research and development outputs. Instead, Article 2(8) concerns the ‘research, testing and development’ activity that *precedes* ‘their being placed on the market or put into service’ (i.e., their internal deployment or public release).¹⁶⁴ In other words, Article 2(6) concerns a provider’s research and development activities *before* an AI model or system is put into service.¹⁶⁵
56. This means that, rather than being an exception to Article 2(1)(a), Article 2(8) is in fact fully consistent with the ‘scope’ of the AI Act as delineated in Article 2(1).¹⁶⁶ Under Article 2(1), the AI Act comes into play once a provider puts into service (i.e., deploys) an AI system. Section 2.1.1. of this chapter explored how Article 2(1)(a) could be interpreted to include internal deployment of AI systems. Article 2(8) concerns the time before any type of deployment, either internal or external. This is consistent with the text of Recital 25, according to which ‘[t]hat exclusion [i.e., the exception under Article 2(8)] is without prejudice to the obligation to comply with this Regulation where an AI system falling into the scope of this Regulation is placed on the market or *put into service as a result of such research and development activity*’.¹⁶⁷ In other words, Recital 25 clarifies that, while research and development activity before (internal and external) deployment falls outside the scope of the AI Act, the deployment of the resulting AI system will trigger the application of the Act under Article 2(1)(a).¹⁶⁸
57. It remains unclear which precise pre-deployment ‘research, testing or development’ activities are exempted from the scope of the AI Act under Article 2(8).¹⁶⁹ Recital 25 also broadly refers to ‘scientific research and development activity [...] prior to being placed on the market or put into service’.¹⁷⁰ Before an AI model or system is ready for internal and/or external deployment, the typical AI development cycle includes (i) experimentation and planning (for instance, around model architectures, training algorithms, datasets), (ii) pre-training to develop a base model, and (iii) post-training techniques to improve the base model (for instance, fine-tuning, reinforcement learning, and safety training).¹⁷¹ Depending on the interpretation of Article 2(8), the exemption could include all, or only some, of these activities.

¹⁶² AI Act, art 2(8). See also AI Act, recital 25 ([...] it is necessary to ensure that this Regulation does not otherwise affect scientific research and development activity on AI systems or models prior to being placed on the market or put into service’).

¹⁶³ See Section 3.2.

¹⁶⁴ AI Act, art 2(8).

¹⁶⁵ AI Act, art 2(1)(a); see Section 2.1.1.

¹⁶⁶ See Section 2.1.1. Also see the commentary on Article 2, Section 2.4. in this work.

¹⁶⁷ AI Act, recital 25 (‘As regards product-oriented research, testing and development activity regarding AI systems or models, the provisions of this Regulation should also not apply prior to those systems and models being put into service or placed on the market. That exclusion is without prejudice to the obligation to *comply with this Regulation* where an AI system falling into the scope of this Regulation is placed on the market or *put into service as a result of such research and development activity* [...]’) (emphasis added).

¹⁶⁸ Section 2.1.1. offers arguments as to why *internal* deployment could trigger the application of the AI Act.

¹⁶⁹ AI Act, art 2(8).

¹⁷⁰ AI Act, recital 25.

¹⁷¹ Bengio and others (n 2) 30, 33, 34, 41–43. See, similarly, the commentary on Article 2, Section 2.4.1. in this work.

4. Conclusion

58. This chapter aims to clarify the interpretation of the AI Act and the extent to which it could apply to internal deployment. In this spirit, while acknowledging the jurisdictional limits of the AI Act, this chapter put forward several interpretative pathways supporting the inclusion of internal deployment within the scope of the AI Act (Section 2) and also examined potential objections and exceptions to such inclusion (Section 3).
59. Specifically, this chapter explored and stress-tested the following four arguments in favour of including the internal deployment of AI systems within the scope of the AI Act (Sections 2.1.1.–2.1.4.). *First*, Article 2(1)(a) (‘providers [...] putting into service AI systems’) could be interpreted to include within the scope of the AI Act those EU-based or foreign AI developers that make an internal AI system available for their own use in the EU.¹⁷² *Second*, Article 2(1)(a) (‘providers placing on the market [...] AI systems’) could be interpreted to include within the scope of the AI Act those EU-based or foreign AI developers that make an internal AI system available for use on the EU market in a business-related context, including for use in the business of the developer itself.¹⁷³ *Third*, Article 2(1)(b) (‘deployers of AI systems that have their place of establishment or are located within the Union’) could be interpreted to include within the scope of the AI Act those EU-based AI developers or EU-based staff of a foreign developer that make use of an internal AI system.¹⁷⁴ *Fourth*, Article 2(1)(c) (‘output produced by the AI system is used in the Union’) could be interpreted to include foreign AI developers that make available in the EU (e.g., to EU-based staff) the output produced by an internal AI system deployed abroad.¹⁷⁵ This chapter also explored how the AI Act could cover the internal deployment of GPAI models, either through these models’ integration into AI systems under Article 2(1) and Recital 97, or through an extensive interpretation of the expression ‘placing on the market’ in Article 2(1)(a).¹⁷⁶
60. The application of the AI Act to internally deployed AI models and systems, however, remains an open question. After exploring potential arguments in support of including internal AI models and systems within the remit of the AI Act, this chapter analysed arguments against their inclusion as well as potential exceptions or carve-outs.¹⁷⁷ Specifically, this chapter examined the following two arguments. *First*, Article 2(6) (‘AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development’) could be interpreted to exempt from the scope of the AI Act internal AI models or systems that an AI developer trains and makes available internally for the exclusive purpose of carrying out scientific research and development (i.e., research and development that is not product-oriented and that is undertaken in controlled environments, for example ‘in vitro’, for monitoring and diagnostics).¹⁷⁸ *Second*, Article 2(8) (‘[a]ny research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service’) could be interpreted to exempt an AI developer’s research and development activities *before* an AI model or system is put into service, either internally or externally.

¹⁷² See Section 2.1.1.

¹⁷³ See Section 2.1.2.

¹⁷⁴ See Section 2.1.3.

¹⁷⁵ See Section 2.1.4.

¹⁷⁶ See Sections 2.2. and 2.1.1.

¹⁷⁷ See Sections 3.1. and 3.2.

¹⁷⁸ See Section 3.1.

61. If looked at in isolation, all these provisions of the AI Act may appear as conflicting and even disjointed. That friction, however, is only superficial. Observing all the examined provisions contextually and systematically, the joint interpretation of Articles 2(1), 2(6), 2(8), 3(1), 3(3), 3(4), 3(9)–3(12) and 3(63) and Recitals 12, 13, 21, 25, 97, 109 and 110 can not only be brought to coherence, but these provisions can in fact be viewed as complementary to each other.¹⁷⁹

62. Figure 1 visually represents how these provisions complement and intersect with each other.




Figure 1. Visual representation of the scope of the AI Act in relation to AI models and systems deployed internally within AI providers, under a joint interpretation of Articles 2(1), 2(6), 2(8), 3(1), 3(3), 3(4), 3(9)–3(12) and 3(63) and Recitals 12, 13, 21, 25, 97, 109 and 110 AI Act.

- In yellow: internal AI models and systems arguably falling *within* the scope of the AI Act under Article 2(1)(a)–(c) and Recital 97.¹⁸⁰
- In blue: the activities preceding internal deployment¹⁸¹ and AI models or systems specifically developed and deployed for the sole purpose of scientific R&D,¹⁸² arguably falling *outside* the scope of the AI Act under Article 2(8) and Article 2(6) respectively.

¹⁷⁹ Also see the commentary on Article 2, para 9 in this work.

¹⁸⁰ See Sections 2.1.1.–2.1.4. and 2.2.

¹⁸¹ See Section 3.2.

¹⁸² See Section 3.1.

63. In summary, the deployment of an AI system, either *within* or outside of an AI provider, triggers the application of the AI Act.¹⁸³ This is clarified by Article 2(1)(a)–(c), which sets the clear-cut rule that making available or using an AI system or its outputs in the EU triggers the application of the AI Act and, in doing so, does not differentiate between internal systems and external systems.
64. GPAI models trigger the application of the AI Act once they are integrated into an AI system (whether internal or external) that is put into service or once they are independently placed on the market.¹⁸⁴ This is clarified by Article 2(1)(a) and Recital 97, according to which a GPAI model counts as being placed on the market once it is integrated into an AI system. System integration is necessary or at least beneficial for most of the possible internal uses of a GPAI model, which in turn narrows down the number of cases in which an internal GPAI model is not covered by the AI Act.
65. Internal AI models or systems specifically developed and deployed for the exclusive use of scientific research and development fall outside of the scope of the AI Act.¹⁸⁵ This exception is laid out in Article 2(6), according to which the AI Act ‘does not apply to AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development.’ The scope of this exception is narrow. It arguably only applies to AI models and systems that providers have trained and deployed internally for the exclusive purpose of scientific research and development. While the AI Act does not define ‘scientific research and development’, it is possible to infer from Recital 25 and the REACH Regulation that scientific R&D is different from ‘product-oriented’ R&D. By contrast to purely scientific R&D, any AI model or system that engages in R&D that is product-oriented does not fall within the Article 2(6) exception.
66. Finally, research and development activities that providers undertake before ‘putting into service’ or ‘placing on the market’ an AI model or system (for instance, experimentation and planning) do not trigger *per se* the application of the AI Act.¹⁸⁶ This is clarified by Article 2(8), according to which the AI Act ‘does not apply to any research, testing or development activity regarding AI systems or AI models *prior to* their being placed on the market or put into service.’

¹⁸³ See Figure 1, yellow box ‘After deployment’ and Sections 2.1.1.–2.1.4.

¹⁸⁴ See Figure 1, yellow box ‘After deployment’ and Section 2.2.

¹⁸⁵ See Figure 1, blue box ‘Scientific R&D exception’ and Section 3.1.

¹⁸⁶ See Figure 1, blue box ‘Before deployment’ and Section 3.2.