



Inter-Parliamentary Union
For democracy. For everyone.



Case studies for AI in parliaments

October 2026

Drawn from the May 2026 meeting of the IPU
Centre for Innovation in Parliament's
Parliamentary Data Science Hub, hosted by the
Chamber of Deputies of Italy

Table of contents

[Introduction](#) 2

[The state of implementation of AI tools in parliaments](#) 4

[Bill drafting and amendments](#) 6

[Normative assistant — Chamber of Deputies of Chile](#) 7

[Regulatory impact assistant — Chamber of Deputies of Chile](#) 10

[Quorum classification assistant — Chamber of Deputies of Chile](#) 13

[Admissibility assistant — Chamber of Deputies of Chile](#) 16

[Amendment-writing machine — Chamber of Deputies of Italy](#) 19

[Legislative amendment grouper — Chamber of Deputies of Brazil](#) 21

[Amendment vote-ordering and similarity-detection system — Chamber of Deputies of Italy](#) 23

[Committee work and plenary proceedings](#) 26

[Bill-path determination assistant — Chamber of Deputies of Brazil](#) 27

[Automatic voice-to-text transcription of plenary and committee sessions — Senate of Spain](#) 29

[Document and records management](#) 31

[Content categorization according to the EuroVoc thesaurus — Senate of Spain](#) 32

[Generative AI and AI-assisted development for legislative tracking — Chamber of Deputies of Italy](#) 34

[Automated translation of legislative and regulatory acts — Verkhovna Rada of Ukraine](#) 37

[MP support and research](#) 40

[Ulysses Chat — Chamber of Deputies of Brazil](#) 41

[Argumentative assistant — Chamber of Deputies of Chile](#) 44

[AI assistant for search, analysis and work with legal acts — Verkhovna Rada of Ukraine](#) 47

[Plenary proceedings and transcription](#) 50

[Parliamentary questions and legislative support documents — National Assembly of Mauritius](#) 51

[Multilingual support for parliamentary documentation — National Assembly of Mauritius](#) 55

[Public engagement and accessibility](#) 58

[Plain-language bill summarizer — Chamber of Deputies of Brazil](#) 59

[Internal administration, IT and security](#) 60

[AI-assisted re-engineering of a parliamentary legacy application — Chamber of Deputies of Italy](#) 61

[Digital assistant based on RAG techniques to assist in the modernization of COBOL codebases — Chamber of Deputies of Italy](#) 63

[Facial recognition and vehicle licence plate recognition for access control to administrative buildings and territories — Verkhovna Rada of Ukraine](#) 65

Introduction

This collection brings together 21 case studies documenting how artificial intelligence (AI) is being applied across parliaments today, drawing on the direct experience of six parliaments and chambers: the Chamber of Deputies of Brazil, the Chamber of Deputies of Chile, the Chamber of Deputies of Italy, the National Assembly of Mauritius, the Senate of Spain and the Verkhovna Rada of Ukraine. It offers a practical, first-hand view of AI adoption in parliaments, building on the [Use cases for AI in parliaments](#) published by the Inter-parliamentary Union (IPU) in 2025. The authors also wish to thank the Senate of Italy for their support.

The case studies are drawn from the May 2026 in-person meeting of the IPU Centre for Innovation in Parliament's Parliamentary Data Science Hub, hosted in Rome by the Chamber of Deputies of Italy.

Each case study follows a consistent structure (with some or all of the subsections listed below) so that readers can quickly compare approaches and identify what is most relevant to their institution's priorities:

- Overview
- Business problem
- Solution
- Human interaction
- Results and impact
- Risk management
- Lessons learned
- Next steps
- Data
- Architecture and technologies

The case studies are organized into functional areas that reflect the breadth of parliamentary work where AI is already delivering value: bill drafting and amendments; committee work and plenary proceedings; document and records management; member support and research; public engagement and accessibility; and internal administration, IT and security.

Applications range from amendment-grouping and plain-language bill summarization to legacy-system modernization, facial-recognition access control and generative AI assistants for legislative research.

For senior management and ICT leadership, this document is intended as a practical reference: a way to identify comparable use cases, understand the governance and risk-management choices other parliaments have made, and inform decisions about where AI can most usefully be introduced in their own institution's operations.

The functional areas covered by these case studies are set out in the table below (the case studies are presented under their primary functional area in the document itself):

Functional area	Brazil	Chile	Italy	Mauritius	Spain	Ukraine
Bill drafting and amendments	✓	✓	✓			✓
Committee work	✓		✓			
Document and records management	✓		✓		✓	✓
HR and internal administration	✓		✓			
IT infrastructure and cybersecurity			✓			✓

Functional area	Brazil	Chile	Italy	Mauritius	Spain	Ukraine
MP support and research	✓	✓	✓			✓
Plenary proceedings and transcription				✓	✓	
Public engagement and accessibility	✓		✓			✓
Other (physical security)						✓

The state of implementation of AI tools in parliaments

Taken together, these case studies highlight a set of recurring themes in how parliaments are putting AI into practice, even though the tools themselves vary widely in scope and sophistication:

- **The most mature applications run on public data:** The Normative Assistant and Regulatory Impact Assistant of the Chamber of Deputies of Chile draw only on published legislation, decrees and regulations rather than confidential or personal information. This keeps privacy risk low and makes the underlying knowledge base easier to build, curate and keep current.
- **Adoption matters as much as the technology:** Often, the harder challenge has been not building the model but persuading MPs and staff to trust and use it. Embedding an assistant into channels people already use, such as – in the case of Chile – the internal Members’ Portal and WhatsApp, has proved more effective than launching it as a stand-alone tool that people must remember to open.
- **It is important to consider where AI adds the most value:** Well-defined, high-volume tasks that demand complex analysis are more suited to AI deployment than those involving open-ended reasoning. In Italy, the Chamber of Deputies’ amendment-ordering and similarity-detection system automatically classifies and sequences large volumes of amendments – work that previously consumed extensive staff time but follows clear, established procedural rules.
- **What AI is deliberately not being asked to do matters just as much:** At the Chamber of Deputies of Italy, staff were careful to keep their tool focused on procedural sorting rather than admissibility or legal interpretation, which remain the responsibility of the Parliamentary Research Service and the Budget Service. Across the case studies, AI consistently supports rather than replaces judgement on matters of legal or political substance. It is clear from these case studies that maintaining “humans in the loop” is critical for trust – and necessary to ensure a safe and robust process.

These examples suggest that parliamentary AI is maturing fastest where the underlying problem is procedural, the data is open and the tool integrates with existing MP and staff workflows. Where decisions are more complex, legal or political, human knowledge outweighs AI efficiencies. The case studies also highlight the risk that, the longer a parliament takes to implement its own AI solution to a well-defined problem, the more likely it is that unofficial, shadow-AI tools will be developed.

Sharing is learning

A common thread runs through many of these case studies: the most durable AI applications often draw on ideas developed elsewhere. Cooperation ranges from situations in which one institution adapts another’s tool, to arrangements involving joint development, or pilots that become regional resources. Each works best when every party’s contribution is acknowledged.

The clearest illustration is the European Parliament’s “ArchiBot” tool, which was adapted by the Chamber of Deputies of Chile as “ArchileBot” and then [extended through a regional exchange to Argentina, Mexico, Peru and Uruguay](#). In this case, a single transfer ended up becoming a shared resource across Latin America.

Without cooperation, parliaments risk duplicating effort and losing the institutional memory of who solved a problem first. The real acceleration often comes from early, informal exchanges between technical teams – a process that is easy to overlook once a tool is live.

Inter-parliamentary networks exist to support this knowledge exchange. The IPU Centre for Innovation in Parliament, with its thematic and regional hubs, helps seed and share early-stage ideas between institutions before duplication sets in. This collection draws on such an exchange and creates new knowledge for other parliaments to learn from.

Future editions of these case studies could ask parliaments to note where their solution drew on prior work, capturing provenance within the case study itself. A mature parliamentary AI ecosystem might well be measured by how effectively knowledge moves between institutions.

Bill drafting and amendments

Normative assistant

Parliament/chamber: Chamber of Deputies of Chile

Functional area(s): Bill drafting and amendments; MP support and research

Date: July 2026

Overview

The Chamber of Deputies of Chile has developed a normative assistant to help MPs and staff quickly identify the legislation applicable to a given topic. Available through the Members' Portal and WhatsApp, the application uses semantic search to retrieve relevant legal provisions at the article level from natural language queries. It supports legislative drafting, legal research and other day-to-day parliamentary activities by providing fast and accurate access to the legal framework.

Business problem

MPs and staff frequently need to identify legislation applicable to a specific topic when drafting bills, analysing existing legislation or preparing legal and technical documents. This process often requires users to search across multiple legal databases and review lengthy legal texts, making it time-consuming and dependent on individual legal expertise. Existing keyword-based search tools may also fail to retrieve relevant provisions when users are unfamiliar with the exact legal terminology or citation.

The assistant was developed to address this challenge by enabling semantic search across the legal corpus. Instead of relying on exact keywords, users can formulate questions in natural language and retrieve the most relevant legal provisions at the article level.

The objective was to reduce the time spent locating applicable legislation, improve consistency in legal research and make legal information more accessible to parliamentarians during their daily work.

Solution

The solution is an internal assistant that quickly finds the legislation applicable to a specific topic. Instead of searching through multiple legal databases or relying on exact keywords, users can describe their question in everyday language (e.g. "What legislation regulates personal data protection?", "What are the legal requirements for responsible pet ownership?" or "What regulations govern recycling and waste management?").

The assistant identifies the most relevant legal provisions and presents the corresponding articles together with their source, allowing users to review the original legislation directly. This makes it easier to verify information and understand the legal basis for a particular issue without having to manually search through extensive legal texts.

The application is available through the Chamber's internal Members' Portal and via WhatsApp (for MPs and their staff), making it easily accessible during day-to-day parliamentary work. It is intended as a support tool rather than a replacement for legal analysis: users remain responsible for interpreting and applying the legislation, while the assistant helps them locate the relevant legal information more efficiently.

Human interaction

The application fits naturally into existing workflows by supporting the initial stage of legal research, replacing time-consuming manual searches with a faster way to identify relevant legislation. The AI recommends relevant laws, but does not interpret them, assess legal implications or recommend policy decisions.

Users are responsible for reviewing the retrieved legislation, verifying its relevance, and deciding how to apply it in legislative drafting, legal analysis or parliamentary discussions. Human judgement remains essential throughout the process, with the assistant serving as a tool to improve speed and access to information rather than replacing legal expertise.

Results and impact

The assistant is still in the early stages of deployment, and the main achievement so far has been the establishment of regular use among staff and MPs.

One of the greatest challenges has been not the technology itself, but encouraging adoption: making potential users aware of the tool, helping them understand what it can do, demonstrating how to use it effectively, and building confidence in using a parliamentary AI assistant instead of general-purpose commercial AI tools. A growing group of users have now incorporated the assistant into their daily work, particularly when researching the legal framework applicable to a policy issue or when preparing legislative initiatives.

Early feedback has been positive, highlighting the convenience of accessing authoritative legal information through natural-language queries, as well as the value of having responses grounded in the official Chilean legal framework. The Chamber is currently collecting user feedback and monitoring usage patterns to guide future improvements. As adoption continues to expand, it expects to establish more robust indicators on user satisfaction, time savings, and the impact of the assistant on legislative research and drafting workflows.

Risk management

The main risks identified were: inaccurate or incomplete retrieval of legal provisions; AI hallucinations; responses to questions outside the application's intended scope; and the possibility that users might treat the assistant's responses as legal advice rather than legal references.

These risks were addressed through a combination of prompt engineering and human oversight. The assistant was given explicit instructions to: answer only questions related to Chilean legislation; avoid speculating when sufficient information was not available, and acknowledge when it could not answer a query, instead of generating potentially incorrect information. It was also instructed not to provide legal interpretations, recommendations or policy opinions, but instead to focus on identifying the applicable legal framework and citing the relevant provisions whenever possible.

Users remain responsible for interpreting and applying the legislation in their parliamentary work. As the application operates exclusively on public legal information and does not process confidential parliamentary or personal data, privacy risks are minimal.

Lessons learned

One of the main lessons learned was that a reliable AI legal assistant depended not only on the AI model, but also on having high-quality, well-organized legislative information. The application would not have been possible without having Chilean legislation available in digital, structured and machine-readable formats.

The team also learned that maintaining those legal sources was an ongoing responsibility. As legislation evolves, the knowledge base must be kept up to date to ensure that users can rely on the assistant's responses.

Our advice to other parliaments is to invest early in the quality, organization and governance of their legislative information. A well-curated legal corpus provides the foundation on which reliable AI applications can be built.

Next steps

Future work will focus on systematically collecting user feedback and measuring user satisfaction to guide the assistant's continuous improvement. It will also focus on

continuously updating the assistant's legal knowledge base to reflect new legislation and amendments, ensuring that its responses remain up to date. Another priority is to integrate the application more seamlessly into the day-to-day workflows of MPs and staff, making it easier, more intuitive and more convenient to use as part of their regular legislative work. The long-term objective is for the assistant to become a natural component of parliamentary research rather than a stand-alone digital tool.

Data

The application draws on publicly available Chilean legal and regulatory sources, including laws, decrees, regulations, resolutions, instructions, circulars and municipal ordinances. These documents constitute the official legal framework that governs public administration and legislative activity in Chile.

The assistant does not process confidential parliamentary information or personal data.

Architecture and technologies

The application is built on the GPT-4o language model and uses a retrieval-augmented generation (RAG) architecture to search for and retrieve information from a curated corpus of current Chilean legislation and regulations.

The assistant is available through two channels: a desktop version is integrated into the Chamber of Deputies' internal Members' Portal, while a mobile version is delivered through WhatsApp to support quick legal queries during parliamentary activities. Both interfaces connect to the same AI services and legal knowledge base.

Regulatory impact assistant

Parliament/chamber: Chamber of Deputies of Chile

Functional area(s): Bill drafting and amendments

Date: July 2026

Overview

The Chamber of Deputies of Chile has developed the regulatory impact assistant to support the technical analysis and drafting of legislation. The application reviews draft bills against the existing legal framework in Chile to identify potential contradictions, complementary provisions and regulatory gaps. By providing a comprehensive assessment of how a proposed bill interacts with current legislation, the assistant helps MPs and staff improve the technical quality, coherence and consistency of legislative proposals before they advance through the legislative process.

Business problem

The complexity of the Chilean legal framework has made it challenging for MPs and staff to assess how new legislative proposals interact with existing legislation. Draft bills must be reviewed against a large and constantly evolving body of laws to identify potential inconsistencies, overlaps, complementary provisions or regulatory gaps. Conducting this analysis manually is a time-consuming exercise that depends heavily on specialized legal expertise, making it difficult to perform consistently under the time constraints of the legislative process.

As legislative activity has grown, there has been a need to strengthen the technical quality of bills while providing more efficient support to parliamentary offices. Staff require tools that can rapidly identify relevant legislation and highlight legal issues that may affect the coherence and implementation of proposed laws. Improving the consistency of legislative drafting at an early stage also helps reduce the need for subsequent amendments and contributes to a more robust legislative process.

Solution

The assistant was integrated into the Chamber of Deputies' parliamentary digital platform, where it is available to MPs and staff as part of their daily legislative work.

Users submit the text of a draft bill, either in full or in part, and the assistant compares the proposal against the current legal framework in Chile. While the tool can analyse short excerpts, more complete drafts generally produce more comprehensive results. The assistant identifies existing legislation that is related to the proposal, highlighting provisions that may complement, overlap with or contradict the draft bill. It also points out potential regulatory gaps and explains how similar legal issues have been addressed in previous legislation.

Rather than simply listing relevant laws, the assistant provides a structured analysis describing the relationship between the proposed bill and the existing legal framework, together with references to the legal sources that support its findings. The analysis helps parliamentary users understand the broader legal context of a proposal and identify issues that may require further consideration before a bill is introduced or discussed. This allows MPs and their staff to strengthen the technical quality and legal coherence of legislative proposals while reducing the time required to carry out complex legal research manually.

Human interaction

The assistant is primarily used during the preparation and review of draft legislation. It fits into the existing legislative drafting process by providing rapid access to relevant legal analysis before a bill is introduced or amended. Users interact with the assistant conversationally, asking follow-up questions, requesting additional analysis or refining the scope of the review as their work progresses.

The assistant recommends relevant legislation, identifies potential contradictions, complementary provisions and regulatory gaps, and provides a structured assessment of how a draft bill relates to the existing legal framework. However, it does not determine whether a proposal should proceed; nor does it replace legal or policy analysis.

Users are responsible for evaluating the relevance and accuracy of the findings, interpreting their implications and deciding whether changes to the draft are appropriate. Final legislative decisions, legal judgement and political considerations remain entirely with MPs and staff.

Results and impact

The assistant is currently available to MPs and staff through the parliamentary portal. Adoption has been gradual, with users primarily relying on it when preparing or reviewing draft legislation. Early experience has shown that the quality of the analysis is strong when the assistant is used for its intended purpose: assessing how a draft bill interacts with the existing Chilean legal framework.

Risk management

The main risks identified were the possibility of hallucinations and the potential for users to rely excessively on AI-generated outputs. To reduce the risk of inaccurate or irrelevant responses, the assistant was designed with a narrow scope: it performs a single, well-defined task, i.e. analysing draft legislation against the existing Chilean legal framework.

This focused design is reinforced through carefully engineered prompts that guide the model to produce structured analyses based on the relevant legal context, rather than attempting to answer general-purpose questions. The Chamber also recognizes that generative AI cannot replace professional legal judgement. Users are therefore required to review and validate every output before using it in their legislative work, and they remain fully responsible for any information incorporated into draft legislation or other parliamentary documents. The assistant is intended to support legal analysis, not to make legal or legislative decisions.

Privacy was not considered a significant risk in this use case, as the assistant operates exclusively on publicly available legislation and regulations rather than confidential parliamentary information or personal data.

Lessons learned

One of the main lessons learned relates to user experience rather than analytical performance: because the tool is currently presented as a conversational assistant, some users initially approached it as a general-purpose chatbot, requesting tasks beyond its purpose, such as drafting official documents, providing information about the status of bills or answering procedural questions. This highlighted the importance of aligning the user interface with the assistant's specific function – and of clearly communicating what the assistant is (and is not) designed to do, so that users understand both its capabilities and its limitations.

The team learned that this challenge could be addressed not only through effective user onboarding and training, but also by redesigning how users interact with the application. For this reason, the next stage of development will embed the assistant directly into the legislative drafting platform, allowing it to analyse bills as they are drafted rather than relying on free-text chat interactions.

More broadly, the Chamber found that integrating AI into existing parliamentary workflows – rather than offering it as a stand-alone application – was more effective in terms of encouraging adoption. Our advice to other parliaments is to begin with a narrowly defined use case, design the application around a genuine institutional need, and pay as much attention to user adoption and workflow integration as to the AI technology itself.

Next steps

As mentioned above, the next step is to integrate the assistant directly into the Chamber's legislative drafting platform, enabling it to analyse each provision of a bill as it is being drafted rather than through a stand-alone chat interface. This will provide context-aware analysis throughout the drafting process, improve usability, and encourage broader adoption by embedding the assistant into existing parliamentary workflows. Beyond this application, the Chamber will continue developing specialized AI assistants for other parliamentary functions, as well as sharing its experience, methodologies and lessons learned with other legislatures interested in AI adoption.

Data

The application draws on publicly available Chilean legal and regulatory sources, including laws, decrees, regulations, resolutions, instructions, circulars and municipal ordinances. The assistant does not process confidential parliamentary information or personal data.

It operates exclusively on public legal documents, making privacy and information sensitivity risks minimal.

Architecture and technologies

The application is built on the GPT-4o language model and uses a retrieval-augmented generation (RAG) architecture to search for and retrieve information from a curated corpus of current Chilean legislation and regulations. The assistant is available through a desktop version, integrated into the Chamber of Deputies' internal Members' Portal.

Quorum classification assistant

Parliament/chamber: Chamber of Deputies of Chile

Functional area(s): Bill drafting and amendments

Date: July 2026

Overview

The Chamber of Deputies of Chile has developed the quorum classification assistant to support constitutional review during the legislative drafting process. The application analyses draft bills to determine whether they fall within the constitutional framework, identifies the quorum required for their approval and provides the relevant constitutional provisions and case law that support its assessment. By helping MPs and staff identify constitutional requirements at an early stage, the assistant contributes to a more informed legislative drafting process.

Business problem

Determining the constitutional quorum required for a bill is a critical step in the legislative process in Chile, as it affects both the procedure for its approval and its constitutional validity.

Identifying the appropriate quorum is not always straightforward. It often requires reviewing constitutional provisions, as well as relevant case law from the Constitutional Court and legislative precedents, to determine whether a proposal regulates matters subject to special constitutional requirements.

This analysis is traditionally carried out by parliamentary legal advisers with constitutional expertise and is both time-consuming and complex, particularly when bills address multiple legal issues or introduce amendments across different policy areas. As legislative activity increases, providing timely and consistent constitutional assessments becomes increasingly challenging, creating pressure on parliamentary support services and making it more difficult to deliver early guidance during the drafting process.

Solution

The assistant was developed to address this institutional need by supporting the preliminary constitutional analysis of draft legislation, enabling constitutional issues to be identified earlier in the legislative process while ensuring that final legal interpretation and procedural decisions remain the responsibility of parliamentary experts.

The assistant is integrated into the Chamber of Deputies' legislative drafting platform, where it supports MPs and staff: as users draft or revise legislative provisions, the assistant analyses each individual article from a constitutional perspective, rather than assessing the whole bill. The assistant determines the constitutional quorum likely to apply to the proposal and explains the reasoning behind its assessment. It identifies the relevant constitutional provisions, references applicable case law and legislative precedents, and indicates whether these support or challenge the proposed classification. Rather than providing only a conclusion, the assistant presents a structured explanation that allows users to understand the constitutional basis for its recommendation.

Human interaction

The assistant is integrated directly into the legislative drafting platform, allowing users to access AI-supported analysis within their existing workflow, without requiring a separate tool or process.

Human judgement remains essential throughout the process: users review the assistant's analysis, consider the broader legislative context, validate its relevance and make the final decisions regarding the wording of provisions and the procedural approach to be followed.

The tool is designed to enhance legal analysis and provide timely guidance, while maintaining human responsibility for legislative decisions and constitutional interpretation.

Results and impact

The assistant has completed an internal validation process with positive results, including testing by specialized users with legal expertise. As the legislative drafting platform is not yet in production, the assistant has not yet been used by MPs or staff in their day-to-day legislative work. Measurable impacts from real-world adoption are therefore not yet available.

Early validation results indicate that the tool can provide relevant support in the analysis of quorum requirements by organizing constitutional criteria and applying them consistently at the article level. The expected impact is a reduction in the time required for legal research and preliminary constitutional analysis during the drafting and revision of legislative provisions, allowing parliamentary teams to focus more on substantive legislative considerations.

In practice, once deployed, the assistant is expected to improve the accessibility and consistency of constitutional analysis within the drafting workflow, particularly by helping users to identify potential quorum requirements earlier in the legislative process, while preserving human review and decision-making as essential components of the legislative process.

Risk management

The main risks identified for the assistant relate to the accuracy and reliability of AI-generated analysis, as well as the potential for users to rely excessively on automated recommendations in a legally sensitive context.

To address these risks, and given that quorum classification involves constitutional interpretation, the assistant was designed as a decision-support tool rather than an autonomous decision maker. It was also developed to provide analysis based on validated constitutional and legislative information sources, and its performance was reviewed through an internal validation process involving specialized legal users.

Human oversight remains a fundamental safeguard: users are responsible for reviewing the assistant's outputs, applying their legal judgement and making the final decisions regarding legislative drafting and procedural considerations.

Finally, privacy was assessed as a low-risk area in this use case, as the assistant operates exclusively on publicly available legislation and regulations rather than confidential parliamentary information or personal data.

Lessons learned

One of the main lessons learned is the importance of designing AI applications around existing parliamentary workflows rather than creating stand-alone tools. Integrating the assistant directly into the legislative drafting platform ensures that it becomes part of users' daily activities and it is used within its intended scope.

Another key learning was the importance of involving legal specialists throughout the development and validation process. Their expertise was essential to assess the quality and relevance of the assistant's outputs and validate the accuracy of its responses. For other parliaments considering similar applications, an important recommendation is to establish specialized validation processes.

AI tools can provide significant value in parliamentary processes, but their effectiveness depends not only on the technology itself, but also on reliable institutional knowledge, including real-world examples of quorum assessments and expert review of the generated outputs.

Next steps

The next step is to deploy the assistant as part of the Chamber of Deputies' legislative drafting platform, and to evaluate its performance in real parliamentary workflows.

Based on user interaction, feedback and satisfaction levels, the institution will assess opportunities to refine and adapt the tool, with the objective of providing a seamless and intuitive user experience. These improvements will focus on ensuring that the assistant effectively supports users' needs and integrates naturally into their existing drafting processes.

Based on the experience gained, the institution will also explore opportunities to extend AI-supported assistance to other areas of the legislative process.

Data

The assistant draws on publicly available legal and constitutional sources. Its primary knowledge base includes decisions of the Constitutional Court of Chile – accessed through the publicly available repository of the Library of the National Congress – as well as current Chilean legislation, with particular emphasis on the Constitution.

The assistant uses these authoritative legal sources to support its constitutional analysis and generate recommendations regarding the quorum requirements applicable to individual legislative provisions.

As the application relies exclusively on public legal documents and does not process confidential parliamentary information or personal data, privacy and data sensitivity risks are considered limited.

Architecture and technologies

The application is built on OpenAI's text-embedding-3-large embedding model and uses a retrieval-augmented generation (RAG) architecture to search for and retrieve information from a curated corpus of publicly available legal sources, including the Constitution of Chile, current legislation and regulations, and decisions of the Constitutional Court of Chile. The assistant is available through a bill-drafting system as a built-in feature rather than a free-form chat interface, allowing users to request a quorum assessment without leaving their drafting workflow.

Admissibility assistant

Parliament/chamber: Chamber of Deputies of Chile

Functional area(s): Bill drafting and amendments

Date: July 2026

Overview

The admissibility assistant was developed by the Chamber of Deputies of Chile to support the drafting and review of legislative proposals. Integrated into the legislative drafting platform, it analyses draft legislative provisions to identify potential constitutional issues that may affect their admissibility. Drawing on constitutional doctrine and other authoritative legal sources, the assistant evaluates whether a proposal addresses matters that may be regulated by law and whether it complies with constitutional requirements governing the initiation of legislation, and raises any other constitutional issues that could affect its admissibility. It then provides a structured legal assessment to help MPs and staff identify potential constitutional concerns before a bill is formally introduced.

Business problem

Assessing the constitutional admissibility of legislative proposals – i.e., determining whether a proposed provision complies with constitutional requirements – is a fundamental step in the drafting process within the Chamber of Deputies of Chile. This often requires consulting multiple legal sources including the Constitution, constitutional doctrine, jurisprudence and parliamentary precedents.

This analysis can be time-consuming and requires specialized legal expertise. As a result, the quality and efficiency of the analysis may vary depending on the experience of the drafting team and the time available to conduct legal research.

The Chamber identified an opportunity to provide earlier and more consistent access to constitutional analysis directly within the legislative drafting process: rather than replacing legal judgement, the objective was to support MPs and staff by providing structured legal analysis based on authoritative legal sources, enabling potential admissibility issues to be identified before a bill is formally introduced.

Solution

The assistant was integrated directly into the Chamber of Deputies' legislative drafting platform as a built-in feature. Rather than interacting with a free-form chatbot, users draft a legislative provision and request an admissibility assessment by clicking a dedicated button within the drafting interface.

The assistant reviews each legislative provision individually and provides a structured legal assessment highlighting potential constitutional issues that could affect its admissibility. The analysis considers the constitutional framework governing legislative admissibility and identifies relevant legal principles, constitutional doctrine and jurisprudence that support its conclusions. Where appropriate, it explains the reasons why a provision may be considered admissible or inadmissible and identifies the constitutional rules involved. The assessment is presented while the bill is still being drafted, allowing MPs and staff to review the analysis, refine the proposed text where necessary and identify potential constitutional concerns before the bill enters the formal legislative process.

The assistant is intended to support legal analysis and improve consistency in the drafting process, while leaving all legislative and legal decisions to the users.

Human interaction

The assistant is used by MPs and staff involved in the drafting and review of bills. It is integrated directly into the legislative drafting platform, allowing users to access AI-supported legal analysis within their existing workflow, without requiring a separate tool or process.

The assistant analyses individual legislative provisions and provides a structured assessment of potential constitutional issues that may affect their admissibility, together with the legal reasoning and relevant constitutional sources supporting its conclusions. Human judgement remains essential throughout the process: users review the assistant's analysis, assess its relevance in the broader legislative context and make the final decisions regarding the wording of legislative provisions and the admissibility of the proposal.

Results and impact

The assistant has completed an internal validation process with positive results, including testing by specialized users with legal expertise. As the legislative drafting platform is not yet in production, the assistant has not yet been used by MPs or staff in their day-to-day legislative work. Measurable impacts from real-world adoption are therefore not yet available.

Early validation results indicate that the assistant can provide relevant support by organizing constitutional doctrine, jurisprudence and other authoritative legal sources into a structured admissibility assessment for each legislative provision. The expected impact is a reduction in the time required for legal research and preliminary constitutional analysis during the drafting and review of legislative proposals, enabling parliamentary teams to focus on legislative policy and technical drafting.

Once fully deployed, the assistant is expected to improve the consistency and accessibility of constitutional admissibility analysis by making specialized legal knowledge available directly within the drafting workflow. It is also expected to help identify potential constitutional issues earlier in the legislative process, supporting better-informed drafting while preserving human review and decision-making as essential components of the legislative process.

Risk management

The main risks identified for the assistant relate to the accuracy and reliability of AI-generated legal analysis, as well as the potential for users to rely excessively on automated recommendations in a constitutionally sensitive context.

To address these risks, and given that assessing the admissibility of legislative proposals involves legal interpretation, the assistant was designed as a decision support tool rather than an autonomous decision maker, and it was developed to provide analysis based on constitutional jurisprudence. Its performance was reviewed through an internal validation process involving specialized legal users, ensuring that the assistant's recommendations were consistent with established constitutional principles.

Human oversight remains a fundamental safeguard: users are responsible for reviewing the assistant's analysis, applying their legal judgement and making the final decisions regarding legislative drafting and admissibility.

Finally, privacy risks were considered limited in this use case, as the assistant operates exclusively on publicly available legal sources and does not process confidential parliamentary information or personal data.

Lessons learned

One of the main lessons learned is the importance of designing AI applications around existing parliamentary workflows rather than creating stand-alone tools. Integrating the assistant directly into the legislative drafting platform ensures that it becomes part of users' daily activities and it is used within its intended scope.

Another key learning was the importance of involving legal specialists throughout the development and validation process. Their expertise was essential to evaluate the quality and legal soundness of the assistant's assessments and to ensure that its conclusions were consistent with the constitutional framework governing legislative admissibility.

For other parliaments considering similar applications, an important recommendation is to establish specialized validation processes. In this case, the assistant was built around the constitutional rules governing legislative admissibility, with its outputs systematically reviewed by legal experts to ensure that the analysis remained aligned with the applicable constitutional framework.

Next steps

The next step is to deploy the assistant as part of the Chamber of Deputies' legislative drafting platform and evaluate its performance in real parliamentary workflows. Based on user interaction, feedback and satisfaction levels, the institution will assess opportunities to refine and adapt the tool, with the objective of providing a seamless and intuitive user experience. These improvements will focus on ensuring that the assistant effectively supports constitutional analysis while integrating naturally into existing legislative drafting processes. Based on the experience gained, the institution will also explore opportunities to extend AI-supported legal assistance to other areas of parliamentary work.

Data

The assistant draws on publicly available constitutional and legal sources. Its primary knowledge base is Chapter V of the Political Constitution of the Republic of Chile, which establishes the constitutional framework governing the National Congress and the legislative process. In particular, the assistant relies on Articles 63 and 65, which define, respectively, the matters that must be regulated by law and the rules governing legislative initiative, including matters subject to the exclusive initiative of the President of the Republic.

The assistant uses these constitutional provisions, together with other authoritative legal sources, to assess whether individual legislative provisions comply with the constitutional requirements for admissibility. It also considers constitutional restrictions applicable to certain types of proposals, such as laws granting amnesties or general pardons.

As the application relies exclusively on publicly available constitutional and legal information and does not process confidential parliamentary information or personal data, privacy and data sensitivity risks are considered limited.

Architecture and technologies

The application is built on the GPT-5.2 language model and uses a retrieval-augmented generation (RAG) architecture to search and retrieve information from a curated corpus, with a focus on Chapter V of the Political Constitution of the Republic of Chile and the constitutional rules governing legislative admissibility. The assistant is available through a bill-drafting system as a built-in feature rather than a free-form chat interface, allowing users to request an admissibility assessment without leaving their drafting workflow.

Amendment-writing machine

Parliament/chamber: Chamber of Deputies of Italy

Functional area(s): Bill drafting and amendments; MP support and research

Date: June 2026

Overview

The Chamber of Deputies has developed an AI-assisted amendment-writing machine to support MPs and parliamentary group staff in drafting amendments to bills under consideration.

The application enables users to work directly on the text of a bill, which is automatically structured into its constituent legislative elements: articles, headings, paragraphs, subparagraphs, and provisions amending existing legislation.

Rather than drafting amendments from scratch, users can select the relevant provision and specify the intended modification – e.g. inserting, deleting or replacing words, sentences or entire provisions. The system then generates the corresponding amendment in accordance with the formal drafting rules and parliamentary syntax of the Chamber of Deputies.

The application incorporates a generative AI chatbot that provides access to official documentation related to the bill, including research dossiers prepared by the Parliamentary Research Service and memorandums submitted during committee hearings. The application has been operational in a production environment since July 2024 and is currently used by MPs and parliamentary group staff.

The inception and development of the first version of this tool was a joint project with the Italian Senate, in the context of the “Polo Informatico Parlamentare”, a common initiative fostering inter-chamber cooperation and projects (more information can be found [on the website of the Senate of Italy](#), and use cases can be found [on the IPU website](#)).

Business problem

The drafting of amendments is a core component of parliamentary activity. It requires not only political intent but also a high degree of technical precision. Amendments must comply with the formal drafting rules and parliamentary conventions of the Chamber of Deputies.

This can be a demanding task for MPs and parliamentary group staff, particularly when dealing with lengthy or technically complex bills. A further challenge lies in accessing and analysing the information necessary to draft well-informed amendments, which is typically contained in research dossiers and memorandums submitted during committee hearings.

Solution

The tool is integrated into the Chamber of Deputies’ digital legislative environment. It automatically breaks down the text of a bill into its constituent legislative elements. Users can create amendments by intervening on entire legislative provisions or by working directly on the text of the bill.

The application then generates the corresponding amendment proposal using the formal drafting rules and parliamentary syntax of the Chamber. The application also incorporates a generative AI chatbot providing access to official documentation related to the bill, including research dossiers and memorandums submitted during committee hearings.

Human interaction

The application supports users throughout the amendment-drafting process, while leaving full responsibility for the content, validation and submission of amendments to MPs and

parliamentary group staff. The substance and policy objective of the proposed amendment remain entirely under the user's control.

The chatbot provides answers based exclusively on official documentation related to the bill and includes references to the underlying source documents. Human judgement remains central throughout the process.

Results and impact

The application has made the amendment-drafting process more accessible and efficient for MPs and parliamentary group staff. By allowing users to work directly on the text of a bill, it enables the rapid generation of formally compliant amendment proposals without requiring users to draft the wording from scratch. This reduces the risk of drafting errors, improves consistency in the formal presentation of amendments and streamlines the overall drafting process.

Risk management

The principal risks concern the reliability of chatbot answers. Generated amendments are presented for review and may be edited before confirmation. The application supports drafting but does not submit amendments autonomously. The chatbot is designed to work only on the official documents linked to the bill; any other questions are blocked by a robust guard system.

Lessons learned

A key lesson is that AI delivers the greatest value when embedded within a specific parliamentary workflow.

A second lesson is the importance of combining drafting assistance with seamless access to authoritative parliamentary documentation.

The experience also highlights the importance of maintaining meaningful human control, with substantive choices and political responsibility remaining entirely with MPs and parliamentary group staff.

Next steps

Future developments may focus on enhancing the system's ability to assist users in identifying the legislative provisions most relevant to a specific policy objective. Another promising area is the automated generation of consolidated legislative texts incorporating amendments adopted during the legislative process.

Data

The application draws on legislative texts and parliamentary documentation already available within the Chamber of Deputies' information systems. The primary source is the bill under consideration, acquired in Extensible HyperText Markup Language (XHTML) format and automatically transformed into a structured representation. The application also uses official parliamentary documents including research dossiers and memorandums submitted during Committee hearings.

Architecture and technologies

The application is implemented as a .NET web application with a modular service architecture and Oracle-based persistence layer. The amendment-generation service is implemented separately in Python and exposed through application programming interfaces (APIs). The solution is containerized and deployed within the Chamber's Kubernetes environment, allowing integration with existing parliamentary systems and workflows.

Legislative amendment grouper

Parliament/chamber: Chamber of Deputies of Brazil

Functional area(s): Bill drafting and amendments; Committee work

Date: May 2026

Overview

At the Brazilian Chamber of Deputies, legislative consultants support MPs throughout a bill's passage, from committee stage until it reaches the plenary session. During the committee stage, many amendments are presented for each bill, and legislative consultants are responsible for analysing them and advising MPs on whether to accept or reject each one.

The Ulysses legislative amendment grouper is an AI system that automatically groups parliamentary amendments by semantic similarity to facilitate this analysis, reducing the manual effort required for parliamentarians and legislative consultants to identify related or redundant amendments in large volumes of legislative bills.

Business problem

Budget bills must be processed within a short time frame set by the National Congress. During this period, thousands of amendments are tabled by parliamentarians for each bill. The MP responsible for the bill, and the legislative consultant supporting them, face a significant challenge in deciding whether each amendment should be accepted or rejected, and in providing justification for that decision. Although some amendments may appear similar, differing by only a single word or even a comma, they can represent opposing ideas. On the other hand, accepting some amendments means rejecting others, even if they do not relate to the same article. This complex analysis, carried out under significant time pressure, presents a high risk of error, which increases the pressure on both the legislative consultant and the MP.

Solution

This Ulysses module was developed to automatically group parliamentary amendments by semantic similarity, helping parliamentarians and legislative experts analyse related proposals efficiently through visualizations such as word clouds and spatial distributions. The system allows users to adjust the number of groups, assign amendments to predefined groups and provide feedback to continuously improve performance over time. In summary, the AI system:

- automatically reads and interprets proposed amendments, and groups them using semantic similarity analysis
- provides intuitive graphical and list-format visualizations of amendment groups, including word clouds and 3D spatial distributions, allowing users to quickly understand relationships between amendments, adjust grouping parameters and assign amendments to predefined thematic groups
- establishes a feedback-based learning mechanism that captures user insights on AI-generated groups, enabling iterative refinement of semantic analysis algorithms and ensuring progressively more accurate categorization aligned with specialized legislative knowledge

Human interaction

As soon as the legislative digital platform receives an amendment, that amendment is read and interpreted by the AI system. In a parallel interface available only to legislative consultants, all amendments are presented alongside their respective bill. Consultants can adjust parameters to reorganize the original groups as the analysis evolves, and can create

and assign labels to each group. This ensures human oversight and accountability throughout the process.

Results and impact

Hundreds of parliamentary amendments can be analysed for a single bill, allowing the legislative consultant to issue a reliable opinion to the parliamentarians involved in the discussions – work that previously took teams several days can now be done in minutes, enabling simulations.

In addition, conflicting amendment groups can be systematically identified, and the analysis process is standardized, thus reducing bias.

Risk management

The main risk was that users (i.e. legislative consultants) would not accept the system. This was addressed through a strong engagement strategy, gathering their feedback quickly so that they could see their requested adjustments implemented within a short time frame.

Lessons learned

The final adjustments to the AI system were implemented relatively quickly because the manager of the Consultant Unit directed all consultants to test the system thoroughly and provide feedback. This broad engagement made the difference.

Next steps

The Chamber of Deputies is open to supporting other parliaments that want to learn more about this work.

Amendment vote-ordering and similarity-detection system

Parliament/chamber: Chamber of Deputies of Italy

Functional area(s): Bill drafting and amendments; Committee work

Date: June 2026

Overview

The Chamber of Deputies has developed an AI-based application to support the management of amendments throughout the legislative process. The system is fully integrated into the online drafting environment and assists staff in two key activities: automatically determining the order in which amendments must be considered and voted upon, and identifying amendments that are identical or substantially similar in content.

The system combines parliamentary procedural rules with artificial intelligence techniques capable of analysing amendment texts and identifying both textual and semantic similarities. Based on the characteristics of each amendment and its relationship with the bill under consideration, the application automatically generates the proposed voting order and highlights groups of identical or similar amendments.

The information produced by the system supports several parliamentary services. Committee and Assembly secretariats use it to prepare amendment dossiers, and the Parliamentary Research Service and the Budget Service use the identified relationships between amendments when assessing procedural and financial admissibility. MPs also benefit indirectly through the official amendment dossiers generated by the parliamentary information system.

The system has been in a production environment since 2021 and has become an established component of the Chamber's legislative support infrastructure. While it automates a number of complex analytical and procedural tasks, all legal, procedural and financial decisions remain entirely under human responsibility.

The inception and development of this tool was a joint project with the Italian Senate, in the context of the "Polo Informatico Parlamentare", a common initiative fostering inter-chamber cooperation and projects. The Senate of Italy developed a similar tool for amendment management, launched in 2016 (more information can be found [on the website of the Senate of Italy](#), and use cases can be found [on the IPU website](#)).

Business problem

Managing amendments is one of the most demanding activities in parliamentary legislative work. During the examination of a bill, committees and the plenary may receive hundreds or even thousands of amendments submitted by MPs and parliamentary groups. These amendments must be processed according to complex procedural rules governing their voting order. At the same time, staff must identify amendments that are identical or substantially similar, as these may require consistent procedural treatment and comparable admissibility assessments.

Traditionally, both activities relied heavily on manual analysis by staff. This process was time-consuming and could easily lead to human errors, especially when dealing with large volumes of amendments.

The Chamber of Deputies identified an opportunity to support staff through an AI-based solution capable of automatically calculating the voting order of amendments, and identifying

identical and similar amendments, preserving full human control for procedural and legislative decisions.

Solution

The Chamber of Deputies developed an AI-based application integrated into the parliamentary drafting and legislative platform. The application analyses the text of amendments in relation to a bill and automatically performs two tasks:

- Generating the voting order of amendments according to procedural parliamentary rules by combining parliamentary procedural rules with information extracted from the amendment text, automatically classifying amendments according to their legislative effect and applying the relevant ordering criteria
- Identifying proposals that are identical or highly similar

The application supports committee secretariats, assembly staff, the Parliamentary Research Service and the Budget Service, providing a common information base for legislative and procedural activities.

Human interaction

The application operates as a decision-support tool and is fully integrated into existing parliamentary workflows. Staff review and validate the results. They have the possibility, through dedicated functions, to manually correct errors or imperfections that the system generates.

The Parliamentary Research Service and the Budget Service remain fully responsible for all admissibility assessments. Human oversight remains an essential component of the process.

Results and impact

The application has significantly improved the efficiency of amendment management within the Chamber of Deputies.

Tasks that previously required extensive manual review can now be completed more quickly and consistently: staff can process large volumes of amendments while reducing the risk of identical or related proposals being overlooked.

The system has also strengthened consistency in admissibility assessments by making equivalent amendments immediately visible to the services responsible for procedural and financial evaluations.

Since its introduction in 2021, the system has become a routine and trusted component of the Chamber's legislative support infrastructure.

Risk management

The principal risk identified was the possibility that the application might provide incorrect output. This could happen especially in the early stages, when amendments have not yet been drafted and may contain spelling errors. To address this risk, the system was designed as a support tool that can be re-launched at different stages of the amendment life cycle. All outputs must be reviewed and validated by staff before being used in official parliamentary documentation or proceedings.

Lessons learned

A key lesson learned from the project is that AI can provide significant benefits when applied to well-defined parliamentary processes governed by clear procedural rules. Rather than attempting to automate complex legal or political judgement, the project focused on supporting activities that require extensive analysis and repetitive processing. This approach facilitated user acceptance and allowed measurable operational benefits to be achieved without altering established parliamentary procedures.

Next steps

Future developments may focus on enhancing the system's ability to identify not only identical and textually similar amendments, but also amendments that pursue the same legislative objective while using different wording. Future research may explore more advanced semantic analysis techniques capable of understanding the underlying legislative intent of a set of amendments.

Data

The application processes legislative documents already available within the parliamentary information system. Its primary inputs are amendment dossiers containing the text of amendments and related metadata as well as the bill to which the amendments refer. The system operates exclusively on parliamentary legislative documents and does not require access to external personal data sources.

Architecture and technologies

The application is integrated with the Chamber of Deputies' drafting environment and is exposed as a REST-based service. The architecture combines rule-based processing and machine-learning techniques. The solution uses neuro-linguistic programming (NLP) techniques and semantic language models based on word embeddings, implemented as a Python-based service using Flask, SpaCy, Scikit-Learn, Gensim, Pandas and NumPy, deployed in a containerized environment communicating through REST application programming interfaces (APIs).

Committee work and plenary proceedings

Bill-path determination assistant

Parliament/chamber: Chamber of Deputies of Brazil

Functional area(s): Committee work

Date: May 2026

Overview

At the Brazilian Chamber of Deputies, each bill is presented by an MP. The Clerk's staff therefore carry out a complex analysis before determining the following:

- Whether the bill meets the necessary constitutional conditions to follow the legislative process in committees
- All other bills that should be annexed to the bill (owing to similarity)
- The type of procedure applicable to the bill in the committees
- The priority that will be applied to the bill
- The sequence of committees that the bill should follow

This AI system performs all five analyses for the Clerk's staff before each step of a bill's passage through the committees.

Business problem

This Ulysses module uses AI to support the analyses carried out by the Clerk's staff (the body responsible for the legislative process in the Chamber of Deputies) on each bill received in full text, following the work of legislative consultants and presentation of the final draft by the parliamentarian. These analyses indicate whether the bill meets the conditions required to proceed through the legislative process in committees, as well as determine the terms of that process.

Before the system was implemented, the process of distributing propositions to thematic committees – known as “initial dispatch” – required extensive manual and technical analysis by the Clerk's staff. This activity requires deep knowledge of the legislative process and, although it is crucial to the institution's work, is susceptible to inconsistencies, delays and work overload. Furthermore, the growing volume of bills makes the process increasingly challenging in terms of timeliness and accuracy.

The number of bills presented has increased significantly since 2025: in one case, a single MP presented 400 bills in one month.

Solution

Once each bill is officially submitted to the digital legislative platform, an AI system performs five analyses with the following aims:

- Determining whether the bill meets the constitutional conditions required to follow the legislative process in committees, exposing any possible constitutional conflicts, together with their justification
- Identifying all other bills already in progress in the Chamber of Deputies that should be annexed to the bill in question, whether owing to similarity or because they complement each other within the same proposal; rather than being processed independently, they can be attached to the bill under analysis
- Suggesting the precise type of procedure applicable to the bill in the committees, i.e. whether, owing to its characteristics or under the applicable rules, the bill must first be evaluated by thematic committees before being voted on in plenary, or whether it can become law once it has a conclusive opinion from the thematic committees through which it must pass:

- bill subject to conclusive discussion by committees
 - bill subject to discussion in plenary
 - bill subject to internal discussions in committees
- Suggesting the priority to be applied to the bill as it moves through the committees (ordinary time, priority (urgent) or special) according to the Chamber's rules of procedure
- Suggesting the sequence of committees that the bill should follow, taking into account the complex blend of subjects addressed in the bill and the previous analyses

Users select which suggestions to accept and can edit those that are not complete, so that the digital legislative platform uses the final version to produce a digital dispatch, signed by the Clerk, setting out the bill's entire path from the first committee to the plenary discussion.

Human interaction

MPs may or may not use consultants to support the drafting of the bill. Regardless of the approach taken, the MP adds the bill to the digital legislative platform. Once the bill is there, the Clerk's staff begin a set of complex analyses that will shape the path the bill will follow.

Results and impact

The stand-out benefits observed to date include the following:

- Increased operational efficiency of the Clerk's team
- Reduction of the initial processing time for bills, contributing to greater agility in the legislative process as a whole
- Greater precision in the distribution of bills, based on consistent technical criteria
- Support for decision-making by technical teams, which can count on automated recommendations while retaining their role as supervisors and validators of the process
- Standardization of actions regardless of the origin of the bill
- Process scalability, allowing the volume of bills to keep pace with growing demand (in 2026, there has been an increase of approximately 40% in the number of proposals compared with the same period in 2025) without compromising on quality or distribution deadlines

Risk management

The high level of complexity and the need for deterministic analysis meant a single type of AI could not handle the task. The system was split into multiple modules, according to the type of AI used.

Lessons learned

The business unit (the Clerk's staff) joined forces with the IT team. The Clerk's office provided high-level specialists with expertise of the legislative process and the Constitution to take part in this work. This decision made a difference.

Next steps

The Chamber of Deputies is open to explaining to other parliaments how this AI system was developed.

Automatic voice-to-text transcription of plenary and committee sessions

Parliament/chamber: Senate of Spain

Functional area(s): Plenary proceedings and transcription

Date: May 2026

Overview

The Senate of Spain uses this AI system to generate transcripts of plenary and committee sessions. These transcripts are displayed in real time on the Senate website and serve as a draft for the official proceedings record. The draft transcript is then manually reviewed and replaced on the website with the revised official text.

The transcribed text is synchronized with the video displayed on the website, allowing users to search the text and access the specific moment in the video where the speech was made.

Business problem

Previously, the transcript was not available on the website until the day after the parliamentary session. The editors of the official record also had to transcribe the entire text from the audio recordings.

Solution

The system transcribes speech to text in near-real time. It has been fine-tuned by training it with data from previous sessions to avoid errors such as specific parliamentary jargon or the names and surnames of members of the Senate. Editors can now correct errors based on the AI-generated draft.

Human interaction

Although the draft transcript is displayed immediately, it is manually reviewed and replaced on the website with the revised official text the following day.

Results and impact

Citizens can now view and search the transcripts of parliamentary sessions as they are taking place. The work of drafting the minutes has changed completely. Previously, the bulk of the work was transcription. Now, most of the time is spent on revision, drafting and copy-editing.

Risk management

Initially, the transcripts were not very accurate because a general-purpose Spanish speech-to-text engine was used. The engine therefore had to be fine-tuned with data from previous sessions. This has helped to avoid errors with specific parliamentary terminology or the names of Senate members.

Additionally, the text of the automatic transcripts has to be clearly marked on the website to make it clear that they are drafts that will not be official until the following day.

Lessons learned

It is important to perform fine-tuning of the engine to avoid errors with specific words related to the Chamber.

Next steps

The system is being expanded to perform automatic speaker recognition and to automatically split the video for each intervention.

Data

Plenary transcripts serve as the input data.

Architecture and technologies

The automatic speech recognition system transforms audio signals into structured text. It is designed using a hybrid approach that combines deep learning models, inference frameworks and proprietary modules from the vendor, as detailed below:

- **Base models (OpenAI Whisper):** The system uses the Whisper Large v3 model, which employs a transformer encoder-decoder architecture. This model is open-source and licensed under the Apache License 2.0.
- **Execution backend (CTranslate2):** The models are not executed using the standard PyTorch environment, but rather using CTranslate2. This is a specialized inference engine for transformer models written in C++ and Python. It implements advanced techniques such as weight quantization, neural-layer fusion and cache optimization, maximizing throughput on both CPU and GPU and drastically reducing response latency. This framework is open-source and licensed under the MIT License.
- **Diarization integration:** The backend optimizes audio processing. Instead of processing the entire audio linearly, it takes as direct input the time segments previously defined by the diarization pipeline. This avoids sequential processing as well as the unnecessary transcription of silent or noisy segments.

Document and records management

Content categorization according to the EuroVoc thesaurus

Parliament/chamber: Senate of Spain

Functional area(s): Document and records management

Date: May 2026

Overview

Documentalists at the Spanish Senate use an AI-based decision-tree system to categorize documents according to the EuroVoc thesaurus and other internal standard categorization systems.

The categorization process, which was previously done manually, is now performed using a series of categories proposed by the AI system. This reduces categorization time, improves quality and prevents errors. Human supervision is performed on all categories suggested by the system.

Business problem

The categorization work was carried out by several interns and reviewed by documentalists, which produced bottlenecks and led to a lack of quality and homogeneity in the results.

Solution

A system has been implemented that can suggest EuroVoc thesaurus labels for a given text or document. This allows staff to understand the subject matter of the text.

Human interaction

During the document cataloguing process carried out by the documentalists, the AI system proposes a series of EuroVoc categories/labels for validation. The documentalist can confirm the categories proposed by the AI, and remove them or add their own.

No category is assigned to a document if there has been no prior human review.

Results and impact

Work that used to take days can now be done in hours thanks to this tool. The quality of the work is also higher because the AI system's category suggestions allow for standardized labelling criteria across all users; whereas previously, the criteria varied widely and depended on the individual.

Risk management

There have been several cycles of trial and error to identify the optimal number of categories proposed by the tool. Problems have also been detected in the training data. These issues have introduced noise and biases into the model because many different people have been working on this over the years, and each used their own personal criteria for categorization, resulting in non-homogeneous training data.

Lessons learned

Choosing the AI algorithm best suited to the needs was costly, and initial approaches failed to accurately predict categories owing to insufficient training data. The classic AI approach for classifying EuroVoc thesaurus subjects was not viable without adequate training data for each category/descriptor. On the other hand, the approach using only a large language model was unsuitable for the specific requirements in this case. The most effective solution was a hybrid AI model that combined the best of both worlds. Ultimately, a hierarchical

classification approach was employed, combining the use of embeddings (using the RoBERTa-base-bne language model) with a threshold descent strategy that utilizes 358 models to classify 1,339 EuroVoc descriptors.

Next steps

The Senate intends to use the same system for all documents and texts used in the Chamber – not just for the collections for which it was initially designed.

Data

The model has been trained and optimized for the categorization of parliamentary questions, but it can be extended to all types of documents.

Architecture and technologies

First, the text is embedded using the RoBERTa-base-bne model. This is a RoBERTa-based model trained on 570 GB of data from the National Library of Spain. The model consists of 12 layers, returns a 768-dimensional vector and has a total of 125 million parameters. A more detailed description of this model and how it was trained can be found [on this webpage](#), while [this paper](#) describes the model in more detail.

These embedded models are then used with Extreme Gradient Boosting (XGBoost), an open-source machine learning library that utilizes gradient-boosted decision trees. The documentation can be found [on this webpage](#).

The models were trained on a dedicated set of 244,698 questions labelled by Senate documentalists over several legislative sessions. The questions were labelled using the EuroVoc thesaurus, the INE¹ thesaurus and a thesaurus developed by the Senate itself. More than one descriptor was used for each question. As a result, the models are multi-class and multi-label, meaning they can predict different descriptors and predict one or more descriptors.

These models consist of 100 decision trees. Instead of training a single model that can predict all descriptors, several models were trained, each capable of classifying only a limited number of them. This was done by taking advantage of the hierarchical structure of thesauri; each descriptor has a parent descriptor and different child descriptors, so a model was trained for each descriptor to classify its child descriptors.

¹ [Instituto Nacional de Estadística](#) (National Statistics Institute)

Generative AI and AI-assisted development for legislative tracking

Parliament/chamber: Chamber of Deputies of Italy

Functional area(s): Document and records management; Public engagement and accessibility; IT infrastructure and cybersecurity

Date: June 2026

Overview

The Chamber of Deputies has implemented an AI-powered system to collect, organize and analyse legislative data.

The system automates the aggregation and analysis of information from diverse sources, providing a public website where citizens and experts can monitor the legislative process, the status of delegated legislation and the implementation rate of parliamentary directives issued to the Government.

The portal integrates a conversational AI chatbot named “Norma”, which interprets natural language questions and dynamically queries databases to generate illustrative tables and graphs. The system also generates periodic reports in print-ready formats, leveraging generative AI to provide textual descriptions and analyses of the underlying charts and tables.

The platform was developed using AI-assisted development paradigms (commonly referred to as “vibe coding”), which significantly accelerated the production of application components and facilitated the rapid creation of functional prototypes.

The platform is scheduled for full deployment to the production environment in October 2026.

Business problem

Modern parliamentary administrations generate vast amounts of heterogeneous data across the life cycle of a bill. Tracking this data presents a dual challenge: information regarding legislative production and executive compliance is often scattered across distinct, non-communicating databases, requiring intensive manual labour by technical staff; and while legislative data is technically public, its dense technical structure creates a barrier for citizens, journalists and non-expert researchers.

This scenario presents a unique opportunity to transform raw, siloed administrative data into actionable public knowledge.

Solution

To address these challenges, the Chamber of Deputies has developed an integrated, AI-driven data ecosystem that automates the collection, organization and public dissemination of legislative analytics.

The core components include: an “extract, transform, load” (ETL) pipeline that unifies heterogeneous sources for internal monitoring of the Observatory on Legislation and Public Policies; the Norma conversational interface, a public-facing web portal integrating an advanced AI chatbot that allows users to query complex legislative tracking databases using natural language; and a generative reporting engine that uses large language models (LLMs) to automatically generate textual descriptions and narrative analyses of data trends.

Human interaction

Human–AI interaction is architected around two distinct modalities. For internal institutional users, the system automates data aggregation and drafts preliminary narrative analyses for periodic reports, leaving qualitative validation, procedural assessment and final publication entirely under human responsibility. For external users – including citizens, journalists and researchers – interaction is facilitated through the Norma chatbot, which interprets user intent and dynamically outputs structured tables and charts.

Results and impact

Although scheduled for full production release in October 2026, the platform has already delivered substantial operational impacts. The development team achieved a major boost in engineering velocity through AI-integrated development paradigms, allowing for the rapid deployment of functional prototypes and streamlined validation loops with diverse stakeholders.

The system has successfully automated the extraction and harmonization of fragmented data from heterogeneous repositories, significantly reducing the manual overhead previously required by the Observatory.

Risk management

To ensure reliability within a high-stakes constitutional environment, the project implements a risk-management framework centred on human oversight and architectural transparency. A strict human-in-the-loop paradigm is enforced: validation of aggregated legislative data and review of narrative commentary remain strictly within human competence.

The Norma chatbot includes a transparency feature (“*Cosa ho fatto*”, or “What I did”) that explicitly details the step-by-step logic and database queries used to synthesize each response.

Lessons learned

The adoption of AI-assisted development tools proved highly effective, demonstrating that a lean team can rapidly build complex, production-ready enterprise applications when augmented by generative engineering tools. A key lesson is that technological velocity must be coupled with continuous stakeholder alignment. The iterative co-design process with domain experts from the Documentation Services and parliamentary oversight bodies proved essential for translating nuanced parliamentary procedures into precise algorithmic rules.

Next steps

Following the production launch in October 2026, the platform’s road map will focus on data expansion, deeper open-data integration and infrastructure optimization. The next phase will involve expanding the data ingestion pipeline to include new institutional data sets. The Norma chatbot will be significantly upgraded, with future efforts focusing on migrating the LLM infrastructure on-premises.

Data

The system ingests and processes granular data sets tracking the life cycle of legislative activity and executive compliance, including: enacted laws and evolution of those laws; legislative process metrics; amendment tracking; voting records; decree-laws and conversion legislation; delegated legislation; and implementing measures and government compliance tracking.

Architecture and technologies

The system architecture includes a decoupled multi-node backend with dedicated REST application programming interfaces (APIs) for internal administrative requests and public website queries deployed on independent server nodes.

The presentation layer is split into an internal management system for parliamentary officials and a public web portal for civic engagement.

The conversational layer powering Norma integrates Minerva, an advanced LLM developed by Sapienza University of Rome and trained on Italian institutional, cultural and legal data sets.

Automated translation of legislative and regulatory acts

Parliament/chamber: Verkhovna Rada of Ukraine

Functional area(s): Bill drafting and amendments; Document and records management; Public engagement and accessibility

Date: June 2026

Overview

The Verkhovna Rada of Ukraine is developing and implementing an AI-based automated translation system for legislative and regulatory acts. The application supports the process of making Ukrainian legislation accessible in foreign languages, primarily in the context of European integration and the alignment of the legal framework with European Union (EU) standards. The system is designed to improve the efficiency of the legislative process and to enhance the international transparency of parliamentary work.

Business problem

Against the backdrop of an active European integration process and a situation of martial law, the Verkhovna Rada of Ukraine is facing a significant increase in the volume of legislative and regulatory acts requiring translation into foreign languages – primarily English, as the principal language of communication with EU institutions and international partners.

The traditional translation process is highly labour-intensive: it requires a substantial number of qualified translators, considerable processing time and continuous alignment of specialized legal terminology. Given that parliament adopts dozens of laws each month – including as part of harmonization with the EU acquis – ensuring timely, high-quality translation manually has become virtually impossible. A further challenge is maintaining terminological consistency across different acts: the same legal term may be rendered differently by different translators, which hampers foreign partners' understanding of legislation and reduces the quality of international legal engagement.

Solution

A foundational tool for the automated translation of texts and documents has been deployed and is in active use by staff of the Verkhovna Rada Secretariat in their day-to-day work. The system enables rapid translation between Ukrainian and English in both directions. Users interact through a straightforward web interface featuring two panels: one displaying the source text and the other the translation, allowing the result to be reviewed immediately and edited as needed. A dedicated file translation function supports “.txt” and “.docx” formats, enabling staff to work with complete documents rather than isolated text fragments.

Following automated translation, users may further refine the output using an intelligent post-editing tool that clarifies phrasing and improves stylistic coherence. The system also draws on a shared legal terminology glossary, which an administrator can update to maintain consistency across translations. All translated files are stored in a cloud environment and accessible to users as a searchable list, allowing previous translations to be retrieved and reused.

Human interaction

The AI system performs the initial translation of documents automatically, without human involvement. However, the final decision on quality and publication always rests with a

human – a qualified translator-editor within the Secretariat. The translator-editor reviews the machine-generated output, makes corrections where terminology is ambiguous or contextually complex and approves the final version. The system does not publish any translation without human sign-off. In this way, AI handles the routine, labour-intensive work, freeing subject-matter experts to focus on substantive editing and terminological standardization.

Results and impact

At the pilot stage, the following outcomes are anticipated:

- Significant reduction in translation turnaround time – from several days to a matter of hours – through automation of the initial translation
- Improved terminological consistency across legislative and regulatory acts, enabled by a centralized glossary and a single underlying language model
- Reduced workload for Secretariat translation staff, allowing resources to be redirected to the most complex and high-priority documents
- Greater accessibility of Ukrainian legislation for international audiences, particularly EU institutions in the context of the accession process for Ukraine
- Development of a structured translation repository serving as a reference resource for future translations and terminological research.

The project also directly supports Strategic Goal 15 – the approximation of Ukraine to the EU – by enabling the technological integration of legislative information with EU information systems.

Risk management

The following risks and mitigation measures have been identified:

- Accuracy of legal text translation: Legal documents contain highly specialized terminology and complex syntactic structures. The system uses a language model fine-tuned on a corpus of Ukrainian legal texts and EU acquis translations, and mandatory review by a translator-editor provides an additional safeguard against errors.
- Confidentiality and data protection: Processing of legislative texts takes place within a secure cloud environment in accordance with the Verkhovna Rada's cybersecurity requirements and NIS2 standards. Documents with restricted access classification are excluded from automated processing.
- Overreliance on automated outputs: A mandatory human editing procedure has been implemented. The system is configured to flag low-confidence translation segments for priority attention by the editor.
- Terminological divergence from EU acquis: The terminology database is regularly updated in line with current EU glossaries.

Lessons learned

During the planning and piloting phases, it became clear that the key determinant of success was not so much the choice of base language model as the quality and volume of specialized legal data used for fine-tuning.

Parliaments considering similar projects are advised to proceed as follows:

- Begin by building a high-quality terminological glossary before technical implementation of the system.
- Involve translators and legal experts in system design from the outset – their practical knowledge is critical to defining priorities and functional requirements.
- Plan for phased implementation, testing initially on a single document type (e.g. primary legislation) and only scaling to other categories once quality has been validated.

- Avoid attempting to remove the human from the process entirely – the editor’s role changes but does not disappear, and this must be built into the operating model from the start.

Next steps

Following the successful completion of the pilot on primary legislation, the system is planned for extension to secondary legislation, resolutions and international agreements. In the medium term, integration with the open data portal via the “Legislation of Ukraine” application programming interface (API) is envisaged, to enable automatic publication of translations alongside original texts. Consideration is also being given to sharing this experience with parliaments of other EU candidate countries within the framework of bilateral parliamentary cooperation.

Data

The system draws on the texts of legislative and regulatory acts from the Legislation of Ukraine database in machine-readable format. The corpus consists of existing EU acquis translations, legal terminology glossaries (Ukrainian–English) and audit logs of translation edits for ongoing model fine-tuning.

Documents with restricted access classification are excluded from automated processing.

Architecture and technologies

The solution is built on cloud infrastructure (the AWS platform is under consideration in line with the Verkhovna Rada’s overall cloud strategy), utilizing large language models (LLMs) fine-tuned on legal texts. The system integrates with the automated database and the open data portal via the “Legislation of Ukraine” API to receive and publish translations. Integration with the Verkhovna Rada’s unified automated document management system is also planned.

MP support and research

Ulysses Chat

Parliament/chamber: Chamber of Deputies of Brazil

Functional area(s): Committee work; Document and records management; MP support and research; HR and internal administration

Date: May 2026

Overview

Ulysses Chat is a digital assistant based on generative AI platforms from major technology providers. It was developed by the Brazilian Chamber of Deputies to answer questions, using institutional knowledge as context. It draws on the Chamber's official legislative and administrative data and documents, following an extensive data curation process to ensure accurate interpretation of rules and production of reliable results, under an AI governance framework. Available on desktop and mobile, it responds to text or voice commands.

Business problem

The Brazilian Chamber of Deputies intranet has about 12,000 users, including parliamentarians, the Chamber's staff and outsourced collaborators, who need information analysis and research across a wide variety of dimensions and subjects. Many of these internal users started using AI on their own, generating similar data and documents regularly without formal editing or curation.

A second risk was that generative AI accounts could potentially be linked to internal roles and positions, allowing tech providers to map individual users' behaviour and return biased responses based on their personal profile, regardless of their actual work tasks.

A third risk was the lack of transparency and traceability regarding the prompts used to support decisions and tasks.

Solution

Ulysses Chat combines several generative AI platforms with multiple internally developed AI agents that communicate with one another. It was built under AI governance rules, following extensive data curation work – carried out by individual business units, each responsible for its own data – to ensure data quality across a wide range of internal data sets, so that users no longer need to upload official data and documents themselves. All data and documents were incorporated through a data governance process designed to identify data quality issues and the adjustments needed. To this end, data owners and data stewards from each business unit tested, over multiple cycles, the business rules and information embedded in each chat interpretation.

In addition, this approach delivers a significant reduction in token costs compared with a scenario in which each of the more than 12,000 users had to repeatedly upload practically the same documents and data.

One of the product requirements was to avoid easily associating internal professional users' credentials with their personal accounts (for those users who also have a personal generative AI account outside the Chamber). This requirement was met by storing each user's credentials internally, so that people are not exposed to the AI platforms that support Ulysses Chat.

Ulysses Chat provides an AI interaction environment with agents specialized in the following areas:

- The Chamber's services, covering the operating rules of, and activities carried out by:

- units providing internal services and services to citizens across a wide variety of business areas (including the radio station, TV station, news agency, social media, police department, medical department, engineering, building administration, finance and accounting, procurement and contract oversight, library, archive, museum, research, training and capacity-building, information technology, auditing, legal affairs, human resources and events)
- on the legislative side, units supporting parliamentary offices, legislative committees and consultancies assisting in the drafting of laws, together with units supporting legislative procedural rites in both committees and plenary sessions
- Internal legislation, internal standards and official bulletins and manuals
- Information about parliamentarians' activities in the legislative process, including speeches, votes, authorship and reporting of bills, attendance, and trips and expenses
- Legislative bills, including their full content, proposed amendments, substitutes, attachments and orders issued at each stage of the legislative process
- Committees (permanent, temporary, ongoing and closed), including their meeting histories, witnesses and reports, from establishment to closure
- Information about the location and contact details of people within the organization chart
- News produced by internal communication channels
- Contracts established with the Chamber

Human interaction

Immediately after Ulysses Chat launched, a wide set of presentations was made to business units to identify ideas for improvement. The platform is evolving and continuously incorporating new agents and functionalities. The identified benefits include the following:

- Cost savings through the optimization of information traffic to and from the platforms that support Ulysses Chat
- Reduction in the time users would otherwise need to search for data and to develop and continually update prompts for institutional data, including correlating information on topics they are not yet familiar with
- More reliable results through retrieval-augmented generation (RAG) performed under specific rules agreed with business managers and data managers – this is achieved through data governance, in which data managers and curators, working with the Personal Data Protection Officer and the IT unit, take part in various AI governance processes

Results and impact

The implementation journey involved a complete review of the quality of data and documents, bringing together the IT unit and business units. Several agreements and awareness meetings were held first to secure the sponsorship needed to engage the many people affected by the changes that were required to make the product reliable.

Successive testing cycles brought together more than 150 people from dozens of business units to assess the quality of responses (correctness, completeness and relevance). Their feedback guided many actions to improve data quality and the functioning of the AI agents. These evaluation cycles also increased the institution's maturity in building more robust data and document management processes.

Other, specific tests were also carried out, this time involving groups of specialized people. These tests included the following:

- Ethical tests, covering privacy problems, discrimination, emotional involvement, offensive speech and the amplification of vulnerability conditions

- Usability and accessibility tests, covering voice problems, loss of limb mobility, and total or partial vision impairment
- Robustness and cybersecurity tests
- Plain language tests

Following these tests, a large pilot was conducted over four weeks, expanding the sample of testers to 250 people and moving into the real environment (Camaranet, the Chamber's intranet). Access to that environment was restricted to those authorized for testing.

The product was then launched through several events aimed at audiences grouped by role, so that the examples used were relevant to each context. In compliance with the Chamber's AI governance policy, Ulysses Chat is systematically monitored through a continuous process in which each piece of user feedback is classified and analysed.

Lessons learned

The strong partnership between the business units and the IT unit – involving about 250 people in total – made a difference. An established AI governance policy and a data governance policy offered a solid foundation for many of the requirements.

Next steps

The Chamber of Deputies is open to providing further information to other parliaments that are interested in this work.

Argumentative assistant

Parliament/chamber: Chamber of Deputies of Chile

Functional area(s): MP support and research

Date: July 2026

Overview

The Chamber of Deputies of Chile has developed the Argumentative Assistant to support parliamentary debate and legislative preparation.

The application generates structured legal arguments based on the current Chilean legal framework, allowing users to explore different political perspectives on a given issue. Every response is grounded in existing legislation and includes references to the official legal sources used to support the arguments.

The assistant helps MPs and staff both prepare their own positions and anticipate alternative arguments that may arise during legislative debate, contributing to more informed and evidence-based parliamentary deliberation.

Business problem

Parliamentary debate requires MPs and their advisers to develop well-founded arguments within limited time frames. Preparing for legislative discussions often involves reviewing existing legislation, identifying supporting legal provisions and considering how the same issue may be viewed from different political perspectives.

While legal information is publicly available, transforming it into coherent lines of argument remains a time-consuming and largely manual task. The Chamber of Deputies identified an opportunity to explore whether AI could support this process by helping users structure legal arguments more efficiently. Rather than replacing political judgement or policy decisions, the objective was to provide a tool that could assist users in exploring different perspectives, identifying relevant legal foundations and preparing for legislative debate.

Solution

The Chamber of Deputies introduced the Argumentative Assistant as a conversational AI tool available through the Members' Portal, where MPs and staff can interact with it through a chat interface. Users can ask the assistant to generate legal arguments on any public policy topic. Examples of potential prompts include: "Provide arguments on intellectual property from a centre-right perspective" or "Develop arguments supporting stricter environmental regulation from a progressive perspective".

The assistant generates structured arguments grounded in the current Chilean legal framework and includes references and links to the official legal sources used to support each response. The assistant also remembers the context of the conversation, allowing users to refine their requests through follow-up questions and progressively develop more complete lines of argument.

A distinctive feature of the application is that users can choose the political perspective from which the arguments are developed. This allows them not only to prepare arguments supporting their own position but also to anticipate alternative viewpoints that may emerge during parliamentary debate. Responses can be shared directly by email from within the application, and users can provide feedback on the quality of the assistant's answers to support its continuous improvement. The assistant serves as a starting point for analysis and

discussion, while the final interpretation, political judgement and use of the arguments always remain the responsibility of the user.

The initiative formed part of the Chamber's broader strategy to evaluate practical AI applications that could support the daily work of MPs and staff.

Human interaction

The primary users are MPs and parliamentary advisers involved in preparing legislative debates, speeches and policy discussions. The assistant fits into the research and preparation phase of their work, providing a first draft of arguments that users can refine.

Rather than replacing human judgement, the application supports an iterative process. Users can request additional arguments, explore alternative political perspectives, ask follow-up questions or narrow the discussion to specific legal aspects.

The AI automates the time-consuming task of structuring arguments and identifying legal foundations, but it does not determine the political position to adopt or recommend legislative decisions. Human judgement remains essential throughout the process.

Users are responsible for evaluating the relevance and quality of the arguments, verifying that they are appropriate for the parliamentary context and deciding how (or whether) to use them in legislative work. The final decision always remains with the MP or parliamentary adviser.

Results and impact

The Argumentative Assistant is still in the early stages of adoption, and the Chamber of Deputies is currently focused on expanding its regular use among MPs and staff.

The Chamber is currently collecting user feedback and monitoring adoption to better understand how the assistant is used in practice. As usage becomes more widespread, it expects to introduce more systematic measures of user satisfaction, time savings and the overall impact on the preparation of parliamentary work.

Risk management

One of the main considerations for this application was ensuring transparency around its use of political perspectives. The assistant is intentionally designed to generate arguments from a user-selected ideological viewpoint rather than to produce politically neutral responses. To avoid misunderstandings, the selected political perspective is always explicit and fully controlled by the user, ensuring that any ideological framing is intentional and transparent.

Another risk identified was the possibility that users might rely too heavily on AI-generated content without applying their own judgement. To mitigate this, the application was designed as a decision support tool rather than a decision-making system. Human oversight remains essential, and users are expected to review, interpret and adapt the generated arguments before using them in parliamentary work, rather than relying on them without critical assessment.

Prompt engineering was also used extensively to define the assistant's scope and behaviour, instructing it to avoid speculation, refrain from inventing legal information, distinguish clearly between legal sources and generated reasoning, and acknowledge when it lacks sufficient information to answer confidently.

Finally, because the assistant relies exclusively on publicly available legal information, it does not process confidential parliamentary records or personal data, resulting in minimal privacy and confidentiality risks.

Lessons learned

One of the main challenges has been encouraging users to incorporate the tool into their daily work – by demonstrating its practical value, explaining its capabilities and building

confidence in an assistant that is deliberately grounded in official legal sources (rather than relying on general-purpose AI systems).

An important lesson from the initial deployment is that users first need to understand the assistant's specific purpose (unlike a general-purpose AI tool, it is designed to support a very specific task). Early interactions showed that users often asked questions unrelated to its intended function, making user education an essential part of the deployment process. Increasing adoption therefore depends not only on promoting the tool, but also on ensuring that users understand how and when it should be used.

Another important lesson is that defining a narrow scope for the AI assistants makes them more reliable and easier to govern. Limiting the application to a well-defined task made it possible to provide more precise instructions, better control its behaviour and reduce the likelihood of responses outside its intended purpose.

More importantly, specialization led to more consistent and higher-quality outputs, as the assistant could focus on performing one task well rather than attempting to address many different objectives simultaneously.

Next steps

The immediate priority is to expand the assistant's adoption among MPs and staff, encouraging its use in real parliamentary work and gathering evidence on its practical value. As more users incorporate the tool into their daily workflows, the Chamber will systematically collect feedback, monitor user satisfaction and evaluate whether the assistant effectively addresses the needs it was designed to support.

These findings will guide future development, helping prioritize improvements based on real user experience and ensuring that the application continues to evolve in line with parliamentary needs.

Data

The application draws on publicly available Chilean legal and regulatory sources, including laws, decrees, regulations, resolutions, instructions, circulars and municipal ordinances. The assistant does not process confidential parliamentary information or personal data. It operates exclusively on public legal documents, making privacy and information sensitivity risks minimal.

Architecture and technologies

The application is built on the GPT-4.1 language model and uses a retrieval-augmented generation (RAG) architecture to search for and retrieve information from a curated corpus of current Chilean legislation and regulations. The assistant is available through a desktop version, integrated into the Chamber of Deputies' internal Members' Portal.

AI assistant for search, analysis and work with legal acts

Parliament/chamber: Verkhovna Rada of Ukraine

Functional area(s): Bill drafting and amendments; Document and records management; MP support and research; Public engagement and accessibility

Date: June 2026

Overview

The Verkhovna Rada of Ukraine is developing and implementing an AI assistant for working with legal acts. The application uses natural language processing and machine learning technologies to help MPs and staff efficiently search for, analyse and process legislation.

The system is aimed at improving the quality of the lawmaking process and reducing the labour intensity of routine analytical and legal tasks.

Business problem

The corpus of current Ukrainian legislation is extremely large and continuously growing: every year the Verkhovna Rada adopts hundreds of laws and introduces thousands of amendments to existing acts. MPs, their assistants and staff employees spend a significant portion of their working time manually searching for and analysing legal acts, verifying the currency of their versions and identifying relationships between various legislative provisions.

Although the existing search tools in the “Legislation of Ukraine” database offer robust functionality, they require users to know the exact document details or specific search queries. This creates a significant barrier: without legal training or experience working with the database, users can struggle to find the provision they are looking for. At the same time, in the context of European integration, there is a growing need for rapid analysis of the compliance of national legislation with the requirements of the EU acquis.

The AI assistant offers the possibility of overcoming these limitations: it allows users to ask questions in natural language, receive consolidated analytical responses with references to specific provisions, identify contradictions between acts, and track changes in legislation automatically.

Solution

The AI assistant is an interactive conversational interface integrated with two document sets: “Primary Laws” (800+ documents) and “Translations of EU Acts” (600+ documents). Users can submit queries to the assistant in natural language – for example, “What provisions govern the procedure for conducting land auctions?” or “Are there any contradictions between the Law on Local Self-Government and the latest amendments to the Budget Code?” – and receive a structured response with references to specific articles and versions of the relevant acts. The system can summarize content and classify provisions by thematic categories. It is also able to store chat history and index files in a repository. The interface is accessible via the internal staff web portal, and the assistant’s outputs are always accompanied by precise source references, enabling users to independently verify any information provided.

Human interaction

The AI assistant functions as an intelligent search and analytical tool but does not replace human legal judgement. The system provides information, identifies connections and formulates preliminary conclusions. However, all substantive decisions regarding the interpretation of provisions, the preparation of opinions, and legislative initiatives remain exclusively with the human.

The assistant always indicates sources so that users can independently verify every statement. The system is configured to display the confidence level of a response and to explicitly flag cases where a matter is debatable or requires additional expert assessment. This ensures informed, rather than blind, use of AI recommendations.

Results and impact

During the pilot phase, the following results are expected to be achieved:

- A significant reduction in time spent searching for, and conducting initial analysis of, legal acts – from hours to minutes for standard queries
- Improved accessibility of the legislative database for MPs without specialized legal training, as a result of the natural language interface
- Reduced workload for the analytical staff regarding routine queries, allowing their expertise to be focused on complex systemic issues

The implementation of the project supports Strategic Goal 1 (enhancing legislative capacity and optimizing the lawmaking process) and Strategic Goal 2 (parliamentary openness and accessibility of legal information).

Risk management

The following risks and mitigation measures have been identified:

- Accuracy of legal responses: Inaccurate or outdated information may lead to erroneous decisions. To mitigate this risk, the system is connected to the current “Legislation of Ukraine” database in real time. Each response contains references to specific articles for verification.
- Overreliance on AI: There is a risk that users may treat the assistant’s responses as a definitive legal opinion. Mandatory staff training is conducted. The system displays confidence levels and recommends involving legal experts in complex cases.
- Confidentiality of internal information: User queries may relate to sensitive legislative initiatives. A secure cloud infrastructure and role-based access control are in place.
- Model bias: The algorithm may favour certain interpretations of legal provisions. Responses are regularly audited and there is a user feedback mechanism for model correction.

Lessons learned

The key conclusion of the preparatory phase was the understanding that technical implementation is only part of the project – equally important are the formulation of knowledge base requirements with input from practising legal experts, and the identification of real use case scenarios prior to the commencement of development.

Recommendations for other parliaments:

- Start with narrow, well-defined use cases (for example, searching for provisions in a specific area of law) rather than attempting to cover the entire body of legislation at once.
- Involve end users – MPs and committee staff – in testing at early stages: their feedback fundamentally reshapes development priorities.

- Invest in the structuring and machine-readability of the legal database prior to AI implementation – the quality of input data determines the quality of the system’s responses.
- Do not underestimate change management work: convincing experienced lawyers to trust an AI tool takes time and requires tangible results to be demonstrated.

Next steps

Following a successful pilot, an expansion of the assistant’s functionality is planned: expanding the knowledge base to the full corpus of legal acts in the “Legislation of Ukraine” database, adding a comparative analysis module for EU member states’ legislation, integrating with the “Draft Laws” system for real-time analysis of submitted draft laws, and connecting to the “Legislation Online” web portal to provide public access to the assistant’s core functions.

The possibility of adapting the solution for the parliaments of other partner countries – within the framework of parliamentary cooperation programmes – is also under consideration.

Data

The system relies on two document sets – “Primary Laws” (800+ documents) and “Translations of EU Acts” (600+ documents) – in machine-readable format. The system does not process restricted-access documents.

Architecture and technologies

The solution is based on a cloud architecture that uses retrieval-augmented generation (RAG) technologies to ensure accurate responses with attribution to specific sources. The deployment platform is AWS cloud services, in accordance with the overall IT strategy of the Verkhovna Rada. The system integrates with the “Legislation of Ukraine” database via application programming interface (API), as well as with the unified automated document management system and the “Draft Laws” system. The interface is implemented as a web application with support for secure user authentication.

Plenary proceedings and transcription

Parliamentary questions and legislative support documents

Parliament/chamber: National Assembly of Mauritius

Functional area(s): Plenary proceedings and transcription

Date: August 2026

Overview

The National Assembly of Mauritius is implementing a generative AI-powered Parliamentary Questions Module to assist MPs and parliamentary officers in drafting parliamentary questions. The system analyses previous questions, parliamentary records, legislation and ministerial portfolios to generate drafting suggestions and relevant precedents. It also performs automated compliance checks against the standing orders and rules of the National Assembly to improve the procedural admissibility and quality of questions before submission. All AI-generated outputs remain subject to review and validation by parliamentary officers.

Business problem

Parliamentary questions are a fundamental oversight mechanism through which MPs hold the executive accountable. Yet preparing these questions is one of the most time-sensitive and resource-intensive processes within the National Assembly of Mauritius, requiring extensive research across legislation, standing orders and rules, ministerial portfolios, previous parliamentary questions and answers, Hansard records, committee reports and other parliamentary documents. Officers responsible for supporting this process must also ensure that every question complies with the procedural requirements of the National Assembly before it is admitted.

The increasing volume of parliamentary information accumulated over decades, combined with tight submission deadlines and limited specialist resources, has made the research and validation process increasingly demanding. Identifying relevant precedents, avoiding duplication, verifying ministerial responsibility and ensuring compliance with the standing orders and rules all require significant manual effort and institutional knowledge. This creates pressure on parliamentary officers while increasing the risk of inconsistencies or delays during periods of high parliamentary activity.

Solution

The solution to this business problem was not to replace parliamentary expertise, but to augment it. The Parliamentary Questions Module was conceived to provide intelligent research assistance, generate drafting suggestions, retrieve relevant precedents and perform automated compliance checks against the standing orders and rules of the National Assembly. By reducing repetitive research tasks and accelerating document analysis, the solution enables parliamentary officers to focus on procedural judgement and quality assurance while maintaining full human oversight over every parliamentary question submitted.

The module has been integrated into the National Assembly's existing parliamentary questions workflow. Once a parliamentary question is submitted by an MP and uploaded to the platform, it is analysed through an interactive dashboard that brings together all relevant information in a single workspace for parliamentary officers.

The system automatically identifies the appropriate ministerial portfolio, classifies and subedits the question using the National Assembly's parliamentary lexicon, and checks it

against the standing orders and rules. Where potential procedural issues are detected, the dashboard immediately flags the question and highlights the specific areas of possible non-compliance, enabling officers to review and address them before the question is processed further.

The platform also supports collaborative working. Library officers and Clerks' assistants can enrich each question by adding research, references and supporting information, while the Clerk can request additional background information, contextual analysis and summaries from a large language model (LLM) to support procedural consideration. The system provides decision support throughout the process, but all recommendations are advisory, with the final determination on admissibility and approval remaining solely with the Clerk and authorized parliamentary officers.

Human interaction

The Parliamentary Questions Module assists parliamentary officers involved in the parliamentary questions process, including Clerks' assistants, Library officers and the Clerk. Human judgement remains central to the process: parliamentary officers evaluate the AI's recommendations, verify the supporting research, determine whether any procedural concerns have been satisfactorily addressed, and decide on the admissibility of the question. The Clerk retains sole authority for procedural interpretation and final decisions, ensuring that the AI functions solely as a decision support tool and not as a decision maker.

Results and impact

The Parliamentary Questions Module is now in the final production stage and has transformed the way parliamentary questions are processed within the National Assembly. By consolidating research, procedural checks and supporting information into a single workspace, the platform has significantly reduced the time required for parliamentary officers to analyse and prepare questions. Automated identification of ministerial portfolios, classification using the parliamentary lexicon and compliance checks against the standing orders and rules have improved consistency while reducing repetitive manual tasks.

In practice, the interactive dashboard immediately flags potential procedural issues or areas of possible non-compliance, allowing officers to address them at an early stage of the workflow. The platform also facilitates collaboration by enabling Library officers and Clerks' assistants to contribute research and supporting material directly to each parliamentary question, creating a comprehensive digital case file for review.

The result is a more efficient, consistent and transparent parliamentary questions process. Parliamentary officers can focus on procedural analysis, legal research and decision-making rather than routine information retrieval, while the Clerk retains full authority over the interpretation of the standing orders and the final determination of the admissibility of every parliamentary question. The application has strengthened institutional knowledge management, improved the quality of parliamentary support services and enhanced the National Assembly's capacity to respond effectively to increasing workloads.

Risk management

From the outset, the Parliamentary Questions Module was designed in accordance with a risk-by-design approach, with governance, security and parliamentary integrity embedded throughout its development.

Key risks identified included inaccurate or ungrounded AI-generated outputs, procedural misclassification, overreliance on automated recommendations, inconsistent user adoption, and the protection of parliamentary information. There was also a recognized risk that users might accept AI-generated suggestions without sufficient verification or that new ways of working could face resistance during implementation.

To address these risks, the platform was implemented strictly as a decision support system. All AI-generated recommendations – including portfolio allocation, compliance checks against the standing orders and rules, drafting suggestions, and research outputs – must be reviewed and validated by parliamentary officers before any action is taken.

Human-in-the-loop controls, role-based access, audit trails, continuous quality assurance, structured user training, champion-led change management and a phased roll-out were implemented to support adoption and maintain quality. The Clerk retains exclusive authority for interpreting the standing orders and determining the admissibility of every parliamentary question.

In addition, the project underwent a comprehensive data protection impact assessment (DPIA) and incorporates responsible AI principles, ensuring that parliamentary confidentiality, transparency and accountability are maintained at every stage of the process.

Lessons learned

One of the most important lessons learned was that successful AI adoption in parliament depends far more on governance, user engagement and change management than on the technology itself. Involving parliamentary officers from the design stage, incorporating their feedback throughout development, and aligning the solution with existing parliamentary procedures were critical to achieving user confidence and acceptance.

Another key lesson was that AI should complement – not replace – parliamentary expertise. Maintaining a human-in-the-loop approach, with the Clerk retaining final authority over procedural decisions, ensured trust in the system while preserving parliamentary privilege and institutional integrity. Embedding compliance with the standing orders and rules into the application from the outset also proved more effective than treating it as a later enhancement.

Our advice to other parliaments is to begin with a clearly defined parliamentary business need rather than a technology objective. Invest in high-quality parliamentary data, establish strong governance and risk management from the outset, and develop AI solutions in close partnership with end users. AI delivers the greatest value when it enhances existing parliamentary processes while respecting the unique constitutional and procedural role of parliament.

Next steps

The Parliamentary Questions Module forms part of the National Assembly's broader generative AI programme and is already being expanded to support other core parliamentary functions. The same AI platform is being implemented for parliamentary transcription, translation, research, document summarization, legislative analysis and knowledge retrieval, creating an integrated digital ecosystem for parliamentary business. The National Assembly also intends to share its experience, governance framework and lessons learned with other legislatures through international parliamentary networks, contributing to the responsible adoption of AI across parliaments.

Data

The system draws on previous parliamentary questions, the National Assembly's standing orders, rules and regulations, and ministerial portfolios.

Architecture and technologies

The Parliamentary Questions Module is built on the National Assembly's secure generative AI platform hosted within Oracle Cloud Infrastructure (OCI). It combines large language models (LLMs), retrieval-augmented generation (RAG) and vector search technologies to provide accurate, context-aware responses grounded in authorized parliamentary information, rather than relying solely on general AI knowledge.

The application integrates with the National Assembly's existing digital ecosystem, including the Parliamentary Electronic Document Management System (PEDMS), the National Assembly's Member Sitting Application System, the National Assembly website, Hansard records, parliamentary questions and answers, legislation, committee reports and other official parliamentary publications. It also incorporates the National Assembly's parliamentary lexicon and procedural rules to support portfolio identification, classification, and automated compliance checking against the standing orders and rules. Information published through the parliamentary website and the Member Sitting Application System further enriches the knowledge base available to the platform.

The platform operates within the National Assembly's secure ICT environment, using role-based access controls, audit logging and governance controls aligned with the parliament's information security management system (ISMS). Rather than replacing existing parliamentary systems, the AI application acts as an intelligent layer that enhances research, analysis and decision support while preserving established parliamentary workflows and maintaining full human oversight over all parliamentary decisions.

Multilingual support for parliamentary documentation

Parliament/chamber: National Assembly of Mauritius

Functional area(s): Plenary proceedings and transcription; Public engagement and accessibility

Date: August 2026

Overview

The National Assembly of Mauritius is implementing a generative AI-powered multilingual support application to facilitate parliamentary business in English, French and Mauritian Creole. The application assists with translation, transcription, summarization and language consistency across parliamentary debates, legislative documents and official publications. It supports MPs and parliamentary officers by improving the accessibility, efficiency and quality of multilingual parliamentary records, while preserving human review and editorial oversight.

Business problem

The National Assembly operates in a multilingual and multicultural society, reflecting the country's rich ethnic and cultural diversity. Parliamentary business is conducted in English and French, with increasing efforts to support the introduction of Mauritian Creole to enhance accessibility and democratic participation.

The need to produce timely and accurate translations, transcripts and parliamentary publications across multiple languages was placing growing demands on parliamentary staff.

Solution

The National Assembly introduced a multilingual support application to assist parliamentary officers in preparing and publishing parliamentary business in English, French and Mauritian Creole. The application helps translate parliamentary documents, transcribe debates, generate summaries and ensure consistent use of parliamentary terminology across all three languages.

Users simply upload or select a parliamentary document or debate and choose the required language, and the application produces a draft translation, transcript or summary. It also highlights key parliamentary terms and maintains consistent wording throughout the document. In addition, where required, the application also provides alternative wording and contextual suggestions to improve clarity and consistency across languages. Officers can compare versions side by side, make amendments where necessary and approve the final text before publication.

The application supports the Hansard Unit, translators, parliamentary officers and the Parliamentary Library by reducing the time spent on routine language tasks while improving consistency across parliamentary publications. Every output is reviewed and validated by parliamentary language professionals before it becomes part of the official parliamentary record, ensuring that the authority and integrity of the National Assembly are fully preserved.

The final output is a verified multilingual parliamentary record or document that can be published through the National Assembly's existing systems, making parliamentary information more accessible to MPs and the public while maintaining the accuracy and authority of the official parliamentary record.

Human interaction

The multilingual support application is used by Hansard editors, translators, parliamentary officers, the Parliamentary Library and the Parliamentary Digital Services team. It has been integrated into existing parliamentary workflows and assists staff in preparing multilingual transcripts, translations, summaries and official parliamentary publications.

Human judgement remains essential throughout the process: parliamentary language professionals review, edit and validate all AI-generated outputs to ensure linguistic accuracy, procedural correctness and consistency with parliamentary practice. No AI-generated content is published without human approval.

Results and impact

The multilingual support application is now in the final production stage and has made the process of preparing multilingual parliamentary records and publications significantly more efficient. Tasks that previously required extensive manual drafting and comparison across languages can now be completed more quickly, allowing translators, Hansard editors and parliamentary officers to focus on reviewing and refining content rather than producing first drafts.

The application has improved consistency in the use of parliamentary terminology across English, French and Mauritian Creole and has strengthened collaboration between the Hansard Unit, translators and parliamentary officers through a single integrated workflow. It has also accelerated the preparation of multilingual debates, summaries and parliamentary documents, supporting the timely publication of parliamentary information.

Most importantly, the application has enhanced the accessibility of parliamentary proceedings for the country's multilingual population while preserving the quality and authority of the official parliamentary record.

Risk management

The principal risks identified included inaccurate translations or transcriptions, inconsistent use of parliamentary terminology, overreliance on AI-generated outputs, and the protection of parliamentary information. There was also a risk that linguistic nuances or parliamentary expressions could be misinterpreted across English, French and Mauritian Creole.

To address these risks, the application was designed as a decision support tool rather than an automated publishing system. All translations, transcripts and summaries are reviewed and approved by Hansard editors, translators and authorized parliamentary officers before publication. Parliamentary terminology and language conventions were incorporated into the application to improve consistency, while secure access controls, audit trails and governance measures protect parliamentary information.

The project was also developed in line with the National Assembly's information security management system (ISMS), responsible AI principles and a comprehensive data protection impact assessment (DPIA), ensuring that human oversight remains central throughout the process.

Lessons learned

The key lesson learned was that AI is most effective when it complements parliamentary expertise rather than replacing it. Involving translators, Hansard editors and parliamentary officers throughout the project both improved the quality of the outputs and helped build user confidence.

Our advice to other parliaments is to build the application around a well-curated parliamentary language corpus, establish strong governance and human oversight from the outset, and ensure that AI supports existing parliamentary workflows rather than changing them.

Next steps

The multilingual support application is being expanded as part of the National Assembly's broader generative AI programme. Future developments include enhanced support for live parliamentary transcription, real-time translation, multilingual citizen engagement through the National Assembly website and the Member Sitting Application System, and integration with other AI-assisted parliamentary services. The National Assembly also intends to share its experience and good practices with other parliaments through international parliamentary networks.

Data

The application draws on authorized parliamentary information, including Hansard debates, legislation, bills, parliamentary questions and answers, committee reports, official parliamentary publications, procedural rules, the parliamentary lexicon, and other records maintained by the National Assembly. It also uses multilingual parliamentary corpora in English, French and Mauritian Creole to ensure consistent terminology and language usage. The platform is designed to operate primarily on publicly available or authorized parliamentary information and is hosted within the National Assembly's secure ICT environment, with appropriate access controls and governance measures to protect parliamentary information.

Architecture and technologies

The multilingual support application is built on the National Assembly's secure generative AI platform hosted on Oracle Cloud Infrastructure (OCI). It uses large language models (LLMs), retrieval-augmented generation (RAG) and vector search to provide context-aware translation, transcription and multilingual language support.

The application integrates with the Parliamentary Electronic Document Management System (PEDMS), Hansard, the Member Sitting Application System, the National Assembly website and other official parliamentary information systems. It uses the Assembly's multilingual parliamentary corpus, terminology database and procedural rules to ensure consistency across English, French and Mauritian Creole.

Operating within the National Assembly's secure ICT environment, the platform enhances existing parliamentary systems while maintaining role-based access controls, audit logging and full human oversight over all published parliamentary content.

Public engagement and accessibility

Plain-language bill summarizer

Parliament/chamber: Chamber of Deputies of Brazil

Functional area(s): Public engagement and accessibility

Date: May 2026

Overview

The Brazilian Chamber of Deputies summarizes all bills in plain language for citizens, making these summaries available through its app and portal

Business problem

Each year, Brazilian MPs present about 10,000 bills to the Chamber of Deputies. All these bills are made available through an app and the institutional portal, where they are automatically open to citizen contributions through e-polls. However, legislative text is generally not accessible, which presents a challenge for the average citizen.

Solution

As soon as a new bill is officially added to the digital legislative platform, the AI system summarizes it using plain-language techniques and publishes the summary on the Chamber's app and portal, making it accessible to all citizens.

The final summary is capped at 100 words, even for bills that amend other bills. Once a summary has been reviewed and amended by staff, the "summarized by AI" label is removed.

Human interaction

Citizens can see the summaries immediately, labelled "summarized by AI". The app and portal interfaces also allow citizens to provide feedback on each summary. Any negative feedback is reviewed by a small group of legislative consultants, who adapt the summary as needed. The update is then immediately reflected on the app and portal.

Results and impact

The expectation is that more citizens will understand bills, engage with the e-polls and leave comments.

Risk management

The main risk was the inability to review poor-quality summaries. This was mitigated through an agreement between the communication department, the IT department and the legislative consultants, so that consultants could review all negative feedback and analyse the 10–20 most accessed bills each week.

Lessons learned

Always include business units with the necessary skills to review negative AI outputs.

Next steps

The Chamber of Deputies is open to sharing this case study with other parliaments that would like to learn more.

Internal administration, IT and security

AI-assisted re-engineering of a parliamentary legacy application

Parliament/chamber: Chamber of Deputies of Italy

Functional area(s): Document and records management; Committee work

Date: June 2026

Overview

The Chamber of Deputies of Italy set about completely re-engineering “Relazioni”, a business-critical application used to manage reports and statutory obligations associated with parliamentary documents. The system supports two parliamentary services: the Legislative Drafting and Legal Texts Service, which manages reports linked to parliamentary documents and legislative activities, and the Parliamentary Oversight Service, which monitors statutory obligations arising from legislation, government actions and related parliamentary procedures.

Originally developed using ASP.NET WebForms technology, the application had become increasingly difficult to maintain and evolve, so the Chamber launched a re-engineering initiative using an AI-assisted software engineering tool named Codex to analyse the legacy codebase, map dependencies, support architectural assessment, assist with code generation and refactoring, and troubleshoot issues during the migration.

Business problem

The old application supported an established parliamentary workflow, but its underlying technology had become outdated. It was originally developed using ASP.NET WebForms, characterized by a tight coupling between the user interface, business logic and data access layers. Maintenance activities became progressively more complex and newer developments required greater effort.

The Chamber identified an opportunity to use generative AI as an accelerator for software modernization, aiming to improve code quality and support the migration to a more modular, maintainable and scalable architecture, while keeping developers fully in control of all design and implementation decisions.

Solution

The Chamber of Deputies completely re-engineered a legacy ASP.NET WebForms application named “Relazioni” using modern technologies for front end, back end, application programming interface (API) services, containerized deployment and automated CI/CD pipelines. Codex was used throughout the project as an AI-assisted development tool, supporting the team in understanding the legacy codebase, mapping dependencies, identifying obsolete components and technical debt, generating portions of code, supporting refactoring, and troubleshooting. The project followed an incremental approach rather than replacing the entire application at once.

Human interaction

Codex was used by the development team not as an autonomous developer, but as a support tool. Developers used structured prompts to ask Codex to analyse the legacy repository, identify dependencies, suggest refactoring approaches, generate code skeletons and support troubleshooting. All AI-generated outputs were reviewed, adapted and validated by human developers before being used. Architectural choices, functional validation and final acceptance remained under the responsibility of the development team.

Results and impact

The project enabled the Chamber to transform a legacy application into a more modern, maintainable and scalable system in just a few weeks, while preserving its institutional functions. The use of Codex helped accelerate repetitive and time-consuming development activities. As a result, the development team was able to devote more time to architectural design, user experience, functional validation and quality assurance.

Risk management

The main risks concerned possible regressions, incomplete understanding of legacy business rules and excessive reliance on AI-generated code. The team adopted an incremental and controlled development approach. Codex outputs were never accepted automatically. Each suggestion was reviewed by developers, and tested and validated against the expected behaviour of the original application. Complex tasks were broken down into smaller, verifiable steps.

Lessons learned

The project confirmed that AI-assisted development is most effective when coupled with clear objectives, well-structured prompts and strong human supervision. Codex was particularly useful for repetitive coding tasks, code analysis and troubleshooting, but less reliable when business rules were implicit or highly specific to the parliamentary domain. The key recommendation is to treat generative AI as an accelerator for skilled development teams, not as a substitute for software engineering expertise.

Next steps

Future developments will focus on strengthening the integration of the application within the Chamber's broader digital ecosystem for legislative monitoring and parliamentary oversight, particularly through progressive integration with the new Observatory on Legislation and Public Policies.

Data

The project used Codex during the software development process to support the re-engineering of the existing application. The main inputs were the legacy application codebase, the existing application structure, the data model, functional requirements and technical constraints. Use of AI was limited to software engineering tasks such as code analysis, refactoring support, generation of repetitive code structures, and troubleshooting.

Architecture and technologies

The legacy application was based on ASP.NET WebForms. The re-engineered solution introduced a more modular architecture based on the following:

- A React front end
- A Spring Boot back end exposing API services
- Containerized deployment using Docker
- Automated CI/CD pipelines
- Deployment on Kubernetes

Codex was used to analyse the legacy repository, map dependencies, generate code skeletons, assist with refactoring and support troubleshooting.

Digital assistant based on RAG techniques to assist in the modernization of COBOL codebases

Parliament/chamber: Chamber of Deputies of Italy

Functional area(s): Document and records management; HR and internal administration; IT infrastructure and cybersecurity

Date: June 2026

Overview

The Chamber of Deputies' Technological Innovation and Digitalization Department is exploring the possibility generative AI-based solutions to support the documentation of legacy code written in Common Business-Oriented Language (COBOL).

The objective of this project is to develop a prototype digital assistant for understanding, documenting and reverse-engineering source code and operational workflows, to support re-engineering and innovation activities.

The chatbot uses retrieval-augmented generation (RAG) techniques: a structured knowledge base is built through static analysis of the COBOL code using call graphs, data flows, program dependencies and domain logic. The chatbot isolates and retrieves the relevant chunks and uses them as context to answer the user's questions.

Given the criticality of the information processed, the development and service delivery model is on-premises, meaning all software and data will be installed and managed within the Chamber of Deputies' data centre.

Business problem

Evolving a large COBOL codebase, currently consisting of more than one million lines of code, is an expensive and time-consuming task. Current job market conditions make it very difficult to find COBOL developers, meaning that maintaining the codebase is difficult and costly.

Potential migration to a new platform is currently inhibited by a lack of understanding of the current codebase – the risk is too high. A digital assistant that could improve this understanding would make maintenance activities faster and cheaper, ultimately reducing risks and allowing for migration to newer platforms.

Solution

The envisioned solution consists of an agentic architecture based on different agents:

- One for identifying call graphs, data flows, program dependencies and domain logic
- A second for managing the RAG using the previous elements and prior knowledge about COBOL, as well as specific documents held by the department
- A third, an IBM Granite large language model (LLM), together with output interface agents for interfacing with operators

All of this is deployed on-premises in the department on suitable CPU/NVIDIA-based machines.

Human interaction

Operators use the assistant to gain a comprehensive understanding of the codebase. The interaction is based on a chatbot interface.

Results and impact

The envisioned results are the creation of documentation about the codebase, and a reduction in the time and effort needed to maintain the codebase.

Risk management

Information classification and risk analysis were conducted to evaluate the probability and impact of risks relating to confidentiality, integrity and availability. Processing sensitive parliamentary data and complex payroll systems requires a strict on-premises solution and complete data anonymization.

A decision was made to develop the pipeline and RAG locally to ensure sovereignty, keeping all information entirely on the infrastructure of the Chamber of Deputies.

Personnel accessing sensitive data are expressly authorized, with segregation of duties, limited access to data and program source code, secure access control, and logging and event registration.

Lessons learned

Although the project is ongoing, early lessons learned include the following:

- The importance of code anonymization, as sensitive information is often embedded in legacy code
- The importance of designing systems that can be deployed on local and constrained resources
- The importance of continuous learning about generative AI

Next steps

Following the production launch in November 2027, the road map will focus on extending the platform to a similar codebase dealing with staff salaries (more than 1,500 people on the monthly payroll).

Data

Details on this point are not yet available. Portions of the anonymized codebase will be published as open data sets for future research and benchmarking.

Architecture and technologies

The system architecture includes: virtual machines hosting Docker for containerizing the agents' execution environments using a completely on-premises approach, with machines mirrored to allow for parallelization of requests; and a virtualized mainframe test environment for executing a replica of the codebase in case queries ask for refactoring of the code. A valid LLM for this purpose is one from the IBM Granite family, which has specific training and specialization in COBOL and can be run on-premises.

Facial recognition and vehicle licence plate recognition for access control to administrative buildings and territories

Parliament/chamber: Verkhovna Rada of Ukraine

Functional area(s): IT infrastructure and cybersecurity; Other (physical security of parliamentary staff and infrastructure)

Date: June 2026

Overview

To strengthen physical security under martial law and modernize access control, the Verkhovna Rada of Ukraine is piloting an AI-based system for facial recognition and vehicle licence plate recognition, providing fast and accurate real-time identification of individuals and vehicles entering its buildings and grounds.

Business problem

The Verkhovna Rada is a high-security facility that receives hundreds of MPs (450 in total), Secretariat staff, journalists and accredited visitors daily.

The traditional access control system, based on manual verification of electronic passes and documents, created several operational challenges. During peak hours, queues formed at checkpoints, slowing down parliamentary work and causing inconvenience for MPs and staff. The human factor in manual verification did not eliminate the possibility of errors or abuse.

With the introduction of martial law, physical security requirements for parliamentary infrastructure increased significantly: more reliable identity verification was needed for everyone entering the buildings and grounds of the Verkhovna Rada.

An additional factor was the need to integrate access control with the “Smart Rada” system to confirm the physical presence of Members during votes and sessions. The implementation of AI facial recognition became a response both to wartime security challenges and to the need for digital modernization of the access control system in line with EU standards.

Solution

A pilot version of the facial and licence plate recognition project was developed with support for multi-factor authentication. At this stage, the tool does not have a graphical interface, but a system has been implemented in which an administrator manually uploads images directly to folders in the S3 storage system. The servers then automatically verify and register a person for subsequent identification on cameras.

The project also supports the uploading of pre-recorded video for facial recognition. To enhance reliability, the system will be implemented as a component of multi-factor authentication – facial recognition can be combined with verification of an electronic pass or other identifier. This prevents unauthorized access even in the event of a lost or stolen pass.

Human interaction

In this system, AI acts as the first line of verification, significantly accelerating and standardizing the access control process. The AI automatically identifies individuals and vehicles and monitors the video stream in real time – a task that is physically impossible to

perform manually at the same speed and scale. However, when the system is fully implemented and deployed, humans will remain central to the process in critical situations.

A guard will intervene when the system has low confidence in an identification, when an alarm is triggered, or in non-standard situations. A human will always make the final decision on access in disputed cases. Monitoring centre operators will maintain continuous oversight of the system and respond to incidents.

Results and impact

Implementation of the system is expected to deliver tangible results for the security service and at parliamentary checkpoints.

The most noticeable improvement will be a significant reduction in the average time taken to identify a person entering the buildings or a vehicle entering the grounds: a process that previously required several seconds of manual document or pass verification will now occur almost instantaneously and without stopping the flow of people or vehicles. This should eliminate peak-hour queues and improve convenience for MPs and staff.

From a security standpoint, the system will provide significantly more reliable identity verification than visual pass or ID checking, eliminating the possibility of a person using someone else's credentials.

A centralized audit log of all access control events will facilitate incident investigations and improve accountability.

Risk management

The following risks and mitigation measures have been identified:

- Recognition accuracy and false triggers: Both false positives (access granted to unauthorized individuals) and false negatives (a person or licence plate not recognized) are key concerns. The process includes human intervention in the identification process to mitigate this risk.
- Protection of personal and biometric data: This is a critical issue, as the system will process the biometric data of MPs and staff. Database encryption, strict access controls and compliance with Ukrainian personal data protection legislation are planned.
- Security under martial law: The system is designed with operational recovery plans and backup operating modes in case of power outages or cyberattacks.

Lessons learned

The key lesson is that the quality of the technical infrastructure determines the quality of the outcome. Upgrading video cameras and network equipment at security checkpoints is just as important as the recognition system itself: an algorithm cannot accurately identify a person or licence plate from a low-quality video signal, especially in poor lighting or non-standard camera angles.

It is also important to understand that a human is an integral part of this process, as no software can be trusted with 100% certainty.

Next steps

Implementation of the application is planned for the main administrative building of the Verkhovna Rada, with expansion to all facilities of the parliamentary complex and installation at all checkpoints. Further development plans include integrating the project with the existing access control infrastructure: turnstiles, electronic locks and checkpoints. A centralized audit log accessible to security monitoring centre operators is planned to record all events.

Looking further ahead, deeper integration with the “Smart Rada” system is under consideration for automated quorum monitoring and verification of MPs’ presence during votes.

Parliament is also exploring opportunities to share its experience with other parliaments considering implementing biometric security systems.

Data

The system relies on several categories of data:

- The biometric database contains registered digital templates of faces and vehicle licence plates stored in encrypted form. These are mathematical models of facial and licence plate characteristics – not photographs – making direct reconstruction as images impossible. Biometric data is a special category of personal data under Ukrainian law. It is processed in accordance with Ukrainian personal data protection legislation.
- Real-time video data is processed for identification purposes but is not stored in full. Only events are recorded: a frame from the moment of identification, a timestamp and the verification result.
- The event audit log is a structured cloud-based database recording all access control events: identifications, rejections, alarms and manual guard interventions.

Under martial law, additional security protocols are in place to prevent leakage of data concerning staff.

Architecture and technologies

The following architecture and technologies are in place:

- Video capture layer: The system accepts a real-time video data stream from cameras or an uploaded file in supported formats.
- Processing and identification layer: The facial recognition hardware-software complex processes the video stream on cloud servers to ensure real-time processing with minimal latency. The recognition algorithm compares the obtained biometric template against the database, returns a confidence level and records it in the database.
- Management and monitoring layer: A cloud-based management dashboard provides the ability to retrieve the required information in a sorted format upon request.
- Recovery plan: Backup operating modes are planned for failure scenarios, including autonomous operation of local servers when disconnected from the central network, backup power supply for security checkpoint equipment, and procedures for switching to manual access control in the event of a complete system shutdown.