


RIVISTA DELLO STATO DIGITALE

01



Pubblicazione scientifica in formato digitale su informatica e sfera pubblica

Numero 1 – Anno 2025

I contributi di questa Rivista sono sottoposti alla valutazione di un revisore in forma anonima (single blind peer review), con la sola eccezione della rubrica “Lo Scaffale”.

La Rivista si conforma alle linee guida stabilite dalla Committee on Publication Ethics (COPE), nel rispetto del Codice etico consultabile in:
<https://www.rivistastatodigitale.eu>

Ogni riflessione e ogni suggerimento sono i benvenuti, nello spirito di una comunità scientifica aperta e partecipata. Per ogni informazione in merito all’invio dei contributi è possibile contattare la Rivista all’indirizzo rsd@irpa.eu

Direttore scientifico

Bruno Carotti

Vicedirettori

Paolo Clarizia, Gianluca Sgueo

Comitato scientifico:

Sabino Cassese, Stefano Battini, Enrico Carloni, Lorenzo Casini, Edoardo Chiti, Sveva del Gatto, Stefano Civitarese Matteucci, Fulvio Costantino, Giovanni Gallone, Barbara Marchetti, Marco Macchia, Bernardo Giorgio Mattarella, Enrico Nardelli, Luigi Previti, Giorgio Resta, Stefano Rossa, Aldo Sandulli, Luisa Torchia, Riccardo Ursi, Giulio Vesperini.

Primo redattore - Coordinamento editoriale

Gianluca Buttarelli

Comitato di redazione

Gianluca Buttarelli, Alessia Madeddu, Agostino Sola, Giulia Taraborrelli

IRPA | ISTITUTO DI RICERCHE
SULLA PUBBLICA
AMMINISTRAZIONE

Piazza Venezia, 11 - Roma

Roma, Ottobre 2025



Pubblicata con licenza Creative Commons CC BY 4.0., che richiede l’attribuzione dell’opera. Per conoscere i termini d’uso, si può visitare il sito:
<https://creativecommons.org/licenses/by/4.0/>

Progetto grafico: **Nuvola Studio**

Sommario

Presentazione -----	6
Il progetto: una sfida del tempo presente-----	7
Potere e ideologia nel mondo digitale----- <i>di Luisa Torchia</i>	9
Editoriale: verso mari perigliosi----- <i>di Bruno Carotti</i>	11
Le vaste latitudini della digitalizzazione -----	14
Democrazie e poteri (privati) globali----- <i>di Sabino Cassese</i>	15
Governare le macchine. Governare con le macchine----- <i>di Gianluca Sgueo</i>	23
Il fattore tempo nella regolazione digitale: RegLag e RegTech----- <i>di Diana Pittelli</i>	43
Dentro e fuori le pubbliche amministrazioni "digitali" -----	60
Dalla sovranità digitale all'emersione di una funzione di governo dei dati----- <i>di Agostino Sola</i>	61
Generative AI in the Public Sector under the AI Act: Challenges and Opportunities for Procurement----- <i>di Giulia Fantoni</i>	73
Note minime sul cd. 'Piano Europeo Automotive':----- <i>di Gianluigi Delle Cave</i>	91
I sentieri giurisprudenziali della digitalizzazione -----	114
La giurisprudenza amministrativa sulla digitalizzazione ----- <i>di Paolo Clarizia</i>	115
Lo scaffale -----	129
<i>di Bruno Carotti</i>	

DENTRO E FUORI LE PUBBLICHE AMMINISTRAZIONI “DIGITALI”

Generative AI in the Public Sector under the AI Act: Challenges and Opportunities for Procurement

Giulia Fantoni*

Abstract

The article delves into the challenges and opportunities of integrating Generative AI into the public sector within the regulatory framework set forth by the EU AI Act. GenAI offers promising potential for improving administrative efficiency. However, it also raises legal and operational issues, especially in the context of public procurement. Framing procurement as a critical entry point, this study adopts a rights-based approach anchored in the right to good administration. Following a review of the technical and regulatory foundations of GenAI, the article assesses three procurement models: in-house development, reliance on private suppliers via subscription services, and hybrid solutions. In this regard, it scrutinises the consequences of each of these models in terms of cost-effectiveness, transparency, impartiality, data protection and cybersecurity. Moreover, it touches upon the EU Model Contractual Clauses for AI Procurement, published in March 2025. The article concludes by advocating hybrid (and small-scale) approaches as the most responsible for the deployment of GenAI in the public sector.

Summary

1. Preliminary Considerations: Generative AI in the Public Sector. – 2. Understanding Generative AI: Technical and Regulatory Framing under the AI Act. – 2.1 Technical Foundations: From AI Systems to Generative AI – 2.2 GPAI in the AI Act. – 3. Procurement Pathways for Generative AI. – 3.1 Building In-House. – 3.2 Relying on Private Providers: Subscription-based Services. – 3.3 Hybrid Models: A Middle Way? – 4. Final Remarks and Future Outlook: Embracing Hybrid and Small-scale Solutions for Responsible Procurement.

1. Preliminary Considerations: Generative AI in the Public Sector

The release of ChatGPT by OpenAI in November 2022 made available to the public a tool capable of generating various types of content based on users' inputs. The rapid increase in popularity of Generative Artificial Intelligence (hereinafter Generative AI or GenAI) raised awareness of its potential societal impact¹, prompting a dedicated regulatory response at the European Union level. Notably, the Regulation

* PhD candidate in Administrative Law, University of Trento.

(EU) 2024/1689 (hereinafter, AI Act) dedicates an entire Chapter to this technology². While much discourse on its prospective uses has been focusing on the private sector, GenAI has been gaining momentum in public administrations as well³. In the European Union, a study published by the EU Public Sector Tech Watch in April 2025 mapped 61 existing use cases of this technology in the public sector among 20 different Member States⁴, confirming growing insti-

tutional interest. Additionally, several administrations – including the European Commission⁵ – have implemented or are working on the development of internal GenAI tools. This might be justified by the positive impact GenAI is expected to have on public entities. For example, by 2033, it could generate an annual return of \$1.75 trillion globally in the public sector, including \$29 billion in Italy⁶. Furthermore, initial evidence indicates that Generative AI could boost productivity in

- 1 On the matter, particularly eloquent is the opening paragraph to the G7 Leaders' Statement on the Hiroshima AI Process, which reads: «We, the Leaders of the Group of Seven (G7), stress the innovative opportunities and transformative potential of advanced Artificial Intelligence (AI) systems, in particular, foundation models and generative AI. We also recognize the need to manage risks and to protect individuals, society, and our shared principles including the rule of law and democratic values, keeping humankind at the center». See G7 LEADERS, *G7 Leaders' Statement on the Hiroshima AI Process*, Hiroshima, 30th October 2023. Available at https://www.mofa.go.jp/ecm/ec/page5e_000076.html. See also EUROPEAN COMMISSION, *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions on boosting startups and innovation in trustworthy artificial intelligence*, COM(2024) 28 final, 2024. Available at <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52024DC0028>.
- 2 Chapter V, *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)*. In 2024, the international momentum around AI governance increased, with institutions such as the OECD, the African Union, and the United Nations issuing frameworks focused on transparency, trustworthiness, and the responsible development of artificial intelligence. On the matter, N. MASLEJ, L. FATTORINI, R. PERRAULT, *et al.*, *The AI Index 2025 – Annual Report*, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, 2025 [online]. Available at <https://hai.stanford.edu/ai-index/2025-ai-index-report>. Notwithstanding other policy initiatives in the field, the AI Act of the European Union remains, to date, the first and only binding regulation on the matter at the global level. This reflects the regulatory model of the European Union, which is based on the protection of fundamental rights, in contrast to the deregulatory approach of the United States. See A. BRADFORD, *Digital Empires: The Global Battle to Regulate Technology*, OXFORD UNIVERSITY PRESS, Oxford, 2023; A. BRADFORD, *The False Choice Between Digital Regulation and Innovation*, in *Northwestern University Law Review*, 2024, p. 377-352.
- 3 For instance, in late January 2025, OpenAI announced the launch of ChatGPT Gov – a customised version of its GPT model specifically designed for U.S. governmental agencies. Its launch reflects a broader trend: since 2024, over 90,000 U.S. users across more than 3,500 federal, state, and local entities have reportedly used ChatGPT to support their daily professional activities. See OPENAI, *Introducing ChatGPT Gov*, 2024. Available at <https://openai.com/global-affairs/introducing-chatgpt-gov/>.
- 4 The mapped use cases mainly referred to public services, and occasionally to strategic sectors such as security, economy and construction. Additionally, since GenAI is a relatively new technology, over 60 percent of the use cases were identified as in the planning, development or testing phase. However, 17 use cases have passed the testing phase and are in use in the public sector. See EUROPEAN COMMISSION – DIRECTORATE-GENERAL FOR DIGITAL SERVICES, A. BRIZUELA, M. COMBETTO, S. KOTOGLOU *et al.*, *Analysis of the generative AI landscape in the European public sector*, 2025 [online]. Available at <https://op.europa.eu/en/publication-detail/-/publication/7787b257-167d-11f0-b1a3-01aa75ed71a1/language-en>.
- 5 For instance, in October 2025, the European Commission launched GPT@EC, a GenAI tool developed by the Directorate General DIGIT, based on the earlier GPT@JRC model. Designed to enhance internal productivity while safeguarding institutional data, GPT@EC supports various administrative tasks—such as drafting, summarisation, and coding—while ensuring compliance with EU standards and avoiding reliance on third-party AI providers.
- 6 For more, M. CARRASCO, C. HABIB, F. FELDEN, *et al.*, *Generative AI for the Public Sector: From Opportunities to Value*, Boston Consulting Group, 2023. Available at <https://web-assets.bcg.com/df/1e/9cde767044e5bc1d85f3e788f702/generative-ai-for-the-public-sector-from-opportunities-to-value.pdf>. For a similar study commissioned by the Estonian government in 2024, IMPLEMENT CONSULTING GROUP, *The economic opportunity of generative AI in Estonia*, 2024. Available at <https://implementconsultinggroup.com/article/the-economic-opportunity-of-generative-ai-in-estonia>.

civil servants while lowering their cognitive load⁷. Despite these encouraging premises, there are still major legal and operational challenges to overcome.

Much of the academic discourse on the risks of GenAI has focused on issues such as cybersecurity, privacy and intellectual property rights⁸. Nonetheless, a crucial area remains relatively unexplored: how public administrations could procure this technology. This is not a secondary aspect, since procurement is the gateway through which automation enters the public sector. Deploying effective AI tools requires not only funding, but also adequate expertise – both within public institutions and, often, from external providers⁹.

Given all the above, the present article takes procurement as a critical entry point to understanding the governance of GenAI in the public sector. It adopts a rights-based interpretation of the AI Act, grounded in the Charter of Fundamental Rights – particularly the right to good administration (GA¹⁰) –, and considers how this framework might guide public authorities in adopting this relatively

new technology responsibly. While the scope of the article is intended to be general, particular attention will be paid to the ‘Model Contractual Clauses for AI Procurement’ (MCC-AI) issued by the European Commission in March 2025¹¹.

2. Understanding Generative AI: Technical and Regulatory Framing under the AI Act

Before assessing the challenges and opportunities public administrations face when procuring GenAI, it is essential to outline the main technical features of this technology and its regulatory framework under the AI Act.

2.1 Technical Foundations: From AI Systems to Generative AI

The EU Regulation on AI defines an «artificial intelligence system» as a machine-based system designed to operate with varying levels of autonomy and capable of adapting

7 In 2024, the Swedish Municipality of Uddevalla tested Microsoft’s Copilot M365 to support municipal staff. Survey results showed positive outcomes: 73% found work more enjoyable, 67% noted improved work quality, 57% reported less mental fatigue, and 32% experienced reduced work-related stress. On the matter, T. ANDERSSON, *Uddevalla kommun har under 2024 utvärderat Copilot M365*, MyAI, 2024. Available at <https://my.ai.se/usecases/623>.

8 C. NOVELLI, F. CASOLARI, P. HACKER, *et al.*, *Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity* in *Computer Law & Security Review*, 2024, p. 1-16.

9 AGENZIA PER L’ITALIA DIGITALE, *Piano Triennale 2024-2026 per l’Informatica nella Pubblica Amministrazione – Aggiornamento 2025*, Roma, 2025, p. 35. Available at https://www.agid.gov.it/sites/agid/files/2025-01/Piano_Triennale_per_L_informatica_nella_PA_2024-2026_Aggiornamento_2025.pdf.

10 This approach clearly diverges from a merely literal interpretation and can be classified as both systematic and teleological. It is rooted in the right to good administration (as per Article 41 of the Charter of Nice). The author believes this interpretative choice is not isolated: the AI Act itself, at its Article 1, underlines the need for AI to comply with the fundamental rights enshrined in the Charter. Additionally, the Agenzia per l’Italia Digitale (AgID), in its 2024-2026 Three-Year Plan for IT in Public Administration, explicitly stated procuring AI must respect the principles of effectiveness and efficiency – two key corollaries of the right to good administration. On the matter, *Id.*

11 EUROPEAN COMMISSION, *Updated EU AI model contractual clauses*, 2025 [online]. Available at <https://public-buyers-community.ec.europa.eu/communities/procurement-ai/resources/updated-eu-ai-model-contractual-clauses>. MCC-AI may serve to mitigate (or even eliminate) some of the risks observed in other jurisdictions. For example, in the United Kingdom, concerns have been raised about the lack of clear, actionable guidance for public entities tasked with procuring AI responsibly. Some commentators have gone so far as to describe this situation as one of ‘regulatory hallucination’, underscoring the gap between expectations and available institutional support. On the matter, A. SANCHEZ-GRAELLS, *Responsibly Buying Artificial Intelligence: A ‘Regulatory Hallucination’* in *Current Legal Problems*, 2024, p. 81-125.

post-deployment¹². Within this framework, the Act introduces the notion of ‘General-Purpose AI’ referring to models «including [those] 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 [...]»¹³. The term is frequently used interchangeably with «foundation models», a concept coined in 2021 by the Human-Centered Artificial Intelligence (HAI) Institute at Stanford University¹⁴. The main examples of foundation models currently existing are LLMs, which are trained through self-supervised learning on vast amounts of data¹⁵ and can generate

a wide range of outputs in response to natural language prompts¹⁶. These capabilities emerge from a multi-phase training process¹⁷ composed of data collection¹⁸, data pre-processing¹⁹, training²⁰, fine-tuning²¹, and deployment²². When foundation models are used to generate content – including, but not limited to, text and code – they are considered Generative AI²³.

That said, there are several aspects that public entities should carefully consider when evaluating the adoption of GenAI. While these models often display a certain degree of fluency and assertiveness²⁴, they do not ‘understand’ the content they generate. Rather, they create outputs that are the most statistically probable based on

12 Article 3, par. 1, n. 1, AI Act. The same definition was given also by the ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT (OECD). On the matter, OECD, *Recommendation of the Council on Artificial Intelligence*, Paris, 2024. Available at <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

13 Article 3, par. 63, AI Act.

14 R. BOMMASANI, D.A. HUDSTON, E. ADELI, *et al.*, *On the Opportunities and Risks of Foundation Model* in *ArXiv* [online], 2021. Available at <https://arxiv.org/abs/2108.07258>.

15 In practice, given a textual input, the model learns to predict the most likely output – such as the next word in a sentence – based on patterns found in the training dataset. This process, known as self-supervision at scale, allows models to be trained on massive amounts of text without the need for manual annotation. On the matter, A. SANTOSUOSSO and G. SARTOR, *Decidere con l’IA. Intelligenze artificiali e naturali nel diritto*, Bologna, il Mulino, 2024, p. 34.

16 *Id.* p. 47-74.

17 The five-step training framework for LLMs is drawn from: STANFORD UNIVERSITY IT, *AI Demystified: Introduction to Large Language Models*, 2024. Available at <https://uit.stanford.edu/service/techtraining/ai-demystified/llm>. Instead, for a three-phase framework, Y. LIU, H. HE, T. HAN, *et al.*, *Understanding LLMs: A Comprehensive Overview from Training to Inference*, in *ArXiv* [online], 2025. Available at <https://arxiv.org/abs/2401.02038>.

18 In this phase, a variety of sources are consulted for the purpose of obtaining diverse datasets, including books, scientific articles, websites, and social media platforms. The objective is to ensure sufficient diversification across a range of topics, linguistic registers and formats, thereby enabling the model to develop a broad and flexible understanding of natural language. STANFORD UNIVERSITY IT, *Id.*

19 The process involves preparing raw textual material for the model by removing noise and irrelevant content and tokenising the text into separate units like words or sub-words. Finally, the data is structured into input-output pairs for training. See *Id.*

20 During this phase, the model is exposed to a large volume of textual data, which allows it to learn how to predict the subsequent word in each sentence. See *Id.*

21 Fine-tuning is defined as the adaptation of pre-trained models to more specific tasks. To achieve this, it is necessary to expose the model to a narrower dataset. See *Id.*

22 At this stage, the LLM becomes operational. See *Id.*

23 H. TONER, *What Are Generative AI, Large Language Models, and Foundation Models?*, Center for Security and Emerging Technology [website], 2023. Available at <https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models/>. Additionally, Recital 99 of the AI Act explicitly identifies LLMs as a primary example of GPAI models.

24 NATURE, *Why scientists trust AI too much – and what to do about it*, Editorial – Nature [website], 2024. Available at <https://www.nature.com/articles/d41586-024-00639-y>.

their training data²⁵. For this reason, even when prompted to explain their responses, the justification provided is typically an *ex post* rationalisation, not an accurate reflection of the computational process behind it²⁶. This phenomenon – often referred to as the ‘black box’ problem²⁷ – characterises not only GenAI but machine learning systems more broadly²⁸. In the public sector, such opacity poses serious concerns: it stands in tension with principles of publicity and transparency, as well as with procedural requirements to provide reasoned justifications for administrative decisions²⁹.

Another specific limitation of GenAI are

hallucinations, meaning the generation of factually incorrect or fabricated content³⁰. In parallel, the reliance on web-scale datasets makes these systems prone to reproducing and reinforcing harmful stereotypes and biases³¹. Both these issues raise serious concerns in the public sector, as they potentially undermine the principle of impartiality, which underpins administrative activity and protects individuals from arbitrary or discriminatory treatment³².

Finally, public administrations should be conscious of the critical environmental impact of Generative AI³³. Their development and deployment entail high electrici-

25 A. SANTOSUOSSO, G. SARTOR, *Decidere con l’IA. Intelligenze artificiali e naturali nel diritto*, cit., pp. 5, 34 and 53-54.

26 See *Id.*, p. 71. Due to this intrinsic characteristic of GenAI, some scholars have defined this technology as an «agency without intelligence». On the matter L. FLORIDI, *AI as Agency Without Intelligence: on ChatGPT, Large Language Models, and Other Generative Models*, in *Philosophy & Technology*, p. 1-7, 2023.

27 B. MARCHETTI, *La garanzia dello human in the loop alla prova della decisione amministrativa algoritmica*, in *BioLaw Journal - Rivista di BioDiritto*, 2021, p. 367-385.

28 A. SANTOSUOSSO and G. SARTOR, *Decidere con l’IA. Intelligenze artificiali e naturali nel diritto*, cit., p. 5.

29 On the relationship between transparency and procedural institutions, *Manuale di Diritto Amministrativo*, edited by E. Casetta, F. Fracchia, Milano, Giuffrè, 2024²⁶.

30 Z. Ji, N. Lee, R. FRIESKE, *et al.*, *Survey of hallucination in natural language generation* in *ACM Computing Surveys*, 2023 p. 1-38; R. AZAMFIREI, S.R. KUDCHADKAR, J. FACKLER, *Large language models and the perils of their hallucinations*, in *Critical Care*, 2023, p. 1-2.

31 On the matter, R. GEIRHOS, P. RUBISCH, C. MICHAELIS *et al.*, *ImageNET-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness*, in *ArXiv* [online], 2022. Available at <https://arxiv.org/abs/1811.12231>; M. GLICKMAN and T. SHAROT, *How human-AI feedback loops alter human perceptual, emotional and social judgements*, in *Nature Human Behaviour*, 2025, p. 345–359.

32 According to some authors, the principle of impartiality is the application of the principle of legal equality. See S. LARICCIA, *Il principio di imparzialità delle pubbliche amministrazioni*, in U. ALLEGRETTI, L. AMMANNATI, S. AMOROSINO, *et al.*, *Studi in onore di Giorgio Berti*, JOVENE, Napoli, 2005. Note that the issues of bias and hallucinations also have a direct impact on the administrative activity. If GenAI would be used to support the adoption of administrative decisions, the use of discriminatory models could lead to their annulment for abuse of power. See B.M. ARMIENTO, *Pubbliche amministrazioni e intelligenza artificiale. Strumenti, principi e garanzie*, Napoli, Editoriale Scientifica, 2024.

33 For an overview of the environmental costs and potentials of artificial intelligence, B. MARCHETTI, *I costi ambientali dell’IA*, in *BioLaw Journal - Rivista di BioDiritto*, 2025, p. 525–527; N. RANGONE, *Intelligenza artificiale, tutela dell’ambiente e regolazione europea* in *BioLaw Journal - Rivista di BioDiritto*, 2025, p. 529–548. More specifically, for an analysis of the debate on the environmental costs of artificial intelligence in the United States, see M. MERLER, *I costi ambientali dell’intelligenza artificiale: il dibattito negli Stati Uniti*, *BioLaw Journal - Rivista di BioDiritto*, 2025, p. 591–614. Furthermore, for an investigation of the possible application of the Do No Significant Harm (DNSH) principle to limit the negative environmental impacts of artificial intelligence, see L. DE GAETANO, *Il principio Do Not Significant Harm (DNSH) e i costi ambientali dell’intelligenza artificiale* in *BioLaw Journal - Rivista di BioDiritto*, 2025, p. 571–590.

ty consumption³⁴, intensive water usage³⁵, and substantial CO₂ emissions³⁶. These factors should not be overlooked when considering the principle of cost-effectiveness in administrative procurement and operations³⁷.

Given these considerations, GenAI should not be relied upon for decision-making. However, this technology may still serve as an auxiliary tool to carry out internal administrative tasks such as document drafting, data synthesis, or preparatory operation leading to a final (human) decision³⁸. This position finds support in emerging case law (both national and supranational)³⁹, doctrine⁴⁰, and administrative guidelines⁴¹,

which affirm the need for human oversight and accountability in public decision-making processes.

2.2 GPAI in the AI Act

As already mentioned, the EU AI Act regulates AI broadly, while focusing specifically on General-Purpose AI in its fifth Chapter. The Regulation applies both to public and private providers and deployers of AI operating within as well as outside the EU whose AI systems enter the EU market or affect individuals within it⁴². The Act adopts a risk-based approach, imposing obligations that are evaluated on a case-by-case basis

34 Each query submitted to ChatGPT is estimated to consume approximately ten times more electricity than a standard Google search. See ELECTRIC POWER RESEARCH INSTITUTE, *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*, 2024 [website]. Available at <https://www.epri.com/research/products/3002028905>.

35 Shaolei Ren, a researcher at the University of California – Riverside, estimates that submitting a series of queries (ranging from 5 to 50) to ChatGPT results in the model consuming up to 500 milliliters of water. See ASSOCIATED PRESS, *Artificial intelligence technology behind ChatGPT was built in Iowa – with a lot of water*, 2023 [online]. Available at <https://apnews.com/article/chatgpt-gpt4-iowa-ai-water-consumption-microsoft-f551fde98083d17a7e8d904f-8be822c4>.

36 A 2023 study on the LLM BLOOM, which has 176 billion parameters, estimated that its training phase emitted approximately 24.7 tons of CO₂ considering only dynamic energy consumption. When accounting for the entire lifecycle, including equipment production and operational energy use, total emissions rose to 50.5 tons of CO₂. For comparison, training OpenAI's GPT-3 generated over twenty times more emissions than BLOOM, approximately 500 tons of CO₂. See A.S. LUCCIONI, S. VIGUIER, and A.L. LIGOZAT, *Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model in Journal of Machine Learning Research*, 2023, p. 1-15.

37 For insights on more sustainable procurement of artificial intelligence, L. PARONA, *La ponderazione dei costi ambientali nell'approvvigionamento di sistemi di intelligenza artificiale da parte delle amministrazioni pubbliche*, in *BioLaw Journal - Rivista di BioDiritto*, 2025, p. 549–569.

38 For a concise overview of internal uses of generative AI in public administration, see the case of Germany, where the Land of Baden-Württemberg launched the F13 system in 2024 to support civil servants in drafting documents, summarizing texts, and conducting legal research. See ALEPH ALPHA, *GovTech Campus Deutschland, STACKIT and Aleph Alpha create a platform for AI applications for the German administration: New F13 goes live in Baden-Württemberg*, 2024 [website]. Available at <https://aleph-alpha.com/govtech-campus-deutschland-stackit-and-aleph-alpha-create-a-platform-for-ai-applications-for-the-german-administration-new-f13-goes-live-in-baden-wuerttemberg/>.

39 For the Italian context, among the others: Consiglio di Stato, sezione VI, 8 aprile 2019, n. 2270 and Consiglio di Stato, sezione VI, 4 aprile 2020, n. 881, whose principles have been codified in Article 30, Legislative Decree No. 36/2023 (the so-called new Italian Public Procurement Code) which states: «*Le decisioni assunte mediante automazione rispettano i principi di: a) conoscibilità e comprensibilità [...]; b) non esclusività della decisione algoritmica [...]; c) non discriminazione algoritmica [...]*». More recently, the European Court of Justice clarified that data subjects subject to automated decisions must receive meaningful, understandable information on the logic and criteria used, while balancing this right with trade secrets and third-party rights. See Court of Justice of the European Union, Case C-203/22, *Dun & Bradstreet Austria*, ECLI:EU:C:2025:117.

40 B. MARCHETTI, *La garanzia dello human in the loop alla prova della decisione amministrativa algoritmica*, cit., p. 7.

41 For a comprehensive overview of the guidelines adopted at the European level, EUROPEAN COMMISSION, JOINT RESEARCH CENTRE, *PSTW: Analysis of the generative AI landscape in the European public sector*, 2025 [online]. Available at <http://data.europa.eu/89h/561de3c2-b968-42c8-8f14-fecb572d874f>.

42 Art. 2, AI Act.

and proportionate to the risk level of each AI system, classified as unacceptable⁴³, high⁴⁴, limited and minimal or no risk⁴⁵. Alongside these categories, the Regulation includes an additional one. Its Chapter V – which addresses GPAI separately – distinguishes between models with and without systemic risk. This reflects the unique challenges foundational models present, such as unpredictability and broad societal impact⁴⁶.

Prior to delving into the concept of ‘systemic risk’ linked to GenAI *per se*, it is crucial to acknowledge that eight potential applications of AI systems are categorised as ‘high-risk’ in Annex III of the AI Act. Of these, the most relevant for the public sector are listed in its paragraph 5, which refers to AI systems used to determine access to «essential private and public services and benefits»⁴⁷. This list is limited and does not cover the full range of public administrative functions. Thus, it can be argued that its restricted scope reveals an incomplete and arguably inconsistent classification of ‘high-risk’ systems under the AI Act.

Moreover, Article 6, par. 3 introduces a broad exception: AI systems listed in Annex III can escape high-risk status if they do not pose a «significant risk» to health, safety, or fundamental rights. This includes systems performing narrow procedural tasks, supporting human decisions without replacing them or serving as preparatory steps⁴⁸. This exception appears problematic: it is reasonable to presume that AI systems utilised during preparatory stages are likely to impact decision-making. Therefore, the exemption of these uses from high-risk classification fails to require appropriate safeguards, thereby undermining the right-based intent of the AI Act⁴⁹.

Given all the above, and considering the known risks of machine learning, deep learning, and Generative AI – such as opacity, bias, and hallucinations⁵⁰ – the author believes these technologies should be classified as high-risk by default when deployed by public entities. This should be the case at least when these systems influence⁵¹ or determine decisions that affect individual

43 Chapter II of the AI Act, titled «Prohibited AI practices».

44 Chapter III of the AI Act, titled «High-risk AI systems». When it comes to high risk, article 6 identifies three main areas falling under it. AI systems «intended to be used as a safety component of a product, or the AI system is itself a product» and that must «undergo a third-party conformity assessment» as per Art. 6, par. 1, letters a) and b). The third area includes all the uses listed in Annex III of the Regulation. See also S. WACHTER, *Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond* in *Yale Journal of Law & Technology*, 2025, p. 671-718.

45 EAD., p. 677 and Chapter X of the AI Act.

46 Also S. WACHTER, *Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond*, cit., pp. 10, 677.

47 These include eligibility for assistance (par. 5, lett. a), creditworthiness (par. 5, lett. b), insurance risk assessment (par. 5, lett. c), and emergency call classification (par. 5, lett. d). See Annex III, AI Act.

48 These exemptions are listed in Art. 6, par. 3, letters a) to d), AI Act. In any case, however, the exemption can operate when an AI system referred to in Annex III is used to profile a natural person as per Art. 6, par. 3, concluding sentence, AI Act.

49 This is not the only inconsistency within the Regulation. For instance, one shall notice the different treatment of AI systems used in law enforcement and those employed in other public sector administrations. While the former are explicitly classified as high-risk under Annex III, paragraph 6, the latter are not subject to the same degree of regulatory oversight. This choice is difficult to justify, considering that both domains involve decisions that can significantly affect the rights of individuals.

50 This right-based approach is stated in Article 1 of the AI Act.

51 Therefore, the author believes that GenAI should be considered as high-risk also when used in the cases listed under Art. 6, par. 3, AI Act.

rights or other legally protected interests⁵². This would align with the precautionary logic of the AI Act and guarantee the respect of the right to GA.⁵³

Having established this baseline, it is now possible to turn to a more specific discussion regarding GenAI and systemic risk. According to Article 3(65) of the AI Act, systemic risk refers to «a risk that is specific to the high-impact capabilities of [GPAI] models, having a significant impact on the Union market due to their reach, or [...] actual or reasonably foreseeable negative effects on public health, safety, public scrutiny, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain»⁵⁴. The AI Act introduces a rebuttable presumption of systemic risk where the model has been trained using computing power exceeding 10²⁵ floating point opera-

tions (FLOPs). As of today, only a very limited number of GenAI models meet this threshold⁵⁵.

Regardless of whether a model qualifies as systemic-risk GPAI, all providers of GPAI are subject to a set of obligations⁵⁶. These include maintaining up-to-date documentation explaining the capabilities and limitations of the model⁵⁷ and ensuring that downstream providers who decide to integrate GPAI into their own AI systems have sufficient understanding of its use and risks⁵⁸. In addition, GPAI providers must comply with EU copyright law and make available a «sufficiently detailed» summary of the datasets used to train the model⁵⁹. In addition to this, if the model qualifies as systemic-risk GPAI, further obligations apply: the provider must carry out risk assessments – possibly including adversarial testing⁶⁰ – to identify and

52 This would exclude, for instance, using the GenAI tool as a support to write email and/or refining texts.

53 This proposed ‘by-default’ classification is further supported by a closer reading of the MCC-AI of the European Commission, that will be further recalled in the third paragraph of the present article. Although the Commentary to the MCC-AI distinguishes between clauses for high-risk and light-risk AI systems, it explicitly notes that the latter are largely inspired by the former, with only limited differences. It also states that the MCC-AI Light should only be used in situations where an AI system does not qualify as high-risk but still poses a risk to health, fundamental rights, or safety (p. 5 of the Commentary), while also warning that where no such risk exists, their use may be disproportionate and should be adjusted (p. 4 of the Commentary). Yet, the substance of the two sets of clauses is remarkably similar. This is especially evident in Article 3 on «data governance», which imposes identical obligations on suppliers in both versions – for instance, it requires to ensure that datasets are examined for potential bias (Art. 3, par. 3.1, letter f) and are representative and as error-free as possible (Art. 3, par. 3.2). This convergence is particularly interesting given that the AI Act itself, in Article 10, reserves such data governance obligations exclusively for high-risk systems. The near-equivalence of the MCC-AI, regardless of risk classification, thus supports a presumption of high-risk where AI systems rely on data-intensive techniques such as ML, DL, and GenAI. For the Commentary as well as MCC-AI High-Risk and Light-Risk, see EUROPEAN COMMISSION, *Updated EU AI model contractual clauses*, cit., p. 4.

54 A GPAI model is deemed to present systemic risk either when such impact is determined based on «appropriate technical tools and methodologies» (Art. 51, par. 1, lett. a) or by decision of the EU Commission (Art. 51, par. 1, lett. b). The latter scenario can happen either *ex officio* or following the recommendation of an expert group. See Article 51, par. 2, AI Act.

55 For example, the most advanced models of ChatGPT by OpenAI like GPT4 or GPT 4o might be considered as GPAI with systemic risk. See D. DUSHI, *How is ChatGPT regulated by the EU AI Act: Reflections on higher education*, Global Campus of Human Rights, 2024 [website]. Available at <https://www.gchumanrights.org/preparedness/how-is-chatgpt-regulated-by-the-eu-ai-act-reflections-on-higher-education/>.

56 More precisely, these obligations are set forth under Art. 53, AI Act.

57 Art. 53, par. 1, lett. a), AI Act.

58 Art. 53, par. 1, lett. b), nn. 1-2, AI Act.

59 Art. 53, par. 1, lett. c), AI Act.

60 The term «adversarial testing» refers to a method for evaluating a machine learning model to understand its behaviour «when provided with malicious or inadvertently harmful input». On the matter, GOOGLE DEVELOPER PROGRAM, *Adversarial Testing for Generative AI*, N/A [website]. Available at <https://developers.google.com/machine-learning/guides/adv-testing>.

mitigate systemic risks, maintain records of serious incidents (to be reported to relevant national authorities), implement corrective measures where necessary, and ensure robust cybersecurity safeguards⁶¹.

All the just recalled obligations – directed surprisingly only to providers (and not to deployers) of GPAI⁶² – shall always be taken into account by public entities deciding to procure GenAI-based tools.

3. Procurement Pathways for Generative AI

As outlined in the introduction, public administrations can generally buy GenAI through three main pathways: by developing a model in-house⁶³; by subscribing to private provider services; or by adopting a hybrid approach using existing models customised for specific public needs⁶⁴.

The 2024 G7 Toolkit highlighted procurement as a substantial challenge for public administrations seeking to (partially) automate their work⁶⁵. The document emphasised the necessity for flexibility and internal expertise in addressing the legal and tech-

nical challenges associated with AI adoption, particularly in relation to privacy and discrimination⁶⁶. As possible policy options to tackle the just recalled challenges, the Toolkit recommended crafting specific procurement mechanisms and guidelines to support public entities in selecting AI systems that suit their purpose(s). It is within this context that the Model Contractual Clauses for Artificial Intelligence were developed by the EU Commission in March 2025⁶⁷. These clauses follow a risk-based logic, distinguishing between high- and light-risk AI systems and suggesting corresponding contract provisions. Because the present article focuses on GenAI and adopts the position that its use in the public sector should be treated as high-risk by default, the analysis will concentrate solely on the high-risk contractual clauses (MCC-AI HR). It should be noted, however, that the current model clauses do not explicitly address Generative AI systems that may pose systemic risks.

Given all the above, the present article will proceed with an analysis of procurement models by identifying relevant elements that shall be considered for each scenario.

61 For more details, Art. 55, AI Act.

62 It is also for this reason that an ‘high-risk-by-default’ approach when public administrations plan to introduce GenAI tool in their work become more relevant: as it obligates public entities to comply with the (claimed) rights-based approach of the AI Act.

63 Or by commissioning its development to a third party. In both cases the public entity would be procuring GenAI as a provider by procuring the development of a GenAI model tailored to its specific needs. For more, C. CHRYSOCHOU and E. O’CONNELL, *Beyond the Buzzwords: A Practical Guide to AI Procurement with Model Clauses and GDPR*, EIPA, 2024 [online]. Available at <https://www.eipa.eu/blog/beyond-the-buzzwords-a-practical-guide-to-ai-procurement-with-model-clauses-and-gdpr/>.

64 A report published in 2024 by EIPA identifies a different classification: procuring the development of GenAI services as a provider; as a downstream provider; and by procuring the use of GPAI models by contractors. See EAD.

65 Namely, Challenge n. 2 of the Toolkit. On the matter, OECD and UNESCO, G7 Toolkit for Artificial Intelligence in the Public Sector, Como, 2024. Available at <https://www.g7italy.it/wp-content/uploads/1728922597-g7-toolkit-for-ai-in-the-public-sector.pdf> <https://www.g7italy.it/wp-content/uploads/1728922597-g7-toolkit-for-ai-in-the-public-sector.pdf>.

66 The Toolkit also highlighted a significant disparity in AI literacy between the public and private sectors, which could potentially compromise the effective integration of AI technologies within public administrations. This discrepancy has the potential to result in delays in the acquisition of AI solutions, which may already be obsolete by the time they are adopted by public entities. See Id.

67 See EUROPEAN COMMISSION, *Updated EU AI model contractual clauses*, cit., p. 4.

Following a rights-based approach, each option will be evaluated in the light of the right to GA, particularly regarding cost-effectiveness, transparency, and impartiality. Moreover, as data is the backbone of (Gen)AI⁶⁸ and plays a central role in adapting models to the public sector needs, privacy and cybersecurity risks will also be integrated into the evaluation⁶⁹.

3.1 Building In-House

As clarified, the first option public administrations have when procuring GenAI is building it in-house⁷⁰.

While this approach offers a high degree of personalisation, potentially making it the most adaptable to the specific needs of the adopting public entity, it entails very high

substantial costs. To date, no comprehensive data exists that accurately estimates the economic burden public sector entities would face by opting to develop proprietary GenAI models in-house. However, available data on development costs faced by major tech companies shows that creating these models requires multi-million-dollar investments – costs that are likely beyond the reach of most public administrations⁷¹.

Given the generally low level of AI literacy across public administrations⁷² and the exceptional cost of building a model entry in-house, the most likely scenario could be outsourcing the development of customised AI models to third parties. This approach could involve the use of ‘public data’⁷³ during the pre-training phase of language models. Achieving this would require public sector

68 For this reason, open-data initiatives are to be encouraged. On the matter, Directive (EU) 2019/1024 of the European Parliament and of the Council of 20 June 2019, *on open data and the re-use of public sector information (recast)*. See also Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022, *on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act)*.

69 A reference to cybersecurity can also be found in Article 108, paragraph 4 of the new Italian Public Procurement Code, which requires contracting authorities, «[n]elle attività di approvvigionamento di beni e servizi informatici per la pubblica amministrazione, [...] nella valutazione dell'elemento qualitativo ai fini dell'individuazione del miglior rapporto qualità prezzo per l'aggiudicazione, [di tenere] sempre in considerazione gli elementi di cybersicurezza, attribuendovi specifico e peculiare rilievo nei casi in cui il contesto di impiego è connesso alla tutela degli interessi nazionali strategici». Given that the provision does not explicitly mention privacy, and requires contracting authorities to consider cybersecurity aspects only where «strategic national interests» are involved, it is reasonable to assume that special caution should nonetheless be exercised with regard to both cybersecurity and data protection (also in cases of direct awards), and in any event, even where (Gen)AI-based tools are not used strictly within the context of the recalled national strategic interests. Indeed, the AI Act itself, in Articles 2, par. 7 and 15, underscores the centrality of privacy and cybersecurity. On the matter, see also C. NOVELLI, F. CASOLARI, P. HACKER *et al.*, *Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity*, cit., p. 4.

70 Or asking a third party to build such a model for them. As per Article 3, par. 3 of the AI Act, in this scenario, public entities would still be classified as providers.

71 For example, the training cost of Gemini alone, without counting the salaries of the personnel employed, has been estimated between 30 and 191 million dollars. ChatGPT-4 had a technical creation cost between 41 and 78 million dollars. According to Sam Altman, CEO of OpenAI, the overall cost exceeded 100 million dollars. On the matter, see K. BUCHHOLZ, *The Extreme Cost of Training AI Models*, *Forbes*, 2024 [online]. Available at <https://www.forbes.com/sites/katharinabuchholz/2024/08/23/the-extreme-cost-of-training-ai-models/>. The Chinese company DeepSeek, on the other hand, declared that training one of its latest models costed around 5.6 million dollars. Now, reading the data and assuming an equivalence of costs between the private and public sectors, one must conclude that any public administration that was to choose to build a LLMs in-house would face huge or even unsustainable expenses (both in terms of training the model and in terms of technical expertise required for its development). On DeepSeek, see M.W. ROELOFFS, *What Is DeepSeek? New Chinese AI Startup Rivals OpenAI – And Claims It's Far Cheaper*, *Forbes*, 2024 [online]. Available at <https://www.forbes.com/sites/maryroeloffs/2025/01/27/what-is-deepseek-new-chinese-ai-startup-rivals-openai-and-claims-its-far-cheaper/>.

72 G7 ITALIA, *G7 Toolkit for Artificial Intelligence in the Public Sector*, cit., p. 12.

73 This meaning the data possessed by public administrations.

bodies to make their unstructured open data available to developers as this could enhance the performance and relevance of these models⁷⁴.

However, enabling private providers to train models on public sector data inevitably raises concerns, particularly regarding the protection of public sector knowledge assets. To this end, contractual safeguards would be crucial. Section D of the MCC-AI High-Risk⁷⁵ – notably Article 15 – provides a clear framework: suppliers may only use public organisation datasets for contractually agreed purposes and must delete them upon request unless otherwise specified.

From a data protection point of view, notwithstanding greater control over the training dataset, there would still be risks related to the collection and preparation of the datasets. In fact, the named datasets could unintentionally include sensitive data⁷⁶. From a cybersecurity point of view, there remains the risk of manipulation of the dataset by malicious actors if strong cybersecurity systems are not enforced. Furthermore, the risks related to the deployment include possible unauthorised access to the model by

malicious users, who could extract sensitive data. Therefore, this would require the implementation of robust authentication systems and periodic reviews of the deployment configurations, to identify vulnerabilities⁷⁷.

In any case, even with greater control over training data⁷⁸, it shall always be considered that building in-house would not solve core issues like opacity and hallucinations⁷⁹.

3.2 Relying on Private Providers: subscription-based services

A different scenario arises when public administrations choose to rely on GenAI models provided by private suppliers through subscription-based services. This option typically comes with a limited financial impact⁸⁰; therefore, it may fall below the thresholds requiring open competition under EU procurement law. For instance, Article 50 of the Italian Public Contracts Code allows the direct award of services and supplies – without prior consultation of multiple economic operators – for contracts valued below €140,000⁸¹. Within this framework, several

74 H. CHAFETZ, S. SAXENA, S.G. VERHULST, *A Fourth Wave of Open Data? Exploring the Spectrum of Scenarios for Open Data and Generative AI*, TheGovLab, 2024 [website]. Available at https://static1.squarespace.com/static/6637ad07a322a115424c1a01/t/6638fca1ec01706d012d845b/1715010735768/A+Fourth+Wave+of+Open+Data%3F+Exploring+a+Spectrum+of+Scenarios+for+Open+Data+and+Generative+AI_05.06.24.pdf.

75 EUROPEAN COMMISSION, *Updated EU AI model contractual clauses*, cit., p. 4.

76 I. BARBERÁ, *AI Privacy Risks & Mitigations Large Language Models (LLMs)*, *AI Privacy Risks & Mitigations – Large Language Models (LLMs)*, EDPS, 2025 [online]. Available at <https://www.edpb.europa.eu/system/files/2025-04/ai-privacy-risks-and-mitigations-in-llms.pdf>. These risks persist across all procurement models. However, in-house development or fine-tuning of pre-trained models allows for greater direct control over dataset management and cybersecurity measures. In contrast, subscription-based services provided by private vendors limit the ability of public administration to enforce cybersecurity protocols directly on the underlying infrastructure, making it more challenging to mitigate the corresponding risks.

77 EAD.

78 A control always required for high-risk systems under Article 10 of the AI Act.

79 Paragraph 2.1 of the present article (p. 4-9).

80 For instance, ChatGPT by OpenAI offers a monthly plan called ChatGPT Team with a pricing of 30 dollars per user a month. The option comes with tools for deep researching, analysing data, uploading documents, etc. It promises to offer a secure and collaborative workspace while excluding the users' data from training by default. On the matter see OPENAI, *ChatGPT Pricing*, N/A [website]. Available at <https://openai.com/chatgpt/pricing/>.

81 See Art. 50, par. 1, lett. b), Legislative Decree No. 36/2023.

Italian public administrations have already carried out direct award procedures for subscriptions to GenAI services offered by private companies such as OpenAI or Microsoft⁸².

At first glance, this procurement model appears to be a cost-effective and efficient solution. However, such an approach demands careful preliminary evaluations, particularly regarding the protection of the informational assets of public entities. When proceeding through direct award, it is essential to verify that the contractual terms include adequate safeguards, such as explicit prohibitions on the reuse of user-provided data for improving the performance of the model⁸³. Nonetheless, a residual degree of uncertainty remains concerning the actual enforcement of such commitments by private providers, which requires caution.

From a functional standpoint, reliance on private providers through subscriptions of services generally precludes the possibility of training models on datasets provided by public administrations. This means that the models could not be adapted to the specific institutional needs of each public administration, therefore limiting their effectiveness. Additionally, this solution would weaken some key principles tied to the right to GA – like transparency, impartiality, and

explainability. These values, which are already challenged by the inherent opacity of GenAI systems, would be even more weakened due to the impossibility of any control over the training data⁸⁴. This would also clash with the safeguards on data governance required under Article 10 and 26 of the AI Act.

Privacy and cybersecurity considerations further complicate the picture. A typical point of vulnerability in this form of procurement could occur at the prompting stage. If public officials are not adequately informed about data usage policies, they may input sensitive personal or institutional data. Additionally, even when providers do not store the inputted information, there is still the risk that malicious actors could extract it through reverse engineering⁸⁵. To mitigate these risks, civil servants using GenAI procured through subscriptions should implement robust anonymisation techniques – even at the prompt level – and follow secure model access⁸⁶.

3.3 Hybrid Models: A Middle Way?

The third scenario taken into consideration in the present article involves public admini-

82 Among the others, AGENZIA DEL DEMANIO, *Acquisto di n. 10 licenze “Copilot for Microsoft 365”*, 2024 [website]. Available at <https://www.agenziademanio.it/it/gare-aste/forniture-e-servizi/gara/Acquisto-di-n.-10-licenze-Copilot-for-Microsoft-365>.

83 For example, the already cited ChatGPT Team allegedly states that it does not use inputted data for the training of its models. See OPENAI, *ChatGPT*, cit., p. 15. However, the free version offered by OpenAI does retain user’s data and use them to improve and train models. On the matter, OPENAI, *How your data is used to improve model performance*, N/A [website]. Available at <https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance>.

84 Paragraph 2.1 of the present article (p. 4-9).

85 Even when training data is not memorised, personal information can still be inferred by malicious actors using model inversion attacks, which reverse-engineer the input data to reveal private information. See M. FREDRIKSON, S. JHA and T. RISTENPART, *Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures*, Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 1322–33, CCS 2015, New York, USA, Association for Computing Machinery, 2015.

86 I. BARBERÁ, *AI Privacy Risks & Mitigations Large Language Models (LLMs)*, *AI Privacy Risks & Mitigations – Large Language Models (LLMs)*, cit., p. 15.

strations relying on pre-trained models that are either open source or commercially available, with the option of fine-tuning⁸⁷.

From a cost-effectiveness point of view, this approach can be appealing. Fine-tuning an existing model is generally cheaper than training one from scratch. However, the actual expense depends heavily on the quantity and complexity of the data used for customisation. The more data and the more refined the task, the higher the cost – especially when advanced techniques such as reinforcement learning with human feedback or manual data annotation are involved, both of which require human oversight to improve model performance⁸⁸. In any case, as fine-tuning would involve using data relating to public administration, it is crucial to include a reference to the already recalled Section D of MCC-AI HR (specifically, Article 15) in the procurement contract.

In terms of efficiency, this option allows for

a fair degree of personalisation. Even though LLMs are general-purpose by design, they often require adaptation to understand the specific terminology or context of public administration. This can be achieved by using open data to carry out fine-tuning⁸⁹ or to implement Retrieval-Augmented Generation (RAG). The latter enables the model to access a dedicated database of relevant documents prior to generating a response, which can significantly enhance domain relevance⁹⁰.

When it comes to transparency and impartiality, while fine-tuned models offer some control over the data used, the so-called ‘black box’ problem remains. As already mentioned, hallucinations and bias emerge in all generative models, regardless of the training approach⁹¹.

In terms of privacy and cybersecurity, public bodies may not have access to, or knowledge of the data used to train the model initial-

87 «Fine-tuning is a process in which a pretrained model, such as an LLM, is further trained on a custom data set to adapt it for specialized tasks or domains». For more, D.M. ANISUZZAMAN, J.G. MALINS, P.A. FRIEDMAN, *et al.*, *Fine-Tuning Large Language Models for Specialized Use Cases in Mayo Clinic Proceedings: Digital Health*, 2025, p. 1-13.

88 On the matter, Y. LIU, H. HE, T. HAN, *et al.*, *Understanding LLMs: A Comprehensive Overview from Training to Inference*, *cit.*, p. 6.

89 GPAI models fine-tuned by downstream actors potentially turns them into providers of new models. When such modifications occur, only the changes themselves need to be documented under Article 53 of the AI Act. Article 55 obligations, concerning systemic risk, will in any case apply only in the defined situations under Article 51. The AI Office will offer further guidance on when such obligations are triggered. Regardless of whether the downstream party is reclassified as a model provider, it must still comply with all applicable obligations for AI systems under the AI Act. On the matter, EUROPEAN COMMISSION, *General-Purpose AI Models in the AI Act – Questions & Answers*, 2025 [website]. Available at [https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers#:~:text=To%20capture%20the%20most%20advanced,\(2\)%20AI%20Act.](https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers#:~:text=To%20capture%20the%20most%20advanced,(2)%20AI%20Act.)

90 RAG appears to be a promising solution for generating outputs via LLMs, as it integrates model knowledge with external databases. This enhances output accuracy and allows for continuous database updates. On this matter, Y. GAO, Y. XIONG, X. GAO, *et al.*, *Retrieval-Augmented Generation for Large Language Models: A Survey* in *ArXiv* [online], 2023. Available at <https://api.semanticscholar.org/CorpusID:266359151>. The data sources used for both fine-tuning and the underlying database of a RAG system may include court decisions, regulations, circulars, internal FAQs, historical administrative acts, and documentation related to administrative procedures. Notably, legal scholar Roberto Cavallo Perin emphasised that AI trained on a comprehensive corpus of documents produced over 30 to 40 years of regulatory interpretation can offer, through machine-generated outputs, a retrospective analysis of how specific situations were addressed in a particular historical context – also highlighting the potential shortcomings of the public administrations. See R.C. PERIN, *Il nucleo essenziale del diritto amministrativo*, Conference held as part of the PRIN 2020 project *Re.S.To.R.E. – Recovering the State Towards a Reformed Economy*, University of Trento, 30 January 2025. Additionally, on the possibility of combining fine-tuning with RAG, see K. RANGAN, Y. YIN, *A fine-tuning enhanced RAG system with quantized influence measure as AI judge*, in *Scientific Reports – Nature* 2024 [online]. Available at <https://www.nature.com/articles/s41598-024-79110-x>.

91 Paragraph 2.1 of the present article (p. 4-9).

ly⁹². This is why it is crucial to include a reference to data governance under Article 10 of the AI Act in the procurement contract when using this method. Furthermore, reliance on the original developer for updates can compromise the long-term adaptability of this model. Using RAG could mitigate some of these risks by enabling administrations to control the external knowledge base that the model refers to. However, it does not eliminate risks entirely. Furthermore, if inputs and source documents are stored without adequate safeguards, sensitive data could potentially be exposed, or information could leak from the underlying database without users realising it⁹³.

4. Final Remarks and Future Outlook: Embracing Open-Source and Smaller-Scale Solutions for Responsible Procurement

Based on the analysis conducted above, the first option – developing a GenAI system in-house or outsourcing its development – does not seem to be the most cost-effective approach. In contrast, the second procure-

ment model – subscribing to services provided by private vendors – may appear more economically accessible, even opening to the possibility of setting in place mechanisms such as direct awards. Yet this model generally allows for no (or limited⁹⁴) personalisation and poses challenges in ensuring alignment with public interest values. A hybrid model, therefore, emerges as the most balanced and prudent approach. It combines financial viability with the capacity for meaningful institutional oversight over sensitive data and strategic information infrastructure⁹⁵.

It shall be point out that, regardless of the procurement route chosen, public authorities must always ensure that robust risk management frameworks are established⁹⁶. Administrations must also verify that datasets used in model development are relevant, properly processed⁹⁷, and meet standards of quality, representativeness, and accuracy⁹⁸. Furthermore, they must ensure that deployed GenAI systems are robust and secure from a cybersecurity perspective⁹⁹. As deployers of GenAI – here assumed as high-risk by default – public administrations

92 I. BARBERÁ, *AI Privacy Risks & Mitigations Large Language Models (LLMs)*, *AI Privacy Risks & Mitigations – Large Language Models (LLMs)*, cit., p. 15.

93 EAD.

94 Through prompt engineering.

95 This direction is reflected in the EU GenAI4PA initiative, which encourages the fine-tuning of foundational models supplied by European-based private actors. On the matter, I. ZEPPA, *GenAI4EU for the public administrations*, AI Office, 2025 [online]. Available at https://reform-support.ec.europa.eu/document/download/1fb44fbc-c607-417e-9eb7-c9e72395d4c0_en?filename=Ivana%20ZEPPA%20-%20GenAI4EU_PA.pdf. For more on the initiative, refer to: B2MATCH, *GenAI for Public Administration - Call for Proposals* [website]. Available at <https://www.b2match.com/e/genai4publicadmin>; EUROPEAN COMMISSION, *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions on boosting startups and innovation in trustworthy artificial intelligence*, cit., p. 2.

96 If the model is developed in-house, this responsibility rests with the administration; if it is sourced externally, it must be contractually imposed on the provider. Such a system is required under Article 9 of the AI Act and Article 2 of the Model Contractual Clauses for the Procurement of High-Risk AI Systems (MCC-AI HR). Moreover, as per Article 2.10 MCC-AI HR, risk management must be conceived as an iterative process extending across the full lifecycle of the system.

97 This means they should be adequately annotated, labelled, cleaned as per Article 10, AI Act.

98 in line with Article 10 of the AI Act and Articles 3.1 and 3.2 MCC-AI HR.

99 Article 15 AI Act; Article 8 MCC-AI HR.

must also conduct a fundamental rights impact assessment under Article 27 AI Act¹⁰⁰. Given the knowledge asymmetry between providers and deployers, administrations are advised to require the cooperation suppliers in completing this assessment, as encouraged under Article 12 MCC-AI HR. Moreover, public entities shall always make sure that all the obligations for providers of GPAI models under Article 53 of the AI Act are met¹⁰¹. To avoid the additional regulatory burdens associated with systemic risk under Article 55 of the AI Act¹⁰² – and to strengthen accuracy, data stewardship, and environmental sustainability – public administrations should actively explore hybrid and small-scale deployment strategies. In particular, the deployment of Small Language Models (SLMs)¹⁰³, for instance fine-tuned on pre-trained solutions, should be prioritised. These smaller models could be more transparent, cost-efficient, and tailored to specific administrative tasks, thereby reducing the likelihood of hallucinations¹⁰⁴. In addition, SLMs are significantly less resource-intensive, making them a more sustainable al-

ternative to LLMs¹⁰⁵.

Nonetheless, the adoption of SLMs would still call for the implementation of a robust governance framework. This would include the adoption of clear internal guidelines on permissible use cases¹⁰⁶, comprehensive training for civil servants on the risks and limitations of GenAI, and the creation of interdisciplinary teams capable to ensure a proper human review of the outputs created by the machine. This is because, above all, the public-sector deployment of GenAI tools must always rest on the principle that meaningful human oversight is essential to the responsible use of (Gen)AI, since artificial intelligence, by its nature, does not comprehend its own limitations¹⁰⁷.

100 This should include a detailed explanation of the intended use of the system, its operational frequency and duration, oversight measures, and expected risks.

101 These were already recalled in par. 2.2 of the present article (p. 9-12).

102 IBID.

103 J. BERNABE-MORENO, *The power of small*, IBM, 2025 [website]. Available at <https://www.ibm.com/think/insights/power-of-small-language-models>; J. JOHNSON, *Small Language Models (SLM): A Comprehensive Overview*, HuggingFace, 2025 [website]. Available at <https://huggingface.co/blog/jjokah/small-language-model> and H. BLONDEAU, J. MORONEY, *Government AI: Starting Small with Small Language Models (SLMs)*, Elder Research, 2025 [website]. Available at <https://www.elderresearch.com/blog/government-ai-starting-small-with-small-language-models-slms/>.

104 A. SANTOSUOSSO and G. SARTOR, *Decidere con l'IA. Intelligenze artificiali e naturali nel diritto*, cit., p. 94-95.

105 On the matter, UNESCO, *Small language models (SLMs): A cheaper, greener route into AI*, 2024 [website]. Available at <https://www.unesco.org/en/articles/small-language-models-slms-cheaper-greener-route-ai>. In any case, the author believes that additional efforts should in any case be made to explore additional computationally efficient techniques and energy-conscious infrastructures – including the use of renewable energy sources – to mitigate the environmental impact of GenAI systems.

106 This is already the case in several EU Member States. As of April 2025, the EU Public Sector Tech Watch has mapped 31 existing guidelines focusing on the responsible use of GenAI by civil servants. See EUROPEAN COMMISSION – DIRECTORATE-GENERAL FOR DIGITAL SERVICES, A. BRIZUELA, M. COMBETTO, S. KOTOGLOU *et al.*, *Analysis of the generative AI landscape in the European public sector*, cit., p. 3.

107 J. PIATER, *What is AI really? – AI in Law and Practice: Regional Perspectives on European Rules*, Cross-Boarder Seminar, University of Innsbruck, 3rd April 2025.

IA generative nel settore pubblico secondo l'AI Act: sfide e opportunità per le procedure di appalto

Abstract

L'articolo analizza le sfide e le opportunità legate all'integrazione dell'IA Generativa nel settore pubblico alla luce dell'AI Act europeo. La GenAI presenta un notevole potenziale per migliorare l'efficienza amministrativa, ma solleva anche rilevanti questioni giuridiche e operative, in particolare nell'ambito dell'approvvigionamento pubblico. Considerando la fase del procurement come punto d'ingresso strategico, lo studio adotta un approccio ispirato alla tutela del diritto a una buona amministrazione. Dopo una disamina dei profili tecnici e normativi della GenAI, l'articolo analizza tre modelli di approvvigionamento: sviluppo interno, ricorso a fornitori privati tramite servizi di abbonamento e soluzioni ibride. Analizza le conseguenze di ciascun modello in termini di economicità, trasparenza, imparzialità, protezione dei dati e cybersecurity. Inoltre, prende in considerazione le EU Model Contractual Clauses for AI Procurement, pubblicate nel marzo 2025. L'articolo conclude auspicando l'adozione di modelli ibridi e su scala ridotta come modalità responsabile di introduzione della GenAI nel settore pubblico.

Keywords:

Public Administrations – Generative AI – AI Act – Procurement – Right to Good Administration

Parole chiave:

Pubblica amministrazione – IA generative – AI Act – Procurement – Diritto ad una buona amministrazione