

ParlementAIre

Ornella Vaccarelli¹, Eden Brenot¹, Emmanuel De Salis², Henrique Marques Reis²,
Jacqueline Kucera³, Philippe Meyer³, Hatem Ghorbel² and Jean Hennebert¹

¹*Institute of AI and Complex Systems (iCoSys), School of Engineering and Architecture of Fribourg (HELA-FR), HES-SO University of Applied Sciences and Arts Western Switzerland;*

²*Haute Ecole Arc (HE Arc), HES-SO University of Applied Sciences and Arts Western Switzerland;*

³*Parliamentary Library, Research & Data, Parliamentary Services, Switzerland;*

Abstract

ParlementAIre is an innovative initiative to create a secure virtual assistant using conversational AI to support parliamentarians in their legislative work. Developed with the Parliamentary Library, it leverages advanced natural language processing to provide seamless access to relevant materials and parliamentary documents. By streamlining workflows and enabling informed decision-making, ParlementAIre enhances the efficiency and effectiveness of parliamentarians.

Keywords

legislative processes, conversational AI, document retrieval, natural language processing, augmented governance, virtual assistant, artificial intelligence in parliaments

1. Introduction

In modern legislative systems, the complexity and volume of documentation demand efficient, high-quality processes to support decision-making and governance. Parliamentarians are tasked with managing large amounts of legislative material while addressing pressing social challenges. This requires tools that not only streamline administrative operations but also improve strategic decision making.

The integration of artificial intelligence (AI) into parliamentary workflows has garnered significant attention in recent years, with the potential to transform legislative processes. From improving decision-making through automated document analysis to fostering public engagement through AI-powered chatbots, parliaments around the world are beginning to explore innovative applications of AI technology [1, 2]. The integration of AI into governance, especially in parliament workflows, is accelerating the shift towards more efficient, data-driven, and transparent decision-making processes [3]. The close collaboration of the project with the Parliamentary Library ensures seamless access to authoritative, preformatted documents, bridging the gap between technological innovation and legislative practice. This integration underscores the pivotal role of AI in the management of structured and unstructured data within parliamentary environments [4, 5, 6]. The use of AI-driven systems for legislative work is increasingly recognized for its potential to enhance operational efficiency and foster smarter, more responsive governance.

AI technologies have already demonstrated substantial impact across various domains of legislative work. For instance, AI-driven classification systems, using standardized thesaurus such as EuroVoc, have been implemented to automatically classify large volumes of parliamentary documents across jurisdictions such as Italy, the European Union, and Brazil [7, 8, 9, 10]. These tools not only automate categorization, but also provide deeper insight into legislative trends and priorities. Additionally, AI-enhanced tools for legislative drafting, such as those used in Italy, Brazil, and Chile, assist in creating, revising, and managing bills more efficiently [11, 12, 13].

The use of artificial intelligence (AI) in transcription and translation further enhances the accessibility and transparency of parliamentary processes. The systems developed by parliaments such as Estonia and Israel automate the transcription of legislative sessions and translating historical documents, making them more accessible to the public and supporting the preservation of legislative history [14, 15, 16]. Furthermore, summarization of legislative proceedings, including the conversion of records into podcasts, exemplifies the innovative use of AI to facilitate wider engagement with legislative content [17]. AI-powered chatbots, such as the one being developed as part of the ParlementAIre initiative in Switzerland, represent another key advancement. These tools provide intuitive interfaces for querying parliamentary processes and accessing documents, as demonstrated by similar applications in the European Parliament, Italy, and Chile [18, 19, 20, 21]. These chatbots not only improve user

AI days HES-SO '25 January 27–29, 2025, Switzerland

✉ ornella.vaccarelli@hefr.ch (O. Vaccarelli)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

experience but also improve accessibility to parliamentary information, particularly for citizens and stakeholders. The use of AI in governance is not limited to internal efficiency. It also extends to public engagement and security, as evidenced by AI applications for analyzing public opinion on proposed legislation in Brazil, and ensuring the security of parliamentary systems in Finland [22, 23]. Such applications highlight AI's transformative role in modernizing and securing democratic processes.

In this dynamic and rapidly evolving landscape, Switzerland is positioning itself at the forefront of innovation, harnessing AI to significantly enhance legislative transparency. In collaboration with the University of Bern, the Swiss Parliament has successfully automated the transcription of parliament sessions, including the historic debates of the inaugural session of the National Council in 1848 [24]. This initiative underscores Switzerland's commitment to innovation and the preservation of its legislative heritage.

Building on these advancements, *ParlementAIre* emerges as a pioneering initiative designed to further augment legislative processes in Switzerland. The project focuses on developing a secure conversational AI and virtual assistant specifically tailored to the needs of parliamentarians. By incorporating cutting-edge AI technologies, *ParlementAIre* aims to enhance the efficiency, precision, and accessibility of parliamentary workflows, ultimately empowering parliamentarians to engage more effectively in the legislative process. Specifically, the project seeks to increase the productivity of parliamentarians by integrating conversational AI systems capable of addressing their information retrieval and drafting needs. These systems are powered by a comprehensive knowledge base derived from research mandates executed by the Parliamentary Library. A key focus is prototyping functionalities such as intelligent question-answering, document summarization, contextualized document search, and structured drafting assistance.

2. Augmented Legislative Processes

The integration of AI in legislative workflows offers significant advantages by automating data retrieval, improving efficiency, and ensuring precision in legislative analysis and drafting. Beyond these immediate advantages, the project envisions an *augmented legislative process* where advanced AI tools do not replace human decision making, but complement it by streamlining routine operations. This approach ensures that parliamentary work maintains its rigor and depth while benefiting from enhanced productivity.

Figure 1 illustrates the workflow of the *ParlementAIre* system, showing the interaction between parliamentarians, the Parliamentary Library, and the AI-based tool. This collaboration ensures that all AI-assisted services are grounded in accurate, structured, and contextually relevant data.




Figure 1: Workflow of the *ParlementAIre* System: Parliamentarians (blue), the Parliamentary Library (orange), and the *ParlementAIre* tool (green). The workflow supports four core use cases: (1) Documentary Research, (2) Question/Answer, (3) Research Summary, and (4) Assisted Legislative Drafting. Color coding emphasizes the distinct roles and contributions of each entity.

The workflow begins with parliamentarians identifying a legislative theme and requesting relevant documents from the Parliamentary Library. These documents are curated, formatted, and fed into the AI system, ensuring compatibility with its processing pipeline. The system then provides four key functionalities tailored to legislative work:

1. **Documentary Research** – Efficiently retrieving and organizing relevant documents.
2. **Question/Answer** – Generating concise answers to legislative queries.
3. **Research Summary** – Synthesizing key insights from large set of documents.
4. **Assisted Legislative Drafting** – Proposing structured drafts based on contextual input.

The development of *ParlementAIre* follows an agile methodology tailored to meet the nuanced requirements of legislative work. Each development cycle spans three months, culminating in a Minimum Viable Product (MVP) that is rigorously tested in collaboration with the Parliamentary Library. Continuous refinements are driven by user feedback. In particular, this feedback has guided critical design decisions to enhance both usability and security:

- **Secure Data Management** – Parliamentary documents are processed within a self-hostable system, ensuring compliance with strict data privacy standards.
- **Transparent AI Outputs** – Every system response is accompanied by a traceable reference to the underlying source documents, promoting transparency and trust.

By addressing each of these challenges, *ParlementAIre* aims to provide a transformative tool for the Swiss Federal Parliament. This project represents a significant leap forward in legislative technology, promoting a more informed, efficient, and inclusive legislative process.

3. Technical Framework for Augmented Legislative Processes

The technical foundation of *ParlementAIre* is structured around four interconnected research axes, designed to ensure efficiency, precision, and adaptability to legislative workflows. The architecture integrates a robust back-end powered by LLM-RAG techniques with an intuitive front-end built using Vue.js, creating a seamless user experience for parliamentarians. Moreover, the entire system is self-hostable, enabling deployment in environments completely disconnected from the internet, a key feature for ensuring data security and compliance with sensitive legislative processes.




Figure 2: System Architecture of ParlementAIre: The parliamentarian interacts with the system via a Vue.js interface (frontend), which communicates with the backend through RESTful APIs. The backend processes requests by embedding documents through LangChain and storing them in the Milvus vector database. Contextual retrieval and processing by the LLM using RAG techniques deliver tailored outputs. Depending on the use case selected—Documentary Research, Question/Answer, Research Summary, or Assisted Legislative Drafting—the frontend displays results in an interactive and user-friendly manner.

3.1. Vector Databases

The document processing pipeline ensures that legislative documents are efficiently loaded, segmented, and indexed for rapid semantic search. The key steps are as follows:

1. **Loading PDFs:** Documents are read and extracted using `PyPDFDirectoryLoader`, a tool designed to handle large batches of parliamentary documents.
2. **Text Segmentation:** The extracted text is divided into chunks using a combination of the `MarkdownHeaderTextSplitter` and `RecursiveCharacterTextSplitter`. The `MarkdownHeaderTextSplitter`, a more advanced approach, organizes the text based on the

hierarchical structure of Markdown headers, forming a tree-like structure. In contrast, the `RecursiveCharacterTextSplitter`, a more classic method, is configured with `chunk_size=1500` and `chunk_overlap=100` to ensure compatibility with LLM processing.

3. **Document Embedding:** The segmented text chunks are converted into high-dimensional semantic embeddings using `LangChain` [25], which seamlessly integrates components for embedding generation. The embedding generation is powered by *bge-m3-custom-fr*, a specialized version of the BGE-M3 model [26] tailored for French language processing [27], and accessible through Ollama [28]. This model ensures the creation of semantically rich embeddings that capture contextual nuances critical to legislative content. Other embedding methods are currently under study for the different languages used in Switzerland, including German and Italian.
4. **Vector Storage:** The embeddings are stored and indexed in Milvus [29] vector database, a Kubernetes-ready, self-hostable solution that supports high-speed, in-memory operations. Milvus is optimized for large-scale storage and retrieval of high-dimensional embeddings, ensuring fast and accurate searches.

The embeddings, stored and indexed in the Milvus vector database, enabling efficient similarity search and retrieval of document chunks most relevant to the parliamentarian’s query.

3.2. Similarity Search

Similarity search is performed using the **Hierarchical Navigable Small World (HNSW)** [30] algorithm, a graph-based indexing technique that organizes data points into a hierarchical small-world network. Each data point is connected to its nearest neighbors, enabling efficient approximate nearest neighbor retrieval. To optimize performance, HNSW restricts the number of candidate neighbors during graph construction (`efConstruction`) and limits the maximum number of connections per node (`M`). A heuristic-based priority queue is employed to prioritize the most promising search paths. In our implementation, the parameters are configured as `M=48` and `efConstruction=400`, achieving a balance between computational efficiency and search accuracy in high-dimensional spaces.

To further refine the search results and ensure diversity, the **Maximum Marginal Relevance (MMR)** [31] algorithm is applied. MMR iteratively selects vectors by balancing their relevance to the query and their diversity from already selected results. The MMR criterion is defined as:

$$\text{MMR}(\mathbf{q}, \mathbf{v}_i) = \lambda \cdot \text{Cos_sim}(\mathbf{q}, \mathbf{v}_i) - (1 - \lambda) \cdot \max_{\mathbf{v}_j \in S} \text{Cos_sim}(\mathbf{v}_i, \mathbf{v}_j), \quad (1)$$

where $\mathbf{q}$ is the query vector, $\mathbf{v}_i$ is a candidate vector, and S represents the set of vectors already selected in the result set. The parameter λ controls the trade-off between relevance (similarity to the query) and diversity (dissimilarity from the selected results), with $\lambda = 0.4$ used in this implementation to achieve an effective balance. The similarity metric used in both HNSW graph construction and MMR scoring is **cosine similarity**, defined as $\text{Cos_sim}(\mathbf{q}, \mathbf{v}_i) = \frac{\mathbf{q} \cdot \mathbf{v}_i}{\|\mathbf{q}\| \|\mathbf{v}_i\|}$, where $\mathbf{q}$ is the query vector and $\mathbf{v}_i$ is a stored vector in the embedding space.

3.3. LLM-RAG Integration

The integration of Large Language Models (LLMs) with Retrieval-Augmented Generation (RAG) techniques forms the core of the *ParlementAI/re* system, allowing the delivery of precise and contextually relevant outputs [32, 33].

The integration leverages cutting-edge advancements in language modeling and retrieval systems to provide the following advantages:

- **Scalable Multilingual Processing:** Initial testing has employed open source LLMs such as *Mistral-7B* [34], *Mixtral-8x7B* [35], and *Gemma 1/2-7B* [36], optimized for both computational efficiency and support for multiple languages. This ensures accessibility and inclusivity for diverse legislative contexts.
- **Dynamic Adaptation to Context:** The system dynamically retrieves and processes document chunks according to the query context, ensuring that the output remains relevant and actionable.
- **Computational Efficiency:** By employing RAG techniques, the system reduces the computational overhead associated with generating large volumes of text, focusing LLM computation on the most pertinent information.

3.4. Prompt Engineering

Prompt engineering defines the roles and tasks of the LLM to maximize its performance. A systematic approach, involving iterative testing and refinement, ensures that prompts are optimized for various use cases. These include question-answering, document summarization, and drafting. Task-specific prompts are tailored to accommodate diverse languages, contexts, and user interactions, ensuring that output remains relevant and actionable.

Frontend and Backend Integration

The frontend of the *ParlementAIre* system is constructed using Vue.js, a progressive JavaScript framework, to deliver a user-friendly interface tailored for parliamentarians. The Vue.js-based interface provides:

- **Intuitive Navigation:** Easy access to the four different use cases.
- **Interactive Input:** The interface incorporates responsive input fields, including a text box, drop-down menus, and contextual blocks. These blocks guide users and AI in selecting the most relevant documents by refining themes or criteria, ensuring that queries are processed with greater precision and relevance.
- **Visualization of Results:** Outputs are presented in a clear and user-centric manner, including user queries, generated text outputs, and direct links to retrieved documents. Each response is accompanied by a traceable reference to the underlying source documents, promoting transparency and building trust by enabling parliamentarians to verify the origin and context of the information.

The Vue.js interface communicates with the back-end through RESTful APIs, enabling real-time query handling and response generation, dynamic updates as parliamentarians refine their queries or use different features, and secure interaction with the vector database and LLM-RAG processing pipeline.

Integrated Workflow for LLM-RAG

The complete workflow of the LLM-RAG architecture is illustrated in Figure 2. The workflow is as follows:

- **Frontend Interaction:** Parliamentarians use the Vue.js interface to define themes, enter queries, and select use cases. The interface sends these requests to the back-end.
- **Document Processing:** The backend retrieves relevant documents from the Parliamentary Library, embeds them using LangChain, and stores them in the Milvus vector database.
- **Contextual Retrieval and LLM Processing:** Relevant document chunks are retrieved according to the query context and processed by LLM using RAG techniques.
- **Response Generation:** The backend sends the synthesized response, whether it is an answer, summary, or draft, back to the Vue.js interface for display.

By integrating a dynamic Vue.js interface with a robust back-end architecture, the *ParlementAIre* project delivers a seamless, efficient, and user-centric experience, ensuring it meets the complex demands of legislative work.

4. Conclusions

ParlementAIre represents a bold step forward in leveraging AI to transform legislative processes. By integrating advanced conversational AI tools with parliamentary workflows, the project improves efficiency, precision, and accessibility, enabling parliamentarians to focus on impactful decision making. Through iterative development and close collaboration with stakeholders, *ParlementAIre* is setting a new standard for AI-driven governance.

In addition to that, it is to highlight that the integration of AI into parliamentary processes demands clear regulation and ethical governance. Frameworks such as the EU AI Act [37], with its risk-based rules for safety, transparency, ethical use, and innovation, and the Council of Europe's AI Convention, which emphasizes alignment with human rights, democracy, and the rule of law, provide essential guidance. At the national level, strategies such as the Swiss Federal Administration's AI sub-Strategy, which focuses on building AI competencies, fostering trust and enhancing efficiency, offer practical steps for implementation [38].

The IPU Brief 2024 on Generative AI points out that Parliaments play a critical role in regulating and overseeing AI to ensure its alignment with democratic principles and public interests [1]. They must address the challenges of rapid technological change, by managing risks, fostering collaboration, and ensuring that humans remain central in decision-making processes.

This responsibility extends to capacity building, equipping parliament staff and members with essential skills in AI literacy, adaptability, and collaboration, while embedding equity, fairness, and inclusion into AI applications to reduce bias and promote societal resilience. Furthermore, robust oversight and transparent governance are essential to building trust and protecting data protection.

By taking these steps, parliaments can guide the responsible integration of AI into democratic processes, fostering innovation while protecting fundamental rights, promoting public trust, and ensuring that AI serves the greater good of society.

References

- [1] Inter-Parliamentary Union, Index of parliamentary use cases, Online, 2024. URL: <https://docs.google.com/document/d/1tBRg5CPiW9kBUlrsn5RBpFpaqAVuC9Bk-36w7CkDqf0/edit?tab=t.0>, published by the PU's Centre for Innovation in Parliament as part of the Parliamentary Data Science Hub's project.
- [2] Inter-Parliamentary Union, Using generative ai in parliaments, Online, 2024. URL: <https://www.ipu.org/resources/publications/issue-briefs/2024-04/using-generative-ai-in-parliaments>, guidelines for using generative AI in parliamentary tasks, developed by the IPU's Centre for Innovation in Parliament.
- [3] Inter-Parliamentary Union, Artificial intelligence and strong parliaments, Online, 2024. URL: <https://www.ipu.org/impact/democracy-and-strong-parliaments/artificial-intelligence>.
- [4] J. von Lucke, F. Fitsilis, S. Gagnon, Using artificial intelligence in parliament - initial results from the canadian house of commons (2024).
- [5] F. Fitsilis, J. von Lucke, F. De Vrieze, G. Mikros, M. Palmirani, A. Read, G. Schefbeck, A. Pastor y Camarasa, S. Gagnon, J. Alberto de Oliveira Lima, et al., Guidelines for ai in parliaments, Westminster Foundation for Democracy (2024).
- [6] J. von Lucke, F. Fitsilis, Using artificial intelligence in parliament-the hellenic case, in: International Conference on Electronic Government, Springer, 2023, pp. 174–191.
- [7] IT Department, Chamber of Deputies of Italy, Automatic classification of policy and control acts on government, 2024. URL: https://docs.google.com/document/d/1pxwyZMgiDg_8mZA4d6vvWUPrHQc2Bi4RZCvOHqLUG_Q/edit?tab=t.0, use Case ID: 001.
- [8] IT Department, Chamber of Deputies of Italy, Automatic classification of legislative documents, 2024. URL: <https://docs.google.com/document/d/1vUvsHmlK89nYdrVA0vPhjjqRP3pSew0NjEST9XB3D4/edit?tab=t.0>, use Case ID: 007.
- [9] European Parliament, Classification of input texts with predefined labels from the eurovoc thesaurus, 2024. URL: https://docs.google.com/document/d/1pne-FP5Bsh3hCHWBuNWJpaydYUR1uZqIewRa_LNEpog/edit?tab=t.0, use Case ID: 013.
- [10] Chamber of Deputies of Brazil, Thematic classification of portal content, 2024. URL: <https://docs.google.com/document/d/1pP327BmARRQqeiBpnegi2FBKtarHXLNrtLnEwV4jxE/edit?tab=t.0>, use Case ID: 018.
- [11] IT Department, Chamber of Deputies of Italy, Semi-automated production of verbatim reports of sittings using automatic speech recognition, 2024. URL: <https://docs.google.com/document/d/1fziauOXVZ3jldgWySx-kD1sVLvsFkbrHO5tg-JA6in0/edit?tab=t.0>, use Case ID: 002.
- [12] Chamber of Deputies of Brazil, Retrieval of relevant documents for the lawmaking expert handling an mp's demand, 2024. URL: https://docs.google.com/document/d/16zqT_2F-nGDZvbUEDF-pQDDUWTKSg1I0pZNO5yz2Qk/edit?tab=t.0, use Case ID: 020.
- [13] Chamber of Deputies of Chile, Bill drafting assistant, 2024. URL: <https://docs.google.com/document/d/1Op76mFS0eC0ThCp2mFlf0EAjAHVog6iqjnQAbjbd-uo/edit?tab=t.0>, use Case ID: 048.
- [14] IT Department, Chamber of Deputies of Italy, Automatic transcription of handwritten manuscripts from historical archives, 2024. URL: https://docs.google.com/document/d/12rXzZ0m_ywPEk6gmNneulNmOpNQtw4lOOLaYSKE2-Nc/edit?tab=t.0, use Case ID: 006.
- [15] Riigikogu (The Estonian Parliament), Ai-powered verbatim records system (hans), 2024. URL: https://docs.google.com/document/d/1gim5xsUqu2dM9VKw_BY0Cv0DTgv8H84qlimHDG5x3uc/edit?tab=t.0, use Case ID: 016.
- [16] Parliament of Israel, Audio-to-text parliamentary transcription system, 2024. URL: https://docs.google.com/document/d/1jzOSFxLe2q32T_xUFR6sSCO8jAZXK8PBRd14v09cgXA/edit?tab=t.0, use Case ID: 047.
- [17] Parliament of Finland, Summarizing of parliamentary documents and conversion into podcasts, 2024. URL: https://docs.google.com/document/d/133s1RfnCv3InfkUb12Kk4lz5R_MAVzkBhkwKCXHGTrg/edit?tab=t.0, use Case ID: 044.
- [18] IT Department, Chamber of Deputies of Italy, Chatbot for the parliamentary documentation website, 2024. URL: https://docs.google.com/document/d/1bMGyFQDP_iQm7w6_JXiNEOCN-Y916GCsA0C280pT9mY/edit?tab=t.0, use Case ID: 008.
- [19] European Parliament, Chatbot on security in the european parliament, 2024. URL: <https://docs.google.com/document/d/1ScVFyWyt3M2T3JuMSbogg11RFo0emV4jckg7V9nCGGCU/edit?tab=t.0>, use Case ID: 014.
- [20] European Parliament, Ai-powered customer support chatbot, 2024. URL: <https://docs.google.com/document/d/1XjDCQ8zgMgZiH9QzmmI5vi2s599-z8QVtSnSAfjNxCI/edit?tab=t.0>, use Case ID: 015.
- [21] Chamber of Deputies of Chile, Regulations and parliamentary allowances assistant, 2024. URL: https://docs.google.com/document/d/1NBck0VONz_L-46KFfKSoI6D7cmUNvoOfEAPl4k9igs/edit?tab=t.0, use Case ID: 049.
- [22] Chamber of Deputies of Brazil, Analysis of citizens' opinions on bills, 2024. URL: <https://docs.google.com/>

- document/d/1O3U2x_0t4kv_zDhIjeitYweqB9fHBS2iA3vXnMZAS5M/edit?tab=t.0, use Case ID: 023.
- [23] Parliament of Finland, Narrowing of the attack threshold for parliamentary applications, 2024. URL: https://docs.google.com/document/d/11Z6kC9h80PVUQWgjnXSl_uEPkK1dPk3XKtCJCO5Vpus/edit?tab=t.0, use Case ID: 045.
 - [24] The Swiss Parliament, Handschriftliche protokolle der bundesversammlung der schweizerischen eidgenossenschaft, 2022. URL: <https://parlament.transkribus.eu/>.
 - [25] LangChain, Langchain: A framework for developing applications powered by language models, 2024. URL: <https://www.langchain.com/>.
 - [26] FlagOpen, Bge-m3: A multilingual embedding model, 2022. URL: <https://github.com/FlagOpen/FlagEmbedding>.
 - [27] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, P. Moi, P. Cistac, H. Rousset, R. Louf, M. Funtowicz, et al., Huggingface’s transformers: State-of-the-art natural language processing, arXiv preprint arXiv:1910.03771 (2019).
 - [28] Ollama, Ollama: Ai tools for seamless embedding and model integration, 2024. URL: <https://www.ollama.com/>.
 - [29] Milvus Community, Milvus: Open-source vector database for scalable similarity search, <https://milvus.io/>, 2024. Accessed: November 20, 2024.
 - [30] Y. A. Malkov, D. A. Yashunin, Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs, IEEE Transactions on Pattern Analysis and Machine Intelligence 42 (2020) 824–836. URL: <https://doi.org/10.1109/TPAMI.2018.2889473>. doi:10.1109/TPAMI.2018.2889473.
 - [31] J. Carbonell, J. Goldstein, Use of maximum marginal relevance (mmr) for reducing redundancy and improving ranking in retrieval and summarization, in: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), ACM, 1998, pp. 335–336.
 - [32] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Advances in Neural Information Processing Systems (NeurIPS) 33 (2020) 1877–1901. URL: <https://arxiv.org/abs/2005.14165>.
 - [33] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive nlp tasks, Advances in Neural Information Processing Systems (NeurIPS) 33 (2020) 9459–9474. URL: <https://arxiv.org/abs/2005.11401>.
 - [34] M. AI, Mistral-7b: A high-performance open-source language model, 2024. URL: <https://www.mistral.ai/>.
 - [35] M. AI, Mixtral-8x7b: A specialized model ensemble for scalable language processing, 2024. URL: <https://www.mistral.ai/>.
 - [36] G. AI, Gemma 1/2-7b: Multilingual language models for diverse applications, 2024. URL: <https://www.gemma.ai/>.
 - [37] European Union, Artificial intelligence act (regulation (eu) 2024/1689), Official Journal version of 13 June 2024, 2024. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>, interinstitutional File: 2021/0106(COD).
 - [38] Swiss Parliament, , Strategy for the use of ai systems in the federal administration, ICT sub-strategy commissioned by the Federal Council in accordance with Art. 17 DTIO, 2024.