

**Report on Centre for Public Policy and
Administration Seminar on
'The Potential for AI in the Public Sector'
15 April 2025**

Contents

Foreword.....	2
Fundamental Questions.....	2
Executive Summary	3
Key Themes	3
Recommendations	4
1. Introduction	5
1.1 Definitions.....	5
1.2 Why This Matters Now.....	5
1.3 Context.....	5
2. Key Discussions and Themes	9
2.1 AI and Public Sector Decision-Making	9
2.2 Trust and Ethics in AI.....	10
2.3 Security and Sovereign AI Compute	11
2.4 AI Skills and Workforce Development.....	12
2.5 AI Collaboration and Ecosystem.....	14
2.6 Government Perspectives on AI Adoption	15
2.7 Crystallising Views: Reflections from Discussion Panel on Trust, Accountability, and the Future of Public Sector AI	17
3. Recommendations	19
1. Establish a Permanent Public Sector AI Forum.....	19
2. Create a Public Register of AI Use in Government	19
3. Introduce Mandatory AI Risk Audits in Procurement	19
4. Pilot a Sovereign AI Design and Implementation Platform	20
5. Deliver an AI Literacy and Capability Framework.....	20
Acknowledgements	21
About the CPPA	22
Education	22
Scholarship.....	22
Translation.....	22
Exchange	22
Contact.....	22

Foreword

The Centre for Public Policy and Administration (CPPA) would like to extend its sincere thanks to all who participated in our seminar, *The potential for AI in the Public Sector*, held on 15 April 2025 at Riddel Hall, QUB. The time, insights, questions, and willingness to engage in open dialogue from all involved made this event particularly valuable.

From local and international academics, technical domain experts and policymakers, to industry innovators and digital transformation leaders, each voice added depth, challenge, and urgency to the discussions. For the CPPA, it was a privilege to convene this group of thinkers and practitioners, not only to share fresh ideas, but also the desire to ensure that innovation in public service delivery remains a shared priority.

In our conversations and feedback following the event, a unifying theme emerged - that there is a collective ambition to advance the use of AI and emerging technologies in a way that enhances accountability, strengthens effectiveness, and ensures that services remain fit for the future.

Fundamental Questions

So what kind of government do we want AI to enable? One that is faster and more efficient - or one that remains responsive to public voice, democratic scrutiny, and the unpredictability of real life? Are there tensions between these goals, or are there meaningful synergies?

The rapid adoption of AI holds clear potential to enhance government efficiency, particularly in speeding up the analysis of complex public data. But as AI systems begin to inform or even drive decisions, serious questions emerge such as: Will this shift favour automated logic over lived experience? Could a reliance on automated systems leave less room for public discourse, democratic oversight, and meaningful political participation?

These concerns have captured the attention of academics, who are increasingly focused on how to design AI that supports rather than supplants the democratic process. Contributors at our seminar voiced their support for AI not as a replacement for human judgment, but as a tool to augment it. The focus now must be on exploring how AI can provide richer context, better data integration, and more transparent foundations for decisions, without undermining equity or accountability.

So where are the opportunities? For government, the real value lies in harnessing AI that connects information across public organisations, reduces duplication, and empowers civil servants to make more informed, timely, and citizen-focused decisions. For researchers, the challenge is to develop AI systems that are not just technically efficient, but aligned with democratic concerns and values.

If designed and deployed thoughtfully, AI has the power to improve the lives of citizens, not by replacing the human touch, but by strengthening it.

We hope this report serves as both a record and a springboard. It captures the richness of the conversations that took place, but more importantly, it sets out an agenda for action.

Centre for Public Policy and Administration
Queen's University Belfast

Executive Summary

This report provides an account and analysis of the recent seminar, '**The Potential for AI in the Public Sector**', held on 15 April 2025 at Queen's University Belfast's Riddel Hall¹. The seminar brought together senior policymakers, leading academics, and AI practitioners from Northern Ireland, the wider UK, Ireland, and Europe.

Amid accelerating AI innovation globally, the seminar explored the critical conditions necessary for successful and ethical AI deployment within public administration, including robust governance frameworks, enhanced AI literacy, and strategic infrastructure investment.

Key Themes

The seminar presentations and conversations gravitated around four central themes:

- **Cohesion and Vision in Public Sector AI:** Speakers broadly agreed that while individual initiatives to deploy AI are valuable, they should form part of a coherent, collective vision. Professor Chris Johnson, Chief Scientific Adviser at the UK Department for Science, Innovation and Technology (DSIT), emphasised this, warning, "... without vision, all we will see are barriers." There was consensus that a unified strategy, anchored in clearly defined principles, is critical for strategic coherence and long-term impact.
- **Selective and Purposeful Use of AI:** Several speakers underlined that AI should be introduced selectively rather than ubiquitously across public services. Professor Madalina Busuioc, Professor of Public Governance at VU Amsterdam, cautioned that government bodies must critically assess the appropriateness of AI deployments, saying: "Trust without evidence is not trust."
- **Enhancing AI Literacy:** Seminar presentations highlighted the pressing need to improve AI literacy at every level of public administration. Professor Helen McCarthy and the Artificial Intelligence Collaboration Centre's (AICC) Director David Crozier argued vigorously for raising standards of AI understanding among policymakers, the private sector and the adult population more widely. They stressed that better AI literacy is essential to ensure that decision-makers comprehend the algorithms that will be deployed, avoiding the inadvertent harms witnessed in poorly governed AI initiatives elsewhere.
- **Trust, Transparency, and Accountability:** The seminar featured weighty discussion around pragmatic advocacy for assertive AI deployment, despite known imperfections, when considered in tandem with rigorous testing and accountability. Many argued that delaying AI implementation awaiting perfection would be detrimental, with one speaker asserting: "A ship is safe in its harbour, but that's not what it's for." More transparency and auditing - designed to obviate the temptation to implement opaque AI systems - would speed up the process towards AI adoption.

¹ <https://www.qub.ac.uk/research-centres/cppa/seminar-series/public-sector-ai-seminar/>

Recommendations

Based on insights from speakers and subsequent panel discussions, this report recommends five high-impact actions to accelerate AI adoption responsibly and effectively:

1. **Establish a Permanent Public Sector AI Forum**

We suggest the creation of the Public Sector AI Forum to act as a channel for views from across the island drawn from policymaking, academia and private sector innovators. This forum would monitor and publish reporting across five pillars: data; skills; infrastructure; public service transformation; and ethics. The forum should aim to share practical lessons, unblock delivery challenges, and feed into AI policy developments at a national level.

2. **Create a Public Register of AI Use in Government**

Every devolved and local government department should publish active and pilot AI use cases, detailing system purpose, training data, risk status, and performance review. This would ensure transparency, foster public trust, accelerate responsible adoption and support academic and third-party validation.

3. **Introduce Mandatory AI Risk Audits in Procurement**

Risk auditing requirements should be embedded into contracts for all AI tools affecting policy or citizen-facing services. This would require independent assessors to evaluate bias, suitability, and potential harm before systems go live.

4. **Pilot a Sovereign AI Design and Implementation Platform**

A sovereign, publicly-governed AI design and implementation platform to enable safe, scalable, and values-aligned AI experimentation and deployment across public services should be piloted for Northern Ireland. The Northern Ireland Executive should strongly advocate for a Sovereign AI Compute presence in the region to support regional autonomy, data protection, and strategic innovation while reducing reliance on external cloud vendors or opaque infrastructure.

5. **Deliver an AI Literacy and Capability Framework.**

There is a strong case now for the definition of minimum AI literacy standards for senior civil servants, policy professionals, and delivery teams. This would involve roll out of short-form professional training across government, focusing on responsible use, procurement, and use-case evaluation.

1. Introduction

Artificial Intelligence is no longer an emerging technology. It is here, embedded deeply within critical sectors from healthcare and security to education and public governance. Governments worldwide recognise AI's transformative potential, alongside the equally significant challenges of transparency, trust, and accountability.

1.1 Definitions

For the purposes of our seminar and this report, we are using the OECD and 'Artificial Intelligence Playbook'² definitions.

'An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.'

1.2 Why This Matters Now

Recent international initiatives, including the **AI Summit in Paris**³ and the ongoing dialogue around the UK's **AI Opportunities Action Plan**⁴, underscore the urgency with which public sector organisations must approach AI deployment. The discussions on 15 April identified significant opportunities for public organisations in Northern Ireland to adopt responsible, targeted AI initiatives, avoiding pitfalls encountered by other jurisdictions.

Professor Chris Johnson encapsulated the urgency succinctly during his keynote, asserting that Northern Ireland faces a stark choice: rapidly develop a local AI vision or risk falling behind. Highlighting examples from jurisdictions like Estonia, he called for Northern Ireland to "own" its AI strategy and "not expect Westminster to solve its AI readiness challenges."

1.3 Context

There has been a great deal of movement in recent months on technology advancement and policy development. International trading relationships are changing too, with the UK government signing off on trade deals with the US⁵, India⁶ and the EU⁷ in recent weeks. Indeed, much of the conversations at our seminar and feedback we received were reacting to, and reflective of, the context outlined below.

The Paris AI Summit in February brought together over 1,000 participants from more than 100 countries. The summit marked a global pivot toward enforceable AI safety standards, international coordination, and sovereign infrastructure investment. A total of 58 countries, including France,

2

https://assets.publishing.service.gov.uk/media/67aca2f7e400ae62338324bd/AI_Playbook_for_the_UK_Government_12_02_.pdf#page=16&zoom=100,96,290

³ <https://www.elysee.fr/en/sommet-pour-l-action-sur-l-ia>

⁴ <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>

⁵ <https://www.gov.uk/government/news/landmark-economic-deal-with-united-states-saves-thousands-of-jobs-for-british-car-makers-and-steel-industry>

⁶ <https://www.gov.uk/government/publications/uk-india-trade-deal-conclusion-summary/uk-india-trade-deal-conclusion-summary>

⁷ <https://www.gov.uk/government/news/pm-secures-new-agreement-with-eu-to-benefit-british-people>

China, and India, signed a joint declaration emphasising the development of AI that is open, transparent, ethical, safe, and trustworthy. Notably, the US and UK did not sign the declaration, citing concerns over its scope and implications for national security.

AI development will see an injection of funding as a direct result of the summit. The EU's €200 billion AI initiative - anchored by €20 billion for four AI gigafactories, alongside the \$400 million 'Current AI' foundation backed by global tech leaders and governments - signals an unprecedented wave of investment in AI infrastructure and public goods. Northern Ireland, with its established expertise in cybersecurity and digital innovation, is well positioned to harness this momentum and attract a share of this capital through targeted research, infrastructure development, and cross-border collaboration.

The AI Opportunities Action Plan published in January outlines a considered route map toward boosting economic growth, creating jobs, and improving people's lives through the adoption of AI. The plan focuses on three main goals: investing in the foundations of AI; pushing for cross-economy AI adoption; and positioning the UK as an "AI maker rather than an AI taker".

The Action Plan sets out bold targets to cement the UK's global leadership in AI, including a 20-fold increase in public compute capacity by 2030, the creation of new AI Growth Zones to fast-track infrastructure like data centres, and expanded access to anonymised NHS data to drive healthcare innovation. The plan also commits to building sovereign AI capabilities, reducing dependence on foreign technologies, and scaling AI talent pipelines through focused education and training - ensuring that AI development aligns with UK values and delivers tangible public benefit.

Many of these issues were addressed directly in our seminar. In particular, there may be an opportunity for Northern Ireland to pilot innovative applications of AI within devolved services. With an integrated health system and access to population-level data, the region can act as a controlled environment for ethical AI deployment, particularly in areas like diagnostics, social care forecasting, or justice reform. By aligning local strategy with the national agenda, Northern Ireland can position itself as a trusted hub for responsible, high-impact AI deployment that is then rolled out across the UK.

In February, the government published the 'Artificial Intelligence Playbook'⁸. The 'AI Playbook for the UK Government' is a practical guide designed to help public sector organisations adopt AI responsibly, securely, and effectively. It sets out a clear, principle-based framework, with 10 core principles guiding lawful, ethical, and collaborative AI use - from ensuring meaningful human oversight to aligning AI with organisational goals and user needs.

For public sector leaders in Northern Ireland, we would suggest that the Playbook offers several high-level takeaways that could inform their approach:

1. **Build a Structured AI Adoption Plan:** Leaders are encouraged to define an AI strategy aligned with their organisation's goals, supported by governance boards, communication plans, and sourcing strategies.
2. **Prioritise Responsible Use Cases:** AI should only be applied to well-suited, clearly defined

⁸ https://assets.publishing.service.gov.uk/media/67aca2f7e400ae62338324bd/AI_Playbook_for_the_UK_Government_12_02_.pdf

problems where it offers real value. Appropriate safeguards should be applied in relation to higher risk, fully automated decisions, especially in sensitive domains.

3. **Invest in Capacity and Talent:** A key recommendation is to establish multidisciplinary teams, invest in AI training across all levels, and partner with external experts, including academia and industry.
4. **Design Governance from the Start:** Effective deployment relies on strong oversight. Leaders are advised to form AI governance boards or integrate AI expertise into existing decision-making structures, backed by ethics committees where needed.
5. **Emphasise Transparency and Public Trust:** From data protection and algorithmic explainability to fallback procedures and public accountability, every stage must be transparent and fair to maintain public confidence.
6. **Adopt Secure and Scalable Solutions:** AI solutions must be resilient against threats, with risk mitigation plans, fallback options, and secure procurement strategies in line with Crown Commercial Service guidance.

Again, many of the themes present in the above were reflected in the seminar discussion, not least Professor Paul Miller's talk on cybersecurity. As per the Playbook and the seminar, the common message emerged that success with AI isn't just about deploying algorithms, it's about building trust, ensuring accountability, and designing systems that genuinely serve people and public value.

In Ireland, the **AI Advisory Council published its latest report, "Helping to Shape Ireland's AI Future (February 2025)"⁹**, outlining key opportunities and strategic policy recommendations to accelerate AI adoption while safeguarding the economy, competitiveness, workforce, and society. The report covered six critical areas: AI and skills, AI ecosystem, AI literacy and education, AI sovereignty and infrastructure, biometrics and public services, and AI and Ireland's creative sector.

The Artificial Intelligence (AI) Advisory Council has proven to be a strong contributor to the country's knowledge base and strategic direction. There would be much in this organisation that would merit a similar body being replicated in the north. The membership is drawn from a range of various fields, including academia, business, law, security, social sciences, economics, and civil society. In addition to providing expert and technical advice to the government on AI policy, the role this body engages in in relation to public communications and promoting trustworthy and ethical AI is valuable.

Local government innovation and how AI rolls out directly to the service user is one of the most important areas for reflection. While high level summits, national trade agreements and devolved policy will capture attention, this should not be to the detriment of local government in Northern Ireland there is considerable invention and innovation both organically and as part of the Regional City Deals.

Contributing to this space, of particular note is The Tony Blair Institute's latest report 'Governing in the Age of AI: Reimagining Local Government'¹⁰ presents a compelling case for radical reform of local government through artificial intelligence. The central recommendation is the creation of a new

⁹ [https://www.gov.ie/en/department-of-enterprise-trade-and-employment/press-releases/irelands-ai-advisory-council-report-to-government-helping-to-shape-irelands-ai-future/#~:text=%27Ireland%E2%80%99s%20AI%20Advisory%20Council%20Recommendations%20%E2%80%93%20Helping%20to%20Shape%20Ireland%E2%80%99s%20AI%20Future%20\(February%202025\)%27](https://www.gov.ie/en/department-of-enterprise-trade-and-employment/press-releases/irelands-ai-advisory-council-report-to-government-helping-to-shape-irelands-ai-future/#~:text=%27Ireland%E2%80%99s%20AI%20Advisory%20Council%20Recommendations%20%E2%80%93%20Helping%20to%20Shape%20Ireland%E2%80%99s%20AI%20Future%20(February%202025)%27)

¹⁰ <https://institute.global/insights/politics-and-governance/governing-in-the-age-of-ai-reimagining-local-government>

national institution, the *Devolved AI Service (DAIS)*, which would support councils in adopting shared AI tools, modernising data infrastructure, and scaling innovation across local authorities. Crucially, it positions AI not as a bolt-on, but as the foundation for reimagining how local services are designed and delivered.

Northern Ireland, facing similar fiscal and service delivery pressures, can draw clear lessons from this vision. With just 11 councils and a more compact governance system, Northern Ireland is uniquely positioned to move faster and more collaboratively than England's fragmented local authority landscape. A regional version of DAIS could act as a central hub for AI tools, procurement frameworks, and capability-building. Such an initiative would reduce duplication, cut costs, and allow councils to benefit from shared data, platforms, and standards.

The report's three flagship innovations offer immediate relevance. First, an *AI co-worker* for social workers could help clear care assessment backlogs, automate administrative tasks, and free staff for higher-value work. This is especially pertinent in Northern Ireland, where adult social care remains under strain. Second, a *Local Navigation Assistant*, a personalised digital front door to council services, could transform citizen interactions, tackling service confusion, benefit underclaiming, and excessive contact centre demand. Finally, an *AI planning assistant* would bring speed and consistency to local development planning, a major issue as Northern Ireland navigates its way towards housing expansion and infrastructure growth.

The report concludes, as did many of our seminar presenters including Professor Madalina Busuioic, that technology alone is not enough - councils must invest in people, data, and leadership. Importantly, councils here are progressing on their AI journey by testing and piloting proven tools, convening partnerships with local universities, and adopting shared digital standards.

Finally, we are mindful of **Queen's University Belfast's 'Strategy 2030'**¹¹. As Northern Ireland navigates this pivotal decade of AI-enabled transformation, Queen's University Belfast is committed to shaping a better future through technology, research, and innovation. Strategy 2030 sets out an ambitious and interdisciplinary agenda, placing AI, secure connected intelligence, and global partnerships at the heart of its mission.

¹¹ <https://www.qub.ac.uk/home/Filestore/Fileupload,1118456,en.pdf>

2. Key Discussions and Themes

The seminar's detailed discussions focused on several interconnected themes. Each of these discussions highlighted distinct viewpoints, practical insights, and critical tensions articulated clearly by leading voices. Below is an expansive and detailed synthesis, faithfully integrating the direct language and arguments presented by the seminar participants.

2.1 AI and Public Sector Decision-Making

The integration of AI into public sector decision-making represents both a transformative opportunity and a serious governance challenge. Across the seminar, a consistent concern emerged around the risk of adopting AI systems that are insufficiently tested, poorly understood, or misaligned with the values and obligations of democratic institutions.

International experience has shown that algorithmic tools - when applied without effective oversight - could lead to serious injustices. As a result it is important to consider adequate testing protocols, source high-quality data, and focus on transparency in the design, procurement and implementation of AI systems. One speaker highlighted that "the slowness that comes from such careful consideration [of testing and oversight] should be seen as a virtue rather than a vice", whereas the broader tone of discussion suggested that proven, balanced and responsible approaches can operate in concert with rapid adoption.

Public decision-making in democratic societies must remain grounded in values such as accountability, equity, and the right of citizens to contest and understand the decisions that affect their lives. AI tools that compromise these principles, whether by obscuring reasoning processes or introducing harmful proxy variables, risk undermining public trust. This is particularly concerning in high-stakes domains such as welfare provision, policing, and healthcare, where the margin for error is small and the impact of mistakes is profound.

Nevertheless, sharp progress can be achieved and the experience of Estonia was cited by one speaker as an example of a compelling model of AI-enabled public sector transformation. Estonia focused on creating the secure digital infrastructure that now underpins almost all government services. Central to this is the X-Road data exchange platform and AI tools like the virtual assistant Bürokratt, which streamline citizen engagement. AI is also embedded in areas such as healthcare diagnostics, judicial case handling, and employment services. With 99% of public services online, Estonia demonstrates how strategic investment, digital literacy, and public trust can combine to deliver efficient, transparent, and citizen-centric governance.

But how can government in the UK and devolved regions reconcile the efficiency benefits of AI with the need for careful, accountable decision-making? AI should be used to support - not replace - human judgment. Governments should use AI to identify aggregate patterns, surface contextual evidence, and improve the timeliness and quality of decision support. Crucially, these systems can be designed around public values from the outset, rather than retrofitted with governance tools after deployment.

This implies a strong role for pre-deployment validation, ethical review, and continuous post-

deployment auditing. It also requires public bodies to become more discerning customers of AI, developing the literacy and institutional expertise to assess when a system is appropriate - and when it is not. A key insight from the seminar was that ill-fitting AI systems are often deployed not because they are the best fit, but because they are available. In conditions of resource pressure, agencies may turn to whatever tools can be procured quickly. This dynamic must be countered by capacity building and clearer public procurement standards tailored to AI.

Ultimately, AI can improve decision-making in the public sector when deployed with oversight and a commitment to democratic governance. As one speaker put it: public administrators must become “demanding and discerning” users of AI. The stakes are too high to approach this with anything less.

Key takeaways: AI should augment, not replace, human judgment in public service decisions. Rigorous pre-deployment validation, ethical review, and ongoing audit processes are essential. Poorly matched or poorly understood AI systems often stem from procurement pressure, not strategic fit. Public administrators must become discerning and demanding customers of AI.

2.2 Trust and Ethics in AI

Trust in AI cannot be assumed, it must be earned. For public bodies considering the use of AI, this principle is not rhetorical but foundational. The ethics of deploying machine learning systems in sensitive, high-stakes domains such as healthcare, welfare, education, and criminal justice often requires a different approach than the one typically adopted in commercial or consumer technology.

Several core questions arose during the course of the seminar, such as: how can AI deliver compelling efficiency gains without amplifying the risks of deeper opacity and undermining accountability in public decision-making? With AI being increasingly labelled as the ‘fourth industrial revolution’, one speaker remarked: “should we allow for the efficiency logics of industrial revolution to shape the way governments function?” Is it ethical to deploy systems that can make faster or even marginally more accurate decisions if those same systems reduce the capacity of citizens to understand or contest how decisions are made? Are we willing to trade procedural transparency and democratic oversight for technical optimisation and economic efficiency?

Participants stressed that AI must not become a vehicle for outsourcing public responsibility to automated systems. Systems trained on historic data risk replicating and entrenching past inequities, particularly when deployed in social contexts where marginalised communities have already borne disproportionate burdens. These harms are not theoretical - they have occurred repeatedly across jurisdictions. The challenge, then, is not only to make AI technically robust, but to ensure it aligns with the ethical obligations of public governance.

Whereas private sector AI is often geared toward maximising engagement, efficiency, or predictive accuracy, public sector applications must also uphold fairness, procedural legitimacy, and the right to redress. Applying off-the-shelf AI tools without adapting them to public values is incompatible with good governance.

To manage this risk, government agencies should ensure fairness and transparency, and invest in tangible practices that make these principles operational. This includes the development of audit

trails, mechanisms for appeal, and clear documentation of system logic. It also requires ethical assurance processes to be built into procurement and design phases, not appended after deployment.

Importantly, explainability is an important consideration. Citizens affected by AI-informed decisions - whether about their benefits, health status, or eligibility for services - will inevitably demand to know not just the outcome, but why it was made. While there are limits to what current AI models can explain, this does not absolve public bodies of the obligation to provide meaningful justification. If a decision cannot be adequately explained or scrutinised, it may not be fit for use in public service as case law can attest.

Trust also depends on competence. Without widespread AI literacy and ethical understanding across government, including among procurement officials and frontline managers, even well-intentioned deployments can go awry. Seminar discussions highlighted the need for structured training programmes, ethical governance frameworks, and shared data and model repositories to support collective learning and reduce duplication of error.

Ultimately, trust is not just about the integrity of the algorithm, it is about the integrity of the institution deploying it. To earn public trust, AI systems must be developed and used in ways that affirm, rather than erode, the principles of democratic accountability. When done well, AI can enhance these values by providing better evidence, reducing arbitrary decision-making, and improving consistency.

Key takeaways: Public sector AI must be built for accountability, not just efficiency. Off-the-shelf commercial tools may not always meet public governance standards. Strong ethical foundations for public sector decision making require audit trails, appeals mechanisms, and meaningful explanations of decisions.

2.3 Security and Sovereign AI Compute

As artificial intelligence becomes embedded in public service delivery, the security and governance of AI infrastructure are no longer peripheral concerns - they are core strategic issues. AI systems, particularly those reliant on large-scale data and high-performance computation, represent new points of vulnerability and new targets for exploitation. This makes questions of sovereignty, control, and trust not only technical but fundamentally political.

At the seminar, participants stressed that the UK and Northern Ireland must move beyond narrow considerations of AI adoption and consider the foundations upon which these systems are built. Who owns the compute infrastructure? Where does the data reside? Who controls access, and under what conditions? These questions are especially pressing given the increasing reliance on outsourced cloud-based AI services, often provided by multinational vendors whose infrastructure may fall outside UK legal or regulatory reach.

The case for sovereign AI compute - infrastructure that is publicly governed, locally operated, and strategically aligned with national priorities - is both defensive and enabling. On one hand, it ensures that sensitive public data remains within accountable systems, subject to UK and devolved data

governance standards. On the other, it enables the development of AI capabilities tailored to public sector needs, including transparency, auditability, and alignment with public values.

Security by design was a recurring theme. Participants warned that AI systems, like all software, are susceptible to attack vectors ranging from data poisoning to adversarial input manipulation. Unlike traditional IT systems, however, the complexity of AI models can obscure such attacks, making them harder to detect and respond to. Ensuring the integrity of training data, model pipelines, and decision outputs requires not only technical safeguards but institutional mechanisms for independent testing, validation, and red-teaming.

In the public sector, the stakes are high. From healthcare allocation to infrastructure planning and law enforcement, the deployment of AI in mission-critical systems introduces new risks of systemic error or malicious manipulation. These risks are amplified in systems with limited visibility or weak operational oversight. Therefore, public bodies must adopt a holistic approach to AI security - one that encompasses the entire lifecycle from data curation to algorithmic inference and post-deployment monitoring.

Northern Ireland, it was argued, is uniquely positioned to act as a testbed for sovereign compute capability. With two research-intensive universities, a high concentration of cyber-security expertise, and an agile scale of governance, the region can lead in developing secure, ethical, and locally governed AI infrastructure. Moreover, the absence of clear regional leadership in the UK on sovereign AI presents an opportunity for Northern Ireland to distinguish itself - both within the UK and as a model for cross-border data governance and innovation.

Crucially, this is not just about defence - it is also about empowerment. Sovereign compute infrastructure can enable Northern Ireland to pilot next-generation public sector AI systems in areas such as healthcare optimisation, real-time infrastructure monitoring, and digital regulatory services. These systems can be built with local needs in mind, shaped by domain experts, and governed in partnership with civil society and academia.

The national AI strategy recognises the importance of compute power, but deployment has lagged. If Northern Ireland is to fully participate in AI-driven transformation, it must act now to invest in secure, high-performance, ethically governed infrastructure. Sovereign AI compute is not a luxury—it is the foundation upon which trustworthy public sector AI will be built.

Key takeaways: AI infrastructure is a strategic asset, not just a technical tool. Sovereign compute capabilities reduce risk, enhance trust, and ensure data governance within UK legal frameworks. Northern Ireland has the capacity to lead in sovereign AI infrastructure, supported by academic and cybersecurity strengths. Public sector AI must be secure and responsible by design, with full lifecycle monitoring.

2.4 AI Skills and Workforce Development

No AI strategy is complete without addressing the fundamental issue of skills. AI technologies are only as effective as the people who procure, deploy, and interpret them. Without a workforce that understands the capabilities and limitations of AI, public sector adoption will stall, systems will be

poorly implemented, and public trust will erode. The seminar underscored this reality: skills are the bridge between aspiration and delivery.

AI literacy (which may be seen more as a fundamental understanding of overarching AI logics, rather than how those are operationalised technologically) is an essential criterion for wider adoption. Across the public sector, from senior civil servants to frontline delivery teams, there is an urgent need for practical, context-aware understanding of AI tools. This includes technical awareness (e.g. what AI can and cannot do), ethical implications (e.g. how to spot risks of bias or opacity), and operational competence (e.g. how to evaluate vendor claims or assess pilot projects). A clear message emerged: AI systems cannot be effectively governed by those who do not understand them.

Seminar participants called for structured capacity-building that goes beyond generic digital skills and speaks directly to the public sector's specific responsibilities. This includes designing continuing professional development (CPD) programmes focused on the responsible use of AI in service design, procurement, and decision-making. Some participants suggested setting baseline AI literacy standards for senior leadership across departments to ensure that executive-level decisions about AI adoption are informed, realistic, and accountable.

The education system also has a critical role to play. Preparing the next generation of public sector professionals requires an overhaul of curricula to embed AI awareness in disciplines well beyond computer science. Law, social policy, public health, planning, and economics must all engage with AI not only as a technical subject, but as a force shaping governance and society. Seminar discussions noted that Northern Ireland's universities have already begun this work, but more systemic integration is needed - particularly to ensure alignment with strategic workforce needs.

Upskilling must also include those already in the workforce. As AI changes job roles and workflows, reskilling and transition planning become essential for maintaining morale and effectiveness. Rather than replace human judgement, well-deployed AI can augment professional roles: freeing up time, surfacing insights, and supporting better decisions. But this only happens when workers are engaged in the design and rollout of new systems. Where change is imposed without consultation or training, resistance and implementation failure often follow.

A further concern raised at the seminar was the unevenness of AI skills across departments and agencies. While some areas have piloted AI solutions and built up internal expertise, others remain at the start of their journey. This poses a coordination challenge: without shared frameworks, peer learning, and centralised support, progress will be fragmented and difficult to scale. Participants proposed developing a cross-departmental skills strategy, supported by centres of excellence and secondment opportunities between academia, government, and industry.

Importantly, discussions highlighted that building AI capability is not only about data scientists and engineers. Ethical governance, policy interpretation, and legal expertise are equally essential to ensuring that AI use remains grounded in democratic norms. The AI-literate public servant of the future will not need to write code, but they must be able to ask the right questions, understand system limitations, and challenge inappropriate uses with confidence.

In summary, public sector innovation through AI will rise or fall on the strength of its people. Strategic investment in AI capability is not a parallel stream, it is the precondition for success. Northern Ireland has an opportunity to lead by embedding AI skills into the fabric of public service at every level: from policy design to service delivery, from classroom to cabinet room.

Key takeaways: AI literacy is a core competency, not a specialist skill. Capacity-building must target civil servants, managers, and frontline staff, not just technical roles. Baseline standards and CPD pathways are urgently needed to support confident and informed AI use. Skills gaps will limit effectiveness and increase risk without structured investment.

2.5 AI Collaboration and Ecosystem

Effective public sector AI does not emerge in isolation. It depends on an ecosystem - a network of organisations, capabilities, standards, and shared infrastructure that together create the conditions for innovation, trust, and scale. Across the seminar, participants consistently returned to this point: collaboration is not an accessory to AI success, but its foundation.

Public sector challenges - whether in healthcare, infrastructure, education, or justice - rarely fall neatly within the remit of a single department. AI has the potential to bridge these silos by drawing insights from complex, overlapping datasets and supporting more joined-up decision-making.

But to unlock this potential, systems must be designed collaboratively from the outset. While there is a clear need for collaboration, each department has a unique remit and sector of operation. For example, the operational principles of health services are different from that of policing. Too often, AI projects are commissioned in isolation, with limited cross-sector input, leading to duplication, fragmentation, and lost learning. Acknowledging this, collaborative approaches could focus on ensuring that there is a high-level dialogue on AI adoption while also ensuring that the unique operational principles of each government sector are respected and well attended to.

Several contributions pointed to the importance of structured, multi-stakeholder collaboration. This includes partnerships between government departments, academic institutions, civic organisations, and private sector innovators. Universities, in particular, were identified as critical partners - offering not only technical expertise but independent scrutiny and a long-term view. But seminar participants were clear: this collaboration must be active and intentional, not simply transactional and incidental. Co-creation must replace consultancy.

The seminar also highlighted the need for shared infrastructure - digital, legal, and institutional - that enables this ecosystem to function. Common data standards, interoperable platforms, and secure environments for experimentation are prerequisites for meaningful collaboration. Without them, even the most well-intentioned partnership can stall at the point of integration. There was widespread support for expanding initiatives that make high-quality datasets, compute power, and model development environments available to multiple partners under clear ethical and governance frameworks.

Northern Ireland has unique advantages in this regard. Its compact scale, strong university presence, and existing innovation hubs position it well to act as a testbed for collaborative AI deployment. But

realising this vision requires more than good will. Participants called for the creation of cross-sector working groups, sustained funding for public-interest pilots, and mechanisms to surface, share, and scale successful local initiatives. Knowledge must not stay trapped in isolated projects or institutions.

Of equal importance is the need to create space for ethical and civic input. AI systems shape the conditions under which rights are exercised and services are delivered. Therefore, collaboration must include voices beyond the technical and managerial elite. Civil society organisations, advocacy groups, and affected communities all have a stake in how public sector AI is developed and deployed. Their involvement is not only ethically justified, it improves system design by surfacing unintended consequences, highlighting context-specific risks, and ensuring public legitimacy.

Finally, participants reflected on the importance of maintaining momentum. One-off collaborations (while valuable) are insufficient in a fast-moving field. What's needed is a long-term ecosystem approach: permanent structures that can convene partners, set strategic direction, fund experimentation, and broker relationships across sectors. This might take the form of a national or regional AI collaboration centre, an annual public sector AI forum, or a cross-government AI board.

The conclusion was clear: no single actor can deliver trustworthy, effective AI for the public sector alone. But together, through deliberate and well-supported collaboration, Northern Ireland can build a distinctive, resilient, and values-led AI ecosystem that serves the public interest.

Key takeaways: No single actor can deliver safe, effective AI in isolation; collaboration is foundational. Shared data, interoperable systems, and common standards are critical enablers. Universities, civic organisations, and government must shift from transactional partnerships to co-creation. Ecosystem momentum requires sustained convening and permanent structures, not one-off initiatives.

2.6 Government Perspectives on AI Adoption

The seminar brought together perspectives from across the UK and Ireland, revealing a diverse set of national approaches to public sector AI. From regulatory design to implementation capacity, the contrast between jurisdictions illustrated both the possibilities and the pitfalls of government-led innovation. What emerged was a clear call for political leadership, strategic clarity, and regionally grounded action, particularly in Northern Ireland.

While there is no one-size-fits-all model for AI adoption in government, comparative insights underscored the importance of coherence. Jurisdictions that have made the most progress (such as Estonia - for further information, see the case study reference in the Appendices) have done so through integrated strategies that link data infrastructure, procurement standards, workforce planning, and citizen trust. Piecemeal initiatives, by contrast, tend to underdeliver, stall, or generate duplication across departments.

Northern Ireland's position was discussed with both realism and resolve. In part reflecting the region's unique character as a post-conflict society with deep social fault lines still existing, delegates acknowledged gaps in local readiness: fragmented leadership, unclear accountability, and under-resourced implementation teams were all cited as barriers to progress. Yet these gaps also create an opening - a chance to craft a distinct local vision for AI that is not simply a derivative of wider UK

strategy, but rooted in the region's institutional strengths, policy needs, and cross-border relationships.

A recurring theme was the need for vision-led adoption. Without a clear articulation of what AI is for (what kinds of problems it should solve, what values it should uphold) governments risk drifting into passive dependence on market-led solutions or simply following the path of least resistance. AI adoption then becomes reactive, driven by vendor offerings or short-term pressures rather than long-term public value.

Participants urged Northern Ireland to assert its own vision: one that sees AI not just as a technical enabler, but as a vehicle for strengthening public service delivery, citizen trust, and policy responsiveness. In this vision, AI is deployed not simply because it is possible - but because it meaningfully contributes to a better, fairer, more effective state.

Crucially, governments must retain control over the governance terms of AI adoption. This means avoiding overreliance on opaque or proprietary systems, ensuring independent oversight, and embedding ethical accountability into procurement and deployment. Across all jurisdictions represented at the seminar, this was seen as a work in progress though the need for progress was urgent.

Data governance emerged as a particularly pressing issue. The fragmentation of public data across departments, legacy systems, and inconsistent standards is a known barrier to AI deployment. But participants also warned against integrating data without clear ethical frameworks. Instead, governments must design data architectures that are secure, transparent, and purpose-built for public good—rather than simply adapted from commercial models.

Seminar discussions also highlighted the importance of regulatory innovation. As new AI use cases emerge, traditional compliance models may struggle to keep pace. Several contributors proposed agile regulatory approaches that support experimentation while safeguarding citizens, such as regulatory sandboxes, adaptive standards, or ethics-by-design audits. Northern Ireland could lead in this space by leveraging its scale and cross-sector relationships to pilot new models of trustworthy AI governance.

Finally, governments must recognise that the legitimacy of AI systems is both political and technical. Citizens must be convinced that AI in public services can improve experiences and support greater equity, rather than threaten the diminution of rights or individual agency. They must be able to challenge decisions, understand how outcomes are reached, and trust that systems are aligned with the public interest. Achieving this requires active, visible government stewardship, not only of technology, but of values.

In sum, AI adoption by government is a leadership issue as much as a technical one. If Northern Ireland is to seize the opportunity ahead, it must move decisively: to clarify strategy, build institutional capacity, and foster collaboration between departments, civil society, academia, and the wider UK and Irish AI ecosystems. Experts at the seminar offered a considered route map, but real progress will depend on sustained, collective political will.

Key takeaways: Vision-led adoption is critical: AI must be used because it delivers public value, not

because it is available. Fragmentation and lack of local strategy leave governments vulnerable to vendor dependency. Northern Ireland can lead by developing a distinct, regionally grounded AI strategy. Regulatory innovation and political leadership are essential to build public legitimacy and policy relevance.

2.7 Crystallising Views: Reflections from Discussion Panel on Trust, Accountability, and the Future of Public Sector AI

The panel discussion for the seminar crystallised many of the conversations across the course of the day. The panel surfaced a compelling and, at times, provocative debate on the role of trust, urgency, and public values in shaping the future of AI deployment within the public sector. What emerged was a rich exchange between pragmatism and caution, and where speakers stressed that confident deployment need not necessarily arrive at the expense of transparency and accountability.

The discussion opened with the question of trust. One view advocated for assertive AI deployment, warning that waiting for perfect systems would delay much-needed reform and reduce public sector effectiveness. Examples from health and national security were cited to support the claim that even imperfect AI can outperform overburdened human systems. As one speaker remarked: "A ship is safe in its harbour, but that's not what it's for." AI, they argued, must be deployed to deliver outcomes, even as its governance evolves.

In contrast, one speaker argued that trust must be earned, not assumed. Systems, when opaque and unaccountable, undermine public legitimacy. Robust testing, transparency, and auditability were essential.

A second point focused on proprietary and black-box systems. One speaker raised concern about the widespread outsourcing of AI to vendors, with public bodies unable to explain, interrogate, or correct algorithmic decisions. The Chief Scientific and Technology Advisor to DSIT countered that while full explainability may be unrealistic, public institutions can use statistical techniques, causal models, and adversarial testing to approximate understanding and embed public values into design. Progress towards AI adoption, they argued, should not be paralysed by idealism.

The discussion also incorporated a broader critique of the foundational logic behind current AI tools. The view was aired that (arguably) mainstream AI - particularly generative models - was ever designed to support civic or democratic outcomes. Built largely for commercial engagement, these systems can lack alignment with public values like care, equity, and procedural fairness. However, there was a note of caution around this with one speaker urging against romanticising human judgement, which itself can be flawed and opaque.

The panel closed on a shared recognition that Northern Ireland must shape its own path and act with purpose and urgency. There was broad agreement that the region cannot wait for Westminster to define its AI strategy. Other jurisdictions are already advancing bold, systems-level AI deployments. For Northern Ireland to lead, it must articulate its own vision, underpinned by values, capability, and responsibility.

Key Takeaways: Trust requires evidence. Deploying AI systems without transparency or

auditability risks public harm and undermines democratic accountability. Public sector AI must be explainable and values-driven. Ownership of AI strategy locally is essential. Northern Ireland must define its own vision and act decisively to develop capability, infrastructure, and leadership. Progress and prudence are not mutually exclusive. AI can be deployed pragmatically - so long as risk, ethics, and equity are built into the process from the start.

3. Recommendations

The following recommendations are drawn directly from the evidence, insights, and stakeholder discussions that emerged throughout the seminar. They are intended to provide practical, high-impact actions that can accelerate the responsible adoption of AI within the public sector. Each recommendation addresses a specific strategic or operational gap identified during the workshop - whether in governance, transparency, capability, infrastructure, or ethical oversight.

Collectively, these proposals represent a coherent roadmap for making AI work in the public interest: not through isolated innovation, but through coordinated leadership, shared standards, and a sustained investment in skills and systems. If adopted, they will help position Northern Ireland as a regional leader in trusted, transparent, and citizen-focused AI governance.

1. Establish a Permanent Public Sector AI Forum

Our first recommendation is **to create a standing AI Forum comprising of policymakers, CPPA representation, academic partners, private sector voices, and civic society observers from across the island to liaise with and to coordinate strategy, share learning, and set standards for public sector AI in Northern Ireland.**

The seminar consistently emphasised the importance of strategic coherence and cross-departmental coordination. Fragmentation was cited as a key barrier to progress. One speaker warned that “without vision, all we will see are barriers.” A permanent forum would institutionalise vision-setting and provide sustained leadership. As highlighted in section 2.5, effective AI depends on ecosystem coordination - not isolated pilots. This forum would also act as a clearinghouse for AI governance, ethics, risk, and procurement expertise, ensuring that departmental efforts align with national and international best practice.

2. Create a Public Register of AI Use in Government

Our second recommendation is **to create and regularly publish a public-facing register of AI tools used across all Northern Ireland Executive departments and agencies.**

Throughout the seminar, transparency was consistently highlighted as a precondition for public trust. A register would help mitigate risks as well as support democratic accountability, enable academic scrutiny, and allow civil society to assess whether AI use aligns with public values. It also complements similar proposals in the UK AI Playbook and the Tony Blair Institute report for building public legitimacy through open disclosure.

3. Introduce Mandatory AI Risk Audits in Procurement

Our third recommendation is **to require all new public sector procurements involving AI systems to undergo a centralised risk assessment, whereby suppliers can consider the effects of bias, explainability, harms, and alignment with ethical standards.**

Speakers at the seminar warned repeatedly about systems deployed “because they are available,” not because they are fit for purpose. AI can encode bias, and harm vulnerable populations if left unchecked. Therefore, we would suggest that a centralised public sector risk assessment tool be developed for bodies to include as part of their procurement processes. This centralised risk assessment tool would take the form of an **AI Risk Declaration and Self-Assessment Disclosure**

Questionnaire to be completed by bidding organisations. We believe this approach would increase rigour, reduce liability, and help public bodies become “discerning customers of AI,” as one speaker put it.

4. Pilot a Sovereign AI Design and Implementation Platform

Our fourth recommendation is **to invest in a regionally hosted, publicly governed AI compute platform to support safe, scalable AI experimentation and application across public services.**

As outlined in section 2.3, compute infrastructure is no longer just a technical issue, it's a matter of security, sovereignty, and capability; and developing sovereign compute should also be accompanied by mechanisms to capture and analyse system data, which will be essential to initiating serious discussions on how AI governance processes should operate in practice. Numerous speakers emphasised the strategic need for UK-based, accountable compute platforms, warning against overdependence on cloud services outside UK jurisdiction. Northern Ireland, with its cybersecurity ecosystem, research capacity at QUB and UU, and scale for agile pilots, is ideally positioned. A sovereign compute pilot would also unlock access to national funds linked to the UK's 20x compute expansion target (AI Action Plan) and place Northern Ireland at the forefront of ethical AI development.

5. Deliver an AI Literacy and Capability Framework

Our final recommendation is **to develop and roll out a tailored AI skills and literacy framework for the Northern Ireland public sector, covering leadership, procurement, policy, ethics, and technical awareness.**

As Section 2.4 makes clear, skills are the bridge between aspiration and delivery. Without sufficient AI understanding among decision-makers, deployments will be misaligned or fail. Multiple speakers called for systemic training across all levels of government, not just for specialists. This framework should include CPD pathways, baseline competency standards, and partnerships with universities for delivery. Importantly, it would also support talent mobility and career development - embedding AI capability as a core public sector competency by 2030.

Acknowledgements

We extend our sincere thanks to all those who contributed their insight and expertise throughout the development of this report. Their thoughtful guidance - shared during and after the seminar - has been invaluable in helping us capture the complexity of the issues discussed, from AI's technical potential to its democratic implications.

While there are many who contributed to this workshop and report, we would like to particularly acknowledge the work of Geoff McGimpsey in compiling this report.

The contributions of speakers, facilitators, and attendees alike played a central role in shaping the recommendations and reflections presented here. This report stands as a collaborative effort, grounded in the spirit of open dialogue and shared purpose that defined the seminar.

About the CPPA

The Centre for Public Policy and Administration is founded on four pillars. Working across each of these pillars, CPPA builds interdisciplinary partnerships locally, nationally and globally to foster our mission:

Education

The Centre builds on QUB's current teaching provision on public policy, and our flagship Msc in International Public Policy. Within CPPA, we seek to broaden the opportunities for PhD research in public policy at QUB, and facilitate the provision of the Continual Professional Development (CPD) programmes for public service practitioners.

Scholarship

Our work advances the University's contribution to the study of public policy as a cross-cutting discipline, wherein the Centre helps situate QUB as an authoritative voice for tackling pressing global public policy and governance challenges of the 21st century.

Translation

CPPA facilitates dialogue, collaboration and co-production between academics and policy practitioners. It promotes opportunities for scholars across QUB to engage with policy professionals and outlets. The Centre provides the means and networks to communicate oftentimes complex, yet important research findings to broader public policy audiences in an accessible and comprehensive manner.

Exchange

The CPPA will maximize the societal and economic impact of publicly funded research by prioritizing knowledge transfer. It will bridge the gap between academics and policymakers, enhancing intelligence gathering processes to mutual benefit. Additionally, the Centre will facilitate regular exchanges between academic and policy communities at local, national, European, and international levels.

Contact

You can reach Centre director Prof Muiris MacCarthaigh by emailing cppa@qub.ac.uk