Small Language Models for Public Administration: Towards Sustainable, Trustworthy, and Transparent AI Systems

Francesco Recanati[1]

[1]*Sapienza University of Rome, Piazzale Aldo Moro 5, 00185 Rome, Italy*

Abstract

This article offers a brief overview of the opportunities and challenges associated with the adoption of artificial intelligence in Public Administration, taking as a reference the *AILO* case developed by the Municipality of Trieste, a specific example of the use of *Large Language Models* (*LLMs*) in a public consultation context. It highlights the main critical issues, including privacy protection, algorithmic opacity, dependence on external providers, and legal compliance. In this context, the paper explores the hypothesis of adopting *Small Language Models* (*SLMs*), trained on domain-specific datasets, which appear to offer advantages in terms of technical sustainability, institutional control, and alignment with the European regulatory framework (*AI Act*, GDPR). The aim is to contribute to the ongoing debate on the integration of AI into public processes by outlining criteria that may guide its development and adoption in a sustainable manner and in accordance with ethical and legal principles.

Keywords

Artificial intelligence in public administration, large language models (LLMs), small language models (SLMs), administrative process automation, AI governance

1. Introduction

Public Administration (PA) is undergoing a profound transformation, driven by growing social expectations and substantial investments in the digitalization of acts and procedures. In Italy, within the framework of national policies, two main objectives are attributed to this process: on the one hand, improving service quality and strengthening operational efficiency and effectiveness; on the other, ensuring greater accessibility, transparency, and accountability towards citizens. This trajectory is rooted in the principles of *good administration* and *impartiality* enshrined in Article 97 of the Italian Constitution — principles progressively elaborated and consolidated in ordinary legislation, most notably in Law 241/1990 on administrative procedure, a cornerstone of Italian administrative law, which introduced efficiency, effectiveness, cost-effectiveness, and transparency among the general criteria governing administrative action [1] — and extensively developed and systematized in legal scholarship. The *Codice dell'Amministrazione Digitale* (Digital Administration Code, Legislative Decree 82/2005), the main legal framework for digital administration in Italy, incorporates and further specifies these principles and legal references, projecting them into the framework of digital citizenship by recognizing specific rights for citizens and imposing binding obligations on public administrations [2]. They are also translated into concrete measures, in a programmatic key, through the *Piano Triennale per l'informatica nella Pubblica Amministrazione 2024–2026* (Three-Year Plan for ICT in the Italian Public Administration 2024–2026) [3].

At the same time, the digital transformation of the Italian PA forms part of the broader European framework of the *Digital Decade 2030* — established by Decision (EU) 2022/2481 of the European Parliament and of the Council [4] — which identifies four strategic dimensions: digital skills, digital public services, the digitalization of businesses, and secure and sustainable digital infrastructures, and provides a common reference for the evolution of administrative and technological systems across the Member States.

Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23–24, 2025, Trieste, Italy
francesco.recanati@uniroma1.it (F. Recanati)
0009-0000-0260-3997 (F. Recanati)

Meanwhile, the rapid spread of *artificial intelligence* (AI) technologies — particularly *large language models* (LLMs) such as *ChatGPT* (OpenAI), *Claude* (Anthropic), or *PaLM* (Google) — is fueling new expectations regarding the automation of an increasing number of low value-added administrative tasks. Among these is, for instance, the drafting of documents and official acts, which could be partially delegated to AI systems, thereby reducing the workload of public employees in repetitive and standardized duties, where human contribution is limited or primarily executive.

These expectations are fully acknowledged in the *Strategia italiana per l'intelligenza artificiale 2024–2026* (Italian Strategy for Artificial Intelligence 2024–2026) [5], which identifies Public Administration as one of the priority sectors for AI adoption. The strategy promotes an ethical and transparent use of AI, aimed at improving the quality of public services and enhancing the relationship with citizens. As stated in the document: "Artificial intelligence can become a central factor in the digital transformation of public administration, thanks to its potential both as a tool to make internal activities more efficient and as a means of delivering services that better respond to citizens' needs" (p. 21).

However, the question remains as to how the strategic guidelines outlined in the Plan will be translated into concrete applications in the coming years. In this respect, a first attempt at operational guidance is represented by the *Bozza di Linee guida per l'adozione dell'IA nella pubblica amministrazione* (Draft Guidelines for the Adoption of AI in Public Administration) [6], published by AgID on 18 February 2025 and submitted to public consultation, which concluded on 20 March 2025. The document seeks to regulate the use, procurement, and development of AI systems in the PA by providing legal references, mandatory technical criteria, prohibitions, recommendations, and methodological principles. Its aim is to support administrations in the conscious and responsible integration of intelligent systems, with particular attention to regulatory compliance and organizational impact.

The introduction of AI systems in Public Administration should, however, also mark a new phase in the digitalization process: one that goes beyond mere dematerialization and procedural automation, moving instead toward forms of "cognitive" automation. This term denotes the use of systems capable not only of executing technical operations, but also of interpreting texts, supporting decision-making processes, and generating content in natural language.

2. The AILO Project of the Municipality of Trieste

A noteworthy example in this direction is the experimental project launched by the Municipality of Trieste in 2024, aimed at exploring the use of *artificial intelligence* in participatory processes. The initiative, promoted by the *Dipartimento Politiche Territoriali* (Department of Territorial Policies) in collaboration with the *Dipartimento Innovazione e Servizi al Cittadino* (Department of Innovation and Citizen Services) and the *Servizio Open Government* [Open Government Service], involved the deployment of a system named *AILO* (*Artificial Intelligence for Local Opinion*). It represents one of the first documented cases in Italy in which a *large language model* (LLM) — specifically OpenAI's *ChatGPT-4*, accessed via API — was integrated into a formal public consultation process.

The consultation focused on the redevelopment of the Barcola waterfront and was conducted through a participatory platform managed by the municipality. It involved more than 2,500 citizens and generated over 8,000 open-ended comments. To analyze this extensive linguistic corpus, an *AI*-based query system was developed by Luca Bandelli, a graduate in Artificial Intelligence from the University of Groningen. The system was designed to interact with *GPT-4* through structured prompts and semantic queries [7].

Municipal staff were able to engage directly with the content of the responses, posing questions such as "What issues are raised by residents of the Roiano neighborhood?" or "How many citizens mention cycling mobility?" and receiving concise, categorized answers. The interface allowed opinions to be filtered by neighborhood, topic, or general orientation, providing a disaggregated thematic and territorial overview useful for identifying local differences and emerging trends.

Within a few days, the system generated a 22-page structured report, including thematic summaries, qualitative analyses, and other linguistic insights [8]. According to the Director of the Department, Lorenzo Bandelli, the same task would have required "many days of work" by municipal staff if

performed manually. The project was ultimately presented publicly to the local community as a concrete example of institutional innovation aimed at strengthening democratic participation and administrative transparency.

3. Large Language Models and Public Administration

The direct adoption of general-purpose *LLMs* — often operated by non-EU entities — in public sector contexts raises legal, ethical, and institutional challenges that are far from marginal. Key concerns include the absence of data localization mechanisms, lack of model auditability, opacity in generative logic, limited interpretability, risks to data confidentiality, and strong dependence on external commercial providers. This condition runs counter to the principle of *digital sovereignty* promoted at the European level (see the European Chips Act) [9].

The *Strategia italiana per l'intelligenza artificiale 2024–2026* (Italian Strategy for Artificial Intelligence 2024–2026) aligns with this perspective. In the section dedicated to Public Administration, it explicitly states: "In line with the general objectives of the national strategy, public administration will not only need to develop procurement expertise for AI solutions available on the market, but must also acquire the capacity and take concrete steps to develop its own solutions, prioritizing platforms developed in Italy and fully compliant with national and EU regulations regarding risk classification and related requirements, with particular reference to the protection of personal data" (p. 23).

An additional layer of complexity arises from the nature of legal language itself: inherently normative, context-sensitive, and deeply embedded in a stratified and multi-level regulatory system that encompasses local, regional, national, and European frameworks. Automated systems must therefore be capable not only of generating formally correct texts, but also of complying with institutional conventions and practices, ensuring consistency with laws, regulations, and official guidelines [10].

Although general-purpose *LLMs* demonstrate remarkable capabilities in natural language generation, their application in the public domain raises a number of substantive concerns, including:

- substantial performance and computational costs, requiring specialized hardware often inaccessible to smaller public entities;
- significant privacy risks, particularly with models developed by non-European companies, with direct implications for GDPR compliance;
- algorithmic opacity, which hinders auditability and transparency;
- a high risk of bias and discrimination in training data;
- and, not least, the absence of legal compliance guarantees — a critical issue for texts intended for binding regulatory contexts.

In other words, unless properly adapted, validated, and governed according to accessible and institutionally legitimate criteria, general-purpose *LLMs* risk becoming technically powerful yet administratively inadequate tools.

4. Small Language Models and Public Administration

In this context, it becomes essential to explore alternative, more contextualized and accessible solutions for AI adoption in Public Administration, solutions that require less complex infrastructure and lower computational power. Among these, particular attention should be devoted to the development and use of *Small Language Models* (*SLMs*): compact models — such as *TinyBERT* (Huawei), *DistilBERT* (Hugging Face), *MobileBERT* (Google), or *TinyLLaMA* (Open LLM) — based on transformer architectures, with a parameter count generally ranging from a few million to around 100 million. These models can be trained on structured collections of legal and administrative texts, tailored to specific application contexts.

By training these models on datasets built from managerial determinations, council resolutions, public notices, technical reports, and procurement documents, local administrations can semi-automate the

drafting of both standardized and complex texts, while retaining full control over data and infrastructure and ensuring that outputs remain legally compliant.

This perspective aligns with emerging principles of *digital self-determination* and *human-centric AI*, emphasizing human oversight (human-in-the-loop and human-on-the-loop), transparency, and the strengthening of local autonomy (see the *Ethics Guidelines for Trustworthy AI*, 2019) [11]. It also fits squarely within the European regulatory framework, which adopts a risk-based approach as outlined in the *Artificial Intelligence Act (AI Act)*, recognizing the specificity of high-risk applications (Chapter III) in public services and requiring strict human oversight and traceability [12].

From a linguistic standpoint, the administrative domain is characterized by formulaic structures, legal-technical vocabulary, and a high degree of intertextuality — features that make it particularly well-suited for training specialized *SLMs*. The recursive and predictable structure of many public documents is a technical advantage for small-scale models, which are not required to "guess" user intent but rather to complete well-established patterns. This synergy between legal language and computational language not only improves model accuracy but also mitigates the risk of semantic drift, misinterpretations, or the textual "hallucinations" typical of general-purpose *LLMs*.

Recent studies show that domain-specific models, when trained on high-quality datasets, can achieve strong performance with limited computational resources [13]. In this scenario, *SLMs* — models that are "tiny but mighty" — emerge as a more technically sustainable, organizationally manageable, and legally coherent alternative, aligned with goals of digital sovereignty and regulatory compliance. Their structural advantage lies in the fact that public bodies can maintain full control over data, models, and outputs. Unlike proprietary *LLMs* — often hosted on extra-European infrastructure and governed by opaque licensing schemes — *SLMs* can be trained and deployed in public or consortial environments, ensuring independence from private vendors and enhanced cybersecurity. In an era where digital sovereignty is increasingly recognized as a condition for democratic legitimacy, this issue is not peripheral but central.

Italian Public Administrations have, for decades, produced vast amounts of highly structured, standardized, and recurring texts: resolutions, determinations, minutes, calls for tenders, reports, convocations. This corpus constitutes a semantic heritage that is already available for training *SLMs*, drastically reducing data acquisition and pre-processing costs.

Another key strength is model auditability: an *SLM* developed in a public setting can undergo independent verification, code audits, bias testing, semantic validation, and source checking. This is crucial in the public sector, where automation cannot come at the expense of transparency, impartiality, and accountability. Unlike proprietary "black box" models, public *SLMs* can be documented, certified, and — if necessary — contested, in line with the principle of algorithmic transparency set out in the *AI Act* (Art. 13).

The compact design of *SLMs* also offers economic and ecological advantages. A model with approximately 100 million parameters can be run on local hardware with affordable computational costs and drastically lower energy consumption compared to multi-billion-parameter models. In a context of growing concern about the environmental impact of digital technologies [14], this represents a strategic lever for promoting the ecological transition of AI in Public Administration.

From a legal standpoint, *SLMs* facilitate the application of the principle of data minimization (Articles 5 and 6 GDPR), can be hosted on local servers to avoid extra-EU data transfers (Article 44 GDPR), and lend themselves to the native implementation of the requirements set forth by the EU *AI Act*, such as technical documentation, risk management, and traceability. In other words, these are models inherently aligned with European law, rather than merely adapted *ex post*.

5. Practical Applications of SLMs in Public Administration

The application areas of *Small Language Models* (*SLMs*) in Public Administration are numerous and, from a technical standpoint, already fully mature. Among the most promising use cases are:

- the automated generation of standard documents (determinations, resolutions, decrees, convocations, minutes) based on legal templates and data provided by staff;
- the pre-filling of calls for tenders and public notices, adaptable to relevant regional or sector-specific regulations;
- the automatic summarization of documents intended for councils and executive boards, with key decision points highlighted;
- the classification and archiving of documents using semantic tagging and intelligent retrieval systems;
- the development of intelligent citizen-facing interfaces (dynamic FAQs, chatbots, automated digital help desks) aimed at improving access to services and the quality of administrative information.

These tools do not replace public employees; rather, they enhance their interpretative, drafting, and decision-making capacities.

6. Conclusions

The introduction of *AI* into Public Administration is not merely a technological evolution; it represents a profound institutional transformation that demands a solid ethical and legal framework. For this reason, the use of *AI* — particularly language models — cannot be guided solely by efficiency criteria; it must be shaped by the principles of transparency, fairness, accountability, and human oversight [15].

Institutions must retain the capacity to operate through the drafting, interpretation, and governance of their own acts, without relying on opaque technologies or unverifiable external models. *Small Language Models* (*SLMs*) appear to meet this need: lightweight, adaptable tools that can integrate *AI* into administrative processes without compromising the principles of good administration, impartiality, legality, traceability, and institutional autonomy.

Ultimately, this is not simply about "adopting AI"; it is about gradually building an administration capable of using it in ways consistent with the ethical and legal principles governing public action — in the public interest and under citizen scrutiny — thereby ensuring the right of citizens to hold the administration accountable.

Acknowledgments

This contribution reflects solely the personal views of the author, who, in addition to being a PhD student at Sapienza University of Rome, also works as an administrative-accounting officer for the Municipality of Trieste. It does not represent an official position of the municipal administration. The AILO initiative is mentioned exclusively as a documented case study, with the aim of offering some general considerations. The author was not directly involved in the project.

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT-4 (OpenAI), Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] Italy. *Legge 7 agosto 1990, n. 241. Nuove norme in materia di procedimento amministrativo e di diritto di accesso ai documenti amministrativi*, and subsequent amendments.

[2] Italy. *Decreto Legislativo 7 marzo 2005, n. 82. Codice dell'amministrazione digitale*, and subsequent amendments.

[3] Agenzia per l'Italia Digitale (AgID). *Piano Triennale per l'Informatica nella Pubblica Amministrazione 2024–2026*.
[4] European Union. *Decision (EU) 2022/2481 of the European Parliament and of the Council of 14 December 2022 establishing the Digital Decade Policy Programme 2030*.
[5] Presidenza del Consiglio dei Ministri – Dipartimento per la Trasformazione Digitale. *Strategia italiana per l'Intelligenza Artificiale 2024–2026*.
[6] Agenzia per l'Italia Digitale (AgID). *Bozza di Linee guida per l'adozione dell'IA nella pubblica amministrazione*.
[7] Redazione (Ed.). AILO: utilizzo pratico della IA al Comune di Trieste, *Digitale Popolare*, 20 August 2024. URL: https://digitalepopolare.it/ailo-utilizzo-pratico-della-ia-al-comune-di-trieste/.
F. Bercic. A Barcola più parcheggi, mobilità e aree verdi: ecco gli esiti del questionario rielaborati con l'IA, *Il Piccolo*, 14 August 2024. URL: https://www.ilpiccolo.it/cronaca/a-barcola-piu-parcheggi-mobilita-e-aree-verdi-ecco-gli-esiti-del-questionario-rielaborati-con-lia-vtooc142.
[8] D. Lettig. L'intelligenza artificiale analizza il sondaggio su Barcola, *RaiNews*, 14 August 2024. URL: https://www.rainews.it/tgr/fvg/articoli/2024/08/lintelligenza-artificiale-analizza-il-sondaggio-su-barcola-73ed571d-9f32-4689-9356-44e75cace2c3.html.
[9] European Union. *European Chips Act*. Regulation of the European Parliament and of the Council, 2023.
[10] G. Sartor. *L'informatica giuridica e le tecnologie dell'informazione*, Giappichelli, Torino 2022, pp. 327–362.
[11] High-Level Expert Group on AI. *Ethics Guidelines for Trustworthy AI*, European Commission, 2019.
[12] European Union. *Artificial Intelligence Act (AI Act)*. Regulation of the European Parliament and of the Council, 2024.
[13] S. Subramanian, V. Elango, M. Gungor. Small Language Models (SLMs) Can Still Pack a Punch: A Survey. *arXiv:2501.05465 [cs.CL]*, 2025. doi:10.48550/arXiv.2501.05465.
[14] E. Strubell, A. Ganesh, A. McCallum. Energy and Policy Considerations for Deep Learning in NLP. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 3645–3650. Association for Computational Linguistics. doi:10.18653/v1/P19-1355.
[15] L. Floridi. *The Ethics of Artificial Intelligence*, Oxford University Press, 2023.