Responsible Municipal Artificial Intelligence in Central and Eastern Europe: Governance, Ethics, and Efficiency in Bulgaria and Slovenia

Nora Dyankova

Sofia University ” St.Kliment Ohridski ”, Bulgaria, ndjankova@uni-sofia.bg

Phd student, Sofia University ” St.Kliment Ohridski ” Sofia , Bulgaria.

Abstract

Artificial intelligence is increasingly promoted as a tool for modernising public administration through faster service delivery, improved resource allocation, predictive analytics, and stronger decision support. Yet at the municipal level, where public institutions interact most directly with citizens, artificial intelligence cannot be evaluated only as a technical instrument or an efficiency measure. This article develops a conceptual and exploratory comparative framework for understanding municipal artificial intelligence through an integrated governance–ethics–efficiency lens. The central argument is that municipal artificial intelligence creates sustainable public value only when performance gains are aligned with institutional capacity, accountability, reliability, fairness, accessibility, and social inclusion. Drawing on literature on artificial intelligence in public administration, intermunicipal cooperation, critical infrastructure governance, smart city fairness, and recent research on artificial intelligence integration in Slovenian public management, the article applies the framework to municipal artificial intelligence readiness in Central and Eastern Europe. Bulgaria is used as the primary case because of its strategic ambition to develop artificial intelligence and the uneven administrative capacity of small and medium-sized municipalities. Slovenia is used as a comparative benchmark because of its more coordinated artificial intelligence and digital governance ecosystem. The article identifies governance capacity, data readiness, staff expertise, procurement competence, ethical oversight, and intermunicipal cooperation as key conditions for responsible municipal artificial intelligence. It concludes that municipalities should adopt pre-deployment impact assessments, accountability protocols, human oversight, shared capacity-building mechanisms, and multidimensional evaluation metrics before deploying artificial intelligence tools in high-impact public services.

Practitioner Points

For public administration practitioners, the main message is that artificial intelligence should be treated as a governance reform rather than as a simple software purchase. Municipalities should begin with the public value problem to be addressed, not with the technology to be acquired. Before deployment, local governments should assess service performance, legal compliance, institutional accountability, data quality, accessibility, cybersecurity, vendor dependence, and the possibility of discriminatory outcomes. The Bulgaria–Slovenia comparison shows that national artificial intelligence strategies are not sufficient unless they are translated into local capacity, procurement standards, staff training, and practical oversight mechanisms. Smaller municipalities should not be expected to develop full artificial intelligence governance capacity alone. They should use intermunicipal cooperation to share procurement templates, audit expertise, model documentation standards, digital skills training, and access to legal and technical

support. Human oversight should remain meaningful in domains such as social services, infrastructure, mobility, and welfare-related decision support. Evaluation should also go beyond speed and cost reduction. Municipal leaders should measure reliability, accessibility, complaint outcomes, staff experience, citizen trust, and distributive effects across neighbourhoods and vulnerable groups.

Keywords

accountability, artificial intelligence, ethics, municipal governance, public administration

1. Introduction

Artificial intelligence has become one of the most visible and consequential elements of contemporary public sector digital transformation. Governments increasingly experiment with machine learning, natural language processing, predictive analytics, computer vision, robotic process automation, and decision-support tools in order to improve administrative performance. These technologies are often associated with faster response times, reduced administrative burden, better forecasting, improved allocation of resources, and more personalised public services. In this sense, artificial intelligence appears to respond to a familiar problem in public administration: governments are expected to do more, to do it faster, and to do it with fewer resources, while also meeting rising citizen expectations shaped by digital services in the private sector.

The municipal level deserves special attention in this debate. Municipalities are the public institutions closest to residents. They manage administrative services, local infrastructure, mobility, utilities, social support, housing-related procedures, permitting, public information, public spaces, and many everyday interactions between citizens and government. Artificial intelligence at this level is therefore not an abstract technological development. It may affect how a resident receives information about social support, how a traffic light responds to congestion, how water leakage is detected, how maintenance priorities are set, or how a digital portal routes service requests. Because municipal services are embedded in everyday life, even apparently routine algorithmic decisions can have visible and unequal consequences.

The promise of artificial intelligence in local government is considerable. A municipal chatbot may answer routine questions outside office hours. Predictive maintenance may allow a city to repair infrastructure before failure occurs. A traffic management system may reduce congestion and emissions. Data analytics may help administrators identify service bottlenecks and allocate staff more effectively. These are not trivial benefits. Local governments in many countries operate under fiscal pressure, face shortages of specialised staff, and must respond to increasingly complex urban problems. In such conditions, tools that improve organisational learning, forecasting, and routine processing are attractive.

However, the same features that make artificial intelligence attractive also make it difficult to govern. Algorithmic systems can be opaque, technically complex, highly dependent on data quality, and difficult to challenge once embedded in administrative routines. Municipalities often lack specialised technical expertise, strong procurement capacity, or internal auditing functions. Smaller municipalities may be especially dependent on private vendors and

standardised tools whose assumptions were not designed for local social conditions. In addition, digital transformation may reproduce existing inequalities when access, literacy, disability, language, income, gender, age, or neighbourhood differences are not incorporated into design and evaluation. A service may become faster for digitally skilled residents while becoming less accessible for others.

The problem addressed in this article is therefore not whether artificial intelligence can improve municipal administration. In some cases, it clearly can. The more important question is under what institutional and ethical conditions artificial intelligence can produce legitimate and sustainable public value at the municipal level. A narrow focus on efficiency is insufficient because speed, cost reduction, and automation do not automatically produce better public administration. An algorithmic system can be efficient but unreliable, innovative but inaccessible, technically impressive but discriminatory, or cost-saving but democratically unaccountable. Public administration cannot evaluate technological innovation by operational performance alone, because public organisations are bound by legality, fairness, proportionality, accountability, and service obligations to the whole population.

This article develops an integrated governance–ethics–efficiency framework for municipal artificial intelligence. The framework is built on a simple but demanding claim: municipal artificial intelligence creates sustainable public value only when efficiency gains are embedded in institutional capacity, ethical legitimacy, reliability, transparency, and social inclusion. Governance, ethics, and efficiency should not be treated as separate policy domains. They are interdependent conditions of successful municipal innovation. Weak governance may turn a potentially useful tool into an accountability risk. Weak ethics may undermine trust and produce exclusion. A narrow efficiency logic may crowd out democratic values and create fragile systems that perform well in demonstration settings but fail under real administrative pressure.

The article contributes to the literature in four ways. First, it shifts attention from artificial intelligence in government generally to artificial intelligence in municipalities, where administrative proximity to citizens makes questions of accessibility, trust, and practical accountability especially salient. Second, it connects debates that are often treated separately: public sector artificial intelligence, intermunicipal cooperation, critical infrastructure governance, and smart city justice. Third, it proposes a compact analytical framework that can be used by researchers and practitioners to examine whether municipal artificial intelligence systems are institutionally governable, ethically legitimate, and operationally useful. Fourth, it extends the framework to the Central and Eastern European context by comparing Bulgaria with Slovenia as a regional benchmark for municipal artificial intelligence readiness.

Accordingly, the main research question of this article is:

Under what governance and ethical conditions can artificial intelligence create sustainable public value in municipal public administration in Bulgaria, and what lessons can be drawn from Slovenia as a Central and Eastern European benchmark?

This main research question is supported by three subquestions:

1. What governance capacities are necessary for municipalities to adopt and supervise artificial intelligence responsibly?
2. How can ethical concerns such as fairness, transparency, accessibility, human oversight, and contestability be integrated into municipal artificial intelligence implementation?
3. How can efficiency gains from artificial intelligence be evaluated together with public value, accountability, and social inclusion rather than through speed and cost reduction alone?

The article proceeds as follows. The next section reviews the relevant literature and identifies the conceptual gap that motivates the framework. The methodology section explains the design of the article as a conceptual synthesis combined with an exploratory Bulgaria–Slovenia comparative case extension. The fourth section presents the governance–ethics–efficiency framework and develops three analytical propositions. The fifth section applies the framework to selected municipal use cases. The discussion section identifies implementation conditions and policy implications for municipalities, including a comparative readiness assessment for Bulgaria and Slovenia and the role of intermunicipal cooperation. The conclusion summarises the argument and outlines directions for future research.

2. Literature Review

2.1 Artificial Intelligence and Public Administration

The public administration literature increasingly treats artificial intelligence as part of a broader transformation in the organisation, delivery, and evaluation of public services. Digital transformation has changed expectations about government responsiveness and has encouraged public organisations to redesign processes rather than merely digitise existing paperwork (Mergel et al., 2019). Artificial intelligence intensifies this transformation because it introduces systems that can classify, predict, recommend, generate language, detect patterns, and support or partially automate administrative choices. Wirtz et al. (2019) identify multiple public sector application areas, including knowledge management, process automation, predictive analytics, and service interaction. De Sousa et al. (2019) similarly show that artificial intelligence research in the public sector has expanded across policy domains, including general public services, economic affairs, and environmental protection.

Much of the initial policy enthusiasm has focused on efficiency. Artificial intelligence promises faster processing, improved information management, better detection of anomalies, and stronger decision support. These benefits are especially relevant in administrations that face high caseloads, limited staff capacity, fragmented data systems, or pressure to provide digital services around the clock. In municipal settings, these pressures are intensified because local governments are expected to deliver visible and responsive services while often having fewer technical resources than central administrations.

At the same time, research on algorithmic administration emphasises that artificial intelligence is not neutral. Algorithmic systems translate data, objectives, assumptions, and organisational priorities into operational decisions. Janssen and Kuk (2016) warn that big data algorithms may strengthen technocratic governance while creating

problems of opacity, bias, and accountability. Veale, Van Kleek, and Binns (2018) argue that public sector machine learning changes the relationship between administrative data, professional discretion, and service delivery. Levy et al. (2021) show that algorithmic systems used in public sector decision-making raise questions across the entire life cycle, including problem formulation, procurement, deployment, oversight, and evaluation. These concerns are not merely technical; they concern the character of democratic administration itself.

The emerging governance literature also stresses that artificial intelligence in the public sector involves cross-sector collaboration. Governments frequently rely on private vendors, universities, consultants, and technology providers. Mikhaylov et al. (2018) argue that public sector artificial intelligence requires collaboration across sectors, but that collaboration creates management challenges. For municipalities, these challenges include vendor lock-in, limited access to training data or model documentation, difficulty auditing proprietary systems, and weak internal capacity to negotiate contracts that protect public values. The governance of artificial intelligence therefore begins before deployment, at the stage of problem definition and procurement.

International governance frameworks increasingly reflect this life-cycle perspective. The OECD AI Principles, adopted in 2019 and updated in 2024, emphasise inclusive growth, human-centred values, transparency, robustness, security, safety, and accountability (OECD, 2024). The European Union's Artificial Intelligence Act introduces a risk-based regulatory structure, with particular obligations for high-risk systems and transparency requirements for certain uses (European Parliament and Council, 2024). The NIST AI Risk Management Framework organises responsible artificial intelligence around governance, mapping, measurement, and management of risks (NIST, 2023). The European ethics guidelines for trustworthy artificial intelligence identify requirements such as human agency and oversight, technical robustness, privacy and data governance, transparency, non-discrimination, societal well-being, and accountability (High-Level Expert Group on Artificial Intelligence, 2019). Together, these frameworks confirm that artificial intelligence governance requires more than technical performance.

Yet general frameworks do not automatically solve municipal implementation problems. Local administrations must translate broad principles into procurement clauses, staff responsibilities, complaint procedures, documentation practices, public communication, and budget decisions. A principle such as transparency must be converted into concrete questions: What must be explained, to whom, at what level of detail, and with what possibility of contestation? A principle such as human oversight must be converted into administrative design: Which official has authority to override an output, how is that override documented, and how are front-line workers trained to avoid automation bias? The gap between general principles and municipal routines is one of the central reasons for the framework proposed in this article.

2.2 Intermunicipal Cooperation and Uneven Capacity

A second body of literature concerns intermunicipal cooperation. This literature is important because the capacity to adopt and govern artificial intelligence varies sharply across municipalities. Larger cities may have digital transformation offices, data units, procurement specialists, cybersecurity staff, and legal departments. Smaller municipalities may have limited staff and rely heavily on external vendors or national platforms. Even when artificial

intelligence appears economically rational, adoption depends on organisational readiness, leadership, coordination capacity, political-administrative culture, and the ability to manage risk.

Intermunicipal cooperation has long been studied as a response to territorial fragmentation, limited scale, and fiscal pressure. Hulst and van Montfort (2007) show that European municipalities have used cooperation as a way to address service delivery challenges while retaining local political identity. Teles and Swianiewicz (2018) emphasise the institutional and governance arrangements that sustain cooperation among local governments. Bel and Warner (2015) examine the relationship between cooperation and costs, showing that expected savings are contingent rather than automatic. Feiock (2007) explains voluntary regional governance through a political contracting logic in which cooperation emerges when expected benefits exceed transaction costs.

These insights matter for artificial intelligence because capacity is itself a governance condition. A municipality that lacks data governance expertise may be unable to assess whether a vendor's model is appropriate. A municipality that lacks legal capacity may not identify risks related to personal data, discrimination, or administrative due process. A municipality that lacks staff training may be unable to integrate an artificial intelligence tool into workflows without creating confusion or overreliance. Intermunicipal cooperation can potentially reduce these capacity gaps through shared procurement templates, joint audit pools, common ethical standards, shared training programmes, and regional centres of expertise.

However, cooperation should not be idealised. It can also create coordination costs, blurred accountability, and uneven influence among participating municipalities. A shared artificial intelligence procurement framework may be efficient, but it may also privilege the needs of larger municipalities. A regional data platform may improve analytical capacity, but it may also obscure who is responsible for errors. The intermunicipal cooperation literature therefore suggests a balanced conclusion: cooperation is a promising mechanism for responsible municipal artificial intelligence, but it requires explicit governance design rather than assumption that shared systems automatically produce shared public value.

2.3 Artificial Intelligence in Critical Infrastructure and High-Impact Services

A third literature concerns artificial intelligence in critical infrastructure and high-impact systems. Municipalities manage or coordinate services in which failures can affect safety, continuity, economic life, public health, and trust in government. Water systems, waste management, traffic systems, energy-related coordination, emergency response, bridges, roads, and digital service platforms are not merely administrative conveniences. They are infrastructures on which residents and businesses rely. Artificial intelligence in such contexts must therefore be evaluated not only for accuracy or speed, but also for robustness, resilience, security, accountability, and performance under real operating conditions.

The NIST AI Risk Management Framework is especially relevant here because it defines artificial intelligence risk as a socio-technical matter that can affect individuals, organisations, and society (NIST, 2023). Its emphasis on mapping context, measuring risks, and managing them through governance is compatible with municipal needs. Similarly, ISO/IEC 42001:2023 treats artificial intelligence management as an organisational system, not only as a technical

compliance checklist (ISO, 2023). These frameworks are not specific to municipalities, but they reinforce the idea that high-impact artificial intelligence requires structured processes, documented responsibilities, continuous monitoring, and risk-based management.

Critical infrastructure applications also reveal the limitations of generic ethical principles. A principle such as fairness may have different operational meaning in a chatbot, a water leakage prediction system, a social services prioritisation model, and an adaptive traffic signal network. In infrastructure contexts, fairness may involve the distribution of maintenance attention across neighbourhoods, resilience for peripheral areas, or the avoidance of data gaps that make some districts less visible to predictive systems. Reliability may involve not only average model accuracy but also performance during extreme weather, data outages, sensor failure, or unusual demand. Accountability may require escalation protocols when an automated alert is ignored or when a vendor-provided model fails.

The literature on high-impact systems therefore supports the claim that municipal artificial intelligence should be assessed in context. The risks of a system depend on the affected service, the population served, the degree of automation, the possibility of human review, the availability of alternative access channels, and the severity of potential harm. Municipalities should not adopt a single artificial intelligence checklist for all uses. They should instead apply a proportionate governance model that is stricter when services are essential, rights-affecting, safety-related, or likely to affect vulnerable residents.

2.4 Fairness, Inclusion, and Smart City Governance

A fourth literature concerns smart cities, fairness, and inclusion. Smart city initiatives often present data-driven technologies as tools for efficiency, sustainability, and innovation. Kitchin (2014) shows that big data and smart urbanism create new forms of real-time analysis and technocratic governance, but also raise concerns about politics, control, privacy, and the way cities are imagined through data. Cardullo and Kitchin (2019) argue that many smart city initiatives frame citizens as users or consumers rather than as rights-bearing participants in democratic urban governance. These critiques are directly relevant to municipal artificial intelligence, because artificial intelligence often operates within smart city infrastructures and digital service platforms.

Inclusion is central because digital innovation can reproduce unequal access to public services. Residents differ in internet access, language proficiency, literacy, disability status, age, income, trust in government, and familiarity with digital interfaces. A service that is technically available online may not be meaningfully accessible to all residents. A chatbot that works well in the dominant language may underserve linguistic minorities. A predictive mobility model may be more accurate in sensor-rich central areas than in peripheral neighbourhoods. A social services decision-support system may rely on data that reflects past administrative surveillance of poor residents rather than genuine need.

The fairness literature on algorithmic systems reinforces this point. Barocas and Selbst (2016) show how data mining can produce disparate impact even when explicit discriminatory intent is absent. Mittelstadt et al. (2016) map ethical concerns associated with algorithms, including opacity, bias, and the difficulty of assigning responsibility. Eubanks (2018) demonstrates how automated decision-making in welfare, housing, and child protection can intensify the

monitoring and punishment of poor and working-class people. In public administration, these concerns are especially serious because residents often cannot opt out of government services in the same way they might avoid a private platform.

Fairness in municipal artificial intelligence therefore requires more than technical bias testing. It requires attention to institutional context, social vulnerability, participatory design, accessibility, complaint mechanisms, and the distribution of benefits and burdens across the urban population. A tool that improves average service speed may still worsen inequality if disadvantaged groups are more likely to receive wrong answers, be misclassified, lack digital access, or find it difficult to appeal an outcome. The literature therefore supports a multidimensional view of artificial intelligence success: municipal systems should be evaluated by their contribution to public value, not only by operational metrics.

2.5 Conceptual Gap

The four literatures reviewed above provide a strong foundation but remain partly fragmented. Public sector artificial intelligence research identifies applications and challenges, but often operates at a broad governmental level. Intermunicipal cooperation explains uneven capacity and shared service arrangements, but has rarely been connected systematically to artificial intelligence governance. Critical infrastructure frameworks emphasise risk, robustness, and accountability, but are often too general for municipal policy design. Smart city and fairness research highlights inclusion and rights, but is not always integrated into operational evaluation frameworks used by local administrations.

The research gap addressed in this article is threefold. First, existing artificial intelligence governance frameworks are often developed at a general governmental or regulatory level and do not sufficiently explain how municipalities with limited administrative and technical capacity can translate these principles into everyday service delivery. Second, research on artificial intelligence in public administration often focuses on central government, large public organisations, or general public-sector management, while municipal implementation remains underexplored. Third, the Central and Eastern European context requires closer attention because municipalities in the region often operate under uneven digital capacity, fragmented data systems, limited procurement expertise, and strong dependence on external vendors. This article addresses this gap by developing a governance–ethics–efficiency framework and applying it to municipal artificial intelligence readiness in Bulgaria, with Slovenia used as a comparative benchmark.

3. Research Methodology

This article combines conceptual synthesis with an exploratory comparative case-study design. The aim is not to test a statistical hypothesis or evaluate the technical performance of a single deployed municipal artificial intelligence system, but to build an analytical framework from existing research and apply it to municipal artificial intelligence readiness in a Central and Eastern European context. Conceptual synthesis is appropriate when a research problem cuts across several bodies of literature that have not yet been fully integrated. In this case, the relevant debates include artificial intelligence in public administration, intermunicipal cooperation, critical infrastructure governance, fairness in smart city policy, and recent artificial intelligence readiness research in public management.

The research design follows four steps. First, the article identifies recurring themes in the public administration literature on artificial intelligence: efficiency, service transformation, accountability, opacity, data quality, vendor dependence, and the changing role of human discretion. Second, it incorporates the intermunicipal cooperation literature in order to account for uneven local capacity and the possibility of shared governance mechanisms. Third, it draws on risk and critical infrastructure frameworks to clarify why high-impact municipal services require stronger oversight than low-risk digital tools. Fourth, it integrates smart city and algorithmic fairness research to ensure that inclusion, accessibility, contestability, and distributive effects are treated as central evaluation criteria rather than secondary ethical concerns.

The exploratory comparative extension focuses on Bulgaria and Slovenia. Bulgaria is selected as the primary case because it has adopted a national concept for the development of artificial intelligence until 2030 and has strategic ambitions in digital public administration, but it also faces challenges related to digital skills, cybersecurity, research and development capacity, and uneven digital uptake. These challenges are particularly relevant for small and medium-sized municipalities. Slovenia is selected as a comparative benchmark because of its more coordinated artificial intelligence and digital governance ecosystem, stronger research networks, and recent NISPAcee research on artificial intelligence integration in Slovenian public-sector management functions (Pevcin, Hržica, & Debelak, 2025). The purpose of the comparison is not to rank the two countries, but to identify governance lessons for responsible municipal artificial intelligence in the region.

The empirical material consists of three sources. First, the article analyses national and European policy documents related to artificial intelligence, digital public administration, open data, digital skills, and municipal governance in Bulgaria and Slovenia. Second, it develops a comparative readiness table that contrasts the two countries across selected dimensions: strategic framework, digital public service maturity, data readiness, administrative and technical capacity, ethical and legal governance, and intermunicipal cooperation potential. Third, it proposes an interview protocol for future fieldwork with municipal and public-sector experts. The protocol is included not as evidence of completed interviews, but as a transparent fieldwork instrument that can be used to validate and extend the framework.

The analytical strategy follows thematic coding. Documentary sources and future interview responses can be coded according to the three dimensions of the proposed framework: governance, ethics, and efficiency. Governance codes include procurement capacity, accountability, data management, vendor dependence, staff expertise, interdepartmental coordination, and intermunicipal cooperation. Ethics codes include fairness, transparency, accessibility, privacy, human oversight, inclusion, and contestability. Efficiency codes include service speed, administrative burden, cost reduction, resource allocation, responsiveness, reliability, and service quality. This coding structure allows the analysis to connect broad artificial intelligence policy frameworks with practical municipal implementation conditions.

The article has three limitations. First, because it is conceptual and exploratory, it does not measure the actual performance of existing municipal artificial intelligence systems. Empirical research is needed to test the propositions developed here across different countries, municipal sizes, service domains, and institutional traditions. Second, the Bulgaria–Slovenia comparison is based primarily on policy documents and secondary literature; it should be complemented by interviews, surveys, and case studies of specific municipalities. Third, the article does not provide

a detailed legal compliance checklist for any specific jurisdiction. Although it considers regulatory frameworks such as the European Union Artificial Intelligence Act, the article's main purpose is analytical rather than doctrinal.

Despite these limitations, the methodology has practical value. Municipalities often face artificial intelligence proposals before they have developed mature governance structures. A conceptual and exploratory comparative framework can help administrators, elected officials, auditors, and researchers ask better questions at the beginning of the process: What problem is being addressed? Who will be affected? What data are being used? Who is accountable? How will errors be corrected? How will exclusion be detected? What counts as success? These questions are not substitutes for technical evaluation, but they define the conditions under which technical evaluation becomes meaningful for public administration.

3.1 Interview Protocol for Empirical Extension

The interview component is designed as a future empirical extension of the framework. It is not presented as completed fieldwork, but as a transparent methodological instrument for validating the conceptual model in Bulgarian municipalities and, where possible, for comparing selected findings with Slovenian experience. The target sample includes five to eight respondents, such as municipal digital transformation officers, information technology managers, e-government specialists, public service managers, procurement officers, social service managers, or experts working on artificial intelligence, open data, and digital public administration. Where possible, the sample should include respondents from both larger municipalities and small or medium-sized municipalities.

The proposed interview protocol focuses on eight themes: current use of artificial intelligence and algorithmic tools; expected benefits; data and technical readiness; institutional and procurement capacity; accountability and oversight; transparency and citizen rights; accessibility and inclusion; and intermunicipal cooperation. Respondents may be asked to assess municipal readiness on a basic scale and to identify the most important barriers to responsible implementation. Responses can be anonymised through codes such as BG1, BG2, BG3, SI1, SI2, and SI3.

Table 1: Interview Themes and Guiding Questions

Theme	**Indicative Questions**	**Link to Framework**
Current use and service domains	Does the municipality currently use artificial intelligence, automation, algorithmic tools, or advanced data analytics? Which service areas could benefit most?	Efficiency and governance
Expected benefits	What benefits are expected: faster services, lower administrative burden, better forecasting, improved communication, or better resource allocation?	Efficiency

Data and infrastructure readiness	Are municipal data sufficiently digitalised, reliable, interoperable, and secure for artificial intelligence use?	Governance and efficiency
Institutional capacity	Does the municipality have staff expertise, procurement competence, monitoring capacity, and political support for responsible implementation?	Governance
Accountability and oversight	Who approves, supervises, explains, corrects, and takes responsibility for artificial intelligence-supported outputs?	Governance and ethics
Transparency and citizen rights	Are citizens informed when automated tools are used, and do they have access to complaint, appeal, or correction mechanisms?	Ethics
Accessibility and inclusion	Could digital or artificial intelligence-supported services exclude elderly people, people with disabilities, linguistic minorities, or digitally disadvantaged residents?	Ethics and efficiency
Intermunicipal cooperation	Could shared infrastructure, standards, training, procurement templates, or audit capacity help smaller municipalities?	Governance and efficiency

3.2 Analytical Matrix for Future Interview Analysis

The interview protocol can be operationalised through an analytical matrix that links expected empirical findings to the governance–ethics–efficiency framework. Since the interviews are proposed as a future empirical extension, the matrix is not presented as evidence of completed fieldwork. Rather, it provides a transparent structure for coding and interpreting future interview data.

Table 2: Analytical Matrix for Future Interview Analysis

Expected Interview Finding	Governance / Ethics / Efficiency Interpretation	Implication for Municipalities
Municipal officials report limited internal expertise for evaluating artificial intelligence systems offered by external vendors.	Governance interpretation: weak technical and procurement capacity may increase vendor dependence and reduce the municipality's ability to assess risks before deployment.	Municipalities should develop procurement templates, require model documentation, and use shared technical expertise through intermunicipal cooperation.
Respondents indicate that municipal data are fragmented	Governance and efficiency interpretation: poor data readiness limits the usefulness of artificial intelligence tools and may reduce	Municipalities should improve data governance, interoperability, data quality, and cybersecurity before

across departments and not always interoperable.	service quality rather than improve it.	adopting artificial intelligence systems.
Respondents express uncertainty about who would be responsible if an artificial intelligence-supported system produced an incorrect or unfair result.	Governance interpretation: unclear accountability may create responsibility gaps between municipal officials, departments, and private vendors.	Municipalities should define accountability protocols before deployment, including responsible units, escalation procedures, complaint mechanisms, and vendor obligations.
Respondents support artificial intelligence for routine administrative tasks but are cautious about its use in social services.	Ethics and efficiency interpretation: artificial intelligence may be useful for low-risk administrative support, but high-impact services require stronger safeguards.	Municipalities should apply proportionate governance: lighter controls for low-risk internal tools and stricter oversight for services affecting rights, welfare, or vulnerable groups.
Respondents emphasise the need for human review of artificial intelligence-generated recommendations.	Ethics interpretation: meaningful human oversight is necessary to preserve professional judgment, accountability, and citizen trust.	Municipalities should ensure that officials can review, question, override, and document artificial intelligence-supported outputs.
Respondents identify elderly people, citizens with low digital literacy, and people with disabilities as groups at risk of exclusion.	Ethics interpretation: digital-first or artificial intelligence-supported services may reproduce inequality if accessibility and alternative service channels are not preserved.	Municipalities should maintain non-digital service channels, provide assisted access, and monitor the effects of artificial intelligence tools on vulnerable groups.
Respondents see intermunicipal cooperation as useful for smaller municipalities.	Governance interpretation: shared capacity can reduce the burden on municipalities with limited staff, legal expertise, procurement capacity, or technical knowledge.	Municipalities should share training, procurement standards, audit capacity, cybersecurity support, data governance standards, and impact assessment procedures.
Respondents expect artificial intelligence to improve speed and reduce administrative burden, but not to replace human decision-making.	Efficiency and ethics interpretation: artificial intelligence is more acceptable when used for augmentation rather than full automation.	Municipalities should frame artificial intelligence as a decision-support tool, especially in high-impact areas, rather than as a replacement for public officials.

4. An Integrated Governance–Ethics–Efficiency Framework

4.1 Governance

Governance refers to the institutional arrangements through which artificial intelligence systems are selected, procured, implemented, supervised, evaluated, and corrected. At the municipal level, governance includes administrative capacity, role clarity, procurement oversight, legal review, data management, auditability, risk management, interdepartmental coordination, and vendor accountability. Governance asks not only whether a system works, but who is responsible for ensuring that it works appropriately in a public service context.

Several practical questions follow. Who authorises the use of an artificial intelligence system? Who defines the service problem and the performance objective? Who verifies the quality and representativeness of the data? Who monitors errors after deployment? Who has the authority to suspend or modify the system? How can citizens challenge or appeal outcomes? How are vendor claims tested? How are staff trained to use outputs without overreliance? How is accountability preserved when a system is shared across municipalities or provided by an external supplier?

Municipal governance must cover the full life cycle. In the pre-deployment phase, municipalities should define the problem, assess alternatives, identify affected groups, evaluate legal and ethical risks, and specify procurement requirements. During deployment, they should establish human oversight, documentation procedures, service continuity plans, and channels for complaints or corrections. After deployment, they should monitor performance, error patterns, differential effects, user satisfaction, accessibility, and changes in model behaviour. Governance is therefore not a single approval moment; it is an ongoing administrative responsibility.

Governance is particularly important because artificial intelligence can blur traditional lines of responsibility. In conventional administration, a decision can often be traced to an official, a rule, a file, or a procedure. In algorithmic administration, responsibility may be distributed across data providers, software vendors, procurement officials, front-line workers, managers, and elected officials. Without explicit governance arrangements, each actor may assume that another actor is responsible. This accountability gap is especially dangerous in municipalities with limited internal expertise.

4.2 Ethics

Ethics concerns the normative conditions under which municipal artificial intelligence can be considered legitimate. It includes fairness, transparency, inclusion, accessibility, human dignity, human oversight, privacy, accountability, and contestability. Ethical legitimacy is not achieved by declaring abstract principles. It depends on whether those principles are translated into service design, procurement, monitoring, staff training, public communication, and citizen rights.

Fairness requires attention to both individual and group effects. An artificial intelligence system may be individually accurate in many cases yet systematically worse for certain neighbourhoods, languages, disability groups, income groups, or age groups. Municipalities therefore need to monitor differential performance and not rely exclusively on average accuracy. Transparency requires that affected residents and municipal staff understand when artificial

intelligence is used, what role it plays, and how decisions can be reviewed. Transparency does not always require full disclosure of model code, but it does require meaningful explanation of administrative use, limitations, and rights.

Inclusion and accessibility are especially important in municipal services. A digital service channel should not become the only practical path to public administration when some residents lack digital access or skills. Chatbots, automated forms, and digital portals should comply with accessibility standards and should offer alternative channels for residents who cannot use them effectively. Public participation also matters. In high-impact applications, municipalities should consult affected communities, front-line workers, civil society organisations, and disability advocates before deployment. Ethical governance must therefore be participatory rather than purely managerial.

Contestability is a core ethical requirement. Residents should have a practical ability to challenge, correct, or appeal outcomes that affect them. This is particularly important in social services, housing-related decisions, fines, eligibility assessment, or prioritisation of scarce public support. A system that is difficult to contest may undermine trust even if it improves administrative efficiency. Ethical legitimacy depends on whether citizens remain subjects of rights rather than objects of automated processing.

4.3 Efficiency

Efficiency captures the performance objectives that often motivate artificial intelligence adoption. Municipalities may seek faster service delivery, lower administrative costs, reduced waiting times, improved resource allocation, better forecasting, fewer errors, or improved responsiveness. Efficiency is a legitimate public value, especially when local governments face fiscal pressure and increasing service demands. Citizens also reasonably expect public services to be timely, reliable, and convenient.

However, efficiency must be defined carefully. There is a difference between internal administrative efficiency and public value efficiency. Internal efficiency may reduce staff workload or processing time. Public value efficiency asks whether the service improvement is meaningful for residents, whether it improves access, whether it reduces burdens on vulnerable groups, and whether it supports policy goals. A chatbot that reduces calls to a municipal office may be internally efficient, but if it gives incomplete answers to residents with complex needs, the public value result may be negative. A traffic optimisation system may improve average travel time, but if it diverts congestion into poorer neighbourhoods, the social result may be unjust.

Efficiency should therefore be treated as necessary but not sufficient. An artificial intelligence system cannot be considered successful if it is fast but unreliable, cheap but discriminatory, innovative but unaccountable, or automated but inaccessible. Efficiency gains are meaningful only when they are embedded in governance and ethics. This does not mean that ethics and governance are obstacles to efficiency. On the contrary, weak governance and weak ethics can undermine long-term efficiency by generating errors, complaints, legal challenges, staff resistance, vendor dependence, public distrust, or system abandonment.

4.4 Interdependence of the Three Dimensions

The three dimensions are interdependent. Governance enables ethical commitments to become operational. Ethics defines the boundaries within which efficiency gains are legitimate. Efficiency provides one reason to adopt artificial intelligence but must be assessed in relation to public value. If one dimension is weak, the others are weakened as well. Weak governance can produce ethical failure because no one monitors exclusion or corrects errors. Weak ethics can damage efficiency because public distrust reduces adoption and increases conflict. Narrow efficiency can undermine governance because pressure for rapid implementation may lead municipalities to bypass risk assessment, staff training, or citizen communication.

The framework therefore supports three analytical propositions. Proposition 1: governance capacity conditions municipal artificial intelligence effectiveness. A system is more likely to produce sustainable value when municipalities have clear accountability, procurement competence, data governance, human oversight, monitoring routines, and capacity to manage vendors. Proposition 2: ethical legitimacy conditions trust and sustainability. A system is more likely to be accepted and maintained when it is fair, transparent, accessible, contestable, and aligned with public service obligations. Proposition 3: efficiency gains are meaningful only when embedded in governance and ethics. Performance improvements should be evaluated together with reliability, rights protection, inclusion, distributive effects, and accountability.

These propositions are intentionally practical. They do not require municipalities to become advanced artificial intelligence laboratories. They require municipalities to treat artificial intelligence as an administrative system that changes responsibilities, risks, and relationships with citizens. This perspective is particularly important because artificial intelligence adoption is often framed as modernisation. Modernisation, however, is not only technological. In public administration it also concerns legitimacy, fairness, competence, and the capacity to serve the whole population.

5. Municipal Use Cases

5.1 Citizen-Facing Administrative Services

Citizen-facing administrative services are among the most visible applications of artificial intelligence in municipalities. Examples include chatbots, automated request triage, digital portals, document classification, appointment scheduling, and systems that help residents identify relevant services. These tools are attractive because many municipal inquiries are repetitive, staff time is limited, and residents increasingly expect digital access outside normal office hours. A well-designed system may reduce waiting times, improve consistency of basic information, and allow staff to focus on complex cases.

From an efficiency perspective, the benefits are clear. A chatbot can answer routine questions about working hours, required documents, forms, deadlines, or service locations. Automated triage can route a request to the correct department more quickly. Digital portals can reduce paperwork and allow residents to track request status. These functions can improve responsiveness, especially in municipalities with high service demand.

Governance challenges arise when the system provides incorrect information, fails to escalate a complex problem, or becomes the de facto gatekeeper to a service. Municipalities need clear responsibility for system content, regular updating, error monitoring, and human escalation. Staff should know when to rely on automated classification and when to override it. Procurement contracts should specify documentation, data protection, service-level requirements, accessibility standards, and the municipality's right to audit or modify the system. If a vendor provides a chatbot model, municipal officials remain responsible for the public service consequences of its use.

Ethical concerns focus on accessibility, transparency, and unequal use. Residents with low digital literacy, limited internet access, disabilities, language barriers, or complex social needs may be disadvantaged by digital-first service design. The municipality should therefore provide alternative access channels and monitor whether certain groups are less likely to use or successfully complete digital services. Citizens should be informed when they are interacting with an automated system and should have an easy route to a human official. In this use case, success is not measured only by the number of inquiries handled automatically. It is measured by whether the system improves access and reduces administrative burden without excluding residents who need public administration most.

5.2 Urban Mobility and Traffic Management

Urban mobility and traffic management are frequently presented as flagship smart city applications. Artificial intelligence may be used for congestion prediction, adaptive traffic signals, parking management, public transport optimisation, route planning, or incident detection. These systems can generate visible efficiency gains by reducing delays, improving traffic flow, lowering emissions, and supporting more responsive infrastructure management.

However, mobility systems also involve distributive choices. Traffic optimisation is not politically neutral. A model may reduce average travel time while increasing traffic through residential streets or peripheral neighbourhoods. Sensor coverage may be better in central districts, making those areas more visible to optimisation systems. Public transport data may reflect existing service inequalities rather than unmet need. A system trained on current mobility patterns may reinforce car-oriented planning rather than support broader goals such as accessibility, safety, and sustainable transport.

Governance therefore requires municipalities to define the policy objective before selecting the technology. Is the goal to reduce congestion, improve bus reliability, protect school zones, reduce emissions, increase accessibility for people with disabilities, or balance economic and residential priorities? Different objectives imply different data, metrics, and trade-offs. Municipalities should not allow technical optimisation to substitute for democratic priority-setting. Oversight should include transport planners, data specialists, elected officials, neighbourhood representatives, and accessibility experts where appropriate.

Ethically, mobility artificial intelligence should be evaluated for infrastructure justice. Municipalities should monitor whether benefits and burdens are distributed fairly across neighbourhoods and social groups. They should communicate how automated mobility systems operate and how complaints can be raised when residents experience negative effects. In this domain, efficiency metrics such as average travel time, vehicle throughput, or incident

response must be complemented by safety, accessibility, environmental, and distributive metrics. The question is not whether artificial intelligence can optimise traffic, but what definition of urban mobility it is asked to optimise.

5.3 Municipal Infrastructure and Utilities

Municipal infrastructure and utilities include water systems, waste collection, street lighting, road maintenance, drainage, energy management, and other services that affect the continuity of everyday urban life. Artificial intelligence may be used for predictive maintenance, anomaly detection, demand forecasting, route optimisation, asset management, and early warning systems. In many of these areas, the efficiency case is strong. Better forecasting can prevent failures, reduce costs, extend asset life, and support more targeted interventions.

The governance requirements are also stronger because these services can be safety-related or continuity-critical. A false negative in anomaly detection may delay response to a serious infrastructure problem. A false positive may waste limited maintenance resources. A demand forecasting model may fail under extreme weather or unusual events. A vendor-provided system may become difficult to audit when it is integrated into operational infrastructure. Municipalities therefore need clear accountability protocols, human supervision, incident response procedures, cybersecurity controls, and documentation of model limitations.

Ethical questions in infrastructure may appear less direct than in social services, but they are still important. Predictive maintenance models may prioritise areas with better historical data, better sensors, or higher economic value. Peripheral neighbourhoods may be underrepresented in data and therefore receive less preventive attention. If a municipality evaluates infrastructure artificial intelligence only by cost savings, it may miss patterns of unequal reliability across the city. Infrastructure justice requires monitoring service continuity, repair times, and investment distribution across neighbourhoods.

This use case illustrates why artificial intelligence governance must be contextual. The main ethical issue may not be individual explainability, as in a benefits decision, but reliability, resilience, and equitable allocation of maintenance attention. The municipality should evaluate whether the system improves public service continuity for the whole population, whether staff understand its limits, and whether governance arrangements are strong enough to manage failure. In infrastructure, responsible artificial intelligence means not only smarter prediction but also accountable resilience.

5.4 Social Services and Decision Support

Social services and welfare-related decision support are among the most sensitive municipal artificial intelligence applications. Systems may be used for eligibility support, risk flagging, case prioritisation, fraud detection, resource allocation, or identification of households that may need intervention. These tools are attractive in contexts of high demand, limited staff, and pressure to allocate scarce resources. However, they also affect vulnerable populations and can produce serious harm if poorly governed.

The efficiency argument in this domain must be treated with caution. Faster processing and better prioritisation may help residents if they reduce waiting times and direct support to urgent cases. But efficiency can become dangerous

when it leads to excessive automation, weak human review, or a presumption that algorithmic outputs are objective. Social services data often reflect unequal patterns of surveillance, reporting, service access, and administrative history. A model may identify poverty, disability, neighbourhood disadvantage, or prior contact with public services as risk signals, thereby reinforcing the very inequalities public administration should address.

Governance in this domain should include strict human oversight, documentation of decisions, appeal mechanisms, fairness testing, data quality review, and clear limits on automation. Artificial intelligence should support professional judgment rather than replace it. Front-line workers should be trained to understand model limitations and automation bias. Residents should know when algorithmic support has influenced an outcome and should have access to meaningful explanation and challenge. Procurement should require information about training data, performance across groups, monitoring procedures, and the ability to suspend the system if harmful patterns appear.

Ethical legitimacy is central. Social services are not ordinary administrative transactions. They often involve people in crisis, people with disabilities, children, older residents, migrants, unemployed people, and low-income households. A wrong classification may delay assistance, increase surveillance, or deny support. In this domain, municipalities should apply the strictest version of the governance–ethics–efficiency framework. Efficiency should never be accepted as sufficient justification. The relevant question is whether artificial intelligence improves professional capacity while protecting rights, dignity, and access to support.

5.5 Cross-Case Lessons

The four use cases demonstrate that municipal artificial intelligence cannot be governed through a single performance logic. Citizen-facing services require accessibility and escalation. Mobility systems require democratic priority-setting and distributive analysis. Infrastructure systems require resilience, safety, and accountability for continuity. Social services require rights protection, human oversight, and strong contestability. In every case, efficiency is relevant, but it takes meaning only in relation to the public purpose of the service.

The cases also show that risk is not only a property of the technology. It is a property of the relationship between the technology, the service domain, the data, the affected population, the degree of automation, and the municipality's governance capacity. A relatively simple rule-based system may be high impact if it controls access to social support. A complex predictive model may be lower risk if it only helps schedule internal maintenance and is reviewed by experienced staff. Municipalities should therefore avoid both technological enthusiasm and technological fear. They should assess each system by context, public value, and institutional readiness.

6. Discussion: Conditions for Responsible Municipal Artificial Intelligence

6.1 Pre-Deployment Impact Assessment

The first policy implication is that municipalities should introduce pre-deployment artificial intelligence impact assessments. Such assessments should not be treated as bureaucratic obstacles, but as tools for clarifying purpose, risk, and accountability before financial and organisational commitments are made. An assessment should identify the

public problem, the expected efficiency gains, the data required, the groups affected, the legal basis, the potential harms, the need for human oversight, the accessibility implications, and the metrics that will be used after deployment.

A municipal impact assessment should also compare artificial intelligence with non-artificial intelligence alternatives. Sometimes a service problem can be solved through better staffing, clearer forms, improved coordination, simpler rules, or basic digitalisation. Artificial intelligence should not be adopted because it is fashionable. It should be adopted when it is proportionate to a clearly defined public value problem and when the municipality has the capacity to govern it responsibly.

6.2 Accountability and Human Oversight

The second policy implication concerns accountability and human oversight. Municipalities should define responsibilities before deployment. This includes political responsibility, managerial responsibility, technical responsibility, front-line responsibility, vendor responsibility, and complaint-handling responsibility. A system should not be deployed unless there is a named municipal unit or official responsible for monitoring and intervention. Vendor accountability should be contractual, but public accountability cannot be outsourced. Residents experience the service as municipal even when the technology is provided externally.

Human oversight must be meaningful rather than symbolic. In high-impact domains, a human official should have authority, time, information, and training to question or override algorithmic outputs. Oversight should be documented so that patterns of override, disagreement, or error can inform future evaluation. Municipalities should also train staff about automation bias, selective adherence, and the limits of prediction. Research on human-artificial intelligence interaction in public decision-making suggests that human oversight does not automatically prevent bias if officials overtrust algorithmic advice or selectively follow it when it confirms stereotypes (Alon-Barkat & Busuioc, 2021).

6.3 Ethics as Practical Governance

The third policy implication is that ethics should be treated as practical governance. Municipalities often endorse values such as transparency, fairness, and inclusion, but the challenge is to operationalise them. Fairness requires data review, performance monitoring across groups, and procedures for correcting unequal effects. Transparency requires public notices, understandable explanations, documentation for staff, and communication about limitations. Inclusion requires accessible design, alternative channels, language support, and consultation with affected communities. Contestability requires complaint procedures, appeal pathways, and human review.

Ethics should also be embedded in procurement. Calls for tender should require information about data sources, model purpose, performance limitations, audit access, accessibility compliance, monitoring, cybersecurity, and procedures for addressing errors. Municipalities should avoid procurement arrangements that prevent meaningful audit or create dependence on opaque vendor systems. Where proprietary tools are used, contracts should still protect the municipality's ability to understand system behaviour, receive documentation, and suspend or modify deployment when necessary.

6.4 Intermunicipal Cooperation as Capacity Building

The fourth policy implication concerns intermunicipal cooperation. Many municipalities will not be able to develop artificial intelligence governance capacity alone. Shared capacity may be especially valuable for smaller municipalities that cannot hire specialised data scientists, algorithm auditors, procurement lawyers, accessibility experts, or cybersecurity professionals. Cooperation can take several forms: joint procurement guidelines, shared registers of artificial intelligence systems, regional audit teams, common training curricula, model documentation templates, shared ethical review boards, and joint negotiation with vendors.

However, cooperation must preserve accountability. A regional platform should not make it unclear which municipality is responsible to residents. Shared systems should include governance agreements that define decision rights, monitoring responsibilities, complaint handling, data stewardship, and cost-sharing. Smaller municipalities should have a voice in shared arrangements so that artificial intelligence tools do not reflect only the priorities of larger cities. Intermunicipal cooperation can support responsible artificial intelligence, but only if it is designed as a governance mechanism rather than merely a cost-saving instrument.

6.5 Comparative Readiness: Bulgaria and Slovenia

The Bulgaria–Slovenia comparison illustrates why municipal artificial intelligence readiness should be understood as a governance-capacity issue rather than only as a technological issue. Both countries operate within the European Union's regulatory and policy environment, including the Digital Decade agenda and the obligations created by the European Union Artificial Intelligence Act. However, their national digital governance contexts create different conditions for municipalities.

Bulgaria has adopted a concept for the development of artificial intelligence until 2030, which identifies infrastructure, skills, trust, innovation, regulation, electronic public administration, and public-sector transformation as relevant priorities. Yet national strategic ambition does not automatically translate into municipal readiness. Small and medium-sized Bulgarian municipalities may face constraints related to staff expertise, data quality, procurement capacity, cybersecurity, and the ability to supervise external vendors. These conditions make governance capacity a central precondition for responsible artificial intelligence adoption.

Slovenia provides a useful comparative benchmark. Its artificial intelligence and digital governance ecosystem places stronger emphasis on coordination, open data, trustworthy artificial intelligence, research capacity, and public-sector experimentation. Recent NISPAcee research on artificial intelligence integration in Slovenian public management functions shows the value of combining qualitative and quantitative data when assessing artificial intelligence readiness and managerial implications (Pevcin, Hržica, & Debelak, 2025). This does not mean that Slovenia has solved municipal artificial intelligence governance. Rather, it provides a regional example of a more coordinated policy environment from which Bulgaria can draw lessons.

Table 3: Comparative Artificial Intelligence Readiness Dimensions in Bulgaria and Slovenia

Dimension	Bulgaria	Slovenia

National artificial intelligence framework	National artificial intelligence concept until 2030 identifies infrastructure, skills, trust, innovation, and regulation as relevant priorities.	National artificial intelligence and digital governance frameworks emphasise trustworthy artificial intelligence, stakeholder coordination, research capacity, open data, and public-sector use.
Digital public service maturity	Digital infrastructure and e-government priorities are present, but uptake and capacity remain uneven, especially outside larger administrations.	Digital infrastructure is comparatively well developed, though digital skills remain a continuing challenge.
Data readiness	Open data and public-sector information reuse are recognised as prerequisites, but local data quality and interoperability remain uneven.	Open data and computational infrastructure are more visible in the national artificial intelligence policy ecosystem.
Administrative and technical capacity	Small and medium-sized municipalities may lack specialised staff, procurement expertise, cybersecurity capacity, and artificial intelligence audit skills.	Smaller municipalities also face capacity constraints, but can draw on stronger research networks and national coordination.
Ethical and legal governance	Trust and regulation are included in national artificial intelligence policy, but municipal-level mechanisms for impact assessment and oversight remain less visible.	Trustworthy artificial intelligence, legal safeguards, and public-sector deployment are more explicitly connected in policy discussions.
Intermunicipal cooperation potential	Highly important for sharing procurement templates, audit capacity, training, and legal expertise across municipalities.	Important for smaller municipalities, but likely supported by a more mature digital governance ecosystem.

The comparison suggests that Bulgaria and Slovenia should not be treated as opposite cases of failure and success. Rather, they represent different configurations of artificial intelligence readiness within the Central and Eastern European context. Bulgaria demonstrates the importance of building basic administrative, data, and human capacity before artificial intelligence systems are deployed widely in municipalities. Slovenia demonstrates the value of coordinated national frameworks, open data infrastructure, stakeholder involvement, and explicit attention to trustworthy artificial intelligence. For both countries, the key lesson is that small municipalities need shared governance mechanisms if they are to use artificial intelligence responsibly.

This comparative extension strengthens the governance–ethics–efficiency framework in two ways. First, it shows that efficiency gains depend on prior institutional conditions such as data quality, procurement competence, and staff

capability. Second, it shows that ethical artificial intelligence cannot be reduced to abstract principles; it requires practical instruments such as impact assessments, complaint procedures, human review, accessible service channels, and shared audit capacity. In the Bulgarian context, Slovenia can be used as a benchmark not because it provides a ready-made model, but because it illustrates how strategic coordination and public-sector artificial intelligence discussion can support more responsible implementation.

6.6 Multidimensional Evaluation

The fifth policy implication is that municipalities should evaluate artificial intelligence systems using multidimensional metrics. Speed and cost reduction matter, but they should not dominate the evaluation. Municipalities should also measure reliability, error rates, accessibility, user satisfaction, staff experience, complaint frequency, appeal outcomes, differential performance across groups, service continuity, cybersecurity incidents, and distributive effects across neighbourhoods. Evaluation should be repeated over time because systems, data, populations, and administrative contexts change.

Multidimensional evaluation can also protect efficiency itself. Systems that generate many complaints, require constant correction, or produce public distrust may appear efficient in narrow metrics but costly in practice. A tool that saves staff time in one department may shift burdens to another department or to residents. A system that performs well in pilot conditions may fail when scaled. Evaluation should therefore examine total public value, not only immediate administrative savings.

6.7 Toward a Municipal Artificial Intelligence Governance Cycle

The discussion can be summarised as a governance cycle. First, the municipality defines the public problem and considers alternatives. Second, it conducts a pre-deployment impact assessment. Third, it designs procurement and governance requirements. Fourth, it deploys with human oversight, documentation, and communication. Fifth, it monitors performance, fairness, accessibility, and complaints. Sixth, it revises, suspends, or terminates the system when evidence shows that it does not serve public value. This cycle should be proportionate to risk: lighter for low-impact internal tools and stricter for high-impact services.

Such a cycle makes artificial intelligence a normal object of public administration rather than an exceptional technological project. It connects innovation to planning, procurement, legality, human resources, citizen engagement, auditing, and evaluation. This is important because the greatest risk for municipalities may not be that they reject artificial intelligence. It may be that they adopt artificial intelligence without the administrative routines needed to govern it. Responsible innovation requires both technological openness and institutional discipline.

7. Limitations and Future Research

This article has developed a conceptual and exploratory comparative framework rather than a full empirical evaluation of deployed municipal artificial intelligence systems. The Bulgaria–Slovenia comparison is based primarily on policy documents and secondary sources, while the interview protocol and analytical matrix are included as instruments for

future validation. The article therefore does not claim to measure the actual performance of artificial intelligence systems in Bulgarian or Slovenian municipalities. Future research should conduct interviews with municipal officials, compare municipalities of different sizes, and examine specific service domains such as citizen information portals, traffic management, infrastructure maintenance, and social service decision support.

Quantitative research could develop indicators for multidimensional evaluation. For example, researchers could measure whether artificial intelligence systems reduce processing time while also tracking complaint rates, appeal success, accessibility outcomes, staff workload, cybersecurity incidents, and differential effects across social groups. Longitudinal studies would be especially useful because artificial intelligence systems may drift over time as data patterns, citizen behaviour, staff practices, and vendor updates change. What appears successful after a pilot may look different after several years of routine use.

Future research should also examine the politics of municipal artificial intelligence. Technology adoption is shaped by elected officials, administrative leaders, vendors, consultants, civil society organisations, and residents. Research should ask how priorities are set, whose expertise counts, who participates in design, and how conflicts between efficiency and rights are resolved. The role of procurement deserves particular attention because many artificial intelligence governance choices are made through contracts before citizens or even front-line staff see the system.

Finally, more research is needed on intermunicipal cooperation for artificial intelligence governance. Shared capacity may be one of the most practical ways to support smaller municipalities, but its design is not straightforward. Researchers should examine which cooperative models preserve accountability, reduce transaction costs, support expertise, and avoid domination by larger partners or private vendors. As artificial intelligence becomes more common in local government, the governance of shared systems may become a central issue for democratic public administration in Central and Eastern Europe.

8. Conclusion

Artificial intelligence in municipal public administration should not be understood as a simple digital upgrade. It is a governance and public-value challenge. Municipalities are under pressure to improve efficiency, responsiveness, and service quality, and artificial intelligence can contribute to these goals. However, local governments also have obligations that cannot be reduced to speed or cost reduction. They must preserve legality, fairness, accessibility, trust, accountability, and democratic legitimacy. These obligations are especially important because municipalities provide services that directly shape everyday life.

The article has argued that municipal artificial intelligence creates sustainable public value only when efficiency gains are aligned with governance capacity and ethical legitimacy. Governance concerns the institutional arrangements through which systems are selected, procured, supervised, audited, and corrected. Ethics concerns fairness, transparency, inclusion, accessibility, accountability, human oversight, and contestability. Efficiency concerns performance goals such as faster service delivery, lower administrative burden, improved responsiveness, and better resource allocation. The central point is that these dimensions are interdependent. Efficiency without governance may

become fragile. Governance without ethics may become formalistic. Ethics without operational design may remain symbolic.

The municipal use cases examined in the article demonstrate the practical relevance of the framework. Citizen-facing administrative services require accessibility and escalation to human staff. Mobility systems require attention to distributive effects and democratic priority-setting. Infrastructure and utilities require reliability, resilience, and accountability for failures. Social services decision support requires the strongest safeguards because it affects vulnerable residents and rights-related outcomes. Across all domains, municipalities need to ask not only whether a system works, but for whom it works, under what conditions, with what accountability, and at what social cost.

The Bulgaria–Slovenia comparison reinforces this argument. Bulgaria shows why artificial intelligence strategy must be translated into local capacity, especially for smaller municipalities that may lack staff expertise, procurement capacity, and data governance structures. Slovenia offers a useful benchmark because its artificial intelligence and digital governance environment places greater emphasis on coordination, open data, trustworthy artificial intelligence, and public-sector experimentation. The comparison does not imply that one model can simply be copied by the other. It shows that responsible municipal artificial intelligence in Central and Eastern Europe depends on the ability to connect national strategy with local implementation capacity.

The policy implications are clear. Municipalities should conduct pre-deployment impact assessments, establish clear accountability and human oversight, embed ethics in procurement and implementation, use intermunicipal cooperation to build shared capacity, and evaluate systems through multidimensional public value metrics. These measures are not barriers to innovation. They are conditions for innovation that citizens can trust.

Artificial intelligence will continue to develop, and municipalities will continue to face pressure to adopt new tools. The challenge is not to choose between innovation and caution. The challenge is to build public institutions capable of governing innovation responsibly. Municipal artificial intelligence becomes valuable when it helps local governments serve residents more effectively while remaining fair, accessible, accountable, and democratically legitimate. Efficiency gains alone are not enough. The future of artificial intelligence in municipal public administration depends on whether local governments can transform technological capacity into public value.

References

Alon-Barkat, S., & Busuioc, M. (2021). Human-AI interactions in public sector decision making: Automation bias and selective adherence to algorithmic advice. Journal of Public Administration Research and Theory, 33(1), 153–169.

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. California Law Review, 104, 671–732. https://doi.org/10.15779/Z38BG31

Bel, G., & Warner, M. E. (2015). Inter-municipal cooperation and costs: Expectations and evidence. Public Administration, 93(1), 52–67. https://doi.org/10.1111/padm.12104

Cardullo, P., & Kitchin, R. (2019). Being a citizen in the smart city: Up and down the scaffold of smart citizen participation in Dublin, Ireland. GeoJournal, 84, 1–13. https://doi.org/10.1007/s10708-018-9845-8

De Sousa, W. G., De Melo, E. R. P., Bermejo, P. H. D. S., Farias, R. A. S., & Gomes, A. O. (2019). How and where is artificial intelligence in the public sector going? A literature review and research agenda. Government Information Quarterly, 36(4), 101392. https://doi.org/10.1016/j.giq.2019.07.004

Eubanks, V. (2018). Automating inequality: How high-tech tools profile, police, and punish the poor. St. Martin's Press.

European Commission. (2025a). Bulgaria 2025 Digital Decade Country Report. European Commission.

European Commission. (2025b). Slovenia 2025 Digital Decade Country Report. European Commission.

European Parliament and Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union.

Feiock, R. C. (2007). Rational choice and regional governance. Journal of Urban Affairs, 29(1), 47–63. https://doi.org/10.1111/j.1467-9906.2007.00322.x

Government of the Republic of Slovenia. (2026). National Artificial Intelligence Strategy of Slovenia 2030 (NsUI 2030). Government of the Republic of Slovenia.

High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.

Hulst, R., & van Montfort, A. (Eds.). (2007). Inter-municipal cooperation in Europe. Springer. https://doi.org/10.1007/1-4020-5379-7

International Organization for Standardization. (2023). ISO/IEC 42001:2023: Information technology — Artificial intelligence — Management system. ISO.

Janssen, M., & Kuk, G. (2016). The challenges and limits of big data algorithms in technocratic governance. Government Information Quarterly, 33(3), 371–377. https://doi.org/10.1016/j.giq.2016.08.011

Kitchin, R. (2014). The real-time city? Big data and smart urbanism. GeoJournal, 79, 1–14. https://doi.org/10.1007/s10708-013-9516-8

Levy, K., Chasalow, K. E., & Riley, S. (2021). Algorithms and decision-making in the public sector. Annual Review of Law and Social Science, 17, 309–334. https://doi.org/10.1146/annurev-lawsocsci-041221-023808

Mergel, I., Edelmann, N., & Haug, N. (2019). Defining digital transformation: Results from expert interviews. Government Information Quarterly, 36(4), 101385. https://doi.org/10.1016/j.giq.2019.06.002

Mikhaylov, S. J., Esteve, M., & Campion, A. (2018). Artificial intelligence for the public sector: Opportunities and challenges of cross-sector collaboration. Philosophical Transactions of the Royal Society A, 376(2128), 20170357. https://doi.org/10.1098/rsta.2017.0357

Ministry of Transport, Information Technology and Communications. (2020). Concept for the development of artificial intelligence in Bulgaria until 2030. Government of Bulgaria.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2), 1–21. https://doi.org/10.1177/2053951716679679

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Organisation for Economic Co-operation and Development. (2024). OECD AI principles. OECD.AI Policy Observatory.

Pevcin, P., Hržica, R., & Debelak, K. (2025). Navigating the automation–augmentation paradox: A case study of artificial intelligence integration in public management functions. NISPAcee Journal of Public Administration and Policy, 18(2), 106–135. https://doi.org/10.2478/nispa-2025-0016

Teles, F., & Swianiewicz, P. (Eds.). (2018). Inter-municipal cooperation in Europe: Institutions and governance. Palgrave Macmillan.

UNESCO. (2022). Recommendation on the ethics of artificial intelligence. UNESCO.

Veale, M., Van Kleek, M., & Binns, R. (2018). Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, 1–14. https://doi.org/10.1145/3173574.3174014

Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial intelligence and the public sector: Applications and challenges. International Journal of Public Administration, 42(7), 596–615. https://doi.org/10.1080/01900692.2018.1498103