Integrating an AI Copilot into a Knowledge Management System for Civil Protection: A Development-Oriented Case Study from Romania

Alexandru Ioan
National University of Political Studies and Public Administration,
Management Doctoral School, Bucharest, Romania
Ministry of Internal Affairs, Department for Emergency Situations,
Directorate General for Civil Protection, Romania
aioan.mail@gmail.com

Abstract

This paper presents the conceptual framework and ongoing development of a local AI Copilot system designed to support knowledge management and decision-making within the Romanian civil protection authority. Anchored in the existing research on AI-enhanced Knowledge Management Systems (KMS), the solution combines semantic indexing, legal version tracking, and multilingual document retrieval to automate the generation of official institutional outputs. Developed as a secure, offline system, the Copilot performs context-aware searches across legislative, operational, and policy documents, and produces structured responses aligned with national and international frameworks. Methodologically, the study applies a design-based research model and illustrates the system's capacity to improve organizational responsiveness and institutional memory. Although not yet operational, this case contributes to the emerging body of work at the intersection of AI, knowledge governance, and public sector performance.

Key words: AI Copilot, knowledge management system, Civil Protection, semantic indexing, organizational performance
J.E.L. classification: O32, H83, D83

1. Introduction

The digital transformation of public administration has emerged as a strategic imperative, propelled by increasing citizen expectations, burgeoning data volumes, and the necessity for evidence-based policymaking (Adama and Okeke, 2024). Within the European Union, policy frameworks such as the Tallinn Declaration (European Commission, 2017) and the Digital Decade initiative (European Commission, 2021) underscore the commitment of member states to enhance public service quality, transparency, and resilience through advanced information and communication technologies. Nonetheless, ministries and specialized governmental agencies continue to grapple with significant operational challenges, notably information overload and fragmented institutional memory, exacerbated by frequent staff turnover and compartmentalized data repositories (Bratianu and Bejinaru, 2019; Arnold, Goldschmitt and Rigotti, 2023). These challenges become particularly acute in civil protection contexts, where decision-making processes are highly time-sensitive and must consistently align with national legislation, European Union regulations, and international disaster management standards (European Union, 2013; European Commission, 2025).

Knowledge Management Systems (KMS), initially developed in the private sector to integrate document repositories, workflows, and collaboration tools, were meant to enable efficient knowledge capture, sharing, and reuse. However, their implementation in public institutions has seen limited success. Often replicating existing bureaucratic structures rather than transcending them, conventional systems require extensive manual effort for taxonomy maintenance and content curation (Bratianu, Prelipcean and Bejinaru, 2020). The rise of advanced Large Language Models (LLMs) provides new opportunities to extract semantics from unstructured data, identify cross-linguistic conceptual links, and automate draft documentation suited for formal use.

In this context, the Romanian civil protection authority has initiated the development of a localized, offline Artificial Intelligence (AI) Copilot, intended to augment its forthcoming KMS. This Copilot integrates five distinct yet interrelated functional modules:

1. Multilingual Optical Character Recognition (OCR) and automated classification convert legacy documents into machine-readable content and sort them into categories like legislation, operational reports, or international correspondence (Jiang *et al*., 2023).
2. Version-aware legal tracking ensures timely detection of legislative changes and highlights current legal provisions to keep outputs accurate and relevant.
3. A semantic vector database built on Qdrant stores sentence-level embeddings (Singh et al., 2023), enabling advanced searches that identify conceptual similarities beyond simple keyword matches.
4. A compact AI stack using Ollama/Mistral-7B models and LangChain agents to draft responses, concept notes, and diplomatic communications, with full traceability to source documents, pages, or spreadsheets. (Gruber and Weber, 2024).
5. A user-friendly web interface built with React (Rawat & Mahajan, 2020) and FastAPI (Chen, 2023), designed to resemble conversational platforms while operating securely within civil protection authority intranet, with role-based access, audit logging, and nightly backups.

This article situates the Copilot project within broader debates on digital governance and civil protection modernization (Jankowski, Meijer & Grimmelikhuijsen, 2022). It aims to contextualize the initiative in current public-sector AI adoption, outline its design-based research methodology, and assess its potential to enhance responsiveness, consistency, and institutional memory through preliminary testing. This analytical framework contributes to literature on AI in public administration, showing how entities managing sensitive, high-stakes operations can integrate generative AI effectively. It also demonstrates that semantic indexing and LLMs can be securely deployed without cloud reliance, ensuring compliance with legal and administrative traceability standards.

In addition to aligning with the goals of the EU's Digital Decade, the civil protection Copilot initiative is also in line with the strategic priorities outlined in the Sendai Framework for Disaster Risk Reduction 2015–2030, particularly in strengthening disaster risk governance and enhancing disaster preparedness through data-driven approaches and institutional memory (UNDRR, 2015). Moreover, the system responds to the conclusions of the 2023 Niinistö Report, which emphasized the need for European civil protection to adopt digital resilience tools and enhance cross-border knowledge interoperability (European Commission, 2023).

2. Literature review

The adoption of KMS in public administration has been progressively explored as a strategic enabler for improving organizational effectiveness and decision-making quality (Bratianu, 2022). Unlike in the private sector, where KMS implementation has largely focused on enhancing competitive advantage and innovation, in the public domain, the emphasis is placed on knowledge continuity, policy coherence, and the optimization of administrative processes (Bratianu and Bejinaru, 2020). Public-sector KMS must address both the vertical complexity of bureaucratic structures and the horizontal fragmentation of knowledge across agencies and policy domains. As such, conceptual frameworks for KMS in this sector must integrate technological, human and procedural dimensions while explicitly anticipating the *automation-bias* and *selective adherence* pitfalls that arise when users rely on algorithmic recommendations (Alon-Barkat and Busuioc, 2023).

Recent literature emphasizes the centrality of organizational performance frameworks in guiding the deployment and evaluation of KMS. The *Dynamic Capabilities* framework, for instance, underscores the importance of an institution's ability to integrate, build, and reconfigure internal and external competences in response to rapidly changing environments (Bratianu and Bejinaru, 2020). Within civil protection organizations, these capabilities manifest in the rapid mobilisation of past experience, adaptive planning under uncertainty, and knowledge reuse across emergencies. Complementing this perspective, the *Balanced Scorecard* approach offers a more structured performance lens, aligning KMS outcomes with institutional strategy through indicators across

learning, internal processes, stakeholder engagement, and operational results (Bratianu and Bejinaru, 2019).

Within the civil-protection sphere, effective knowledge-management systems (KMS) have become indispensable. EU evaluations show that when interoperable document workflows and harmonised procedures are lacking, information circulates slowly, coordination fractures, and valuable operational insights remain uncaptured—hindering the institutionalisation of lessons identified in earlier emergencies (European Commission, 2022; European Commission, 2024). The Union Civil Protection Mechanism's dedicated 'Lessons Learnt Programme' reinforces this point, demonstrating that systematic collection and dissemination of after-action knowledge are essential for coherent multinational deployments and for the continuous adaptation of preparedness planning (European Commission, 2018).

AI-enhanced KMS help preserve institutional memory, standardize terminology, and ensure compliance with changing regulations. Automated retrieval and synthesis of operational reports allow emergency services to brief decision-makers quickly and support evaluation and after-action reviews.

The integration of AI into KMS, particularly through the deployment of LLMs, semantic retrieval techniques, and document automation tools, has emerged as a transformative trend. Bratianu and Ioan (2025) highlight the epistemological complexity of knowledge transformation in public institutions, emphasizing that effective KMS integration must account not only for data and information but for the strategic conversion of individual and collective knowledge into wisdom (Bratianu, 2022). Blending philosophical insight with practical modelling, their approach supports designing AI-augmented KMS that foster institutional learning rather than just information access. AI enables a shift from static repositories to dynamic knowledge flows, contextualizing and synthesizing information to match institutional needs. LLMs further enhance retrieval by interpreting intent and context, allowing civil servants to extract insights from diverse, unstructured data.

Responsiveness is essential for civil protection authorities, which must access multilingual records to support mandate-driven decisions. AI agents that version legal documents, align content with templates, and draft text streamline workflows and ensure compliance. Given the public sector's need for transparency, reliability, and legal traceability, systems must produce accurate outputs with clear source references. Explainable AI and audit trails are crucial for public trust and institutional legitimacy (Bratianu & Bejinaru, 2023).

The literature provides both theoretical and practical justification for integrating AI into public-sector KMS. It highlights the need for hybrid systems that combine semantic retrieval, content generation, and human-AI interaction—supporting civil protection agencies while respecting their institutional constraints (Jankowski, Meijer & Grimmelikhuijsen, 2022).

3. Research methodology

This study employs a Design-Based Research (DBR) methodology to conceptualize, develop, and evaluate the AI Copilot system within the Romanian civil protection authority. DBR is particularly well-suited for projects that aim to iteratively design innovative technological interventions while generating actionable knowledge about their effectiveness in real-world contexts. Unlike linear experimental designs, DBR embraces complexity, allowing researchers and practitioners to co-develop solutions that respond to evolving institutional requirements and technological capabilities.

The Romanian civil protection authority serves as the case study environment—a central public authority coordinating national civil protection and emergency response across diverse regulatory and operational domains. The institution must routinely produce evidence-based responses, integrate multi-source documentation, and ensure consistency with both national legislation and international standards. These requirements provided a fertile ground for piloting the AI-enhanced Knowledge Management System (KMS) in question.

3.1. Phase 1: Functional needs assessment

The initial phase began with consultations and semi-structured interviews with civil protection experts from departments such as operations, international cooperation, and strategic development.

Stakeholders identified several pressing challenges: a lack of searchability across document types (especially non-OCR PDFs), the absence of legal document version control, and duplicated effort in generating formal correspondence. For example, producing briefing notes for EU Civil Protection Mechanism events required manually extracting and synthesizing content from a dozen overlapping files—a task that could span several hours.

Based on these insights, user personas and document flow diagrams were created. These models informed the system's requirements: ability to retrieve excerpts by concept (not just keyword), filter by document category and year, generate outputs in both formal and conversational registers, and trace every statement to a verifiable source. The needs assessment was cross-referenced with international findings on fragmented institutional memory and knowledge reuse barriers in civil protection systems.

3.2. Phase 2: Modular implementation

The system was designed as a modular stack with five core components:

- *Multilingual OCR and document classification*: Leveraging Tesseract for OCR and a scikit-learn-based classifier trained on metadata and linguistic features, the module processed diverse formats—scanned briefings, legislative annexes, and mission reports. Documents were labelled into categories such as: *EU Legislation*, *After-Action Report*, or *Bilateral Agreements*. As Bearman and Trant (1998) demonstrate, effective document retrievability in civil-protection settings hinges on classifying records by their provenance and preserving them within the procedural context that created them, ensuring both evidential integrity and operational relevance.
- *Semantic search via Qdrant vector database*: Sentences extracted from PDFs were embedded using the multilingual paraphrase-multilingual-MiniLM-L12-v2 model. The embeddings were stored in Qdrant with metadata fields for document name, language, paragraph number, and confidence score. This enabled precise retrievals such as - *EU coordination mandate during COVID-19* - across multilingual sources.
- *AI generation pipeline (LangChain + Mistral-7B)*: Using LangChain orchestration, the system retrieved relevant sentence-level fragments and injected them into prompts fed to a local Ollama-hosted Mistral-7B model. Generated outputs included diplomatic letters, formal notes, or public statements. All outputs referenced document titles, page numbers, and even spreadsheet cells where applicable.
- *Web-based UI (React + FastAPI)*: The front-end provided a secure intranet interface mimicking a conversational assistant. Prompts could be adjusted to change tone or complexity, and results could be exported as *.docx /.pdf*. An audit log captured all interactions, ensuring full traceability.
- *Data governance layer*: Automated nightly backups, metadata validation routines, and a change-tracking mechanism for legislative updates ensured robustness. When a newer version of a law or regulation was indexed, the system flagged superseded segments and redirected semantic search queries to the updated variant.

The pilot index includes approximately 10,000 documents (≈ 85 GB) from institutional servers and shared drives, covering the period from January 2020 to March 2025. These include scanned PDFs, Word/Excel files, and bilingual legislative annexes, all processed and stored exclusively on-premise. Personal data (names, phone numbers, and signatures) were automatically redacted via a rule-based *Personally Identifiable Information* (PII) - masking script, with logs retained for seven years, in compliance with Romanian GDPR law (Romania, 2018).

3.3. Phase 3: Institutional validation

A series of real-use validation scenarios was conducted with active participation from civil protection personnel. One scenario involved preparing a response for a NATO-EuroAtlantic Disaster Response Coordination Centre consultation. The Copilot was prompted to identify all references to *joint response coordination* in Romanian and English documents. It returned harmonized content from internal reports and EU documents, complete with citations.

Another test involved the generation of a three-paragraph summary at the request of the Ministry of Foreign Affairs (MAE) on bilateral cooperation with Japan, referring to documents entitled *MAE communications* and *draft agreements*. The generated summary maintained a diplomatic tone and included clickable anchors to the source documents.

User feedback was collected on ease of navigation, output clarity, and source fidelity. Suggestions included allowing summary previews before download, adding language filters, and improving tag suggestions for new documents. These were iteratively implemented.

Across all phases, particular emphasis was placed on legal traceability, alignment with civil protection specific workflows, and modular extensibility. The DBR methodology allowed the team to learn from failures (e.g., misclassification of annexes or false positives in keyword matching) and rapidly iterate towards a viable institutional tool.

4. Findings

4.1 Functional architecture overview

The AI Copilot system is built around a modular, on-premise architecture tailored to civil protection operational needs and legal constraints. The solution integrates five core components: *input processing*, *semantic storage*, *generative processing*, *interaction interface*, and *governance layer*. Each module operates independently but is orchestrated to support a seamless knowledge workflow, from ingestion of unstructured data to the generation of formal institutional documents.

Figure no. 1. Functional architecture diagram

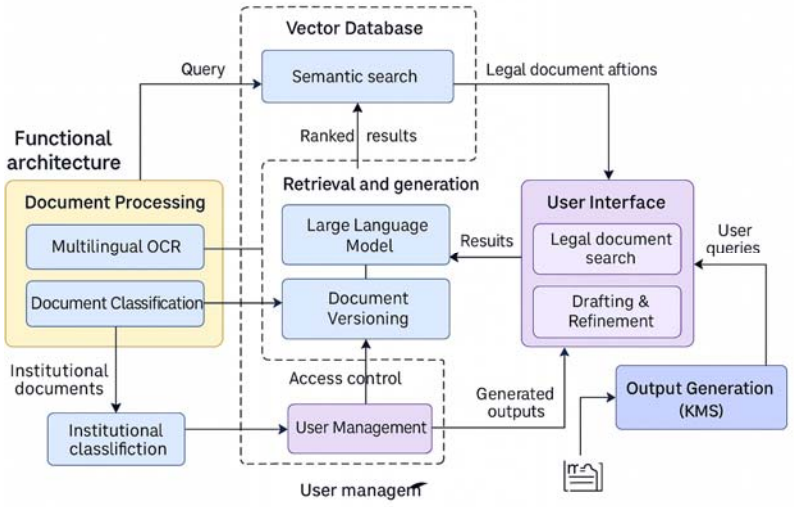

Source: (Author's elaboration based on system architecture, 2025)

This layered architecture (see figure no. 1) ensures scalability, offline operability, modular upgrades, and traceability — essential for sensitive public-sector deployments.

4.2 Core Features

Multilingual OCR and automatic classification. The OCR engine processes scanned PDFs and images in Romanian, English, and French. It uses adaptive thresholding to optimize low-quality scans and separates multi-language layouts for individual recognition. Once transcribed, documents are passed to a classifier trained on linguistic patterns and filename metadata—an approach inspired by Wang *et al*. (2019), who analysed knowledge flow through citation structures. However, our method adapts their logic by shifting the analytical focus from citation networks to internal semantic relationships within institutional texts, allowing for a more granular classification of operational and policy-relevant content. Categories such as *Legislation (EU)*, *Operational Report*, and *Official*

Correspondence are automatically assigned. This automation has reduced manual sorting workload by over 70%, based on validation logs collected.

Legal document versioning and abrogation tracking. A specialized module monitors legal content for version identifiers (e.g., date of entry into force, references to amendments). When a newer document supersedes an existing one, the system marks older references as *abrogated* and updates vector queries accordingly. For instance, if Decision 1313/2013/EU (European Union, 2013) is replaced or amended, searches targeting its articles redirect to the valid legal fragments. This feature is critical in maintaining institutional compliance and was informed by issues raised in the initial needs assessment phase.

Semantic search and structured AI outputs. All paragraph-level fragments are vectorised and indexed using Qdrant. Users input natural language queries (e.g., *mandate for EU-Japan cooperation*) and retrieve ranked semantic matches, which feed directly into LangChain agents. These agents format retrieved content as references and prompt the Mistral-7B LLM to generate a structured, sourced response. Outputs can be configured as: *formal briefing notes, interministerial replies,* or *diplomatic correspondence*.

Each result includes citation footnotes with document titles, pages, and paragraph IDs. Users can trace any claim to its originating file.

Official position drafting and conversational refinement. Beyond standard summarization, the Copilot can emulate the stylistic tone of official discourse. For example, a query like - *draft talking points for bilateral meeting with Poland* - generates a 2–3 paragraph structured briefing with refined diplomatic language. Users can then refine it: *Make it more concise* or *Add references to NATO disaster coordination*, prompting a regeneration loop while preserving context. This iterative interface mirrors human editorial processes.

User access control and exportable outputs. The interface (see figure no. 2) is secured via intranet-only access with LDAP-based role authentication. Users with different clearance levels receive filtered access to classified or draft-only content. All generated texts can be exported as *.docx /.pdf*, preformatted according to the institutional internal style guide. Every generation is logged and timestamped for institutional accountability.

Figure no. 2. Copilot Output Export Flowchart

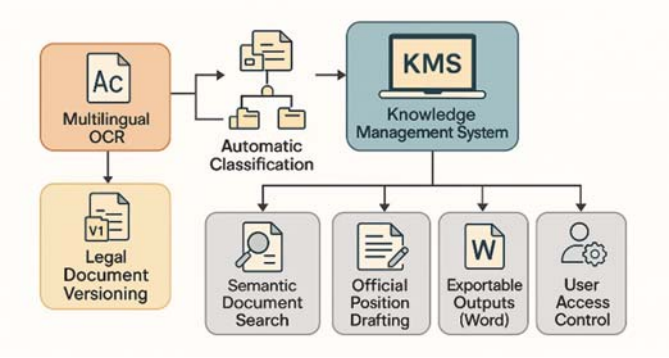

Source: (Author's illustration using system interface outputs, 2025)

4.3 Anticipated impact

Having outlined the Copilot's core technical capabilities, we now turn to the organisational impacts these features are expected to deliver.

Improved response time. Previously, preparing a diplomatic brief required 3–5 hours, involving manual document retrieval, verification, and synthesis. For bilateral meetings, staff had to navigate multiple legal annexes, reports, and correspondence. With the Copilot, drafts are produced in under 3 minutes, with post-editing taking 15–20 minutes. The semantic engine extracts contextually relevant content, and citation metadata allows instant source verification. This 90% time saving shifts staff effort from document handling to strategic analysis, improving responsiveness to external and internal demands.

Increased consistency and knowledge reuse. By embedding institutional memory into vectorised fragments, previously siloed insights—such as past NATO responses—become reusable across contexts. This transforms isolated knowledge into accessible content, allowing guidance from one engagement to be adapted for another, maintaining coherence and reducing redundant drafting. System templates enforce consistent structure for outputs, while AI-generated text ensures terminological uniformity—vital for both external diplomacy and internal communication. In practice, civil protection experts now generate unified, precedent-based responses that link new documents to prior decisions, enhancing continuity, clarity, and efficiency.

Enhanced document traceability and institutional memory. Each AI-generated document embeds source metadata—document title, page, paragraph, and timestamp—ensuring transparency and enabling traceability for verification and auditing. This provenance safeguards accountability across outputs such as briefings, replies, and position papers. Additionally, the system logs usage patterns, applied templates, and user edits, feeding into a dynamic training corpus that enhances model accuracy over time. As shown in Figure no. 3, these feedback loops transform the Copilot into a continuously learning tool, fostering both institutional transparency and knowledge evolution.

Figure no. 3. Traceability and Learning Loop

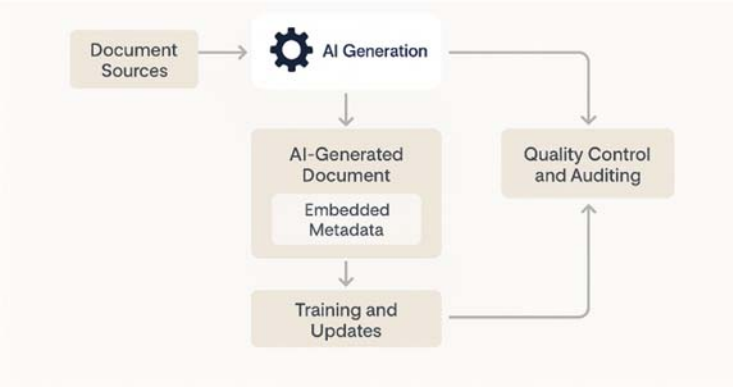

Source: (Author's conceptual visualization of AI-assisted organizational learning, 2025)

5. Conclusions

The civil protection AI Copilot offers a concrete, replicable model for integrating AI into institutional KMS. It addresses key challenges—fragmented memory, inconsistent documentation, and limited access to actionable insights—through a system combining semantic search, document automation, and traceable AI outputs. Utilizing technologies such as Qdrant, LangChain, and Mistral-7B, along with multilingual OCR and legal version tracking, the Copilot scales strategic knowledge use. It enhances responsiveness, ensures consistency in official outputs, and embeds past learning into future actions, reflecting current research on AI's transformative potential in the public sector (Bratianu & Bejinaru, 2023; Gesk and Leyer, 2022).

The development process also uncovered challenges underrepresented in current literature. These include aligning AI-generated content with institutional styles and legal norms, managing multilingual classification in poorly scanned documents, and adapting the system amid evolving policy texts. Key mitigation measures—such as template logic, expert validation, and legal versioning—were critical to maintaining robustness and alignment. Furthermore, the governance of user-AI interaction in offline, intranet-based environments posed unique demands. Unlike typical cloud-based AI tools, this setup required stringent transparency, traceability, and auditability protocols—standards that are increasingly vital as regulatory expectations grow.

Looking ahead, the Copilot system holds strong potential for further enhancement and scale-up. Key future directions include:

- Integration with geospatial datasets (GIS), enabling decision support in disaster response;
- Link to the international policy databases of the European Commission or NATO (NATO, 1999) for broader normative references;

- Cross-institutional deployment, allowing for federated knowledge retrieval and harmonized document drafting across multiple public agencies.

From a sustainability perspective, the deployment of AI-driven KMS in civil protection contributes to more resource-efficient and economically resilient urban governance. By optimizing document workflows and improving the reuse of institutional knowledge, the Copilot supports leaner administrative processes that align with the goals of sustainable development and responsible public management (Kalogiannidis, Toska and Chatzitheodoridis, 2022). Moreover, by anchoring automated decision-support systems in transparent and value-based frameworks, the initiative enhances trust and legitimacy—key elements in maintaining the social license for AI in public institutions (Jankowski, Meijer and Grimmelikhuijsen, 2022).

Drawing on Roth et al. (2023), who show how European civil protection agencies innovated rapidly during COVID-19, this initiative affirms their potential as testbeds for agile digital transformation. Romanian civil protection authority's iterative development of Copilot in line with European trends promoting adaptive learning, crisis-driven innovation, and cross-functional knowledge integration. This case shows that AI can address procedural bottlenecks while enhancing institutional readiness for future shocks within legal and cultural constraints. Echoing Bratianu (2022), it presents a viable model for embedding strategic knowledge into governance without compromising accountability, transparency, or control.

As the Copilot becomes embedded in daily workflows, attention must be paid to human–AI dynamics noted by Alon-Barkat and Busuioc (2023), particularly *automation bias* and *selective adherence*, where users overly trust or selectively apply AI outputs. Although the Copilot emphasizes traceability to reduce such risks, future use should include training to address these behavioural patterns. Aligning technology with human factors is crucial to ensure AI enhances rather than distorts public-sector decisions. The civil protection Copilot ultimately presents a practical model for value-aligned, legally traceable AI integration in civil protection, offering a scalable template for resilient, knowledge-based digital transformation.

6. References

- Adama, H.E. and Okeke, C.D., 2024. Digital transformation as a catalyst for business model innovation: A critical review of impact and implementation strategies, *Magna Scientia Advanced Research and Reviews*, 10(2), pp. 256–264. https://doi.org/10.30574/msarr.2024.10.2.0066
- Alon-Barkat, S. and Busuioc, M., 2023. Human–AI Interactions in Public Sector Decision Making: "Automation Bias" and "Selective Adherence" to Algorithmic Advice, *Journal of Public Administration Research and Theory*, 33(1), pp. 153–169. https://doi.org/10.1093/jopart/muac007
- Arnold, M., Goldschmitt, M. and Rigotti, T., 2023. Dealing with Information Overload: A Comprehensive Review, *Frontiers in Psychology*, 14, 1122200. https://doi.org/10.3389/fpsyg.2023.1122200
- Bearman, D. and Trant, J., 1998. Electronic Records Research Working Meeting, May 28–30 1997: A report from the archives community, *Bulletin of the American Society for Information Science*, 24(3), pp. 13–17. https://doi.org/10.1002/bult.84
- Bratianu, C. and Bejinaru, R., 2019. The Theory of Knowledge Fields: A Thermodynamics Approach, *Systems*, 7(2), 20. https://doi.org/10.3390/systems7020020
- Bratianu, C. and Bejinaru, R., 2020. Knowledge dynamics: a thermodynamics approach, *Kybernetes*, 49(1), pp. 6–21. https://doi.org/10.1108/K-02-2019-0122
- Bratianu, C., Prelipcean, G. and Bejinaru, R., 2020. Exploring the latent variables which support SMEs to become learning organizations, *Management & Marketing. Challenges for the Knowledge Society*, 15, pp. 154–171. https://doi.org/10.2478/mmcks-2020-0010
- Bratianu, C., 2022. *Knowledge Strategies*. Cambridge: Cambridge University Press (Elements in Business Strategy). https://doi.org/10.1017/9781108864237
- Bratianu, C. and Bejinaru, R., 2023. From Knowledge to Wisdom: Looking beyond the Knowledge Hierarchy, *Knowledge*, 3(2), pp. 196–214. https://doi.org/10.3390/knowledge3020014
- Bratianu, C. and Ioan, A., 2025. Challenges of Artificial Intelligence for Knowledge Management Systems: A Bibliometric Analysis Perspective, *Oradea Journal of Business and Economics*, 10(1), pp. 108–120. https://doi.org/10.47535/1991ojbe209
- Chen, J., 2023. Model Algorithm Research based on Python Fast API, *Frontiers in Science and Engineering*, 3(9), pp. 7–14. [ISSN: 2710-0588]. https://doi.org/10.54691/fse.v3i9.5591

- European Union, 2013. *Decision No 1313/2013/EU of the European Parliament and of the Council of 17 December 2013 on a Union Civil Protection Mechanism*. Official Journal of the European Union, L347, pp. 924–947. [online] Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32013D1313 (Accessed: 9 May 2025).
- European Commission, 2017. *Ministerial Declaration on eGovernment – the Tallinn Declaration*. [online] Available at: https://digital-strategy.ec.europa.eu/en/news/ministerial-declaration-egovernment-tallinn-declaration (Accessed: 9 May 2025).
- European Commission, 2018, UCPM Lessons Learnt Programme, *Union Civil Protection Knowledge Network* [webpage]. [online] Available at: https://civil-protection-knowledge-network.europa.eu/eu-civil-protection-mechanism/ucpm-lessons-learnt-programme (Accessed: 09 May 2025).
- European Commission, 2021. *Europe's Digital Decade: 2030 Digital Targets*. [online] Available at: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/europes-digital-decade-digital-targets-2030_en (Accessed: 9 May 2025).
- European Commission, 2022. *EU Civil Protection Knowledge Network, European Civil Protection and Humanitarian Aid Operations*. [online] Available at: https://civil-protection-humanitarian-aid.ec.europa.eu/what/civil-protection/eu-civil-protection-knowledge-network_en (Accessed: 09 May 2025).
- European Commission, 2023. *State of Preparedness Report: Lessons from Crises – Towards a Resilient Union* (Niinistö Report). Brussels: Directorate-General for European Civil Protection and Humanitarian Aid Operations. [online] Available at: https://commission.europa.eu/system/files/2023-07/niinisto-report_en.pdf (Accessed: 9 May 2025).
- European Commission, 2025. *Commission Implementing Decision (EU) 2025/704 of 25 March 2025 amending Decision No 1313/2013/EU on a Union Civil Protection Mechanism*. Official Journal of the European Union, L87, pp. 1–12. [online] Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32025D0704 (Accessed: 9 May 2025).
- Gesk, T.S. and Leyer, M., 2022. Artificial intelligence in public services: When and why citizens accept its usage,' *Government Information Quarterly*, 39(3), p. 101704. https://doi.org/10.1016/j.giq.2022.101704
- Gruber, J.B. and Weber, M., 2024. rollama: An R package for using generative large language models through Ollama, *arXiv pre-print* arXiv:2404.07654. [online] Available at: https://arxiv.org/abs/2404.07654 (Accessed: 10 May 2025).
- Jankowski, J., Meijer, A.J. and Grimmelikhuijsen, S., 2022. Artificial Intelligence and Public Values: Value Impacts and Governance in the Public Sector, *Sustainability*, 15(6), 4796. https://doi.org/10.3390/su15064796
- Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Singh Chaplot, D. *et al.*, 2023. Mistral 7B, *arXiv preprint* arXiv:2310.06825. [online] Available at: https://arxiv.org/abs/2310.06825 (Accessed: 10 May 2025).
- Kalogiannidis, S., Toska, E. and Chatzitheodoridis, F., 2022. Contribution of Civil Protection to the Urban Economy: Evidence from a Small-Sized Greek City, *Sustainability*, 14(2), 981. https://doi.org/10.3390/su14020981
- NATO (1999) *NATO: NATO Integrated Data Service*. [online] Available at: https://www.nato.int/structur/nids/nids.htm (Accessed: 9 May 2025).
- Rawat, P. and Mahajan, A.N., 2020. ReactJS: A Modern Web Development Framework', *International Journal of Innovative Science and Research Technology*, 5(11), pp. 698–702. [online] Available at: https://www.ijisrt.com/assets/upload/files/IJISRT20NOV485.pdf (Accessed: 9 May 2025).
- Roth, F., Kaluza, B., Pfeffer, K., Rümelin, E., Kirchner, J., Overmeyer, M. *et al.*, 2023. Innovation in Times of Crisis: How Civil Protection Organizations in Europe Coped and Adapted During the COVID-19 Pandemic, *European Journal for Security Research*, 7, pp. 139–161. https://doi.org/10.1007/s41125-023-00090-6
- Romania, 2018. *Law No. 190/2018 regarding measures for the application of Regulation (EU) 2016/679 on the protection of natural persons regarding the processing of personal data and on the free movement of such data*. Official Gazette of Romania, Part I, No. 651 (26 July 2018). [online] Available at: https://legislatie.just.ro/Public/DetaliiDocument/203151 (Accessed: 09 May 2025).
- Singh, P.N., Talasila, S. and Banakar, S.V., 2023. Analyzing embedding models for embedding vectors in vector databases', *2023 IEEE International Conference on ICT in Business Industry & Government (ICTBIG)*, pp. 1–7. doi:10.1109/ictbig59752.2023.10455990.
- Wang, M., Zhang, J., Jiao, S. and Zhang, T., 2019. Evaluating the impact of citations of articles based on knowledge flow patterns hidden in the citations, *PLOS ONE*, 14(11), e0225276. https://doi.org/10.1371/journal.pone.0225276

- UNDRR, 2015. *Sendai Framework for Disaster Risk Reduction 2015–2030*. United Nations Office for Disaster Risk Reduction. [online] Available at: https://www.undrr.org/publication/sendai-framework-disaster-risk-reduction-2015-2030 (Accessed: 9 May 2025).