Northern Ireland Assembly

Research and Information Service Briefing Paper

Paper 33/25 September 2024 NIAR 193-2024

Ray McCaffrey

Use of AI in Legislatures

This briefing paper provides information on the use of Artificial Intelligence in legislatures.

1 Introduction

This briefing paper looks at how Artificial Intelligence (AI) is being utilised in parliamentary contexts. It provides examples of parliamentary tasks being performed by AI and discusses the potential pros and cons of its use.

Section 1 provides a brief introduction. Section 2 sets out a definition of AI and provides a basic outline of how it works, section 3 explores the use of AI in a parliamentary context, while sections 4 to 10 look in closer detail at the application of AI in selected legislatures. Section 11 sets out guidelines for the use of AI and section 12 presents some concluding remarks.

Assembly research paper NIAR 45-24, published in May 2024, provides a broader examination of the use and regulation of AI in the UK.

This paper is based on publicly available information, academic articles and information provided by the Assembly Library.

2 Defining AI and how it works

There is no single, agreed definition of AI. However, the Oxford English Dictionary defines Artificial Intelligence as:

> The capacity of computers or other machines to exhibit or simulate intelligent behaviour; the field of study concerned with this. In later use also: software used to perform tasks or produce output previously thought to require human intelligence, esp. by using machine learning to extrapolate from large collections of data.[1]

Generative AI falls within the second part of this definition – replacing human output with software, for example ChatGPT. This ‘machine learning’ goes beyond the AI that people interact with on a daily basis, for example when using Alexa, Siri or chatting with customer service chatbots, whereby “practitioners

[1] Oxford English Dictionary, definition of Artificial Intelligence: https://www.oed.com/dictionary/artificial-intelligence_n?tl=true

develop artificial intelligence through models that can "learn" from data patterns without human direction."[2]

Generative AI describes "any type of artificial intelligence (AI) that can be used to create new text, images, video, audio, or code. Generative AI uses techniques such as Large Language Models (LLMs) to learn from very large datasets in order to generate outputs."[3]

The Parliamentary Office for Science and Technology (POST) explains how AI technologies function:

> All AI technologies are underpinned by an algorithm or a set of algorithms. An algorithm is a set of instructions used to perform tasks (such as calculations and data analysis) usually using a computer or another smart device. AI often involves retraining algorithms with new data.[4]

3 AI in the parliamentary context

Research from 2021 set out the state of play of AI in the parliamentary context, highlighting the advent of two deep learning models - Generative Pre-trained Transformer (GPT) and Bidirectional Encoder Representations from Transformers (BERT) – which "can be generally used for document classification, information extraction and natural language understanding and inference."[5]

[2] McKinsey and Co., What is generative AI? (2 April 2024): https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai

[3] UK Parliament, Guidance on Generative AI for Parliament

[4] United Kingdom Parliament, POST Note, Artificial Intelligence: AN explainer (14 December 2023) https://post.parliament.uk/research-briefings/post-pb-0057/

[5] The Journal of legislative Studies, Artificial Intelligence (AI) in parliaments – preliminary analysis of the Eduskunta experiment, vol. 27, 2021: https://www.tandfonline.com/eprint/QMKQTXB8VKNU7WHHHVC8/full?target=10.1080/13572334.2021.1976947#d1e172

The release and open-access availability in November 2022 has led to a "heightened focus on researching the potential applications, deployment strategies, as well as the associated risks and challenges inherent in AI."[6]

Subsequent research set out a "classification framework for AI-driven value creation in the parliament, based on two parameters":

1. Usage of the AI system: either focus on the core parliamentary functions aiming to solve existing problems, such as legislative drafting as well as parliamentary reporting and editing, analysis of data, and/or to address emerging parliamentary functions that will enable new forms of services and associated value creation potential.

2. Level of AI services: at an intra-parliamentary level, aiming to provide assistance to MPs as well as the scientific and administrative personnel of the parliament. At an inter-parliamentary level, AI services can either be between parliaments and/or between parliaments and civil society.

This is illustrated in more detail in figure 1.

[6] Jorn Von Lucke & Zander Frank, The Open Government Institute, (2024) A few Thoughts on the Use of ChatGPT, GPT 3.5, GPT-4 and LLMs in Parliaments: Reflecting on the results of experimenting with LLMs in the parliamentarian contex: https://dl.acm.org/doi/pdf/10.1145/3665333

Level of AI Services		AI Systems Usage	
		AI to solve existing problems	**AI to enable new forms of value creation**
Intra-Parliamentary	**MPs**	**Potential uses** Parliamentary reporting and editing, analysis of data, legislative drafting, public submissions on a bill, etc. **Relevant AI technologies** NLP, speech recognition, information retrieval, OCR, data analytics, ontology engineering, summarisation	**Potential uses** Encourage new types of evidence-based analysis, reporting and regulation (including multimedia, video, etc.), foresight analysis, anticipatory regulatory activities, etc. Relevant AI technologies Machine learning, computer vision, time-series mining, AI-powered simulation, planning, summarisation
	Scientific and Administrative Personnel	**Potential uses** Parliamentary reporting and editing, analysis of data, legislative drafting, etc. **Relevant AI technologies** NLP, speech recognition, information retrieval, OCR, data analytics, ontology engineering, summarisation, RPA	**Potential uses** Support new types of evidence-based reporting and analysis, support preparatory activities and analysis, etc. **Relevant AI technologies** Machine learning, computer vision, time-series mining, AI-powered simulation, planning, summarisation, RPA
Inter-Parliamentary	**Collaboration with other Parliaments**	**Potential uses** Exchange of best practices, collaborative functions, data exchanges (legislative and financial data), discussion fora, etc. **Relevant AI technologies** Ontologies, knowledge bases, recommender systems, community detection algorithms, information retrieval, dialogue systems	**Potential uses** Co-creation activities, collaborative foresight, anticipatory regulatory activities, etc. **Relevant AI technologies** Knowledge bases, ontologies, recommender systems, NLP, summarisation, MT
	Citizens/Civil Society	**Potential Uses** chatbots for online conversations with citizens about parliament, etc. **Relevant AI technologies** NLP, speech recognition, recommender systems, information retrieval, dialogue systems, reasoning	**Potential Uses** Services to improve parliamentary transparency of decision-making for citizens, AI-moderated online conversations, personalised knowledge sharing, etc. **Relevant AI technologies** NLP, recommender systems and personalisation, text mining, dialogue systems, ontologies

Figure 1: classification framework for AI-driven value creation in the parliament, based on two parameters

3.1 The state of play

In research on AI in public institutions, the Deputy Chief Information Officer of the Brazilian Chamber of Deputies, which is itself a world leader in the use of AI, identified three groups of parliaments in relation to AI:

1. The first group includes parliaments that have already implemented AI projects successfully and are now focusing on making these solutions

sustainable. Their focus is on policy and strategy development and creating requirements for public procurement that include AI services.

2. The second group includes parliaments that are experimenting with AI in a "lab mode" — testing solutions to demonstrate their benefits and gain support and funding from parliamentary leadership.
3. The third group consists of parliaments that have not yet adopted AI, primarily due to the lack of basic IT infrastructure, legislative systems, and data governance.[7]

4 Eduskunta (Finland)

In April 2021 the Committee for the Future in the Finnish Parliament held a hearing with AI, based on the GPT-3 model. This was widely recognised as the first such hearing in any parliamentary context.

Recent academic research provided an analysis of the "Eduskunta experiment":

> There was a thorough preparation process before the actual hearing could take place. The Committee heard several AI scholars and private sector experts beforehand. A couple of experts from academia and the private sector were called in to assist in the preparation, facilitation and reporting for the hearing… The specificities of GPT-3 were analysed and the decision was made to set up two different personality profiles focusing on technological innovation and business possibilities, and environmental policy, respectively. Once these were created, tests questions were used to see how the profiles react. To bridge any difficulties in the interaction with the GPT-3, questions were first directed to the facilitators, who then fed into the AI system them via an internet prompt.
>
> From this point on, the hearing was conducted like almost any regular one. The parliamentarians asked questions to the two different simulated personalities, Muskie and Saara… and these provided written answers.

[7] Popvox Foundation, Representative Bodies in the AI Era, (January 2024): https://static1.squarespace.com/static/60450e1de0fb2a6f5771b1be/t/659d49c8c62d136f72890838/1704806866772/Representative_Bodies_in_the_AI_Era_Vol_1.pdf

After each question and answer session, Committee members had a side discussion with the facilitators to reflect on the responses provided and the reason why. This reflection discussion constituted the only difference compared to a regular committee hearing.

The research highlighted two important lessons from the hearing:

- AI systems are at the gates and parliaments should better prepare for it. Early adopters will have a tough time, but eventually their investments will pay off as they will experience rapidly improving productivity and lower operational costs in the parliamentary workspace.
- when it comes to AI models, there can be no one-size-fits-all.[8]

5 Italian Senate

The Inter-Parliamentary Union's October 2022 Innovation Tracker highlighted the application of AI in the Italian Senate in relation to amendments to legislation. The key points are reproduced at figure 2.

In the Italian Senate, amendments are presented electronically, but not always for the purpose of changing proposed legislation. Often, members – particularly those from opposition parties – propose amendments as a way to advertise alternative policies or to challenge the government. In extreme cases, they may present large numbers of amendments, differing by just a few words, in order to slow down the legislative process – a practice known as "filibustering". This can prove challenging for the Senate staff responsible for analysing and organizing amendments and scheduling voting.

A new AI-powered system for managing amendments was introduced as part of an experiment to handle workload peaks, such as those caused by

[8] The Journal of legislative Studies, Artificial Intelligence (AI) in parliaments – preliminary analysis of the Eduskunta experiment, vol. 27, 2021

filibustering. This system, which uses “text clustering algorithms”, speeds up the detection of groups of similarly worded amendments.

This experimental feature is already demonstrating its value. According to Carlo Marchetti, Head of the Senate’s Information Systems Development Office, the process is almost instant. But a human eye is still needed to review, approve, modify and integrate the system’s results. No decisions are taken without supervision. In this sense, the AI technology is there to assist human beings, not to replace them.

As part of further research into new tools and algorithms to support core parliamentary activities, the Senate is exploring the possibility of detecting not just textual but also semantic similarities, and of identifying related bills that could be similarly affected by the amendments.[9]

Figure 2: application of AI in the Italian Senate

6 Brazilian Chamber of Deputies

The Brazilian Chamber of Deputies is one of the leading parliaments in the use of AI. Its advanced Ulysses Suite “is a comprehensive, integrated array of AI-driven modules, each designed to enhance a particular aspect of the legislative workflow.”[10]

A recent article examined in detail the functions of the Ulysses Suite:

> The uniqueness of the Ulysses Suite lies in its multifaceted functionalities. Each component, while tailored in its specific task, is intrinsically linked to the overarching goal of modernising the legislative operational framework, within the context of a digital transformation strategy. The suite addresses various aspects of legislative work - from automating laborious tasks and processing complex legal texts, to engaging with public feedback and

[9] Inter-Parliamentary Union, Innovation Tracker, (October 2022): https://www.ipu.org/innovation-tracker/story/how-ai-helps-italian-senate-manage-amendments

[10] LegisTech Library, Integrating Artificial Intelligence to Legislative Services - Ulysses Suite in the Chamber of Deputies of Brazil (March 2024): https://library.bussola-tech.co/p/ulysses-chamber-deputies-brazil

enhancing communication channels. This comprehensive approach not only elevates the operational efficiency of the Chamber of Deputies but also fosters a more transparent and participatory legislative environment.

There are eight components to Ulysses, each with a different function. These are summarised in table 1.

Table 1: Ulysses in the Brazilian Chamber of Deputies

Ulysses 1: Content Aggregation and Thematic Organisation in Parliamentary Webpage Ulysses 1 (U1), the first module of the Ulysses suite, is a specialised AI algorithm designed to enhance the organisation and presentation of thematic content of the Office of the Legislative Counsel on the Chamber's digital portal…This algorithm represents a significant step in digitising and streamlining access to legislative information, reflecting a commitment to both innovation and user-centric design. The core functionality of U1 is its ability to categorise and display technical legislative content, including bills, technical studies, news, committees, speeches and pools based on specific themes or subjects selected by users. This categorisation process is dynamic, allowing for real-time reorganisation of content as new information becomes available or as user interests shift.
Ulysses 2 - Automating Legislative Counsel's Requests Distribution Ulysses 2 (U2), as part of the Ulysses suite developed for the Brazilian Chamber of Deputies, serves a crucial role in streamlining the legislative process through its advanced machine learning and natural language processing capabilities. The primary function of U2 is the semi-automated distribution (specialists still can check/revise the U2's suggestions before distribution) of requests to specialised groups of legislative counsels, a task that was previously managed manually with some risk of cognitive bias and inefficiency. The algorithm of U2 is designed to read and interpret textual requests submitted by parliamentary offices. These requests, often laden with complex and nuanced legislative language, are indicative of the diverse and intricate nature of

legislative work. U2's task is to parse this language, understand the underlying request, and determine the most appropriate group of consultants to handle it. This process is crucial for ensuring that each request is addressed by experts most qualified in the relevant area of legislative drafting.

Ulysses 3: Semantic Analysis in Legislative Drafting Requests

Ulysses 3 (U3) is engineered to process a substantial volume of legislative requests efficiently. The Office of Legislative Counselling receives more than 20,000 demands per year, and U3 plays a critical role in managing this influx. Each request, varying from concise to complex legislative documents, requires the algorithm to perform an in-depth analysis to identify semantically similar existing legislative studies and proposals.

The primary function of U3 is to analyse requests submitted by parliamentary offices to the legislative counsel's office. This analysis is conducted using advanced natural language processing (NLP) techniques, allowing the algorithm to interpret the text of the requests in a sophisticated manner. The key here is the semantic analysis – not just searching for keywords, but understanding the context and meaning behind the words. This approach enables U3 to identify technical studies and legislative propositions that are semantically similar to the submitted requests…

A significant challenge for U3 is the brevity and lack of detail in many requests. The algorithm has been fine-tuned to extrapolate and infer detailed insights from succinct inputs, a critical capability given that many requests for new legislation are concise and lack comprehensive detail.

U3 also incorporates a feedback mechanism to continually improve its accuracy and relevancy. The responses it generates are evaluated for relevance by the legislative counsels, and this feedback is used to refine and optimise the algorithm's performance.

Ulysses 4: Analysis and Management of Legislative Amendments

Ulysses 4 (U4), a vital component of the Ulysses suite, is a sophisticated artificial intelligence algorithm designed to enhance legislative drafting. This advanced tool specifically focuses on the critical task of managing and analysing amendments to legislative texts, one of the core activities of legislative work.

At its core, U4 utilises cutting-edge natural language processing (NLP) techniques. This enables the algorithm to read and interpret the language of amendments submitted during the legislative process. The NLP capabilities of U4 are not limited to merely identifying keywords but extend to understanding the context and deeper semantic meanings embedded within the text of the amendments.

One of its remarkable features is its ability to determine the appropriate placement of amendments within the broader context of legislative proposals. This involves a detailed analysis of the text to ascertain where each amendment fits within the structure of a legislative document, ensuring that amendments are contextually aligned with the sections they aim to modify or influence.

Ulysses 5: Transforming Legislative Record-Keeping with AI-Driven Transcription

Ulysses 5 (U5), a component of the Ulysses suite of AI services implemented by the Brazilian Chamber of Deputies, represents a significant leap in the realm of legislative digitization and accessibility.

The primary function of U5 is the automatic transcription of voice to text. This feature is crucial in a legislative environment where proceedings, discussions, and debates form the core of the legislative process. These verbal exchanges, often complex and lengthy, are essential for record-keeping, legal documentation, and public dissemination (social and institutional control of the legislative process). U5's role in converting these oral communications into written form not only enhances transparency but also aids in the accessibility of legislative information to the public (external and internal to the Legislative Power) and other stakeholders.

In developing U5, the emphasis was on integrating it with the existing legislative systems, notably those supporting the traditional stenography or shorthand writing used in legislative assemblies. This integration ensures that the audio feed into the AI is of the highest quality, which is crucial for accurate transcription.

Ulysses 6: Semantic Analysis of the Public's Comments in Legislative Proposals

Ulysses 6 (U6) is the module specifically tailored to analyse and interpret the feedback received from citizens on various legislative proposals, primarily through electronic polls associated with legislative bills.

The core functionality of U6 revolves around its capacity to process and understand the semantic content of citizens' comments. Unlike traditional polling methods that offer binary options of approval or disapproval, the e-poll invites U6 invites citizens to provide detailed commentary on each bill. E-poll for bills is one of the services provided by the institutional portal of the Chamber of Deputies. The commentaries gathered by this service are then submitted to U6 for argumentation analysis and grouping. This qualitative approach allows for a more comprehensive understanding of public sentiment, going beyond mere numerical statistics to grasp the subtleties of public opinion.

When a legislative proposal is introduced into the system, an electronic poll is opened, inviting public participation. Citizens are encouraged not only to vote in favour or against the proposal but also to leave comments elucidating their viewpoints. These comments, which can vary significantly in terms of content, length, and complexity, provide a rich dataset for analysis.

U6 employs sophisticated NLP techniques to analyse these textual inputs. The algorithm is capable of identifying key themes, sentiments, and patterns within the comments. It does this by parsing the language used, understanding the context, and even detecting underlying sentiments and biases. This analysis is not limited to surface-level word matching; instead, it gets into the semantic meaning behind the words, enabling a more detailed understanding of public opinions.

Ulysses 7: Biometric Authentication Integration with Infoleg for Parliamentary Deliberations

Ulysses 7 (U7) is the module specifically developed to integrate with the Infoleg application, used by members of the parliament. This module addresses a critical aspect of legislative procedures: authentication during deliberations in plenary sessions and committee meetings. In light of concerns about the potential for fraud in the virtual and hybrid voting system, like staff voting on behalf of Parliamentarians, the development of U7 becomes crucial.

One of the standout features of U7 is its focus on leveraging biometric authentication, particularly facial recognition, to verify the identity of members participating in various parliamentary activities through the Infoleg App. This technology ensures that the members can engage in deliberations without being physically present in the plenary, thereby enhancing flexibility and participation, in situations permitted by the rules of procedure.

Ulysses 8: Application of Facial Recognition for the Communication and Photos of Parliamentarians

Ulysses 8 (U8) was the module whose primary function revolved around facial recognition technology. The deployment of facial recognition by U8 was confined to a specific context – identifying parliamentarians in photographs taken during public meetings, such as Committees and Plenary Meetings. This targeted application reflects a conscious effort to balance the benefits of AI with the need to respect privacy and avoid intrusive surveillance.

The facial recognition system in U8 operated by scanning photographs and accurately identifying the parliamentarians present in them. This process involved a complex interplay of machine learning algorithms that analyse facial features and cross-reference them with the official database of the Parliament. The sophistication of the system was evident in its ability to handle the nuances of facial recognition – differentiating between individuals who may have similar features, adjusting for variations in lighting and angle, and even accounting for changes in appearance over time.

7 United States Congress

The United States Congress is addressing the potential benefits and drawbacks of AI. In February 2024, the Republican Speaker of the House of Representatives along with the Democratic leader launched a bipartisan Task Force on Artificial Intelligence. The taskforce has a broad remit "to explore how Congress can ensure America continues to lead the world in AI innovation while considering guardrails that may be appropriate to safeguard the nation against current and emerging threats."[11]

The Committee on House Administration (CHA), as well as being represented on the taskforce, is looking at the utilisation of AI in the context of Congress itself. It held a private roundtable on the issue, entitled "Building Artificial Intelligence Guardrails for the People's House", which included external experts, senior staff from the House of Representatives and interested staff from Members' Offices. The roundtable informed the subsequent publication of guardrails for the use of AI in the House and these are reproduced below:

- **Human Oversight and Decision-Making:** Even as AI technologies are utilized for efficiency, human oversight remains crucial. Decisions, particularly those with significant impact, must ultimately rest with human experts who can interpret AI outputs within the broader context and with an understanding of potential limitations and biases.
- **Clear and Comprehensive Policies:** To ensure safe and appropriate use of AI, the development of clear and comprehensive policies is essential, as is the maintenance of accurate AI inventories. Any AI policy should address privacy, data security, and ethical considerations, ensuring that AI tools are known and used in ways that preserve public trust and safeguard sensitive information.

[11] Committee on House Administration, Artificial Intelligence Strategy and Implementation (April 2024): https://static1.squarespace.com/static/60450e1de0fb2a6f5771b1be/t/661fdcd5250a944842c3690e/1713364182472/CHA+AI+Flash+Report+April+2024.pdf

- **Robust Testing and Evaluation:** Before AI technologies are deployed, they should undergo thorough testing and evaluation to assess their reliability, validity, and potential biases. Continuous monitoring and regular reassessment are crucial to adapt to technological advancements and emerging risks. The continuous monitoring and evaluation of data that is entered into or taken out of the AI system is critical.

- **Transparency and Disclosure:** Transparency about the use and capabilities of AI technologies is key to maintaining trust. This includes clear communication about how AI is being used, the data it processes, and the measures in place to protect privacy and security. There should be considerations for disclosing when AI has contributed to legislative drafting or other significant tasks.

- **Education and Upskilling:** To effectively implement AI technologies, ongoing education and upskilling of Members and staff are essential. This includes training on the capabilities and limitations of AI, as well as on the ethical considerations and policy frameworks that guide its use. Building digital literacy and resilience among everyone in the House community will help mitigate risks and foster responsible use of AI.

The CHA also listed some notable achievements from legislative branch agencies and House entities in the period January to March 2024. For example, the Library of Congress:

> …has mapped out work through December 2024 which includes a focus on finalizing their AI Use Case Inventory, creating an AI definition and principles, evaluating current policies, developing new policies, and describing an AI governance model based on the Library's risk frameworks. The Library is also focusing on longer-term staff training and communication recommendations.[12]

[12] Committee on House Administration, Artificial Intelligence Strategy and Implementation (April 2024)

7.1 AI generated content

The United States Congress has to deal with an overwhelming amount of information from a range of stakeholders, and can lack the capacity to adequately process all of this data – what has been described as "a failure of absorptive capacity: the ability of an organization to recognize the value of new, external information, to assimilate it, and to apply it to desired ends."[13]

It has been argued that AI could be used to address a significant portion of routine tasks that in turn would allow staff to focus on higher value work. This viewpoint was set out in an article in Techpolicy.press, which explored how AI generated content was already being used and how it might be used in the future:

> Already, legislators are experimenting with these tools to augment their work. In January, Rep. Ted Lieu (D-CA.) penned an op-ed in the New York Times with the help of an AI, and subsequently introduced an AI-drafted bill. Around the same time, Rep. Jake Auchincloss (D-MA) read an AI generated speech on the House floor. While these examples may feel gimmicky and attention-seeking, AI tools will soon bring real disruptive innovation to legislatures—entailing substantial benefits as well as risks.

The article went on to suggest that:

> Specifically, AI could help improve the speed and efficiency of a range of office tasks, including:
>
> - Drafting routine communications such as op-eds, press releases, social media posts, speeches, dear colleague letters, oversight letters, casework letters, constituent letters, witness questions, and similar materials.

[13] Warner, Justin and Grant Tudor. "The Congressional Futures Office." Paper, Belfer Center for Science and International Affairs, Harvard Kennedy School, May 2019: https://www.belfercenter.org/sites/default/files/2019-06/PAE/CongressionalFuturesOffice.pdf

- Summarizing incoming information such as government documents, hearing transcripts, bulk emails, and interest group communications.
- Leveraging legislative data from Congress.gov and a half-dozen other sources for strategic insights. This could include analysis of potential co-sponsors, insights about past Congresses, or assessments for how likely a bill is to move.[14]

The House of Representatives has established a voluntary AI working group at the direction of the Committee on House Administration. This was accompanied by a "published collection of best practices on safeguarding sensitive data and clearly communicated which commercial AI services were institutionally authorized for use."[15]

The risks of legislatures failing to adapt to emerging AI technologies is real:

> Legislatures can either take steps to keep pace or be overwhelmed by the brute force changes that these technologies are bringing.[16]

For example, experts at the Brookings Institution warned in May 2023 that "AI-enabled technologies will make it much easier to create a large volume of comments (both genuine and fake) that could potentially flood government consultation processes."[17]

Furthermore, in relation to AI drafted bills, "LLM tools have resulted in a surge of Members and staff "drafting" bill text using third-party LLM services. And while these drafts may be formatted to resemble a bill, it is often the case that the LLM-produced text would not have the policy impact that the lawmaker intends

[14] Zach Graves , Daniel Schuman and Marci Harris, Bots in Congress: The Risks and Benefits of Emerging AI Tools in the Legislative Branch, TechPolicy.Press (February 2023): https://www.techpolicy.press/bots-in-congress-the-risks-and-benefits-of-emerging-ai-tools-in-the-legislative-branch/

[15] Popvox Foundation, Representative Bodies in the AI Era, (January 2024):

[16] As above

[17] As footnote 15

— as was recently the case in New York — or could lead to other unintended interactions with United States Code."[18]

The New York case relates to a bill researched and written by AutoGPT which is a more advanced version of ChatGPT. The bill requires landlords to provide tenants with a copy of their lease upon request up to twice a year. Yet critics questioned the value of the bill, including its practical impact and lack of enforcement clauses.[19]

There is a further consequence of the encroachment of AI into the world of legislative drafting. While it may be thought that it will allow the streamlining of processes, the opposite may be true:

> LLM tools are helping lawmakers and staff more easily explore policy ideas and turn them into proposals, which is increasing the volume of drafts submitted to the HOLC. The lawyers reviewing these submissions, however, must still undertake the same review process as always, without any corresponding increase in resources or tools to help them handle the increased workload.[20]

The Congressional Research Service (CRS) has begun exploring the potential application of AI to its work. In a statement before the Committee on House Administration Subcommittee on Modernization in March 2024, the Director of the CRS explained how the service was moving forward with AI generated AI content. Some of the approaches include:

- CRS has established a working group to identify near-, medium-, and long-term opportunities to integrate artificial intelligence (AI) applications into the day-to-day work of the Service
- CRS has initiated a number of AI related activities including:
    - the launch of an Artificial Intelligence Research Portal to provide resources for CRS staff related to the field of AI;

[18] Popvox Foundation, Representative Bodies in the AI Era, (January 2024)

[19] The New York Times, A.I. Wrote a Housing Bill. Critics Say It's Not Intelligent , (July 2023): https://www.nytimes.com/2023/07/14/nyregion/ai-ny-housing-bill.html

[20] As footnote 18

- evaluation of new, generative AI components of a commercial legal research database;
- the issuance of guidance to CRS staff prohibiting the use of Chat GPT or other AI tools in their work for Congress without prior CRS or Library approval
- In collaboration with the Library's Office of the Chief Information Officer, CRS is also exploring the utilization of automated processes to expedite drafting, reviewing, and publishing of bill summaries and the assignment of geographic and organizational subject terms.[21]

8 California legislature

In August 2023 the California legislature adopted the nation's first AI drafted resolution:

> Senate Concurrent Resolution 17 highlights the significant challenges posed by the use of technology, data, and automated systems, including incidents of unsafe, ineffective, or biased systems and unchecked data collection that threatens privacy and opportunities. At the same time, the resolution recognizes the potential benefits of AI, including increased efficiency in agriculture and data analysis that could revolutionize industries.
>
> The resolution affirms the state's commitment to President Biden's vision for safe AI and the principles outlined in the White House Office of Science and Technology Policy's "Blueprint for an AI Bill of Rights." The five principles -- Safe and Effective Systems; Algorithmic Discrimination Protections; Data Privacy; Notice and Explanation; and Human

[21] Committee on House Administration Modernization Subcommittee Hearing, "Legislative Proposals to Support Modernizing the Congressional Research Service and the Use of Federal Data (March 2024): https://www.congress.gov/event/118th-congress/house-event/116991?s=1&r=51

Alternatives, Consideration and Fallback - will guide the design, use, and deployment of automated systems in California.[22]

9 House of Commons

Along with the opportunities presented by AI, there are obvious risks, especially in a setting subject to intense scrutiny. Some of these risks have been highlighted in guidance issued by the House of Commons[23]:

- **Bias and discrimination:** Generative AI models can pick up on and amplify biases present in the data they are trained on, leading to discriminatory or unfair outcomes.
- **Misinformation, factual error and manipulation:** Generative AI can be used to create fake content, such as deepfakes, which can be used to spread misinformation or manipulate public opinion. It can also create factually erroneous content.
- **Intellectual Property infringement:** Generative AI can be used to create content that infringes on the intellectual property rights (IPR) of others, such as using AI to create a logo that is similar to an existing one. Entering parliamentary information into generative AI tools can result in the loss of IPR.
- **Ethical concerns:** Generative AI can raise ethical concerns.

10 Other selected legislatures

[22] Webpage of Claifornia Senator Bill Dodd, Legislature Adopts Sen. Dodd's Groundbreaking AI Resolution (August 2023): https://sd03.senate.ca.gov/news/20230814-legislature-adopts-sen-dodd%E2%80%99s-groundbreaking-ai-resolution

[23] AI guidance produced by House of Commons working group, 2024

The following table provides an overview of AI-related work in selected international legislatures[24]. It is not necessarily intended to reflect best-practice.

Table 2: Use of AI in selected legislatures

Argentina	The Argentinian Chamber of Deputies' "DipLab'' began experimenting with AI in 2019 — identifying processes that could benefit from AI, conducting training sessions, developing AI diagnostic tools with international partners, and creating an AI guide for parliamentary work. Existing AI projects include digitizing transcripts of plenary sessions from 2005 to 2020 and creating an editable transcript database in collaboration with the National Council for Scientific and Technical Research (CONICET). The algorithm indexes legislative information, allowing for effective searching and analysis of public policy trends and legislative changes. Future projects include developing a virtual tool for bill drafting with AI integration, digitizing all parliamentary archives, and launching additional AI-driven projects to improve efficiency and transparency in parliamentary work.
Estonia	The Estonian Parliament, the Riigikogu, has incorporated AI into its operations, using a system named HANS, based on large language models, to transcribe all parliamentary meetings, including plenaries and committees, through automated speech recognition. For plenary sessions, the AI-generated verbatim text is reviewed and corrected by editors before being published online. In committees, the digitally generated transcript is used by staff to manually prepare minutes. HANS has an error rate of around 5%, which can vary depending on the speakers and the meeting context. This implementation of AI has led to a reduction in staff, particularly

[24] Popvox Foundation, International Collaboration and Examples, (January 2024): https://www.popvox.org/representative-bodies-in-the-ai-era/international

	in replacing the Parliament's team of stenographers, who were nearing retirement and difficult to replace due to the nature of the job.
European Union	The European Union has adopted various AI-enabled tools to enhance efficiency and effectiveness including chatbots, which automate question-and-answer processes across multiple sectors, aiding in quick and efficient responses to queries. Additionally, the EP uses an auto-summarizer and Microsoft Editor to provide concise text summaries, facilitating the understanding of complex documents. The EP has also implemented eTranslation, a tool capable of translating documents into multiple languages, an essential function for a multilingual institution like the EP. Another significant integration is the Speech to Text service, which offers real-time transcription, enabling accurate and immediate documentation of spoken words during meetings and discussions.
India	In India, a newly inaugurated parliament building incorporates several AI-powered technologies to enhance operations, accessibility, and security. One innovation is the "Digital Sansad" application that allows Members of Parliament to listen to proceedings in their own languages in real time. This AI-enabled platform leverages automatic speech recognition to accurately transcribe and translate speeches word-for-word. It aims to facilitate participation and understanding across India's linguistic diversity.
Taiwan	Beginning in 2015, Taiwan's legislature integrated AI into its democratic processes through the vTaiwan (virtual Taiwan) and Pol.is platforms. vTaiwan is an online platform designed for public consultation and participatory legislation. Pol.is, an AI-driven tool, plays a crucial role in the vTaiwan process — participants can express their opinions on specific issues, and

	the AI system visualizes these opinions, highlighting areas of consensus and disagreement. This visualization helps legislators and policymakers to understand public sentiment on various issues in a nuanced and detailed manner. This innovative approach to lawmaking has been applied to various issues, including the regulation of online alcohol sales, telemedicine, and fintech regulations.

11 Inter-Parliamentary Union guidelines

The Inter-Parliamentary Union (IPU) has been at the forefront of the discussion on the application of AI in legislatures. As part of its work it is developing, in conjunction with senior technical experts, a set of guidelines on AI governance in parliaments which are due for publication in the second half of 2024. The guidelines will cover the following issues:

- AI applied to parliamentary proceedings
    - Legislative processes
    - Administrative processes
    - Citizen interactions
    - Government oversight
- AI applied to parliamentary proceedings
- Legislative processes
- Administrative processes
- Citizen interactions
- Government oversight
- Ancillary processes for AI governance
    - Information management
    - Information security
    - Risk management
    - AI portfolio management
- AI systems development process

- Human resources
- AI systems certification
- AI outsourcing[25]

In October 2023 the IPU organised an online event on AI in legislative processes. During the broadcast, an official from the Canadian House of Commons "explained that it was becoming increasingly difficult to find, train and retain stenographers, transcribers and other staff performing repetitive tasks, adding that pressure was increasing to deliver services faster and at higher quality. He argued that AI was therefore arriving at the right time, noting the importance of close cooperation between business departments and IT staff in order to find ways that AI could add value in parliamentary processes."[26]

An IPU Issue Brief, published in April 2024, set out a number of considerations that senior parliamentary staff should consider when moving ahead with the implementation of AI technology. These are summarised below:

Expect rapid change: the rapid development of AI technologies will continue, which will improve their precision and capabilities as well as likely reducing their costs. This will bring opportunities but also increased risks, requiring adequate control mechanisms and safeguards. The complexity of GenAI architecture makes it difficult to understand why it produces some results and not others for a given input. This can lead to "hallucinations", whereby GenAI produces inaccurate or misleading results.

GenAI could be overtaken by other, more powerful AI technologies which will be more predictable and coherent. Examples include:

- neuromorphic computing, which aims to replicate the human thought process digitally
- neurosymbolic AI, which combines the statistical, data-driven learning capabilities of neural networks with symbolic reasoning.

[25] Inter-Parliamentary union, Innovation Tracker (May 2024): https://www.ipu.org/innovation-tracker/story/co-creating-guidelines-ai-governance-in-parliaments

[26] Inter-Parliamentary union, Innovation Tracker, (May 2024)https://www.ipu.org/innovation-tracker/story/updates-ai-parliamentary-community

Understand impact and risk: the everyday tasks performed in a parliament – research, drafting bills and amendments, writing speeches etc. – could theoretically be performed by AI, but what level of safeguards would need to be in place before entrusting these tasks to AI?

The level of complexity of controls and safeguards will reflect the importance of the task. For example, AI generated legislation carries significant risk that would have a profound impact on people's lives, whereas other tasks could be more straightforward and carry minimal risk. "Explainability", that is why AI produces certain results, will aid transparency and accountability and result in accountability and ethical use. Parliaments will need to demonstrate to the public that these principles are built-in to their use of AI.

Ensure a strong culture of digital transformation: the adoption of GenAI in parliaments should be subject to the same routines as other technology solutions, namely discovery, learning, experimentation, development of simple use cases, and project initiation. Parliamentary leaders need to engage in continuous dialogue to decide when and how to deploy AI technology. Parliamentary committee hearings are a useful way to introduce members to the pros and cons of AI.

Be realistic about the limitations of AI: Generative AI is not fully reliable. LLMs must be trained to ethical standards and ensure the representation of balanced gender and minority perspectives. This is not straightforward, as these systems can only draw on what already exists. A key risk is that as Generative AI becomes more widespread in its use, it could become self-referential, learning from its own output.

Keep humans in the loop: Human scrutiny and control over Generative AI is crucial. This is particularly the case where AI is used in core parliamentary tasks, such as legislative drafting.

Build capacity through collaboration: The complexity of Generative AI means that individual parliaments may lack the expertise to implement these

systems safely and effectively. The early adopters will have valuable lessons that they can share with other legislatures.[27]

12 Conclusion

The advent of AI, and in particular Large Language Models such as ChatGPT, presents both opportunities and challenges for parliamentary institutions. This paper highlights innovations in legislatures across the world that have applied AI to legislation, committee meetings and public consultations. But AI also reflects the biases and limitations of the information it works with, and human oversight of AI produced content remains a priority.

What can be said with certainty is that AI applications will increasingly become a feature of parliamentary work, and it is important that institutions are prepared to both utilise the opportunities presented by this and also manage the risks that come with its use.

[27] Inter-Parliamentary Union, Using generative AI in parliaments, March 2024

Appendix 1

Ulysses 1: Content Aggregation and Thematic Organisation in Parliamentary Webpage

Ulysses 1 (U1), the first module of the Ulysses suite, is a specialised AI algorithm designed to enhance the organisation and presentation of thematic content of the Office of the Legislative Counsel on the Chamber's digital portal…This algorithm represents a significant step in digitising and streamlining access to legislative information, reflecting a commitment to both innovation and user-centric design.

The core functionality of U1 is its ability to categorise and display technical legislative content, including bills, technical studies, news, committees, speeches and pools based on specific themes or subjects selected by users. This categorisation process is dynamic, allowing for real-time reorganisation of content as new information becomes available or as user interests shift.

One of the key features of U1 is its responsiveness to user-selected topics. Upon choosing a theme, the algorithm instantaneously rearranges all relevant content on the portal, providing a comprehensive and tailored overview of the subject. By focusing on user experience and content accessibility, U1 serves as a bridge between the Chamber and the public, facilitating a more informed and engaged citizenry.

Ulysses 2 - Automating Legislative Counsel's Requests Distribution

Ulysses 2 (U2), as part of the Ulysses suite developed for the Brazilian Chamber of Deputies, serves a crucial role in streamlining the legislative process through its advanced machine learning and natural language processing capabilities. The primary function of U2 is the semi-automated distribution (specialists still can check/revise the U2's suggestions before distribution) of requests to specialised groups of legislative counsels, a task that was previously managed manually with some risk of cognitive bias and inefficiency.

The algorithm of U2 is designed to read and interpret textual requests submitted by parliamentary offices. These requests, often laden with complex and nuanced legislative language, are indicative of the diverse and intricate nature of legislative work. U2's task is to parse this language, understand the underlying request, and determine the most appropriate group of consultants to handle it. This process is crucial for ensuring that each request is addressed by experts most qualified in the relevant area of legislative drafting.

One of the key challenges that U2 addresses is the handling of requests that encompass new or blended themes, which are increasingly common in modern legislative environments. These requests, often straddling multiple domains or presenting new topics, require a sophisticated understanding and categorisation ability that U2 provides.

Before the introduction of U2, the allocation of such requests relied heavily on the subjective judgement of individuals. This method, while functional, was susceptible to inconsistencies and biases, leading to potential inefficiencies and uneven distributions of workloads. U2 revolutionises this process by employing a data-driven approach, significantly reducing the possibility of human bias.

Through its advanced algorithms, U2 presents three percentage-based recommendations, indicating the most suitable legislative counselling groups for each request. This not only simplifies the decision-making process but also introduces

a level of objectivity and precision previously unattainable. As a result, the overall efficiency of the legislative process is significantly enhanced, with quicker and more accurate responses to requests.

Ulysses 3: Semantic Analysis in Legislative Drafting Requests

Ulysses 3 (U3) is engineered to process a substantial volume of legislative requests efficiently. The Office of Legislative Counselling receives more than 20,000 demands per year, and U3 plays a critical role in managing this influx. Each request, varying from concise to complex legislative documents, requires the algorithm to perform an in-depth analysis to identify semantically similar existing legislative studies and proposals.

The primary function of U3 is to analyse requests submitted by parliamentary offices to the legislative counsel's office. This analysis is conducted using advanced natural language processing (NLP) techniques, allowing the algorithm to interpret the text of the requests in a sophisticated manner. The key here is the semantic analysis – not just searching for keywords, but understanding the context and meaning behind the words. This approach enables U3 to identify technical studies and legislative propositions that are semantically similar to the submitted requests.

Such capability is invaluable in legislative drafting, where similar or redundant proposals can be commonplace. By identifying these similarities, U3 aids counsels in drafting new legislative texts more efficiently and effectively. Additionally, it provides a means to notify parliamentarians if their requests closely resemble existing studies or proposals, enhancing the overall quality of the legislative work.

The algorithm must navigate the complexities of the Chamber's rules of procedures. For instance, there's a classification constraint that legislative information regarding the drafting of new documents is confidential until officially submitted. U3, therefore, operates within these constraints, ensuring compliance while performing its analytical tasks.

A significant challenge for U3 is the brevity and lack of detail in many requests. The algorithm has been fine-tuned to extrapolate and infer detailed insights from succinct inputs, a critical capability given that many requests for new legislation are concise and lack comprehensive detail.

By automating and streamlining the initial stages of legislative consultation and drafting, U3 not only saves time and resources, but also significantly reduces the potential for human error and oversight. Its ability to conduct thorough semantic analyses of requests marks a significant leap forward to a more efficient internal service delivery to parliamentarians and staff.

U3 also incorporates a feedback mechanism to continually improve its accuracy and relevancy. The responses it generates are evaluated for relevance by the legislative counsels, and this feedback is used to refine and optimise the algorithm’s performance.

Ulysses 4: Analysis and Management of Legislative Amendments

Ulysses 4 (U4), a vital component of the Ulysses suite, is a sophisticated artificial intelligence algorithm designed to enhance legislative drafting. This advanced tool specifically focuses on the critical task of managing and analysing amendments to legislative texts, one of the core activities of legislative work.

At its core, U4 utilises cutting-edge natural language processing (NLP) techniques. This enables the algorithm to read and interpret the language of amendments submitted during the legislative process. The NLP capabilities of U4 are not limited to merely identifying keywords but extend to understanding the context and deeper semantic meanings embedded within the text of the amendments.

One of its remarkable features is its ability to determine the appropriate placement of amendments within the broader context of legislative proposals. This involves a detailed analysis of the text to ascertain where each amendment fits within the structure of a legislative document, ensuring that amendments are contextually aligned with the sections they aim to modify or influence.

A significant challenge in legislative drafting is the handling of multiple amendments that may be similar or overlapping in nature. U4 addresses this challenge head-on by identifying similarities amongst various amendments. By clustering amendments with shared themes or objectives, the algorithm aids in streamlining the review process. This clustering not only saves time but also prevents potential conflicts or redundancies within the legislative text.

Moreover, U4's functionality is not merely confined to sorting and categorisation. It plays an instrumental role in the semantic clustering of amendments on the legislative proposition as a whole. This feature is crucial for Parliamentarians and advisors, who need to understand how each amendment could potentially alter the original intent or effectiveness of the proposition.

The processing capability of U4 is designed to function in near real-time, allowing for a swift and dynamic response to the influx of amendments typically received during legislative sessions. This timeliness is critical in maintaining the momentum of the legislative process and ensuring that amendments are dealt with efficiently.

Ulysses 5: Transforming Legislative Record-Keeping with AI-Driven Transcription

Ulysses 5 (U5), a component of the Ulysses suite of AI services implemented by the Brazilian Chamber of Deputies, represents a significant leap in the realm of legislative digitization and accessibility.

The primary function of U5 is the automatic transcription of voice to text. This feature is crucial in a legislative environment where proceedings, discussions, and debates form the core of the legislative process. These verbal exchanges, often complex and lengthy, are essential for record-keeping, legal documentation, and public dissemination (social and institutional control of the legislative process). U5's role in converting these oral communications into written form not only enhances transparency but also aids in the accessibility of legislative information to the public (external and internal to the Legislative Power) and other stakeholders.

In developing U5, the emphasis was on integrating it with the existing legislative systems, notably those supporting the traditional stenography or shorthand writing used in legislative assemblies. This integration ensures that the audio feed into the AI is of the highest quality, which is crucial for accurate transcription. It also signifies a significant step in modernising the legislative record-keeping process, merging traditional methods with advanced technology. The resulting system streamlines a stage in the process of recording the debates, which will still go through the final writing process, making it faster, more efficient, and less prone to human error.

Furthermore, U5 is designed to work in real-time, providing almost instantaneous written records of ongoing proceedings. This feature is particularly beneficial for immediate analysis, reporting, and decision-making processes within the legislative body. It also allows for quicker public access to legislative discussions, enhancing transparency and public engagement in the legislative process.

Ulysses 6: Semantic Analysis of the Public's Comments in Legislative Proposals

Ulysses 6 (U6) is the module specifically tailored to analyse and interpret the feedback received from citizens on various legislative proposals, primarily through electronic polls associated with legislative bills.

The core functionality of U6 revolves around its capacity to process and understand the semantic content of citizens' comments. Unlike traditional polling methods that offer binary options of approval or disapproval, the e-poll invites U6 invites citizens to provide detailed commentary on each bill. E-poll for bills is one of the services provided by the institutional portal of the Chamber of Deputies. The commentaries gathered by this service are then submitted to U6 for argumentation analysis and grouping. This qualitative approach allows for a more comprehensive understanding of public sentiment, going beyond mere numerical statistics to grasp the subtleties of public opinion.

When a legislative proposal is introduced into the system, an electronic poll is opened, inviting public participation. Citizens are encouraged not only to vote in favour or against the proposal but also to leave comments elucidating their viewpoints. These comments, which can vary significantly in terms of content, length, and complexity, provide a rich dataset for analysis.

U6 employs sophisticated NLP techniques to analyse these textual inputs. The algorithm is capable of identifying key themes, sentiments, and patterns within the comments. It does this by parsing the language used, understanding the context, and even detecting underlying sentiments and biases. This analysis is not limited to surface-level word matching; instead, it gets into the semantic meaning behind the words, enabling a more detailed understanding of public opinions.

The module's ability to cluster similar arguments is particularly noteworthy. Despite the diversity of expressions and language used by different individuals, U6 can group comments that share common themes or sentiments. This clustering is vital for parliamentarians, as it brings a wide array of opinions into comprehensible thematic categories, highlighting the main arguments for and against a bill.

Furthermore, U6 contributes to transparency and informed decision-making in the legislative process. By systematically analysing and presenting the public's views, it ensures that parliamentarians have a clear understanding of the citizen's views on various issues. This insight can be crucial in guiding legislative decisions, amendments, and discussions.

The insights generated by the module will soon be made accessible to parliamentarians through a user-friendly interface, often in the form of summarised reports or visual data representations. This seamless integration ensures that the valuable insights provided by the public are readily available and can be effectively utilised in the political decision-making process.

Ulysses 7: Biometric Authentication Integration with Infoleg for Parliamentary Deliberations

Ulysses 7 (U7) is the module specifically developed to integrate with the Infoleg application, used by members of the parliament. This module addresses a critical aspect of legislative procedures: authentication during deliberations in plenary sessions and committee meetings. In light of concerns about the potential for fraud in the virtual and hybrid voting system, like staff voting on behalf of Parliamentarians, the development of U7 becomes crucial.

One of the standout features of U7 is its focus on leveraging biometric authentication, particularly facial recognition, to verify the identity of members participating in various parliamentary activities through the Infoleg App. This technology

ensures that the members can engage in deliberations without being physically present in the plenary, thereby enhancing flexibility and participation, in situations permitted by the rules of procedure

The module employs sophisticated algorithms to not only recognize faces but also discern facial movements. This capability is critical in verifying the active presence of members during sessions. By implementing random movements for authentication, such as looking in different directions, U7 adds an additional layer of security, eliminating the possibility of using static images for authentication purposes.

Ulysses 8: Application of Facial Recognition for the Communication and Photos of Parliamentarians

Ulysses 8 (U8) was the module whose primary function revolved around facial recognition technology. The deployment of facial recognition by U8 was confined to a specific context – identifying parliamentarians in photographs taken during public meetings, such as Committees and Plenary Meetings. This targeted application reflects a conscious effort to balance the benefits of AI with the need to respect privacy and avoid intrusive surveillance.

The facial recognition system in U8 operated by scanning photographs and accurately identifying the parliamentarians present in them. This process involved a complex interplay of machine learning algorithms that analyse facial features and cross-reference them with the official database of the Parliament. The sophistication of the system was evident in its ability to handle the nuances of facial recognition – differentiating between individuals who may have similar features, adjusting for variations in lighting and angle, and even accounting for changes in appearance over time.

The utility of U8 extended beyond mere identification. Once the parliamentarians were identified in the photographs, the system automatically tags and labels the images with their names. This functionality serves multiple purposes. Firstly, it streamlines the process of cataloguing and archiving photographic content, a task that, if done manually, would be

labour-intensive and prone to errors. Secondly, and perhaps more significantly, it aids in the rapid dissemination of these images for public relations and communication purposes. Parliamentarians can swiftly access these photographs, which they can then utilise for various forms of outreach and engagement with their constituents, for example in social media. This rapid access to personalised content is particularly valuable in the digital age, where social media and online platforms play an important role in political communication.