Report into the Role of Parliament in Shaping the Future of Responsible AI conference

(28–30 November 2025, Kuala Lumpur)

Hannah Blythyn MS & Adam Price MS

Introduction

Parliaments are often described as "talking shops". In Kuala Lumpur in late November, they were reimagined as something else entirely: global machines for collective intelligence that must now learn to work alongside artificial intelligence or risk being outpaced by it. Democracies worldwide find themselves in a moment when they need to forge a way forward that embraces the opportunities of new technologies but in doing so to not fall short in shaping a values-based approach to ensure the responsible application and utilisation of that technology.

This report sets out what happened at The Role of Parliament in Shaping the Future of Responsible AI conference (28–30 November 2025, Kuala Lumpur), co-convened by the Inter-Parliamentary Union (IPU), Commonwealth Parliamentary Association (CPA) and UNDP, and attended by over 200 parliamentarians and staff from around the world.

Hannah Blythyn MS and Adam Price MS represented the Senedd. We focused on three things:

- Bringing a distinctly Welsh perspective to debates about skills, human judgement and democratic resilience in the age of AI.
- Arguing that parliaments must share not just lessons but infrastructure – culminating in a new Recommendation, added at our delegation's request, on co-designing a shared parliamentary AI platform.
- Connecting global lessons back to the Senedd's own work on AI governance and to the Welsh Government's AI Cymru plan.

1. What the conference was trying to do

The organisers were explicit about their aim: this was not another generic "future of tech" event. It was a working meeting on how parliaments, specifically, should govern AI and use it to augment not undermine democracy.

The agenda was split into two main tracks, plus cross-cutting plenaries:

- **AI policy track** – how to legislate for and oversee AI: from societal transformation and capacity-building to law-making, gender and equality, and international cooperation.
- **AI in parliament track** – the practical side: what AI tools are already being used in parliaments, what works, what fails, what foundations need to be put in place and how to put governance around it.

Across three days, participants moved between big-picture plenaries, detailed case studies (from Mauritius to Chile and Thailand), small-group work and hands-on exercises, including an AI "jailbreaking" game showing how easily safeguards can be circumvented.

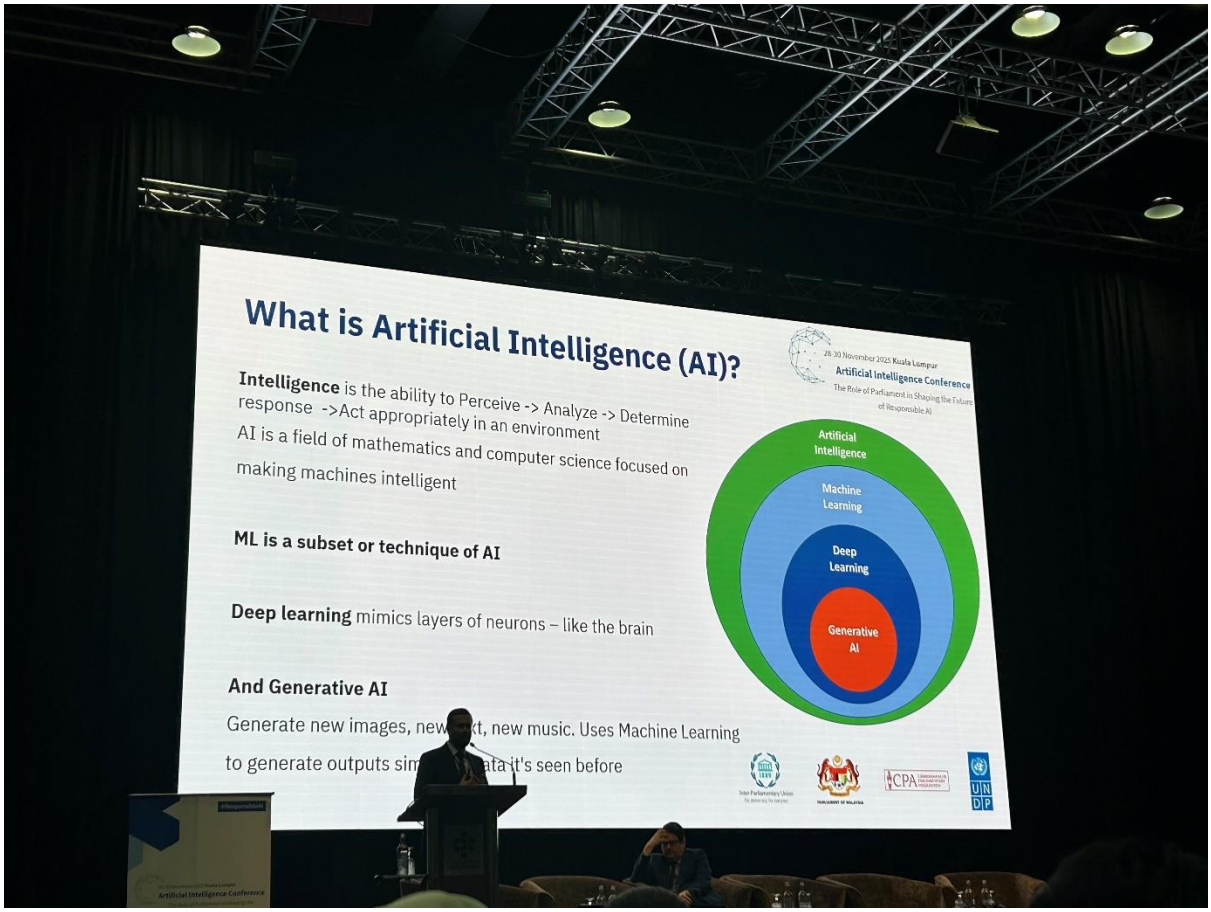

2. Five cross-cutting themes

Although the sessions were diverse, a small set of themes kept resurfacing.

2.1 You cannot skip the plumbing

"Beautiful AI applications require beautiful underlying data." That sentence, repeated in different forms, captured the consensus: parliaments that have not yet digitised their archives, modernised their workflows, considered how their workforce need to be supported and involved, or cleaned their data cannot realistically expect sophisticated AI to perform well.

The message was unglamorous but clear:

- Digitisation (converting paper archives to structured digital data).
- Digitalisation (modern tools for drafting, reviewing and publishing).
- Then AI.

Trying to invert that order, the conference heard, leads to unreliable systems and public mistrust.

2.2 Capacity, not cash, is the real bottleneck

Money matters but lack of skills, confidence and political ownership matters more. Participants from several parliaments complained that IT teams were "running ahead" while members were not engaged; others described AI pilots launched without any consultation, strategy or oversight.

The real gap, echoed across sessions, was **capacity:**

- Members who do not yet understand AI well enough to legislate on it or use it safely.
- Staff who lack training to implement and maintain AI tools.
- Institutions that have not yet created the cross-functional structures, bridging policy, legal, ICT and research, that successful parliaments now rely on.

2.3 Data, sovereignty and a looming AI divide

The most strategic anxiety was about **who owns and controls the underlying data and infrastructure**. Delegates worried that many countries risk becoming permanent consumers of foreign AI, with little control over models trained on their citizens' data.

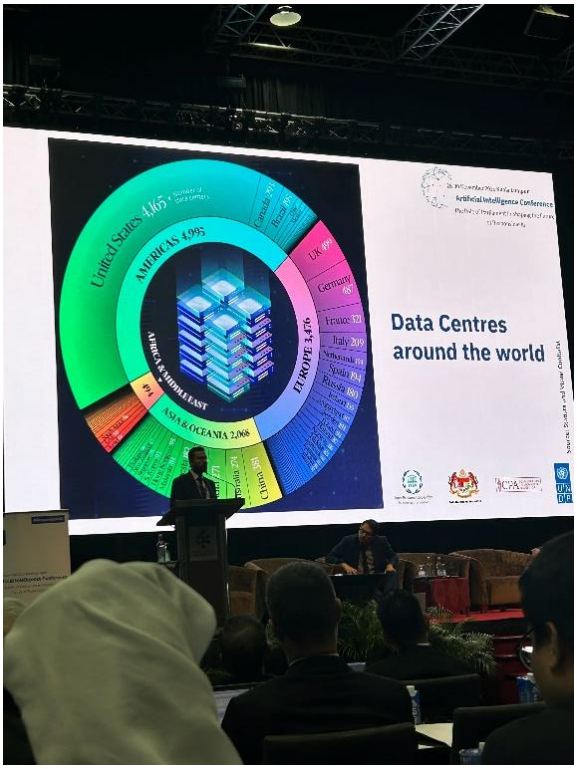

The "bonus insight" in one synthesis session was stark: a global AI divide could harden into **parliamentary inequality**, with richer states able to build sovereign AI systems while poorer ones rely on off-the-shelf tools they neither understand nor influence.

That concern sat behind a recurring interest in shared, open and co-developed tools – precisely the space in which Recommendation 13, championed by our delegation, ultimately landed.

2.4 Democracy's race against time

A second anxiety was temporal. AI systems answer questions in seconds; legislative processes take months or years. In one of the closing discussions, Adam Price suggested that parliaments are machines for collective intelligence but risk losing the "war on time" if people find that AI systems respond to their problems faster than politics does.

The point resonated. Several interventions across the conference stressed that if citizens come to see democratic institutions as slow, opaque and technologically illiterate while AI tools appear quick, personalised and fluent, trust could erode rapidly – especially in moments of crisis, when misinformation and deepfakes spread faster than parliamentary hearings can be scheduled. The conference also sought to emphasise the point that although AI is already with and all round us, the direction of AI is not pre-determined and that as parliamentarians we should approach the application of AI from a citizen not consumer perspective.

2.5 Human judgement must stay in the loop

A more hopeful leitmotif was that AI can act as a "cognitive exoskeleton" for parliaments amplifying human judgement, not replacing it. But that depends on how institutions design and govern its use.

Here, **Hannah's intervention** cut through. In discussion on skills and adoption, she warned against hollowing out human thinking capacity by treating AI as an automatic pilot rather than a tool to support learning and critical reflection.

That warning echoed strongly with the Senedd's own risk register where over-reliance on automated aids and the erosion of higher-order thinking are explicitly flagged as dangers to be managed, not progress to be celebrated.

---

3. Session highlights

What follows is a brief tour through the main sessions, noting where they speak most directly to our work in Wales.

3.1 Navigating societal transformation in the age of AI

This opening policy session set the tone by placing AI inside today's "polycrisis": climate change, geopolitical tension, economic inequality, information disorder. An initial poll showed participants' top worries: inequality, misinformation, "fake news" and data governance dominated the word-cloud.

Speakers warned that generative AI will both intensify existing risks (deepfakes, personalised disinformation, cyber-attacks) and open new ones (autonomous weapons, pervasive biometric surveillance). Parliamentary responses discussed included:

- Setting **red lines** (eg, bans or moratoria on certain uses).
- Risk-based regulation inspired by the EU and other models.
- Stronger transparency obligations on both governments and private providers.
- Dedicated parliamentary mechanisms to monitor cross-border and cross-sector risks.

The underlying message: AI policy cannot be siloed. It touches security, labour markets, health, education and democracy simultaneously. Legislatures need the capacity to see those interconnections.

3.2 Building national capacity

The second policy session drilled into what "capacity" actually looks like. Using a simple framework – pro-profit, pro-people, pro-planet, pro-potential – speakers argued that AI policy needs to pass four tests at once: economic viability, social benefit, environmental sustainability and human development.

A hands-on "jailbreaking" exercise showed how easily a model's guardrails can be bypassed to reveal hidden prompts or generate harmful outputs. The exercise, playful on the surface, underlined a serious point: if parliamentarians do not understand how these systems behave at their limits, they cannot credibly oversee their use in public services or elections.

Key recommendations from group work included:

- **Baseline mapping** of where AI is already used in government, what decisions it shapes and what safeguards exist.
- **Audits of existing law** (data protection, anti-discrimination, cybersecurity) before rushing into new AI-specific Acts.
- **Training for MPs and staff**, with identified "AI champions" in each parliament.

3.3 Law-making and oversight

This session asked a deceptively simple question: what can parliaments do now, with the tools they already have?

Examples ranged from Korea, where 19 separate AI-related bills were consolidated into a single flexible framework, to Bahrain's use of expert hearings on AI governance.

The main conclusions were pragmatic:

- Treat AI not just as a topic for new laws, but as a lens for **using existing oversight tools** – questions, inquiries, hearings, budget scrutiny – to uncover harms early.
- Create or empower specialised committees or cross-party groups with a clear mandate to scrutinise AI across sectors.
- Embed human-rights and equality impact assessment into any new AI legislation.

These points are mirrored almost verbatim in Recommendations 6–8 of the final communiqué.

3.4 Gender & AI

The gender session made a blunt observation: AI systems inherit the data they are trained on, and much of that data carries historical bias. If left unchecked, AI could deepen the under-representation of women and marginalised groups in everything from recruitment to access to public services and diagnosis tools. There was also a focus on the dangers of what is emerging as TF GBV – technology facilitated gender based violence – and how that could be turned on its head with AI potentially being used as a countermeasure to TF GBV, whether that is through regulation, victim support systems or education. The latter could be an opportunity for us here in Wales in line with digital literacy work and RSE.

Participants in the session pressed for:

- Gender-disaggregated data and explicit attention to intersectionality in AI regulation and public-sector deployments.
- Gender responsive AI governance alongside online safety regulation and education.
- Parliamentary inquiries into gender bias and equity in AI systems (later reflected in Recommendation 9).
- More diverse teams designing and auditing AI within both public institutions and vendors.

Session attendees also signed up to taking away the following focused actions:

Within three months:

1. Table a written question or oral question to the relevant Minister on what safeguards exist to prevent gender bias in public-sector AI systems.

Within six months:

2. Advocate for and support measures to ensure that women are equitably included in decision making groups in parliament responsible for internal technology deployment, digital services, AI adoption, in line with equality standards and best practice.

Within twelve months:

3. Initiate or support at least one committee inquiry, to examine gender bias and levels of diversity and equity in AI systems relevant to the jurisdiction.
4. Advocate relevant and appropriate amendments to digital safety laws to address non-consensual synthetic media.

3.5 International cooperation

The international session tried to reconcile two tensions: AI as a global technology and AI as something that lands in very local contexts. Parliaments have different constitutional roles, but face common problems: asymmetric information vis-à-vis tech companies, cross-border data flows, and citizens exposed to platforms regulated elsewhere.

Here, organisations like IPU, CPA and UNDP were presented not just as talking shops, but as **collective bargaining platforms:** places where smaller parliaments can pool expertise, develop common standards and negotiate more effectively with large technology firms. This logic sits behind Recommendations 10–12 on multi-stakeholder dialogue and inter-parliamentary exchange.

3.6 Journeys towards AI adoption in parliaments

This was one of the most practically useful sessions. Parliaments from Europe, Africa, Asia and Latin America described their AI experiments – warts and all.

Some patterns emerged:

- Early pilots tended to focus **on text-heavy, repetitive tasks:** summarising debates, classifying documents, search and translation.
- Success often came from **small, low-budget projects** with clear use cases, not grand strategies. (The European Parliament's first generative AI project was run by a single staff member on a four-digit budget.)
- Long-term ambitions are growing: several parliaments are now exploring bespoke language models trained on parliamentary data, or shared tools for regions with common languages.
- Parliaments need to learn more about AI, what it means and what it can do. A strategic and inclusive approach needs to be taken and guardrails around use, security of data, verification, confidentiality, and transparency.

It was in this context that **Hannah's contribution** landed: she cautioned that AI should support, not supplant, human skill-building – especially among staff and younger members who might be tempted to outsource thinking to machines. That concern closely tracks the Senedd's own risk analysis. It was also clear that where the use of AI has been most successful, elected members have been involved been involved in defining the direction – something that the Senedd should take note of. There are a number of anecdotal examples of different Parliament's use and/or approach to AI which can be shared separately. In short, parliaments approach to AI should not be driven simply by software like Co-pilot but to take AI as a co-pilot to Parliament.

3.7 Building blocks for AI in parliaments

The "building blocks" session provided a kind of operating manual. Drawing heavily on the EU AI Act and the European Parliament's response to it, speakers argued that responsible use requires three pillars:

1. **Governance** – clear policies, risk frameworks, ethics oversight and decision-making structures.
2. **Technical and data capability** – modern infrastructure and integrated data systems.
3. **Organisational capability** – roles, skills, training and culture.

The newly launched **IPU AI Maturity Framework** was introduced as a diagnostic tool for parliaments to assess themselves across governance, technical capability, organisational capability and democratic impact. A live poll revealed that no parliament in the room claimed to be at "advanced practice" or "global standard-setter" level; most clustered at early stages, experimenting with pilots and drafting initial policies.

For Wales, this framework offers a ready-made lens through which the Senedd Commission could review its own AI work.

3.8 Equipping Parliaments for the AI era

This final session in the trio of the parliamentary track encouraging participants to think practically and strategically about the current and future use of AI in their respective parliaments – starting from a baseline and then charting a course to benchmark against. From a standing start of asking the question 'what does digital literacy mean and why does it matter for legislators and parliaments?' Participants established the need to establish and then develop AI literacy (and ongoing upskilling) and a plan for how it could be used within a respective parliament, alongside establishing the guidelines needed to ensure its responsible and ethical application. The discussion included a wide range of areas from safety and confidentiality to clarity on the effective use of AI and how it could aid citizen engagement.

It's important parliaments put these building blocks into place, as one delegate summed up "how can you legislate on AI if you don't understand it yourself?"

4. Our delegation's contributions

4.1 Keeping human thinking at the centre (Hannah Blythyn MS)

In the "Journeys" discussions, Hannah intervened to highlight a risk that is sometimes underplayed in more technical debates: the erosion of human judgement and skill. She warned against the temptation to treat AI as a shortcut that replaces learning rather than supporting it, especially in legislatures where institutional memory and critical thinking are core assets.

This resonated strongly with colleagues from other small and medium-sized parliaments, and chimed with the Senedd's formal risk register, which notes the dangers of "automation bias" and over-reliance on AI tools.

4.2 Framing the democratic risk and proposing a shared platform (Adam Price MS)

In the closing discussions, Adam Price argued that parliaments are, at heart, machines for collective intelligence—but that they risk losing a "war on time" if AI systems provide faster, more personalised responses to citizens than politics does. In that scenario, people may simply route around democratic institutions.

Building on that diagnosis, our delegation pressed for a concrete remedy: **shared parliamentary infrastructure**. Rather than each parliament trying (and mostly failing) to build its own tools from scratch or depend solely on commercial platforms, we proposed that IPU, CPA, UNDP and partners work with parliaments to co-design a **shared parliamentary AI platform**, with a common core and nationally adaptable components.

This idea was taken up enthusiastically in the drafting of the outcome document and now appears as **Recommendation 13**—a late addition to the communique that did not feature in the original draft. It commits parliaments to:

"Participate, through the IPU, the CPA, UNDP and other international partners, in the co-design of a shared parliamentary AI platform… that would address resource gaps, facilitate cross-border collaboration and knowledge-sharing, and support parliamentarians and staff with casework, scrutiny, legislative drafting and policy development."

For a small, digitally-minded parliament like the Senedd, this is potentially transformative: it offers a route to high-quality tools we could not realistically build alone, while maintaining public control, ethical standards and bilingual accessibility.

5. What this means for the Senedd and for Wales

The Senedd Commission and Welsh Government are not starting from scratch. The Commission has already established **an AI Governance Group** and an **AI Opportunities Group**, is piloting Microsoft Copilot tools, and has identified corporate AI risk with detailed sub-risks ranging from bias and hallucination to data privacy, transparency, cyber-security and over-reliance on automation.

Meanwhile, the Welsh Government's AI Cymru plan sets out a vision for ethical, empathetic, enterprising and effective use of AI, underpinned by four strategic pillars: economic growth, educating Wales, equitable delivery and enabling infrastructure.

Against that backdrop, the main implications from Kuala Lumpur are:

1. **Use the IPU AI Maturity Framework to benchmark the Senedd.**
   The Commission should consider formally applying the IPU tool to assess where we stand on governance, technical capability, organisational capability and democratic impact, and to identify gaps to address over the next two years.
2. **Deepen cross-functional governance and Member engagement.**
   The Senedd's AI Governance and Opportunities groups already mirror best practice seen internationally. The next step is to keep Members centrally involved, including through regular briefings, hands-on workshops and identifying cross-party "AI champions" in the Senedd.
3. **Align parliamentary practice with AI Cymru's principles.**
   Conference discussions on fairness, inclusion, bilingual access, public trust and human-centred services fit closely with the Welsh Government's principles of ethical, empathetic, enterprising and effective AI. Parliamentary scrutiny of Welsh AI initiatives can explicitly use those principles as benchmarks.
4. **Prioritise skills and literacy as much as technology.**
   Consistent with Hannah's intervention and the Commission's risk note, the priority should be sustained investment in Member and staff capability—so that AI augments skills rather than replaces them. That means training on how systems work, where they fail, and how to challenge outputs, not just how to prompt them.
5. **Actively shape Recommendation 13's shared platform.**
   Given our role in securing Recommendation 13, there is a strong case for the Senedd to be an active partner in its follow-up:
   - Contributing use-cases (eg, bilingual search, committee briefings, citizen communication).
   - Pressing for strong governance, transparency and open standards.
   - Ensuring smaller and non-metropolitan legislatures like ours are treated as design partners, not just end-users.

6. **Strengthen scrutiny of AI across Welsh public services.**
   The Welsh Government is already deploying AI in health, education, local government and the third sector. The Senedd can now draw directly on the

conference's recommendations—especially 6–9—when designing scrutiny work: mapping use cases, commissioning impact assessments, and holding inclusive national dialogues on values and red lines.

7. **Consider how the Senedd Code of Conduct could and should set out the responsible use of AI**
   Parliaments need to keep pace with the evolution of technology, at the same time of having clear guidelines regarding how it is used – responsibly – by elected Members. The Standards of Conduct Committee is currently considering amendments to the Code of Conduct to clarify Members' use of social media, there is an additional opportunity to lead the way in setting out clear expectations regarding Members use of AI.

---

6. Next steps

We suggest two immediate, practical follow-ups for the Senedd:

1. **Maturity assessment and roadmap.**
   Ask officials to pilot the IPU AI Maturity Framework, then bring back a short roadmap with 12–24 month priorities for parliamentary AI governance, experimentation and utilisation. This should have the involvement of elected Members as well as Commission staff and additionally include engagement with staff representatives.

2. **Engagement on Recommendation 13.**
   Through the IPU and CPA networks, signal the Senedd's willingness to help steer early work on the shared parliamentary AI platform and ensure Welsh (and bilingual) requirements are on the table from the outset.

Kuala Lumpur showed that no parliament yet has all the answers on AI. But it also showed that the ones asking the right questions – about data, skills, ethics and shared infrastructure – are already shaping the future. Our goal, as a small but ambitious legislature, should be to stay firmly in that group.