AI AND WALES: A PRIMER

AI has existed for over 50 years, but new developments in Generative AI (like ChatGPT and DeepSeek) mean that it can now perform multiple functions and create new data.

Generative AI can produce convincing text, but it still presents fiction as fact, and does not truly understand language.

RELIABLE AI CHECKLIST

- ✓ the task is well constrained
- ✓ there is a human collaborator
- ✓ the data is narrow and verified
- ✓ helpfulness, harmlessness, and honesty are built in

CYMDEITHAS DDYSGEDIG CYMRU
THE LEARNED SOCIETY OF WALES

REGULATION

The AI sector is unusually lacking in regulation, standards, and quality assurance.

Trust is more easily lost than gained: since AI development is driven primarily by profit, the public needs to know where AI is used and affects their lives.

The label “artificial intelligence” does not necessarily mean a product is innovative: the term is not protected, so can be misused for commercial purposes.

PUBLIC OPINION

Only **14%** of the Welsh public thinks AI carries more benefits than risks. Disadvantaged groups are least likely to agree that AI will benefit them individually.

Concerns include:

jobs · disinformation · IP · digital divide · privacy

SUCCESS IN WALES USING RELIABLE AI TOOLS

- Improving cancer diagnosis (NHS Wales)
- Locating ancient woodlands (Welsh Gov)
- Ocean modelling for tidal energy (Bangor University)

POLICY AND DEVOLUTION

AI is a reserved matter for UK Government, but crosses over with multiple areas of Welsh Government business.

Welsh Government is exploring where AI can help improve public service delivery across Wales. But barriers remain:

- many services lack even basic automation or digitisation
- poor data quality can prevent useable results with AI
- commercial tools may not align with Welsh Government values.

CONDITIONS FOR ECONOMIC GROWTH

Artificial Intelligence and Wales

The state of the art[ificial intelligence]

How novel is this technology?

Artificial Intelligence is a field of study that emerged in the mid-twentieth century, and refers to any machine that mimics intelligence. Traditional AI is designed to perform a specific task and respond to a particular set of inputs, by learning and making predictions from data.

However, increased public and media attention over the past few years is more than just hype; we are indeed living through an important moment in technological history. What makes this moment important?

- New Generative AI tools can now perform a **wide range of functions** and create new data.
- Machines have now passed the **Turing test** – where a human attempts to distinguish between texts generated by human or machine.
- A huge **surge in investment** means that innovation has been moving unprecedentedly fast, especially in "unicorn" start-ups (worth over $1bn).

Generative AI tools in their current state are by no means perfect. Most notably, large language models like ChatGPT and DeepSeek tend not to know what they don't know, inventing fiction presented (unknowingly) as fact. Moreover, where these tools use the internet for source data, it seems that they are beginning to exhaust what is available, and may have started to enter a downwards spiral of training themselves on their own AI-generated, at times low-quality, web content. However, companies are aware of this risk and have begun to invest in curating and synthesising high-quality data. Large language models are often described as "stochastic parrots", emphasising that while they are able to produce convincing text, they do not possess any *understanding* of language or meaning, instead merely reflecting statistical associations in the training data.

There are therefore a huge number of truly novel AI tools and possibilities emerging at this moment in time. However, investors and policymakers should be wary that the term "artificial intelligence" is not protected, and is thereby vulnerable to misuse for commercial gain, for example to overcharge for a simple algorithm.

When does AI work best?

Currently, AI is still most reliable when the task is well constrained, when there is a human collaborator, and when the training data is focussed on a particular task or domain and verified. It is also useful to build in alignment to key "human values" of helpfulness, harmlessness, and honesty. Tools that meet these criteria are already producing impressive results in Wales, improving cancer diagnosis, finding lost woodlands, and ocean modelling.

A distinct lack of regulation

The AI sector is unusually lacking in regulation, standards, and quality assurance. While engineering in other domains – from cars to kettles – is highly regulated (without stifling innovation), big tech companies have adopted the habit of releasing new generative AI models and letting the public identify the issues.

One reason for this difference is that the internet transcends borders; yet, there are examples of international standards that work. Insurance is a potential risk-based approach to issues caused by AI, but questions remain about who decides when something goes wrong, who is to blame, and what has caused the issue in a black-box technology (where it is not possible to see how the machine makes decisions). The nuclear power sector offers another possible precedent in the form of "safety cases", which outline what safety looks like in precise terms. Another option could be to regulate the *user*, through something akin to an "AI driving licence", supporting the careful use of these potentially harmful technologies.

We are at a juncture where AI development is driven primarily by large commercial companies motivated by profit, rather than social purpose. Social partnerships could be part of the solution: it is vital that the public is aware of where AI is used and affects their lives, and where the data is stored. However, trust is more easily lost than gained, and a high-profile mistake could cause public opinion to turn.

The state of aff[ai]rs in policy

Artificial Intelligence is no longer a niche interest for scientists: it now encompasses multiple disciplines and sectors and it is crucial that this is understood both by academics and by governments. AI is a reserved matter for UK Government, but there are many cross-cutting issues and considerations that impact areas devolved to Welsh Government. The Welsh Government is therefore keen to avoid replicating work taking place in Westminster, and is instead taking a co-ordinating approach across multiple affected areas of the government's work, such as transport, healthcare, ethics, and procurement.

Key AI priorities for Welsh Government include:

- protecting and benefitting all citizens
- supporting public service delivery
- using AI as a catalyst for economic growth.

We know that Wales must avoid wasting resources by replicating work that is already being conducted elsewhere; however, Wales can be ambitious and opportunistic in strategic places.

This webpage offers an overview of Welsh Government AI Policy and Strategy: Artificial intelligence (AI) | GOV.WALES (February 2025).

Protecting citizens

ONS research shows that only 14% of the Welsh public thinks AI carries more benefits than risks (similar to the UK average). Privilege is a dividing line: across the UK, men, non-disabled people, and the 5th IMD quintile (least deprived) are among the groups more likely to agree that AI will benefit them individually.

Members of the public are concerned about AI for a number of reasons. There are fears that it might replace jobs, and this is not without basis. The IMF reports that 40% of jobs are exposed to risk from AI, increasing to 60% in advanced economies. In contrast to the industrial revolution of the nineteenth century, it is now "white collar" jobs that are most at risk. In Wales, the Workforce Partnership Council is working to understand how AI might impact workforce and employment.

Privacy rights are under question, with AI employed in facial recognition for law enforcement, including by South Wales Police. Indeed, especially where AI is used in law enforcement and other serious applications, many citizens are concerned about bias in tools that are often created by undiverse teams, using data skewed heavily towards the majority demographics of the United States.

AI is also playing a role in spreading disinformation, in particular through the generation of fake images and videos which many citizens are unable to distinguish from reality. While fake photographs have existed for nearly as long as real ones, the difference that AI brings is a far greater volume of disinformation, which fact-checkers struggle to keep up with. The critical thinking skills gained through arts and humanities education are more important than ever to combat disinformation.

AI also poses a threat to intellectual property, especially generative AI tools that use the internet as training data. IP policy sits with the UK Government (with the office in Newport, Wales), but the issues are cross-cutting and international. Crucially for Wales, AI poses particular pressures for the Creative Industries, for example when generative AI uses human-made art without credit or payment to create an image. Nearly 200,000 books have been downloaded illegally to train AI. There are also ongoing debates about whether AI can be counted as an inventor – in 2023 the UK Supreme Court judged that it cannot.

Finally, the market domination of big tech companies like OpenAI opens up the possibility of exacerbating the digital divide in Wales by increasing the price of AI tools.

Supporting public service delivery

Welsh Government – working with the Centre for Digital Public Services – has established an AI literacy programme and targeted guidance for the public sector in Wales, focusing on where AI can make the most difference, and understanding the risks. This Welsh guidance is adapted from and supports the UK Government's AI Playbook. Wales's Centre for Digital Public Services has created an online community of practice on Automation and AI and to allow public sector employees to learn AI skills from each other, and avoid reinventing the wheel.

CDPS has also conducted a survey to understand the Welsh public sector's AI maturity and understanding. It found that the approach has been deliberately slow and cautious: there is appetite for

digitisation, and many are exploring the use of AI, but the products available are insufficiently mature, and AI is not yet embedded in business as usual.

There are questions about running before we walk: many public services in Wales would benefit from adopting more basic automation tools before moving on to advanced AI. And in many areas, the quality of data is not yet sufficient to yield useful results with AI. Wales must avoid the temptation to invest in rolling out shiny new AI tools before investing in the solid foundations of data quality.

While multinational corporate AI tools might not be developed within Wales, as a nation we do have control over public spending and procurement. The fact remains that the vast majority of public-facing AI tools have been developed by large US-based commercial companies for the purpose of generating profit, and will be influenced by the biases of those companies and their government. But what might happen if Wales (and other countries) were to link public-sector procurement of technological tools with social purpose, and strong ethical requirements?

AI as a catalyst for economic growth

The UK already has the third largest AI market in the world, after the US and China. In January 2025, the UK Government announced its AI Opportunities Action Plan, with more focus on growth and less on safety. The Welsh Government anticipates that the lion's share of a significant new UK investment by US company Vantage Data Centers will be in Wales.

As is often the case, challenges to growth include infrastructure, money, and skills. Innovate UK funding, STFC's Hartree Centre SME Hubs, plus any new investment through the AI Opportunities Action Plan might go some way towards addressing the issues of infrastructure and money. In the absence of further financial incentives, to attract new companies Wales could look towards an attractive skills offer, favourable working conditions, our culture of national collaboration, and the relative ease of national coordination.

In terms of skills, the Digital Competence Framework is in place for learners in Wales aged 3-16+, but still needs time and effort to have a significant impact on skills. In particular, there is a pressing need to develop AI literacy for school teachers. With a devolved education system, Wales has the opportunity to offer an innovative and progressive digital confidence curriculum across the breadth of the education system, creating citizens and workers equipped to tackle the future. It would be wise for this skills agenda to acknowledge that the majority of the public will be AI users rather than developers, so should include effective prompt engineering, a strong understanding of the specific ethical issues surrounding the *use* of AI tools, critical thinking, and systems thinking. It is possible for AI to be consciously embedded as a tool in teaching and learning, with intelligent assessments that avert cheating; for example, teachers might invite students to put an essay question into a large language model tool, then critique it. In Wales, education inspectorate Estyn has been tasked with conducting a review of AI use in education.

Wales's higher education institutes are core to this highly skilled workforce, and they also produce the world-class research that underpins innovation. Independent analysis commissioned by the Learned Society of Wales shows that research from Wales is driving significant positive impacts both within Wales and internationally. However, university researchers are vulnerable to exploitation by large international tech companies which convert major developments in basic sciences and mathematics into profit without crediting or paying scientists. University researchers are under conflicting requirements to both publish their work and also obtain patents or licenses. Wales and the wider UK would benefit economically from supporting university researchers to develop and retain IP, and incentivising researchers to continue to make advances at their home institution.

Purpose-built machine learning tools for scientific research tend to carry lower risks and yield more accurate and reliable results than large language models, because they usually have a narrow remit, and a quality-assured data set. There are nevertheless concerns within the scientific community that there is insufficient scrutiny of AI ethics and impact by UK research funders compared to EU programmes: scientists must remain mindful of the research sector's long history of causing harm – from eugenics to the atom bomb.

These ambitions for growth and progress have also been in tension with climate targets. Using ChatGPT-4 to write a 100-word email consumes enough energy to fully charge a iPhone Pro Max seven times. However, DeepSeek claims to consume 50–75% less energy than its US counterparts. Indeed, the recent emergence of DeepSeek shows that calls for extremely large and complex systems may need scaling back, and more investment dedicated to the UK's talent for growing software expertise from SMEs.

References

This webpage offers an overview of Welsh Government AI Policy and Strategy: Artificial intelligence (AI) | GOV.WALES (February 2025).

Public awareness, opinions and expectations about artificial intelligence, Office for National Statistics, October 2023.

Gen-AI: Artificial Intelligence and the Future of Work, International Monetary Fund, June 2024.

Broadening the Gains from Generative AI: The Role of Fiscal Policies, International Monetary Fund, June 2024.

Using artificial intelligence at work, Welsh Government, December 2024.

Police to use live facial recognition in Cardiff during Beyoncé concert, The Guardian, May 2023.

How AI-generated disinformation might impact this year's elections and how journalists should report on it, Reuters Institute for the Study of Journalism, University of Oxford, March 2024.

These 183,000 Books Are Fueling the Biggest Fight in Publishing and Tech, The Atlantic, September 2023.

AI cannot be patent 'inventor', UK Supreme Court rules in landmark case, Reuters, December 2023.

This Welsh guidance is adapted from and supports the UK Government's AI Playbook.

Global AI Power Rankings: Stanford HAI Tool Ranks 36 Countries in AI, Stanford University Human-Centred Artificial Intelligence, November 2024.

Wales at the forefront of UK drive to 'turbocharge' AI, Welsh Government, January 2025.

Estyn to review use of Generative Artificial Intelligence (AI) in schools, Estyn, January 2025.

'Thirsty' ChatGPT uses four times more water than previously thought, The Times, October 2024.

About this report

In December 2024, the Learned Society of Wales convened a group of experts from the research, public and public sectors to discuss the current status of Artificial Intelligence for Wales. This report summarises the evidence and expert advice gathered through this forum of experts.

Chair: Rossi Setchi (Cardiff University)

Contributors:

Gary Beauchamp FLSW (Cardiff Metropolitan University)

Anthony Cohn FLSW (University of Leeds and Alan Turing Institute)

Joseph Connor (CarefulAI)

Martina Feilzer FLSW (Bangor University)

Dawn Greenberg (UK Research and Innovation)

Layla Jader FLSW

Richard Jones (Welsh Government)

Mark Lee FLSW (Aberystwyth University)

Chenfeng Li FLSW (Swansea University)

Biagio Lucini FLSW (Swansea University)

Faron Moller FLSW (Swansea University)

Omer Rana FLSW (Cardiff University)

Qiang Shen FLSW (Aberystwyth University)

Mark Taubert FLSW (Velindre University NHS Trust & Cardiff University)

Peter Thomas (Centre for Digital Public Services)

Stephen N Williams FLSW (Queen's University Belfast)

Reyer Zwiggelaar FLSW (Aberystwyth University).