applied sciences

MDPI

Article

A Framework for a Public Service Recommender System Based on Neuro-Symbolic AI

Ioannis Konstantinidis [1], Ioannis Magnisalis [1] and Vassilios Peristeras [1,2,*]

[1] School of Science and Technology, International Hellenic University, 57001 Thessaloniki, Greece; i.konstantinidis@ihu.edu.gr (I.K.); i.magnisalis@ihu.edu.gr (I.M.)
[2] DG DITEC, Council of the EU, 1048 Brussels, Belgium
* Correspondence: v.peristeras@ihu.edu.gr

Abstract

Public service provision is still limited to document-centric procedures that require citizens to submit data and information needed for the execution of a service via documents. This, amongst others, is time-consuming, error-prone and hinders progress towards data-centricity. This study proposes a data-centric framework for a public service recommender system that combines knowledge graphs (KGs) and large language models (LLMs) in a neuro-symbolic AI architecture. The framework expresses public service preconditions as machine-readable rules based on data standards and provides dynamic recommendations for public services based on citizens' profiles through automated reasoning. LLMs are utilized to extract preconditions from unstructured textual regulations and create RDF-based evidence models, while KGs provide validation of preconditions through SHACL rules and explainable reasoning towards semantic interoperability. A prototype use case on students applying for housing allowance showcases the feasibility of the proposed framework. The analysis indicates that combining KGs with LLMs for identifying relevant public services for different citizens' profiles can improve the quality of public services and reduce administrative burdens. This work contributes and promotes the proactive "No-Stop Government" model, where services are recommended to users without explicit requests. The findings highlight the promising potential of employing neuro-symbolic AI to transform e-government processes, while also addressing challenges related to legal complexity, privacy and data fragmentation for large-scale adoption.

Keywords: large language models; knowledge graphs; public service provision; recommender systems; neuro-symbolic AI; no-stop government; data-centric governance

Academic Editor: Luis Javier Garcia Villalba

Received: 9 September 2025
Revised: 8 October 2025
Accepted: 14 October 2025
Published: 20 October 2025

Citation: Konstantinidis, I.; Magnisalis, I.; Peristeras, V. A Framework for a Public Service Recommender System Based on Neuro-Symbolic AI. *Appl. Sci.* **2025**, *15*, 11235. https://doi.org/10.3390/app152011235

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Public service provision constitutes a cornerstone of modern societies, encompassing a wide range of administrative, social, and economic activities crucial for citizens and businesses alike. Governments worldwide have invested heavily in Information Technology (IT) infrastructures to streamline operations, reduce bureaucratic burdens, and improve service quality and accessibility [1,2]. These investments have fueled the evolution of e-Government, aiming to facilitate fully digital and user-centered public services that mirror the efficiency and personalization found in the private sector achieving 100% online public services by 2030 [3,4]. e-Government was envisioned not merely as a means

for digitizing paperwork but as a transformative shift towards proactive governance models [5,6].

1.1. Problem Statement

Public service provision still follows legacy processes, despite decades of digitization efforts and massive investments in e-Government infrastructures [7]. At the core of this issue is the ongoing reliance on document-centric procedures [8]. Legislations, administrative regulations, and eligibility preconditions are primarily encoded in human-readable legislative texts, while citizens are often required to repeatedly submit physical or scanned documents—such as birth certificates, income statements, or proof-of-residency—to demonstrate their qualification for services [9–11]. This document-centric approach not only introduces redundancy and inefficiency but also hinders interoperability and automated processing, as human interpretation is regularly needed to ascertain compliance with public service preconditions [12].

In recent years, the public sector has been exploring ways to adopt emerging technologies, like knowledge graphs (KGs) and artificial intelligence (AI), to tackle the limitations of document-based service provision. Among these innovations are recommender systems, which, while long established in the commercial realm, like e-commerce and media platforms, have not yet been widely applied to public service contexts [13]. The vision of a "No-Stop-Government" (NSG), where government services which are proactively triggered to life changes, offering citizens and businesses timely, relevant support without the need for explicit requests, closely aligns with the concept of recommendation engines that can match user profiles with appropriate services [14]. This shift to NSG offers numerous potential benefits, such as time savings, fewer errors, reduced paperwork, and improved government reputation. To enable such proactive and intelligent service delivery, public administrations must shift from a document-centric model to one that directly leverages data and formalized precondition. By converting public service preconditions into machine-readable format, and by constructing comprehensive citizen profiles grounded in reliable data sources, governments can harness recommender systems to offer personalized and context-aware public services. This approach promises not only to address long-standing efficiency and equity issues but also to open up new possibilities for service innovation and user engagement in the public sector.

Transforming public service provision into a data-centric model, where eligibility preconditions are expressed as formal, machine-readable rules and where "evidence" is sourced from authoritative data repositories, holds promise for mitigating these challenges [15]. However, achieving such a paradigm shift demands methodological, technological, and organizational changes. Without formalizing the preconditions and automating access to high-quality citizen data, public services remain limited to their document-based legacy [16]. In turn, this limits the ability of governments to deliver proactive, personalized, and truly intelligent services, capabilities that recommender systems, if properly integrated, could help unlock.

Previous research [12] explored the shift from document-centric to data-centric public service provision, emphasizing the transformative potential of emerging technologies like KGs and Large Language Models (LLMs). This foundational analysis proposed a conceptual architecture to enable granular data exchange by extracting service preconditions from legislative texts, standardizing evidence types, and integrating authoritative data sources in order to allow the automated execution of public services. This approach demonstrated the feasibility of utilizing LLMs to interpret complex legislative documents and translate them into machine-readable formats, while KGs facilitated the semantic interlinking of service attributes, preconditions, and evidence.

1.2. Objectives and Contributions

The overarching objective of this research is to transcend the document-centric legacy in public service provision by proposing a conceptual framework that leverages machine-readable preconditions, data-centric modeling, and intelligent recommendation capabilities. The objective of this research is to create an environment in which matching citizens to relevant public services can occur proactively, rather than reactively by structuring preconditions into formal knowledge representations, and by enabling automated access to authoritative data sources [10,17].

To achieve this objective, we draw on advancements in KG technologies, which capture and reason over structured information, and LLMs, which can assist in interpreting and extracting rules from complex, unstructured legislative texts. Integrating these techniques, often described as "neuro-symbolic AI", can deliver the best of both worlds: the semantic rigor of formal ontologies and the interpretive flexibility of state-of-the-art natural language processing (NLP) systems [18].

Despite rapid advancements, both KGs and LLMs display critical limitations that paves the way for their combination as neuro-symbolic AI [19] Although KGs can provide structured and explainable representation of knowledge, they are costly to construct and maintain due to manual curation and ontology design and prone to errors, often incomplete, and fail to interpret relationships from unstructured text. Furthermore, as KGs are relied on predefined ontologies, this limits the scalability across different domains. On the other hand, while LLMs are able to understand natural language, they suffer from hallucinations, limited explainability and factual inaccuracies. They fail to reason as the level of complexity increases and their probabilistic nature makes it difficult to explain their results. Therefore, by combining KGs and LLMs, we can minimize the issues posed in each technology, which is the main rationale for the neuro-symbolic AI framework proposed in this research paper.

The contributions of this work are threefold:

- Neuro-symbolic framework for public service recommendation: We propose an architecture that integrates KGs and LLMs, enabling the system to interpret preconditions, maintain semantic models of public services, and align user profiles with applicable services. This approach addresses a significant research gap by applying recommendation logic to public sector provision, where such systems are still underexplored [20].
- Prototype implementation and case study demonstration: We apply the prototype of key pipeline components to a public service use case, showcasing the feasibility of a neuro-symbolic approach to extract preconditions from regulations using LLMs, transform them into Resource Description Framework (RDF) evidence data models while using Shapes Constraint Language (SHACL) rules for validation.
- Pathway Towards NSG: Our work supports the realization of proactive, user-oriented government services through bridging the gap between data sources and service models. These recommender capabilities lay the groundwork for NSG scenarios, where individuals receive targeted, just-in-time assistance based on life events or business changes, without needing to explicitly request services or submit redundant documents.

While the proposed framework is using EU-based interoperability standards and references, it is not limited to this region. Respective standards could be reused or built for or developing governments as well. The source code of the implementation can be found on GitHub at https://github.com/ikonstas-ds/framework-llm-public-services-recsys/tree/main (accessed on 13 October 2025) using KNIME Analytics Platform 5.4.

2. Background and Related Work

2.1. Recommender Systems in the Industry

Recommender systems have undergone significant advancements, evolving from basic filtering methods to sophisticated algorithms tailored for diverse applications. Early systems, such as those used in e-commerce and entertainment, relied heavily on collaborative filtering and content-based approaches to infer user preferences from interaction data or item features [21,22]. These techniques formed the backbone of platforms like Amazon and Netflix, showcasing the effectiveness of personalization in enhancing user satisfaction [23].

The emergence of hybrid systems, which combine collaborative and content-based filtering, addressed limitations such as sparsity and cold-start problems. Hybrid methods like semi-supervised clustering using Gaussian mixture models [24], and demographic filtering have enhanced the accuracy and adaptability of recommendations across domains. A more contemporary exploration of knowledge-based recommenders and their alignment with user needs can be found in a recent survey which outlines advances in incorporating robustness and fairness into these systems [25].

Neuro-symbolic AI methods have showcased solid performance in other domains human–machine interaction and multimodal reasoning like assistive robotics and intelligent rehabilitation [26,27].

Recent work has expanded on the application of semantic web technologies, particularly their integration with LLMs for creating semantically rich representations [28]. This builds on earlier foundations while incorporating state-of-the-art techniques. These approaches align well with the structured and policy-driven requirements of public services, enabling recommender systems to navigate legal complexities and evolving regulations. For example, ontological reasoning has been used to create semantically rich, interoperable data models that provide transparent and explainable recommendations [29].

Advancements in natural language processing (NLP), particularly through LLMs, are revolutionizing the field. These models enable natural language understanding, contextual analysis, and explanation generation, enhancing user interaction and trust. LLMs have demonstrated potential in interpreting user-generated content, extracting domain knowledge from unstructured data, and enabling dynamic, precondition-based recommendations [25].

2.2. Recommender Systems in Public Administration

The evolution of public service provision has been closely tied to the development and maturation of digital government practices [30]. However, persistent bottlenecks remain, as governments still grapple with document-centric workflows that limit the potential for proactive and seamless service delivery. The concept of an NSG aligns closely with emerging e-government models, such as the transition from a one-stop shop to a no-stop shop, where citizens receive services without initiating requests or submitting forms. The no-stop government paradigm shifts the responsibility of service initiation from citizens to public organizations, offering services proactively without explicit citizen actions [31,32]. Countries such as Finland, Austria and Singapore have already piloted proactive services like automatic educational recommendations or child benefits and demonstrating how trustworthy AI in public sector can be achieved through accountability, transparency, and human supervision in proactive service provision [33–35]. Furthermore, another study examined how principles of trustworthy AI are applied in practice within a real-world e-Government context highlighting challenges and opportunities towards trustworthy AI in public administrations [36]. Despite these advancements, significant barriers remain, particularly in bridging the gap between the procedural, document-bound legacy

of public administration and the fluid, data-centric vision of "no-stop" or proactive governance [37], which is the focus of the present framework.

Addressing these issues entails moving toward models of digital government that valorize structured data, machine-readable policies, and interoperable platforms. In this regard, ongoing research and policy discourse underscore the necessity of integrating emerging technologies—such as semantic web standards, KGs, and AI-driven NLP—into the public sector toolkit. Previous research contributed to this discourse by proposing a conceptual framework that uses KGs to model service preconditions and link them to standardized evidence types [12].

2.3. Recommender Systems in Public Service Provision

Recommender systems have predominantly found applications in e-commerce, entertainment, and healthcare, but their role in public service provision remains a relatively underexplored domain [20]. Public service provision involves unique challenges due to the complexity, diversity, and intangibility of services, which distinguish it from other sectors. Despite limited research, several frameworks and solutions have demonstrated potential for enhancing the delivery of public services through recommendation systems.

One prominent application is in healthcare, where frameworks like those proposed at [38], focus on delivering personalized medical services. These systems analyze user data to match patients with appropriate healthcare services, illustrating the adaptability of recommender system principles to the public domain.

In social services, systems leveraging social data are particularly effective. As discussed in [39], social recommender systems that utilize social network insights to enhance decision-making processes. Such systems could be adapted to public services like welfare benefits and employment opportunities.

E-governance and smart city initiatives further exemplify the potential of recommender systems in public service provision. A study on recommender systems for government services in smart cities highlights their ability to address the overwhelming volume of information and services, offering personalized recommendations tailored to the needs of diverse stakeholders [40].

Despite these advances, the deployment of recommender systems in public service provision is not without challenges. Unlike commercial applications, public services may necessitate extensive personal information, emphasizing the need for clarity in data usage and collection processes. Additionally, citizens have the right to understand how their data is utilized and how recommendations are generated. Ensuring transparency in AI systems is crucial to maintain public trust and uphold legal and democratic rights. Clear communication about the system's operations, objectives, outputs, and impacts is essential, especially when these systems affect citizens' daily lives and rights [41]. Systems must account for the intangible and heterogeneous nature of services [42] while ensuring trustworthiness and transparency in recommendations [43]. Ethical considerations and data privacy concerns are particularly critical when handling sensitive user information in public sectors.

2.4. Formal Public Service Modeling

A fundamental barrier to achieving proactive, data-centric public services lies in the difficulty of translating legal and administrative rules—traditionally expressed in natural language documents—into machine-readable representations. As policies and eligibility requirements are typically embedded in dense legal texts, ensuring that these rules can be interpreted, evaluated, and enforced by computational systems demands structured formalization [44].

This effort has led to the introduction of standardized semantic vocabularies and ontologies intended to foster interoperability and facilitate the modeling of public services. For instance, the European Commission has developed the Core Public Service Vocabulary Application Profile (CPSV-AP) and the Core Criterion and Core Evidence (CCCEV) to systematically describe and link public services, the data they require, and their associated precondition [45,46]. By adopting these semantic standards, governments can represent services and eligibility preconditions as linked data, enabling systems to navigate the complexity of administrative processes and dynamically identify which services may apply to a given citizen or business profile.

Furthermore, knowledge representation approaches, such as ontologies and KGs, allow the formal specification of concepts, attributes, and relationships in a structured manner. KGs can capture both static descriptions (e.g., definitions of unemployment, income thresholds) and dynamic data elements (e.g., current employment status) linked to authoritative evidence sources [12]. When combined with reasoning engines, these structured representations enable the automatic evaluation of whether a particular user profile meets the necessary precondition for a service, removing the need for repeated document submission and manual interpretation.

The transition to machine-readable precondition also supports multi-stakeholder settings, where different government agencies must collaborate to deliver integrated services. Interoperability frameworks, semantic standards, and rule-based systems simplify data sharing and conditional checks, promoting integrated decision-making across institutions. This harmonization underpins more advanced functionalities, such as triggering proactive service recommendations when user data changes in a way that satisfies predefined conditions [33].

However, challenges remain. Ensuring the correctness, completeness, and currency of formalized precondition is non-trivial, as policies and regulations can evolve rapidly, requiring continuous maintenance of the underlying knowledge models. Likewise, guaranteeing privacy and adherence to data protection legislation adds layers of complexity, as these systems must handle sensitive user information responsibly and transparently [47]. Despite these difficulties, the convergence of semantic standards, ontological modeling, and automated reasoning holds significant promise for enabling a data-driven paradigm in public service provision, laying essential groundwork for integrating recommender systems that can operate atop this newly machine-readable layer of governmental rules and data.

3. Materials and Methods

3.1. Conceptual Framework

3.1.1. High-Level Vision

The central vision of the conceptual framework is to redefine public service provision by moving away from document-centric processes towards a data-driven, precondition-based paradigm. This shift involves formalizing legal rules and conditions into structured, machine-readable precondition and linking them to evidence data drawn from authoritative and interoperable sources, such as governmental registries, statistical databases, and trusted third-party providers.

To achieve this vision, the framework integrates knowledge representation methods with advanced AI techniques. It leverages semantic models and ontologies to represent public services, eligibility precondition, and user attributes, thus enabling automated reasoning over complex policies and multi-faceted user profiles. Complementing these symbolic approaches, LLMs can assist in extracting and interpreting rules from textual regulation documents, thereby continuously updating machine-readable preconditions as

regulations evolve [48]. This hybrid, neuro-symbolic AI strategy combines the logical rigor of structured knowledge with the flexible interpretive capabilities of cutting-edge language models.

The outcome is an intelligent, proactive recommendation environment, where users benefit from timely suggestions of services they may not even be aware they are entitled to. For instance, when a user's employment status changes to "unemployed", the system can immediately infer eligibility for job placement services without requiring the user to navigate complex administrative layers.

3.1.2. Core Components

The conceptual framework's high-level vision is operationalized through three core components that collectively enable proactive, data-driven recommendation of public services. These components—(1) Precondition Extraction and Formalization, (2) Evidence Modeling, and (3) Profile Construction—are integral for ensuring that a machine-readable, interoperable layer of rules and user attributes underpins the recommender system logic.

1. Precondition Extraction and Formalization

A critical challenge is the translation of legal texts, eligibility conditions, and policy rules—traditionally locked in lengthy documents—into structured, computable preconditions. To achieve this, the framework leverages approaches from the fields of intelligent document processing and semantic interoperability. Building upon research that advocates for representing regulations and processes in a machine-readable format, these preconditions are formalized using ontologies, vocabularies, and rule-based systems. By mapping natural language policy statements to formal models, the framework ensures that changes in regulations can be more easily integrated, and that eligibility checks become a matter of automated reasoning rather than manual interpretation.

2. Evidence Modeling

Once preconditions are formalized, the system must link them to authoritative data sources—"evidence"—that can confirm whether these preconditions are met. Evidence modeling involves defining data schemas and reference ontologies aligned with existing standards, such as the CPSV-AP, to represent the attributes and conditions necessary for service eligibility. By structuring evidence as linked data entities, drawn from trusted registries or other verified information hubs, the framework enables seamless and automated verification. Instead of citizens submitting repetitive documents, the system retrieves and interprets relevant data from known sources, ensuring consistency, reducing administrative burden, and improving both efficiency and accuracy.

3. Profile Construction

The final core component focuses on the creation and maintenance of rich, machine-readable user profiles for citizens and businesses. These profiles aggregate and integrate data points sourced from various public registers—such as demographic data, employment status, income level, or business attributes—into a structured, privacy-compliant representation. By aligning user attributes with semantic precondition defined for each service, the system can continuously update and refine eligibility determinations. When profile data changes (e.g., loss of employment, new academic enrollment, or registering a new business), triggers are activated to match the updated profile against existing precondition, thereby identifying which public services become newly relevant. In this manner, the framework fosters a proactive service environment where recommendations emerge naturally from ongoing alignment between evidence, precondition, and user context.

4. Public Service Recommendation

The Public Service Recommendation component is the culminating part of the conceptual framework, combining the outputs of Evidence Modeling and Profile Construction to generate personalized service recommendations. This component operationalizes the system's capability to proactively align users with relevant public services based on their eligibility and contextual changes. It performs the following key functions:

- Integration of Evidence and Profiles: It synthesizes the validated evidence data and dynamic user profiles to match users against preconditions defined for various public services.
- Service Matching: Utilizing semantic reasoning and matching algorithms, the component evaluates user eligibility for services, ranking them based on relevance and timeliness.
- Proactive Recommendations: As user profiles or policy preconditions evolve, the component dynamically updates recommendations, ensuring users are informed of newly available services without requiring manual intervention.
- Explainability and Justification: The system provides explanations for recommendations, detailing the eligibility criteria met by the user through the generated RDF SHACL rules. This fosters transparency and builds trust in the automated decision-making process.

The core components are illustrated in Figure 1.

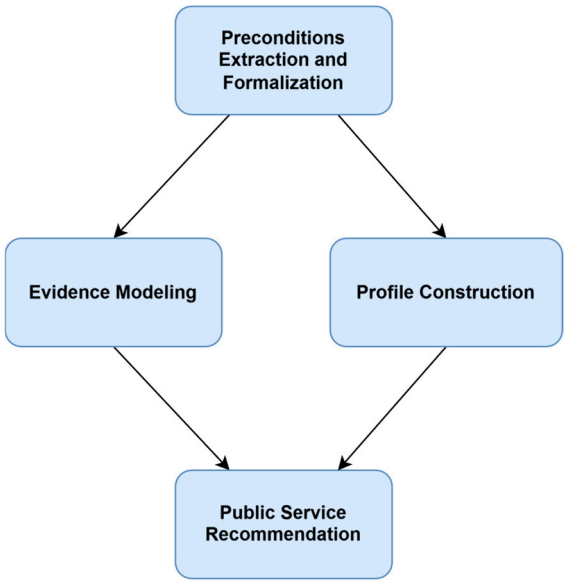


Figure 1. Core component process.

Realizing the full potential of the conceptual framework's core components demands the synergy of multiple enabling technologies. At the heart of these technologies lie KGs and LLMs, which, when integrated, can offer a complementary blend of structured semantic reasoning and flexible natural language understanding.

- *Knowledge Graphs (KGs)*: KGs provide a semantically rich, graph-based structure for representing and interlinking eligibility precondition, evidence sources, public service descriptions, and user attributes [49]. Their graph-based nature enables dynamic queries and inference over interconnected data, supporting real-time evaluation of conditions as user profiles change. By leveraging ontologies and vocabularies tailored to the public sector (e.g., CPSV-AP), KGs facilitate semantic interoperability

and integration with legacy systems, while also providing a flexible backbone for incorporating evolving rules and datasets [50]. As a result, KGs serve as a foundational infrastructure for reasoning about complex eligibility precondition, allowing the recommender system to rapidly identify relevant services.

- Large Language Models (LLMs): While KGs excel at handling structured data and explicit logic, LLMs offer powerful capabilities for understanding and extracting insights from unstructured or semi-structured content, such as policy documents, legal texts, and administrative guidelines [19]. State-of-the-art LLMs, trained on extensive corpora, can interpret regulations expressed in natural language and assist in mapping them to formal precondition, bridging the gap between human-authored rules and machine-readable representations. This is particularly beneficial when updating the system with new or revised regulations, as LLMs can help parse textual amendments, identify the relevant conditions, and suggest corresponding updates to the KG. By doing so, LLMs reduce the manual effort involved in maintaining and scaling the system, ensuring that the framework remains adaptive to changing policy landscapes.
- Neuro-Symbolic Integration: The combination of LLM-driven natural language understanding and KG-based symbolic reasoning constitutes a neuro-symbolic integration approach. KGs provide the structural rigor and interpretability needed for robust public-sector applications, while LLMs offer the flexibility and language comprehension skills necessary for handling messy or evolving regulations. This hybrid approach can support continuous improvement of the recommender system's knowledge base—new rules identified by LLMs can be validated, refined, and integrated into the KG through semi-automated workflows. Over time, the system becomes more adept at handling exceptions, rare conditions, and complex eligibility scenarios, all while maintaining an explainable and trustworthy decision-making process.
- Supporting Tools and Standards: Beyond KGs and LLMs, a range of supporting technologies and standards underpin the framework. Rule engines and inference tools allow for automated reasoning over semantic rules and policies. Data access and integration platforms ensure that evidence sources remain authoritative and up to date, while privacy-enhancing technologies and identity management solutions help safeguard sensitive user data. By assembling these tools into a coherent technology stack, the framework can confidently deliver proactive, precondition-driven recommendations in alignment with policy objectives and user rights.

3.2. Proposed Architecture and Methodology

3.2.1. Neuro-Symbolic AI Integration

The integration of neuro-symbolic AI paradigms within public service recommendation systems represents a transformative approach to achieving proactive and personalized governance. This integration unifies the structured factual knowledge of KGs with the interpretive and generative strengths of LLMs, creating a synergy that addresses their respective limitations [51]. While LLMs such as GPT-4o excel at understanding and generating natural language, their black-box nature and occasional inaccuracies in factual knowledge present challenges. Conversely, KGs explicitly store rich, structured knowledge but are inherently difficult to construct and adapt to new information. The neuro-symbolic approach leverages the complementary strengths of these technologies to deliver more robust and interpretable solutions for public service provision.

KGs provide the structured and semantically rich foundation needed to encode complex relationships between public services, eligibility criteria, required evidence, and citizen or business profiles. By formalizing domain-specific semantics through linked data

models and ontologies, such as the CPSV-AP, KGs enable interoperability across disparate governmental systems, automated reasoning to validate intricate eligibility conditions, and transparency through clear representations of relationships. These capabilities establish KGs as a logical and reusable backbone for public service recommendations, particularly in scenarios where adherence to formal definitions is crucial.

LLMs contribute unparalleled capabilities in processing and interpreting unstructured or semi-structured textual data. These models can automate the parsing of voluminous policy documents and legal texts to identify eligibility conditions, dynamically incorporate regulatory changes to reduce manual intervention, and translate human-authored legal language into machine-readable formats. Their ability to generalize across diverse contexts and tasks makes them indispensable for managing the complexity and variability of public sector documentation. However, their reliance on implicit knowledge often necessitates external augmentation from structured sources such as KGs to enhance inference accuracy and interpretability.

The combination of KGs and LLMs creates a neuro-symbolic framework that overcomes the individual limitations of each approach. KGs enhance LLMs by providing external, factual knowledge during pre-training and inference, improving the accuracy and reliability of generated outputs. Simultaneously, LLMs augment KGs by facilitating tasks such as KG embedding, completion, and the extraction of new facts from unstructured data. This bidirectional relationship enables synergized reasoning, where LLMs and KGs collaboratively inform and refine each other. For example, when a new policy mandates changes to eligibility conditions for a public service, an LLM can identify the updates, propose modifications to the KG, and ensure immediate alignment with user profiles through automated reasoning. The interplay of these technologies supports proactive recommendations and real-time compliance monitoring, achieving a balance between flexibility and rigor.

3.2.2. Prompting Strategy for Precondition Extraction and Evidence Modeling

In order to ensure a consistent extraction of public service preconditions from regulations, we utilized a prompting strategy with LLMs. The approach combined the following techniques:

- Retrieval-Augmented Generation (RAG): While the framework assumes that for each public service the related regulations are known, the texts are preprocessed using chunking and semantic search, ensuring that the model received only relevant sections of text such as a relevant article describing the preconditions. Appropriate context is given to the prompts for inclusion and exclusion criteria of correct results. For instance, the title of the public service is provided and the prompt guides the LLM to focus on the actual preconditions needed and not on evidence or certificates.
- Few-shot prompting: Representative examples of canonical evidence data models and correct triples are provided to guide the model toward more consistent results for structured triples.
- Prompt chaining: Multi-step prompts are used to improve the robustness of the results. This approach allows the hierarchical structuring of the evidence data models while also properly refining the output and minimizing hallucinations. Furthermore, it can improve the handling of more complex and nested preconditions.
- Language handling: Although the source regulations are in Greek, we provided prompts in English language due to the superior reasoning capabilities of LLMs in this language. At the same time, English language will facilitate further semantic interoperability and proper naming convention of the data model classes and properties.

- Rule Formalization and Validation: Together with prompt chaining and the derived evidence data model, last prompts are used to transform the preconditions into RDF SHACL shapes using self-evaluation and correction to further increase the consistency of the final outcome.

3.2.3. Standards and Data Models

The shift toward data-centric public service provision is grounded in the establishment of robust standards and data models that facilitate interoperability, automate validation processes, and enable service recommendations. A critical element of this transition is leveraging evidence data models, which serve as the foundation for validating preconditions and aligning user profiles with applicable public services. Standards such as the CPSV-AP and the CCCEV provide the structural framework for representing these data models in a machine-readable and interoperable format.

CPSV-AP structures public service descriptions, linking them with associated preconditions and criteria, thus ensuring interoperability across administrative domains. CCCEV complements this framework by establishing a connection between procedural requirements and evidence. By formalizing these requirements, CCCEV ensures that the data necessary to validate eligibility can be accurately represented and reused across systems. Together, these vocabularies enable the systematic representation of complex public service rules.

An important addition to this framework is the use of canonical evidence as a contextual guide for the LLM to generate accurate evidence data models in RDF KGs [52]. LLMs play a crucial role in this process by interpreting the modeling best practices for these canonical examples and generating RDF-based evidence models. Leveraging outcomes of existing research on canonical evidences, such as "Birth Evidence" or "Marriage Evidence", can serve as a basis derived from real-world e-Government initiatives and are application profiles of CPSV-AP and CCCEV standards. Such models can provide valuable insights into the structure and attributes that evidence data models should encompass. The use of canonical evidence data models is not intended as a fixed template applied on a specific public service. Instead, it serves to guide the LLM in understanding best practices for designing hierarchical evidence data models and subsequently allowing to properly model evidence data models of any domain.

Once the RDF evidence data models are generated, SHACL rules are formulated to perform validation. These rules ensure that the data satisfies the preconditions for a given public service [44]. For example, SHACL rules might verify that the attributes of an RDF-based evidence data model include required fields such as date of birth, in compliance with predefined criteria. These SHACL rules can be generated through the use of LLMs, which interpret policy documents and procedural requirements to construct validation logic.

In this framework, the interaction between LLMs and SHACL validation processes ensures that public service recommendations are both accurate and context-aware. The LLM interprets unstructured policy and regulatory texts to generate RDF data models, while SHACL rules validate these models to enforce compliance with service preconditions. This iterative process enhances the reliability of public service provision by reducing manual errors and aligning data models with regulatory standards.

Building on these standards, the framework ensures that extracted preconditions and evidences are not modeled in isolation as "data silos" but are semantically based on recognized interoperability assets, primarily CPSV-AP and CCCEV. Rather than defining new ontologies, the framework reuses and extends these standards through canonical evidence data models, which constitute as domain-specific application profiles of CPSV-AP, CCCEV, and related eGovernment Core Vocabularies. Reusing these standards in this

way, preserves semantic consistency across jurisdictions and domains in the public sector. and, subsequently, semantic interoperability allowing also data exchange among public administrations.

Figure 2 presents a simplified example of the KG stemmed from the housing allowance use case, indicating how different entities can link between citizen profile graph and public service nodes.

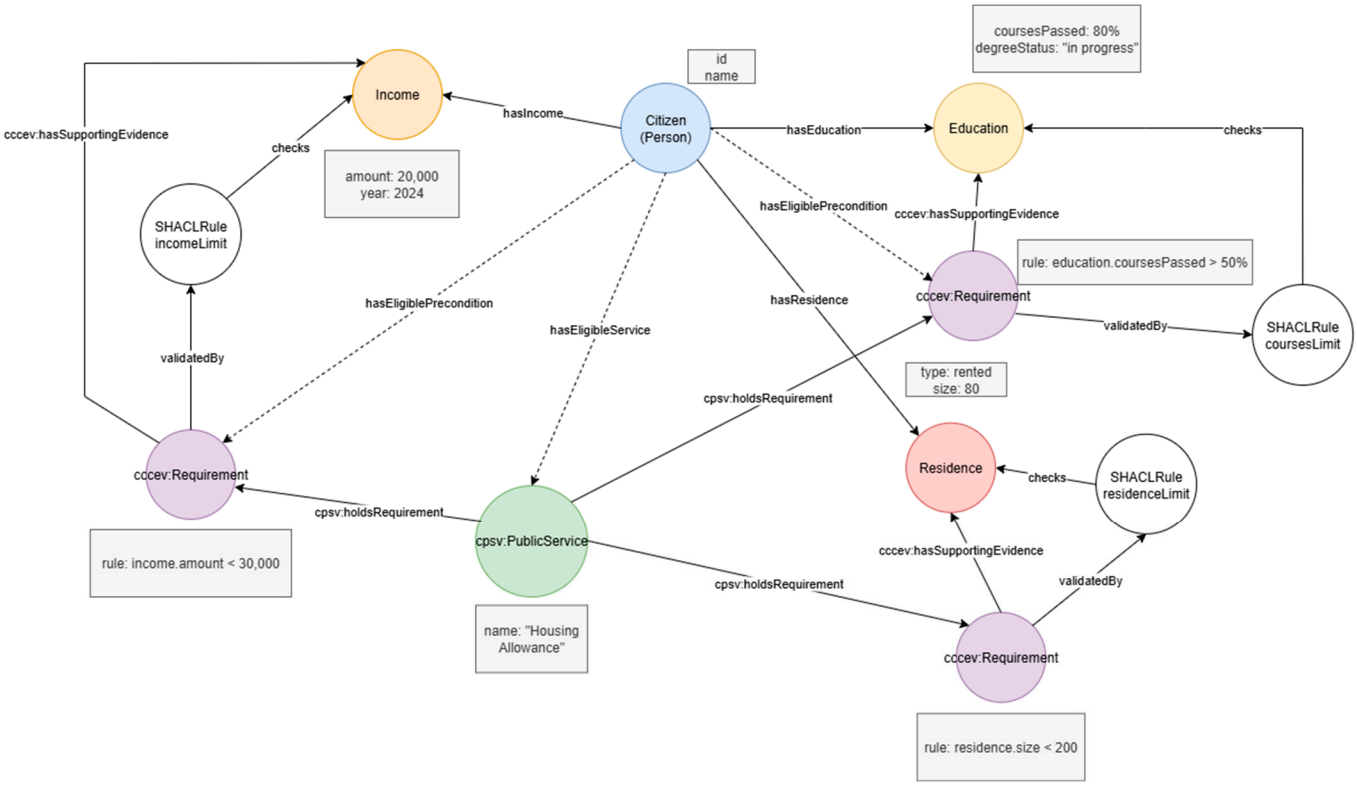


Figure 2. Simplified example of KG depicting the relationships between public service, preconditions (requirements), evidence and citizen profiles and the related SHACL rules.

3.2.4. System Workflow

The technical process of the proposed framework is split into two main phases:

- Phase 1: Create Graph of Public Services and Integrate Preconditions as SHACL shapes.
- Phase 2: Create Citizen Graph and Validate Preconditions.

Phase 1: Create Graph of Public Services and Integrate Preconditions as SHACL shapes

The purpose of this phase is to collect the preconditions from regulations and create the graph of public services that includes basic information of each public service and the extracted preconditions as SHACL shapes.

Step 1. Gather PDF Regulations for Public Services

Preconditions of public services are described within the regulations and policy documents of the related public service. These documents typically exist in unstructured PDF format that require intelligent document processing to extract the required information. The regulations serve as the primary source of information from which service eligibility precondition and evidence requirements are derived.

Step 2. Chunking for Large Texts

In many cases of extensive length of regulations, a chunking method can be employed to split the text into manageable segments. This is an important preprocessing step

that improves the accuracy and effectiveness of the LLM to extract the preconditions using a Retrieval-Augmented Generation (RAG) approach through a vector store to retrieve the most relevant sections having preconditions. To further improve the results, GraphRAG approaches can be employed that can provide further context through relationships indicating promising results in the retrieval phase [53].

Step 3. Extract Parts That Mention Preconditions

Having retrieved the relevant part containing preconditions within that can also fit in the context length of the LLM, we can utilize an LLM to identify and extract only the list of texts describing preconditions for each public service. The preconditions are still kept in their raw unstructured format, but any noise of text is minimized that can further improve the LLM accuracy in the next steps. For example, one precondition could be in its raw format "the applicant must be an adult".

Step 4. Extract Evidence Data Model from the Preconditions

Once the preconditions are extracted and cleaned, the next step is to extract the evidence data models associated with these preconditions. Such evidence data models represent the granular citizen data required to validate the precondition. Using the previously mentioned example of an applicant that must be an adult, the granular data needed in this case is the date of birth of the citizen. This extraction involves using sophisticated prompts to generate the final evidence data model in a structured format like RDF turtle which is a simplistic RDF serialization format that reduces any noise for the LLM as well. The model identifies the relevant evidence data from the preconditions and other data that are implicitly mentioned in the preconditions. Few-shot learning is also applied to improve accuracy of results, while prompt chaining techniques can also be employed to fix any inconsistencies in the LLM results and ensure higher quality of the evidence data models [54].

Step 5. Define Preconditions in RDF SHACL Rules

In this step, the system combines the extracted evidence data models and the retrieved preconditions to be used as context for the LLM. We then prompt the LLM to generate the related SHACL shapes that represent the preconditions as constraints in the derived RDF KG of public services and preconditions. Therefore, any life changes in actual citizen data will automatically pass through the constraint checking mechanism of the SHACL shapes and subsequently provide recommendations to citizens for possible eligibility on public services.

In many cases, eligibility preconditions can include nested or exception rules. To conceptually address this complexity, the framework's prompting strategy with prompt chaining can guide the LLM to generate hierarchical or nested SHACL rules. For instance, a rule may specify that a residence must not exceed 200 m^2, unless it is located in a municipality with fewer than 3000 inhabitants. Figure 3 demonstrates an example of nested SHACL rule.

```
ex:ResidenceShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:size ;
        sh:datatype xsd:integer ;
        sh:maxInclusive 200 ;
        sh:message "Residence size must not exceed 200 m²"
    ] ;
    sh:property [
        sh:path ex:location ;
        sh:node ex:LocationShape ;
    ] .

ex:LocationShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:population ;
        sh:datatype xsd:integer ;
        sh:maxInclusive 3000 ;
        sh:message "Exception: residences in areas with population < 3 000 are exempt"
    ] .
```

Figure 3. An example of nested RDF SHACL shapes for expressing complex preconditions.

Phase 2: Create Citizen Graph and Validate Preconditions

This phase is focused on retrieving the required evidence data of the citizen in order to validate preconditions and provide recommendations for possible eligibility on applying to public services.

Step 6. Use the Evidence Data Model to Retrieve Data for Each Citizen from Registries

The aim of this step is to create the KG of citizen data that are based on the extracted evidence data models, allowing the automated validation of preconditions. It is known that, in most cases, citizen data in each country are scattered across different base registries of public administrations and there is no central registry containing all the data [12]. Therefore, to be able to request the required data from the specific base registry storing the data, a centralized "registry of registries" needs to be established. This registry of registries will act as a data catalog that is a KG containing the metadata and the related data models of each base registry. We can execute queries in SPARQL, which is an RDF query language, against this data catalog to locate the relevant data sources for each type of citizen evidence required based on the evidence data models. Once the relevant registries are identified, the system sends requests to these registries to retrieve the specific data items needed for each citizen. To properly query the data, each base registry needs to securely expose a standard endpoint like a REST API or a SPARQL endpoint that will also ensure semantic interoperability between the different base registries. The retrieved data is then pushed into the citizen graph, MyGraph, a knowledge graph that consolidates citizen data from various registries. MyGraph serves as the unified data source against which eligibility rules will be validated.

Step 7. Validate preconditions with SHACL and provide recommendations

Once MyGraph is created, the RDF SHACL rules are automatically applied by checking whether the citizen data meet the preconditions defined for each public service. The rules automatically determine eligibility by comparing the citizen's data against the predefined precondition, effectively automating what would traditionally be a manual eligibility assessment process. After validation, the system generates recommendations for

public services that the citizen may be eligible for. These recommendations are personalized and proactive, delivered to citizens without the need for them to initiate the application process. This proactive service provision improves the accessibility and efficiency of public service delivery, reducing the administrative burden on both citizens and government agencies.

These steps of the proposed framework are illustrated in Figure 4.

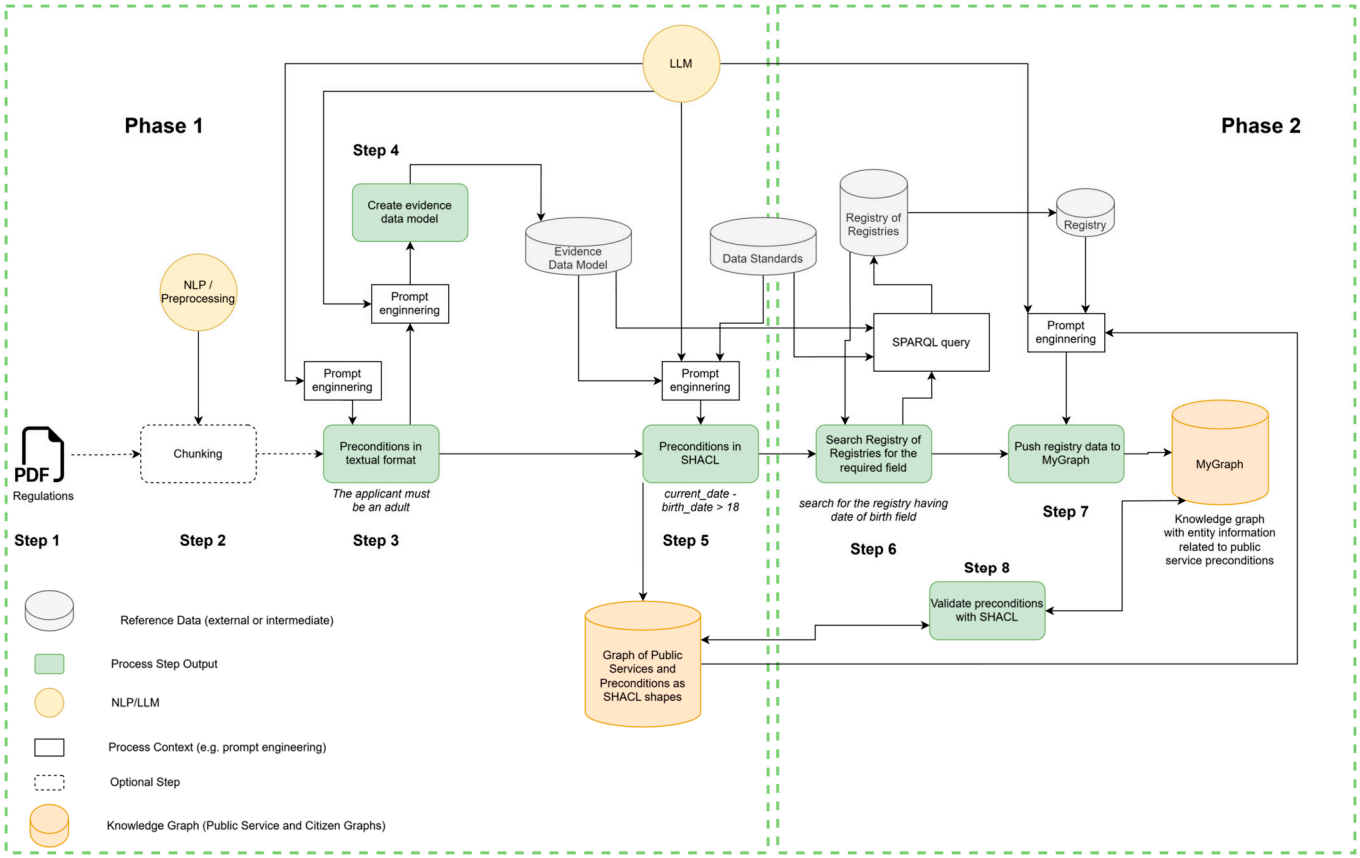


Figure 4. Technical process of the proposed framework.

3.2.5. Evaluation Plan

While the current framework focuses mainly on the conceptual design and implementation of the prototype, an evaluation plan is proposed to assess the framework's effectiveness in future extensions. Evaluation methods are based on standard practices in legislative text parsing and information extraction [48,51].

Accuracy of LLM Extraction: The extraction of public service preconditions from regulations will be assessed on annotations conducted by domain experts. Typical metrics like precision, recall and F1-score will evaluate the general accuracy and hallucinations of the extracted outcome provided by the LLM and follow established practices in recent LLM evaluation research [55–57]. Manual assessment performed in the current prototype phase already constitutes an initial qualitative baseline. Further comparison will be conducted between simple and complex nested eligibility preconditions for proper SHACL rules generation.

Prompt Optimization: Different prompting strategies and combinations will be compared in terms of performance and quality of results in order to optimize accuracy and consistency.

Baseline Comparison: Two reference methods will be compared with our framework:

(i) Direct LLM prompting and RAG on regulation text without integration of KG. This setup represents the standard semantic engine approach, where the citizen data is

used to query on the public service data formed as regulation texts. Already compared to the proposed framework, this approach assumes having access to the whole citizen profile data instead of only using the service-related data as mentioned in our framework following the data minimization principles in public service provision.

(ii) GraphRAG method that will form the legislative texts as KG and will use the whole citizen profile data to query the graph and recommend eligible public services [53].

Scalability and Transferability: Apart from housing allowance, further examples will be compared on different levels of complexity and domains like educational grants or unemployment benefits and on different sizes of data to evaluate the scalability of the framework.

4. Results: Use Case and Implementation

In this section, we evaluate the feasibility of the proposed framework, focusing on a specific use case of public service provision in Greece where a university student can apply for housing allowance. Each public service is captured into regulations entailing the list of preconditions needed for the applicant to be eligible to apply. The regulation for the service to apply for housing allowance includes a specific article (Article 3) that lists the terms and preconditions [58].

The reference data for the related public service is provided by the Greek governmental entity MITOS [59]. From there, we can retrieve the public service title and description, the name of the related regulation and subsequently call the API of the Greek National Printing House (NPH) that provides the related regulation of the public service in PDF format (step 1).

As the length of this regulation is short with approximately 7500 tokens, we skipped the chunking step that retrieves a related reduced segment of the text (step 2). For this case study, we utilize the LLM of OpenAI (San Francisco, United States), GPT-4o, which has shown state-of-the-art results in LLM leaderboards for information extraction [60]. We prompt the LLM to extract the list of all the preconditions required for a citizen to apply to this public service procedure. We use English language as it has better reasoning capabilities and performance in this language. We also prompt the model to focus on the actual preconditions needed and not on any evidence or certificates (step 3).

Then, in order to extract the evidence data model (step 4), we first provide the extracted list of preconditions and instruct the LLM to provide it in plain triples in the form of "concept1 > has_relationship > concept2" to reduce any noise and increase accuracy and we provide the following main guidelines:

- Use generalized entities (e.g., Person, Income, Residence, Education) and their direct relationships.
- Exclude any calculated or aggregated properties, focusing only on fundamental master data. It is crucial to consider evidence data in the most granular level, so that they can be aligned with the retrieved citizen data of MyGraph.
- Represent relationships in a way that allows for external calculations (e.g., family income derived from individual incomes via Person > has_parent > Person). This is important for the proper generation of the SHACL rules to validate preconditions.
- Group related concepts under broader entities. For example, if there are multiple concepts related to education (like the concept variables CourseResult, EducationProgram), they should be grouped under a single Education entity to reflect their hierarchical relationship.
- Avoid precondition-specific eligibility criteria in the evidence data model itself. The purpose of this instruction is to ensure that no eligibility criteria are integrated into the evidence data model.

Additionally, we use few-shot learning by providing examples of canonical evidence data models from existing research that are also based on CCCEV and CPSV-AP [53]. These include the Birth Canonical Evidence and Secondary Education Diploma Canonical Evidence models as examples for structuring classes, properties, and relationships. These examples can help the model to understand patterns on how to define the different classes for the required evidence data model. Finally, in this step, we apply prompt chaining to fix any inaccuracies of the data model and ensure a proper hierarchy and to generate the data model in RDF turtle format.

In the final step of Phase 1 (step 5), we combine the extracted preconditions and the evidence data model to prompt the LLM to generate the related SHACL rules that automatically validate the eligibility of the given citizen to the public service preconditions. Prompt chaining can also be applied to further improve the robustness of the results.

The results (Phase 1) of the different steps of the case study are presented in Table 1. The implementation has been conducted on KNIME Analytics Platform data science platform and the related code is available on GitHub [61].

Table 1. Example of processing steps for a specific public service regulation.

Step	Description	Outcome
Step 1	The full text of the regulation in PDF format	PDF of Greek Government Gazette Issue B 2993/2017—decision number 140832/Z1
Step 2	List of extracted preconditions in textual format	1. The annual family income of the previous year must not exceed thirty thousand (30,000) euros, increased by three thousand (3000) euros for each dependent child beyond the first. 2. The student must reside in a rented house due to their studies in a city other than their main residence, where neither they nor their parents have full ownership or usufruct of another residence. 3. The lease must be valid for at least six (6) months within the academic year for which the allowance is requested. 4. The parents of the student or the student themselves must not be owners or usufructuaries of residences (self-used or rented) exceeding two hundred (200) square meters in total, with exceptions for residences in municipalities or communities with a population of less than three thousand (3000) inhabitants. 5. The student must have successfully passed half of the courses of the previous academic year for which the allowance is requested. 6. The allowance is granted for as many years as the years of study of the respective school or department, according to their regulations, including the time during which they received the allowance while studying in another department or school. 7. The allowance is granted to all undergraduate students of a family, provided that all the above conditions are cumulatively met. 8. Students who are pursuing a second degree, regardless of the method of their admission to higher education, are not entitled to the housing allowance. 9. Students who have exceeded the duration of studies required for obtaining their degree, according to the indicative study program, are not entitled to the housing allowance. 10. Students residing in student dormitories or provided with accommodation by their school are not entitled to the housing allowance.
Step 4	Extracted evidence data model	The full evidence data model is presented in Appendix A.1
Step 5	SHACL rules that validate preconditions	The full SHACL rules are presented in Appendix A.2

While this case focuses on Greek student housing allowance, the framework is designed to be domain-agnostic. The same process of extracting preconditions, identifying the evidence data model and generating the RDF SHACL rules can apply across different domains like unemployment benefits and family allowance, where using prompting strategy and steps we can identify the underlined eligibility preconditions for a given citizen towards proactive public service recommendation. Even in this examined case of Greek housing allowance, the canonical models used were on a different domain for birth and marriage certificates.

5. Discussion

The fusion of KGs and LLMs within a neuro-symbolic AI framework lays the groundwork for transformative advancements in public service provision. By aligning the structured reasoning of KGs with the adaptive capabilities of LLMs, governments can deliver services that are accurate, personalized, and timely. This aligns with the vision of a NSG, where services anticipate and fulfill citizen needs proactively and equitably.

In addition, the framework leverages LLMs to extract complex eligibility rules from legal texts and transform them into structured, computable formats. This capability eliminates the need for manual interpretation of regulations and ensures that service preconditions remain accurate and up-to-date, even as policies evolve. To that end, the use of KGs in this framework ensures the automatic flexibility to changes in evidence data models and citizen data.

Building on these improvements, the framework enables a proactive approach to public service recommendations, aligning closely with the vision of NSG. By leveraging semantic reasoning and integrated user profiles, the system anticipates citizen needs based on life events or contextual changes, such as a change in employment status or family composition. This proactive functionality allows governments to recommend services in a timely and relevant manner, reducing the burden on citizens to identify and apply for services themselves.

Furthermore, while certain canonical evidences were used in this prototype implementation, the fact that these evidences are from a different domain (personal life events) than the domain tested (student housing allowance) strengthens the level of adaptation of the framework in other domains.

5.1. Limitations and Future Challenges

Despite its potential, the framework faces significant limitations and challenges that must be addressed for widespread adoption and success. The current framework has been validated by a prototype that presents the core components of precondition extraction, evidence data modeling and generation of SHACL rules to provide public service recommendations in a proactive manner. Although the pipeline is not yet end-to-end operational, the implemented use case and proposed architecture verify the feasibility of this neuro-symbolic method and form a strong basis for further developing and operationalizing an end-to-end public service recommendation engine.

- Data Fragmentation and Quality Issues: Public sector data is often dispersed across administrative silos, with inconsistencies in format, completeness, and accuracy. These issues pose significant challenges for interoperability and integration that need to be further addressed in Phase 2 of the framework. The reliance on high-quality, standardized data also limits the framework's applicability in regions where digital infrastructure and data governance are less developed.
- Legal and Policy Complexities: Eligibility criteria and service regulations not only vary significantly across jurisdictions but also often involve deep and complex data

models that are embedded within dense legal texts. This complexity poses challenges for scalability and necessitates continuous updates to accommodate changes in regulations. The dynamic nature of policies, combined with their intricate structures, further complicates the integration of these updates into existing systems. One solution could be the representation of regulations as KGs that can address these challenges more effectively within the framework. KGs provide a flexible and structured representation of legal rules and eligibility conditions, enabling systems to dynamically adapt to policy changes. This approach streamlines the process of incorporating new or revised regulations while ensuring consistency and scalability across jurisdictions. Although the proposed framework can theoretically express nested and exception logic through SHACL rules, the current implementation has been limited to less complex rules. Future work will focus on developing full support for more nested dependencies and exception handling.

- Inaccuracies in AI Systems: Employing technologies like LLMs in public administration can introduce both technical and governance risks. While LLMs provide advanced capabilities for interpreting legislative texts, their outputs are not always precise, that can lead to hallucinations or biases. Errors in parsing or formalizing preconditions could lead to incorrect recommendations, impacting citizen trust and service delivery. The "black-box" nature of LLMs also raises concerns about explainability, particularly in high-stakes applications.
- Privacy and Trust Concerns: The framework's reliance on sensitive citizen data introduces challenges related to privacy, trust, and compliance with regulations such as GDPR [47]. Ensuring robust data protection while maintaining transparency in how personal data is used and stored remains a critical concern. Without public trust in these systems, adoption may face significant resistance.
- Data Governance: As the data is retrieved from the different base registries to a central KG, this could lead to data governance issues like data ownership, data sharing and reuse. Further analysis can be conducted on the use of federated KGs where each data-owning authority keeps its own KG.
- Changes in Regulations: Modifications and updates in regulations have not been addressed in the current framework. Further analysis needs to be conducted on applying NLP and LLMs to store the different versions of regulations over time and the dependencies between them.

Nevertheless, the framework helps at tackling certain risks and concerns. As the public service preconditions are expressed in the form of SHACL rules in the KG, the recommendations are grounded in an explainable and human-readable logic. This approach allows having humans-in-the-loop, where providing explainable recommendations through the generated SHACL rules both introduces less risks and accountability issues while the citizens are encouraged to carefully verify the accuracy of AI-generated recommendations. Furthermore, as the citizen information stored in the KG is minimized according to the public service precondition requirements, the framework adheres to data minimization principles.

5.2. Future Directions

The limitations identified provide a basis for future research and development. Academic scholars and practitioners can explore several paths to build upon the proposed framework and address the challenges it presents.

One critical area of focus is improving data integration and interoperability across fragmented public sector datasets. Future research should investigate the development of federated data models and KGs that respect the autonomy of individual registries while enabling seamless collaboration across agencies. Additionally, standardized semantic

vocabularies or documentations of APIs are essential for harmonizing data formats, ensuring consistent evidence modeling, and enabling cross-agency interoperability. Exploring decentralized technologies, such as blockchain [62], could provide secure and tamper-proof mechanisms for data sharing, further strengthening the trust and reliability of public sector data integration.

Another promising direction lies in advancing the interpretation of legislative texts, a key component of the framework. LLMs hold immense potential for extracting eligibility criteria from complex legislative documents, but their domain-specific accuracy can be enhanced through fine-tuning on curated datasets of administrative texts. Combining LLMs with neuro-symbolic approaches, where structured reasoning capabilities of KGs complement the interpretive flexibility of LLMs, could further improve the reliability of these systems. Research into benchmarking methods, for assessing and validating AI-driven interpretations against human-expert evaluations would also contribute to their refinement.

Establishing a secure infrastructure for obtaining citizen graphs from base registries is a critical challenge, particularly in Phase 2 of the framework. Ensuring robust communication channels between registries and the system requires implementing advanced encryption protocols, secure APIs, and resilient authentication mechanisms to safeguard sensitive citizen data from potential breaches or misuse. Future research should explore the design and deployment of such secure infrastructure, emphasizing interoperability and compliance with data protection regulations. Additionally, the feasibility of integrating open-source LLM solutions that can be self-hosted within the business owner's environment merits further investigation. Self-hosted LLMs provide a way to enhance data privacy and security by eliminating dependencies on external, third-party resources, while also offering greater control over the models' behavior and updates. Research should evaluate the performance, scalability, and cost-effectiveness of these solutions, alongside their compatibility with the broader framework.

The framework's applicability can be broadened by expanding its use cases and testing it in diverse contexts. Exploring its implementation in sectors such as healthcare, disaster management, and education would demonstrate its versatility and scalability. Furthermore, cross-border service provision presents a unique opportunity to evaluate the framework's ability to operate within regional frameworks like the European Union, where interoperability between countries is paramount.

Finally, the ethical and inclusive deployment of AI in public service systems should be a priority. Establishing ethical guidelines that ensure fairness, accountability, and inclusivity in AI-driven decision-making is essential for fostering equitable governance. Future research should examine potential risks, such as algorithmic bias or systemic inequities, and propose strategies to mitigate these challenges. Additionally, adaptive and resilient system designs capable of dynamically incorporating policy updates and responding to changes in citizen profiles are critical for ensuring the long-term sustainability and relevance of the framework.

6. Conclusions

This research paper presented a framework for transforming public service provision into a data-centric model, leveraging the synergy of KGs and LLMs. By formalizing service preconditions and evidence into machine-readable formats and dynamically aligning them with citizen profiles and evidence data models, the framework addresses inefficiencies inherent in document-centric workflows. This approach enables proactive, personalized public service recommendations, aligned with the vision of an NSG.

KGs and LLMs display critical limitations if considered separately. In many cases, KGs require manual construction and curation and are prone to errors, often incomplete,

and fail to interpret relationships from unstructured text. On the other hand, LLMs can suffer from hallucinations, limited explainability and factual inaccuracies. They fail to reason as the level of complexity increases and their probabilistic nature makes it difficult to explain their results.

The integration of neuro-symbolic AI combines the structured reasoning capabilities of KGs with the interpretive flexibility of LLMs, ensuring both transparency and adaptability to evolving legal and policy landscapes. Key contributions include the formalization of eligibility preconditions, the creation of dynamic citizen graphs, and the automated validation of service preconditions to provide recommendations, all of which improve accessibility and reduce administrative burdens. Overall, this research paper proposes a framework with a prototype that automates extraction of public service preconditions and recommendations of eligibility, showcasing technical feasibility and providing a solid path towards a full-scale implementation.

While challenges such as data fragmentation, legal complexities, and privacy concerns remain, the proposed framework offers a compelling foundation for future research and practical application, paving the way for more efficient, equitable, and citizen-focused public service delivery.

Author Contributions: Conceptualization, I.K., V.P. and I.M.; methodology, I.K., V.P. and I.M.; software, I.K.; validation, I.K., V.P. and I.M.; formal analysis, I.K., V.P.; investigation, I.K.; resources, I.K.; data curation, I.K.; writing—original draft preparation, I.K.; writing—review and editing, V.P., I.M. and I.K.; visualization, I.K.; supervision, V.P.; project administration, V.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data are partly available within the manuscript. Source code can be found at https://github.com/ikonstas-ds/framework-llm-public-services-recsys/tree/main (accessed on 13 October 2025).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
CCCEV	Core Criterion and Core Evidence
CPSV-AP	Core Public Service Vocabulary Application Profile
GenAI	Generative AI
IT	Information Technology
KG	Knowledge Graph
LLM	Large Language Models
NLP	Natural Language Processing
NSG	No-Stop-Government
RDF	Resource Description Framework
SHACL	Shapes Constraint Language

Appendix A

Appendix A.1

Full evidence data model for the MITOS use case (Section 3).

```
@prefix : <http://example.org/schema#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

:Person a rdfs:Class .
:Income a rdfs:Class .
:Residence a rdfs:Class .
:Education a rdfs:Class .
:Accommodation a rdfs:Class .
:Location a rdfs:Class .

:has_income a rdf:Property ;
    rdfs:domain :Person ;
    rdfs:range :Income .

:has_property a rdf:Property .

:year a rdf:Property ;
    rdfs:domain :Income ;
    rdfs:range rdfs:Literal .

:amount a rdf:Property ;
    rdfs:domain :Income ;
    rdfs:range rdfs:Literal .

:type a rdf:Property ;
    rdfs:domain :Income, :Residence, :Accommodation, :Location ;
    rdfs:range rdfs:Literal .

:has_residence a rdf:Property ;
    rdfs:domain :Person ;
    rdfs:range :Residence .

:location a rdf:Property ;
    rdfs:domain :Residence ;
    rdfs:range rdfs:Literal .

:size a rdf:Property ;
    rdfs:domain :Residence ;
    rdfs:range rdfs:Literal .

:lease_duration a rdf:Property ;
    rdfs:domain :Residence ;
    rdfs:range rdfs:Literal .

:has_relationship a rdf:Property ;
    rdfs:domain :Person ;
    rdfs:range :Person .

:role a rdf:Property ;
    rdfs:domain :Person ;
```

```
    rdfs:range rdfs:Literal .

:has_education a rdf:Property ;
    rdfs:domain :Student ;
    rdfs:range :Education .

:institution a rdf:Property ;
    rdfs:domain :Education ;
    rdfs:range rdfs:Literal .

:program a rdf:Property ;
    rdfs:domain :Education ;
    rdfs:range rdfs:Literal .

:courses_passed a rdf:Property ;
    rdfs:domain :Education ;
    rdfs:range rdfs:Literal .

:study_duration a rdf:Property ;
    rdfs:domain :Education ;
    rdfs:range rdfs:Literal .

:degree_status a rdf:Property ;
    rdfs:domain :Education ;
    rdfs:range rdfs:Literal .

:has_accommodation a rdf:Property ;
    rdfs:domain :Student ;
    rdfs:range :Accommodation .

:population a rdf:Property ;
    rdfs:domain :Location ;
    rdfs:range rdfs:Literal .
```

Appendix A.2

Full SHACL rules for the MITOS use case (Section 3) expressing the preconditions eligibility recommendations.

```
@prefix sh: <http://www.w3.org/ns/shacl#> .
@prefix ex: <http://example.org/schema#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

ex:PersonShape
    a sh:NodeShape ;
    sh:targetClass ex:Person ;
    sh:property [
        sh:path ex:has_income ;
        sh:node ex:IncomeShape ;
    ] ;
    sh:property [
        sh:path ex:has_residence ;
```

```
        sh:node ex:ResidenceShape ;
    ] ;
    sh:property [
        sh:path ex:has_relationship ;
        sh:node ex:RelationshipShape ;
    ] ;
    sh:property [
        sh:path ex:has_education ;
        sh:node ex:EducationShape ;
    ] ;
    sh:property [
        sh:path ex:has_accommodation ;
        sh:node ex:AccommodationShape ;
    ] .

ex:IncomeShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:year ;
        sh:datatype xsd:gYear ;
    ] ;
    sh:property [
        sh:path ex:amount ;
        sh:datatype xsd:decimal ;
        sh:maxInclusive 30000 ;
    ] ;
    sh:property [
        sh:path ex:type ;
        sh:in ("salary" "business" "investment") ;
    ] .

ex:ResidenceShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:location ;
        sh:node ex:LocationShape ;
    ] ;
    sh:property [
        sh:path ex:size ;
        sh:datatype xsd:integer ;
        sh:maxInclusive 200 ;
    ] ;
    sh:property [
        sh:path ex:lease_duration ;
        sh:datatype xsd:duration ;
        sh:minInclusive "P6M"^^xsd:duration ;
    ] ;
    sh:property [
        sh:path ex:type ;
        sh:in ("rented") ;
    ] .
```

```
ex:RelationshipShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:role ;
        sh:in ("parent" "spouse" "child") ;
    ] .

ex:EducationShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:institution ;
        sh:datatype xsd:string ;
    ] ;
    sh:property [
        sh:path ex:program ;
        sh:datatype xsd:string ;
    ] ;
    sh:property [
        sh:path ex:courses_passed ;
        sh:datatype xsd:integer ;
        sh:minInclusive 0.5 ;
    ] ;
    sh:property [
        sh:path ex:study_duration ;
        sh:datatype xsd:duration ;
    ] ;
    sh:property [
        sh:path ex:degree_status ;
        sh:in ("completed" "in-progress") ;
    ] .

ex:AccommodationShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:type ;
        sh:in ("apartment" "house") ;
    ] .

ex:LocationShape
    a sh:NodeShape ;
    sh:property [
        sh:path ex:population ;
        sh:datatype xsd:integer ;
        sh:maxInclusive 3000 ;
    ] .
```

References

1. Dunleavy, P.; Margetts, H.; Bastow, S.; Tinkler, J. *Digital Era Governance: IT Corporations, the State, and e-Government*; Oxford University Press: Oxford, UK, 2006. https://doi.org/10.1093/acprof:oso/9780199296194.001.0001.

2. Nations, U. United Nations E-Government Survey 2020. Available online: https://www.un.org/en/desa/2020-united-nations-e-government-survey (accessed on 13 October 2025).
3. European Commisson. eGovernment Benchmark 2023|Shaping Europe's Digital Future 2023. Available online: https://digital-strategy.ec.europa.eu/en/library/egovernment-benchmark-2023 (accessed on 13 October 2025).
4. European Commission. Decisions (EU) 2022/2481 of the European Parliament and of the Council of 14 December 2022 establishing the Digital Decade Policy Programme 2030. *Off. J. Eur. Union* **2022**, *2022*, 4–26.
5. Janowski, T. Digital government evolution: From transformation to contextualization. *Gov. Inf. Q.* **2015**, *32*, 221–236. https://doi.org/10.1016/j.giq.2015.07.001.
6. Gil-Garcia, J.R.; Dawes, S.S.; Pardo, T.A. Digital government and public management research: Finding the crossroads. *Public Manag. Rev.* **2018**, *20*, 633–646. https://doi.org/10.1080/14719037.2017.1327181.
7. Mureddu, F.; Osimo, D.; Kenny, A.; Upson, M.; Peristeras, V. Fostering a Data-Centric Public Administration: Strategies, Policy Models and Technologies. *Public Adm. Inf. Technol.* **2022**, *38*, 217–244. https://doi.org/10.1007/978-3-030-92945-9_9.
8. Shaheen, R.; Kasi, M. Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies, a Case of USA. *Eur. J. Technol.* **2021**, *5*, 1–15. https://doi.org/10.47672/ejt.641.
9. Susha, I.; Janssen, M.; Verhulst, S. Data collaboratives as "bazaars"?: A review of coordination problems and mechanisms to match demand for data with supply. *Transform. Gov. People, Process Policy* **2017**, *11*, 157–172. https://doi.org/10.1108/TG-01-2017-0007.
10. Wimmer, M.A. *Once-Only Principle Good Practices in Europe BT—The Once-Only Principle: The TOOP Project*; Krimmer, R., Prentza, A., Mamrot, S., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 61–82, ISBN 978-3-030-79851-2.
11. Wimmer, M.A.; Tambouris, E.; Krimmer, R.; Gil-Garcia, J.R.; Chatfield, A.T. Once only principle: Benefits, barriers & next steps. In Proceedings of the 18th Annual International Conference on Digital Government Research, Staten Island, NY, USA, 7–9 June 2017; Part F1282; pp. 602–603. https://doi.org/10.1145/3085228.3085296.
12. Konstantinidis, I.; Kapantai, E.; Michailidis, A.; Deligiannis, A.; Berberidis, C.; Magnisalis, I.; Peristeras, V. From document-centric to data-centric public service provision. *Digit. Gov. Res. Pract.* **2024**, *5*, 1–27. https://doi.org/10.1145/3676279.
13. Anthopoulos, L.G. *Understanding Smart Cities: A Tool for Smart Government or an Industrial Trick?*; Springer: Berlin/Heidelberg, Germany, 2017; Volume 22.
14. McBride, K.; Hammerschmid, G.; Lume, H.; Raieste, A. Proactive Public Services: The new standard for digital governments. In *Proactive Public Services-the New Standard for Digital Governments, White Paper*; Nortal Hertie School of Digital Governance: Berlin, Germany, 2023; pp. 1–44.
15. Loutas, N.; Peristeras, V.; Tarabanis, K. The Public Service ontology: A formal model for describing domain-specific semantics. *Int. J. Metadata, Semant. Ontol.* **2011**, *6*, 23–34. https://doi.org/10.1504/IJMSO.2011.042462.
16. Publications Office of the European Union. Shaping Europe's Digital Future. **2020**. https://doi.org/10.2759/48191.
17. European Commission. Regulation (EU) 2018/1724 of the European Parliament and of the Council of 2 October 2018 Establishing a Single Digital Gateway to Provide Access to Information, to Procedures and to Assistance and Problem-Solving Services and Amending Regulation (EU) No. *Off. J. Eur. Union* **2018**, *2018*, 1–38.
18. Wan, Z.; Liu, C.-K.; Yang, H.; Li, C.; You, H.; Fu, Y.; Wan, C.; Krishna, T.; Lin, Y.; Raychowdhury, A. Towards Cognitive AI Systems: A Survey and Prospective on Neuro-Symbolic AI. *arXiv* **2024**, arXiv:2401.01040. Available online: http://arxiv.org/abs/2401.01040 (accessed on 13 October 2025).
19. Schneider, P.; Schopf, T.; Vladika, J.; Galkin, M.; Simperl, E.; Matthes, F. A Decade of Knowledge Graphs in Natural Language Processing: A Survey. In Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing, Virtual, 20–23 November 2022; Long Papers; He, Y., Ji, H., Li, S., Liu, Y., Chang, C.-H., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2022; Volume 1, pp. 601–614. Available online: https://aclanthology.org/2022.aacl-main.46/ (accessed on 13 October 2025).
20. El Gharbaoui, O.; El Boukhari, H.; Salmi, A. The transformative power of recommender systems in enhancing citizens' satisfaction: Evidence from the Moroccan public sector. *Innov. Mark.* **2024**, *20*, 224.
21. Adomavicius, G.; Tuzhilin, A. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Trans. Knowl. Data Eng.* **2005**, *17*, 734–749. https://doi.org/10.1109/TKDE.2005.99.
22. Ricci, F.; Rokach, L.; Shapira, B. Introduction to recommender systems handbook. In *Recommender Systems Handbook*; Springer: Berlin/Heidelberg, Germany, 2010; pp. 1–35.
23. Kaur, B.; Sharma, R.; Rani, S.; Gupta, D. Recommender system: Towards classification of human intentions in e-shopping using machine learning. *J. Comput. Theor. Nanosci.* **2019**, *16*, 4280–4285.

24. Zhang, Y.; Liu, X.; Liu, W.; Zhu, C. Hybrid Recommender System Using Semi-supervised Clustering Based on Gaussian Mixture Model. In Proceedings of the 2016 International Conference on Cyberworlds (CW), Chongqing, China, 28–30 September 2016; pp. 155–158.
25. Li, Y.; Liu, K.; Satapathy, R.; Wang, S.; Cambria, E. Recent Developments in Recommender Systems: A Survey. *Comp. Intell. Mag.* **2024**, *19*, 78–95. https://doi.org/10.1109/MCI.2024.3363984.
26. Xing, Z.; Ma, G.; Wang, L.; Yang, L.; Guo, X.; Chen, S. Toward Visual Interaction: Hand Segmentation by Combining 3-D Graph Deep Learning and Laser Point Cloud for Intelligent Rehabilitation. *IEEE Internet Things J.* **2025**, *12*, 21328–21338. https://doi.org/10.1109/JIOT.2025.3546874.
27. Xing, Z.; Meng, Z.; Zheng, G.; Ma, G.; Yang, L.; Guo, X.; Tan, L.; Jiang, Y.; Wu, H. Intelligent rehabilitation in an aging population: Empowering human-machine interaction for hand function rehabilitation through 3D deep learning and point cloud. *Front. Comput. Neurosci.* **2025**, *19*, 1543643. https://doi.org/10.3389/fncom.2025.1543643.
28. Teh, S.-A.; Haw, S.-C.; Santoso, H.A. Hybrid-based Research Article Recommender System. *Int. J. Membr. Sci. Technol.* **2023**, *10*, 1587–1606. https://doi.org/10.15379/ijmst.v10i2.1830.
29. Confalonieri, R.; Guizzardi, G. On the multiple roles of ontologies in explanations for neuro-symbolic AI. *Neurosymbolic Artif. Intell.* **2025**, *1*, NAI-240754.
30. Cordella, A.; Bonina, C.M. A public value perspective for ICT enabled public sector reforms: A theoretical reflection. *Gov. Inf. Q.* **2012**, *29*, 512–520. https://doi.org/10.1016/j.giq.2012.03.004.
31. Lindgren, I.; Scholta, H. Untangling the Relationship Between Public Service Automation and No-Stop Government. In Proceedings of the Electronic Government: 22nd IFIP WG 8.5 International Conference, EGOV 2023, Budapest, Hungary, 5–7 September 2023; Proceedings; Springer-Verlag: Berlin, Heidelberg, 2023; pp. 83–94.
32. Scholta, H.; Mertens, W.; Kowalkiewicz, M.; Becker, J. From one-stop shop to no-stop shop: An e-government stage model. *Gov. Inf. Q.* **2019**, *36*, 11–26. https://doi.org/10.1016/j.giq.2018.11.010.
33. Dan, S.; Kalliokoski, J.; Shahzad, K. Trust in public sector AI: The case of Finland's Aurora national program. In *Trust and Artificial Intelligence*; Routledge: London, UK, 2024; pp. 242–257.
34. Innes, K.; Salter, P. Creating a Digital State: From a One-Stop Shop to a No-Stop Shop. CapX. 2021. Available online: https://capx.co/creating-a-digital-state-from-a-one-stop-shop-to-a-no-stop-shop (accessed on 31 August 2025).
35. Bund Länder Städte Gemeinden e government; Finanzen, B.; Austria, D. Austria's e-Government Strategy 2023. 2023. Available online: https://www.digitalaustria.gv.at/dam/jcr:39390f05-12d8-4aec-a4e0-1c8c5b580918/Austrias%20E-Government%20strategy-17072023_EN.pdf (accessed on 13 October 2025)
36. Berman, A.; de Fine Licht, K.; Carlsson, V. Trustworthy AI in the public sector: An empirical analysis of a Swedish labor market decision-support system. *Technol. Soc.* **2024**, *76*, 102471. https://doi.org/10.1016/j.techsoc.2024.102471.
37. Guintchev, P.; Joosten, J.J.; Fernández, S.S.; Adamson, E.S.; Sánchez, A.S.; Heredia, M.S. Specification languages for computational laws versus basic legal principles. *arXiv* **2025**, arXiv:2503.09129. Available online: http://arxiv.org/abs/2503.09129 (accessed on 13 October 2025).
38. Ooi, K.-N.; Haw, S.-C.; Ng, K.-W. A Healthcare Recommender System Framework. *Int. J. Adv. Sci. Eng. Inf. Technol.* **2023**, *13*, 2282–2293. https://doi.org/10.18517/ijaseit.v13i6.19049.
39. Oubalahcen, H.; El Ouadghiri, M.D. Recommender Systems for Social Networks: A Short Review. In Proceedings of the 6th International Conference on Networking, Intelligent Systems\& Security, Larache, Morocco, 24–26 May 2023; Association for Computing Machinery: New York, NY, USA, 2023. https://doi.org/10.1145/3607720.3607725.
40. Lin, R. Recommender System for Government Service in Smart Cities. In Proceedings of the 2022 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), Niagara Falls, ON, Canada, 17–20 November 2022; pp. 771–774.
41. Challenges of Recommender Systems in the Public Sector. Suomi.fi. 2023. Available online: https://kehittajille.suomi.fi/guides/responsible-ai/define-a-data-policy/challenges-of-recommender-systems-in-the-public-sector (accessed on 31 August 2025).
42. Baltaci, I. *Service Recommendation: Recommendation Systems for Service*s; AV Akademikerverlag: London, UK, 2015.
43. Bouzekri, E.; Canny, A.; Fayollas, C.; Martinie, C.; Palanque, P.; Barboni, E.; Deleris, Y.; Gris, C. A list of pre-requisites to make recommender systems deployable in critical context. In Proceedings of the 2nd Workshop on Engineering Computer-Human Interaction in Recommender Systems, co-located with the 9th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EnCHIReS@ EICS 2017), Lisbon, Portugal, 26 June 2017; Volume 1945, pp. 42–55.

44. Heimsbakk, V.; Torkelsen, K. Using the Shapes Constraint Language for modelling regulatory requirements. *arXiv* **2023**, arXiv:2309.02723. Available online: https://arxiv.org/abs/2309.02723 (accessed on 13 October 2025).
45. Group, C.P.S.V.W. Core Public Service Vocabulary Application Profile (CPSV-AP). Interoperable Europe Portal—European Commission. 2024. Available online: https://ec.europa.eu/isa2/solutions/core-public-service-vocabulary-application-profile-cpsv-ap_en/ (accessed on 13 October 2025).
46. Group, C.P.S.V.W. Core Criterion and Core Evidence Vocabulary (CCCEV). Interoperable Europe Portal—European Commission. 2024. Available online: https://interoperable-europe.ec.europa.eu/collection/semic-support-centre/solution/e-government-core-vocabularies/core-criterion-and-core-evidence-vocabulary (accessed on 13 October 2025).
47. Voigt, P.; dem Bussche, A. *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed.; Springer: Cham, Switzerland, 2017; Volume 10, pp. 10–5555.
48. Holzenberger, N.; Van Durme, B. Connecting Symbolic Statutory Reasoning with Legal Information Extraction. In Proceedings of the Natural Legal Language Processing Workshop, Singapore, 7 December 2023; Preoțiuc-Pietro, D., Goanta, C., Chalkidis, I., Barrett, L., Spanakis, G., Aletras, N., Eds.; Association for Computational Linguistics: Singapore, 2023; pp. 113–131. Available online: https://aclanthology.org/2023.nllp-1.12/ (accessed on 13 October 2025).
49. Hogan, A.; Blomqvist, E.; Cochez, M.; D'Amato, C.; de Melo, G.; Gutierrez, C.; Kirrane, S.; Gayo, J.E.L.; Navigli, R.; Neumaier, S.; et al. Knowledge Graphs. *ACM Comput. Surv.* **2021**, *54*, 1–37. https://doi.org/10.1145/3447772.
50. Promikyridis, R.; Tambouris, E. Using Knowledge Graphs to provide public service information. In Proceedings of the 23rd Annual International Conference on Digital Government Research, Virtual, 15–17 June 2022; pp. 252–259. https://doi.org/10.1145/3543434.3543585.
51. Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; Wu, X. Unifying Large Language Models and Knowledge Graphs: A Roadmap. *IEEE Trans. Knowl. Data Eng.* **2024**, *36*, 3580–3599. https://doi.org/10.1109/TKDE.2024.3352100.
52. Karunarathne, T.; Kontopoulos, E.; Konstantinidis, I.; Guzmán Carbonell, A.R. A Canonical Evidence-based Approach for Semantic Interoperability in Cross-border and Cross-domain e-Government Services. In Proceedings of the 15th International Conference on Theory and Practice of Electronic Governance, Guimarães, Portugal, 4–7 October 2022; Association for Computing Machinery: New York, NY, USA, 2022; pp. 131–139. https://doi.org/10.1145/3560107.3560299.
53. Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitansky, D.; Ness, R.O.; Larson, J. From local to global: A graph rag approach to query-focused summarization. *arXiv* **2024**, arXiv:2404.16130.
54. Wu, T.; Terry, M.; Cai, C.J. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April–5 May 2022; Association for Computing Machinery: New York, NY, USA, 2022. https://doi.org/10.1145/3491102.3517582.
55. Bang, Y.; Cahyawijaya, S.; Lee, N.; Dai, W.; Su, D.; Wilie, B.; Lovenia, H.; Ji, Z.; Yu, T.; Chung, W.; et al. A Multitask, Multilingual, Multimodal Evaluation of Chat GPT on Reasoning, Hallucination, and Interactivity. In Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, Bali, Indonesia, 1–4 November 2023; Long Papers; Park, J.C., Arase, Y., Hu, B., Lu, W., Wijaya, D., Purwarianti, A., Krisnadhi, A.A., Eds.; Association for Computational Linguistics: Bali, Indonesia, 2023; Volume 1, pp. 675–718. Available online: https://aclanthology.org/2023.ijcnlp-main.45/ (accessed on 13 October 2025).
56. Liu, X.; Yu, H.; Zhang, H.; Xu, Y.; Lei, X.; Lai, H.; Gu, Y.; Ding, H.; Men, K.; Yang, K.; et al. AgentBench: Evaluating LLMs as Agents. *arXiv* **2023**, arXiv:2308.03688. Available online: https://arxiv.org/abs/2308.03688 (accessed on 13 October 2025).
57. Lv, X.; Lin, Y.; Cao, Y.; Hou, L.; Li, J.; Liu, Z.; Li, P.; Zhou, J. Do Pre-trained Models Benefit Knowledge Graph Completion? A Reliable Evaluation and a Reasonable Approach. In Proceedings of the Findings of the Association for Computational Linguistics, ACL, Dublin, Ireland, 22–27 May 2022; Muresan, S., Nakov, P., Villavicencio, A., Eds.; Association for Computational Linguistics: Dublin, Ireland, 2022; pp. 3570–3581. Available online: https://aclanthology.org/2022.findings-acl.282/ (accessed on 13 October 2025).
58. Greek Government Gazette, Decision No. 140832/Z1; *Determination of the Procedure and Supporting Documents for the Granting of the Housing Allowance to Students of Higher Education Institutions*; Hellenic Republic Ministry of Education, Research and Religious Affairs: Athens, Greece, 2017.
59. MITOS 2023. National Registry of Administrative Public Services. 2023. Available online: https://en.mitos.gov.gr/index.php/Main_Page (accessed on 13 October 2025).
60. Şahinuç, F.; Tran, T.T.; Grishina, Y.; Hou, Y.; Chen, B.; Gurevych, I. Efficient Performance Tracking: Leveraging Large Language Models for Automated Construction of Scientific Leaderboards. In Proceedings of the 2024 Conference on Empirical Methods

in Natural Language Processing, Miami, FL, USA, 12–16 November 2024; pp. 7963–7977. https://doi.org/10.18653/v1/2024.emnlp-main.453.

61. Konstantinidis, I. Code for "A Framework for a Recommender System on Data-Centric Public Service Provision Using Large Language Models" 2025. GitHub Repository. 2025. Available online: https://github.com/ikonstas-ds/framework-llm-public-services-recsys (accessed on 13 October 2025).
62. Konstantinidis, I.; Siaminos, G.; Timplalexis, C.; Zervas, P.; Peristeras, V.; Decker, S. *Blockchain for Business Applications: A Systematic Literature Review BT—Business Information Systems*; Abramowicz, W., Paschke, A., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 384–399.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.