Senedd Cymru
Pwyllgor yr Economi, Masnach a Materion Gwledig
Ymchwiliad: Deallusrwydd Artiffisial ac Economi Cymru
Cyf: AI06
Ymateb gan: Dr Morgan Jones

-

Welsh Parliament
Economy, Trade, and Rural Affairs Committee
Inquiry: AI and the Welsh Economy
Ref: AI06
Evidence from: Dr Morgan Jones

The Role of Artificial Intelligence in Wales: Opportunities, Challenges, and Strategic Directions

Morgan Jones*

Artificial Intelligence (AI) is reshaping industries worldwide. Wales faces unique opportunities and challenges as a small nation navigating this transformative landscape. While the theoretical foundations of AI, including the latest algorithms and architectures, are globally accessible, the infrastructure required to scale and deploy state-of-the-art models often limits smaller countries like Wales. This disparity underscores the importance of Wales focusing on innovation, strategic collaborations, and specific applications of AI that align with its economic, cultural, and social goals.

1 The Current State of AI Development

AI is the simulation of human-like intelligence in machines, enabling them to learn, reason, solve problems, and make decisions. **AI is not new**, the Percepton Algorithm that could perform image classification tasks has been around since the 1950s [15]. Even before this, control theory has been used since the dawn of the industrial revolution, endowing machines with the power of feedback. Feedback allows machines to use real-time sensor data to intelligently adapt their actuation automatically, these systems have the ability to overcome uncertainty and achieve their objectives. Arguably the first mathematical analysis of an autonomous controlled system was James Maxwell's analysis of the flyball governor system, dating all the way back to 1868 [11].

Since the public release of ChatGPT in 2022, much of the AI hype has centred around Large Language Models (LLMs). While LLMs are certainly important it should be emphasized that **AI encompasses far more than LLMs** and investment and research should not only focus on LLMs. LLMs have proven effective in tasks like language generation and creative media production but have a tendency to hallucinate and provide misleading outputs. Their reliance on human oversight to correct small defects and ensure coherence often limits their broader applicability. In fact Machine Learning (ML) techniques as a whole (methods that involve learning directly from data) "often **do not work**. They can be constructed haphazardly, deployed indiscriminately, and promoted deceptively" [14]. Wales should take **heed of the lessons of the AI winter** [4], a period in the 1980s marked by inflated expectations followed by disillusionment, highlighting the importance of tempering current enthusiasm with a focus on foundational challenges and the development of diverse methods. After all, the work of [7] shows that while increasing model size, data, and compute leads to performance improvements, these gains follow a power-law relationship with **diminishing returns**. This suggests that LLM performance will plateau with scale, contrary to the hype that assumes rapid, indefinite improvement.

Reinforcement Learning (RL) is a ML paradigm where an agent learns to make sequential decisions by interacting with an environment, receiving feedback in the form of rewards or penalties. Like LLMs, this field has also attracted significant attention lately. DeepMind used RL to create a programme, AlphaGo, that beat the best human Go player in the world as well as using RL to discover new protein folding structures, AlphaFold [6]. Open source RL environments, such as those of DeepMind [19] and OpenAI [2], play a crucial role in the landscape of AI and robotics. These environments provide benchmarks, allowing for easy comparison between different RL algorithms, and also serves as initial training grounds for agents to learn in simulation and then transfer what they have learnt to reality (Sim-to-Real [22]). However, these environments are often a **bad approximation of reality** [16]. On top of this it appears that

*M. Jones is with the Department of Automatic Control and Systems Engineering, The University of Sheffield, Amy Johnson Building, Mappin Street, Sheffield, S1 3JD. e-mail: morgan.jones@sheffield.ac.uk

for many RL benchmarks simple random greedy algorithms perform competitively with the state of the art RL algorithms [9, 10]. Although benchmarks based on ATARI games, chess and Go are interesting and make for impressive demonstrations, to bridge the reality gap and allow for new smart robotics we need to construct rich, multi-dimensional models grounded in physics as well as data. These models need to be capable of accurately capturing the nuanced relationships within complex uncertain realistic environments.

The recent advancements in LLM and deep learning has largely been due to the Transformer architectures and large-scale computational resources, with industry leaders such as OpenAI and DeepMind driving breakthroughs. This new architecture was first developed by eight Google scientists in the seminal paper "Attention Is All You Need" [20]. What is important to note here is that it is still an open question if the Transformer architecture is truly the best Neural Network (NN) architecture, however current experimentation seem to indicate this is the case. With the rise of deep learning **access to data and computational resources have become vital**. This has caused a situation where many of the frontier models are being developed by large silicon valley companies with deep financial pockets, a vast number of GPUs and access to an abundance of data. Academia is no longer at the forefront of producing the latest generation of ML models. For small research groups, or small countries like Wales, to stay at the cutting edge of AI we need to go beyond the Transformer architecture to reach a new paradigm where less data and computation is required while having safety and performance certifiable.

2 Strategic Considerations for Wales

While the advancements in AI are remarkable, the narrative that AI is a solved problem, destined to replace humans in the job market and render traditional algorithms obsolete, is flawed. Critical open questions remain about the efficiency of Transformer architectures and the limitations of backpropagation for training NNs. Wales should not assume the topic of AI has been completed, there is still much to discover. It is sensible that some of the available welsh resources should be given to prepare for the future development of Artificial General Intelligence (AGI), a world of abundance and universal income brought on by the existence of AI agents that can complete all tasks humans can complete. However, it is important that the **majority of resources should be given towards foundational research** that strives to advance beyond current AI methods or develops specific methods to solve problems of direct interest to Wales rather than pondering hypothetical AGI possibilities. Moreover, although LLMs have proven to be extremely impressive, issues still persist and most developments happen solely in silicon valley. Therefore it is important that Wales **diversifies AI research efforts** to not only concentrate on LLMs.

Having argued for the diversification of research efforts away from LLMs and the difficulty to be competitive with silicon valley companies it is still important that Wales has access to the latest advance LLMs, or any other state-of-the-art AI tool. Therefore, it is necessary that the Welsh government **maintain a good relationship with world leading AI companies** to keep all models accessible to Wales. This could involve partnerships of co-sponsored research projects, tax incentives, access to reliable energy to build data centres and access to data itself. It should be noted that ChatGPT performs relatively well at welsh language tests [13] perhaps due to the new partnership between the welsh government and OpenAI [21]. Such partnerships should be encouraged and expanded.

Wales should prioritize fostering novel ideas and applications that leverage the nation's unique strengths. For example, AI can support Welsh-language preservation, renewable energy optimization, and rural economic development, areas where **Wales can lead rather than follow**. Despite these opportunities, Wales must address several challenges to maximize AI's potential. Small and Medium-sized Enterprises (SMEs), which dominate the Welsh economy, often lack the resources and expertise to adopt AI solutions. Overcoming these barriers will require government support through training initiatives, subsidies,

and partnerships with universities and research institutions.

3 AI Applications Aligned with Welsh Strengths

Wales's unique economic and cultural landscape offers fertile ground for impactful AI applications. Key opportunities include:

- **Energy:** Wales is currently transitioning away from fossil fuel energy production. Wind now makes a significant proportion of energy production. AI can be used to decide where it is optimal to construct new wind farms and what should be the formations of these farms. Using fluid dynamics it is possible to enhance wind energy production by continuously adjusting individual turbine blades, however this is a notoriously difficult optimal control problem [17] where data driven AI methods are required. If the proposed Swansea Bay tidal lagoon [12] ever gets approved then operations and design could be enhanced using AI. Autonomous drones could be used to map out the terrain and landscape of Wales to identify areas for geothermal, wind, or other means of energy generation.
- **Transport:** Wales has a complex road system with many difficult to navigate non-uniform rural roads. There is a real possibility that Wales misses out the the coming driverless car revolution unless significant effort is made to collect the required welsh road training data. Alternatively, Wales could choose to invest in an intelligent rail system. It is much easier to make trains driverless due to the predictability of the tracks. Wales should pioneer driverless trains or at the very least use AI to reduce signalling errors. Moreover, AI paired with autonomous drones might improve track fault identification that could decrease the number of delays.
- **Smart Homes**: Some areas of Wales have a large density of packed terrace houses. AI could be used to control communal heating systems by learning residents habits to perform optimal zone heating while ensuring a fair distribution of heat across the residents. Moreover, the abandoned coalfields of South Wales have great potential to heat homes, not through coal use but through the geothermal heat stored in the mine water [3]. Not only can AI be used to help us efficiently extract this underground heat for communal heating but also installing residential batteries can be used to minimize consumer peak energy demands [5]. AI could be used to learn Welsh consumer habits as well as predict renewable energy generation in order to optimally charge and discharge residential batteries.
- **Smart Environment:** There is opportunity to construct and leverage widespread sensor networks to monitor key factors like air and water quality, traffic patterns, pedestrian footfall, and flood risks. In collaboration with academics we can integrate this data with advanced machine learning models. This can provide actionable insights to support policy-making, optimize resource allocation, identifying emergent diseases and enhance real time flood alerts. The idea of an "Urban Laboratories" is not new [8], there already exists several in the UK, notably UCL. These centres should be integrated with existing governmental services, for instance in flood alerts. Flooding is a very localised phenomena and weather modelling can only go so far to predict which areas are likely to experience it. Local sensors can supplement these predictions. AI can be used to do optimal sensor location placement and integration of sensor data with existing models and alert systems. South Wales is known to have high levels of Radon [1], currently households are oblivious to short term spikes and little mitigation is employed, AI can help here.
- **Manufacturing**: Wales traditionally has had a large industrial sector. This sector could be revitalised with the introduction of new automation technologies. Wales could pioneer smart robotics and AI use in manufacturing, driving up productivity and economic growth.
- **Tourism and Heritage**: Augmented Reality (AR) applications powered by AI could enrich tourist

experiences at iconic sites like Eryri National Park and Conwy Castle. AI can also help monitor and preserve heritage sites by predicting structural deterioration.

- **Economic Development and Healthcare in Rural Areas**: AI can optimize local supply chains, connecting farmers directly with markets and consumers. Enhanced weather prediction models could support agriculture and mitigate the risks of extreme climate events. AI-powered telemedicine platforms and predictive healthcare analytics can bridge gaps in medical access for Wales's ageing and dispersed population, improving care outcomes and reducing costs. There could be a national AI powered genomics project that improves personalized medicine in Wales based on genomic data.

4 Building a Skilled Workforce

A key enabler of AI adoption is a workforce equipped with the necessary skills. Welsh educational institutions, such as Cardiff University and other universities in Wales, could play a pivotal role in helping to foster and teach AI. These institutions should prioritise integrating AI, machine learning, and data science into their curricula and outreach programs, providing a diverse range of educational opportunities tailored to varying levels of expertise. Beyond technical expertise, training programs should emphasise AI literacy for non-technical sectors such as agriculture and manufacturing, empowering workers to harness AI-enhanced tools effectively.

For those new to AI, no-code platforms offer an accessible entry point, enabling individuals to utilise AI tools without requiring technical backgrounds [18]. For professionals aiming to deepen their understanding, education should cover user-friendly software frameworks like TensorFlow, which bridge the gap between simplicity and functionality. Moreover, providing opportunities to learn programming languages such as Python enables individuals to customize and optimize AI systems for specific needs. Finally, Wales should provide a pathway for the next generation of scientists to emerge that push forward the frontier of AI. These scientists should have a profound understanding of the underlying mechanisms and algorithms of ML techniques which is over and above the knowledge of how to simply implement techniques on readily available software.

5 Conclusion

Wales has the opportunity to establish a unique position in the global AI landscape. To achieve this, it is essential for Wales to avoid being overly influenced by the current excitement surrounding LLMs. The Welsh government should approach all AI-related hype with caution, avoiding premature shifts in focus toward speculative concepts such as AGI. Instead, efforts and resources should prioritize practical AI applications that address immediate societal and economic challenges. At the same time, the Welsh government must continue fostering strategic partnerships with leading AI companies worldwide. However, innovation should never be outsourced, Wales must remain committed to advancing the frontiers of AI through its own path.

Biography

Dr. Morgan Jones was born in Caerphilly, Wales in 1994. He received the MMath degree in Mathematics from The University of Oxford, England, in 2016 and the PhD degree from Arizona State University, USA, in 2021. Since 2022, he has held the position of lecturer in Machine Learning and Control Theory in the Department of Automatic Control and Systems Engineering at the University of Sheffield. He also holds the role of Cross Theme Lead in Supporting Digital Education at the Center of Machine Intelligence. His research primarily focuses on the intersection between Machine Learning and Control Theory with particular attention paid to surface reconstruction, the estimation of reachable sets, attractors, and regions of attraction for nonlinear ODEs, extensions of the Dynamic Programming (Reinforcement Learning) framework to non-separable cost functions, System Identification and Randomized Numerical Linear Algebra.

References

[1] J. Appleton and J. Miles. Radon in wales. 2005.

[2] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[3] G. Farr, S. Sadasivam, I. A. Watson, H. R. Thomas, D. Tucker, et al. Low enthalpy heat recovery potential from coal mine discharges in the south wales coalfield. *International Journal of Coal Geology*, 164:92–103, 2016.

[4] J. Hendler. Avoiding another ai winter. *IEEE Intelligent Systems*, 23(02):2–4, 2008.

[5] M. Jones and M. M. Peet. Solving dynamic programming with supremum terms in the objective and application to optimal battery scheduling for electricity consumers subject to demand charges. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1323–1329. IEEE, 2017.

[6] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

[7] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

[8] A. Karvonen and B. Van Heur. Urban laboratories: Experiments in reworking cities. *International Journal of Urban and Regional Research*, 38(2):379–392, 2014.

[9] C. Laidlaw, S. J. Russell, and A. Dragan. Bridging rl theory and practice with the effective horizon. *Advances in Neural Information Processing Systems*, 36:58953–59007, 2023.

[10] H. Mania, A. Guy, and B. Recht. Simple random search provides a competitive approach to reinforcement learning. *arXiv preprint arXiv:1803.07055*, 2018.

[11] J. C. Maxwell. I. on governors. *Proceedings of the Royal Society of London*, (16):270–283, 1868.

[12] S. Petley and G. Aggidis. Swansea bay tidal lagoon annual energy estimation. *Ocean Engineering*, 111:348–357, 2016.

[13] G. PRYS and D. B. JONES. First welsh language evaluations of openai's gpt large language models. *Language and Technology in Wales: Volume II*, page 38.

[14] I. D. Raji, I. E. Kumar, A. Horowitz, and A. Selbst. The fallacy of ai functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 959–972, 2022.

[15] F. Rosenblatt. Perceptron simulation experiments. *Proceedings of the IRE*, 48(3):301–309, 1960.

[16] E. Salvato, G. Fenu, E. Medvet, and F. A. Pellegrino. Crossing the reality gap: A survey on sim-to-real transferability of robot controllers in reinforcement learning. *IEEE Access*, 9:153171–153187, 2021.

[17] R. J. Stevens, D. F. Gayme, and C. Meneveau. Effects of turbine spacing on the power output of extended wind-farms. *Wind Energy*, 19(2):359–370, 2016.

[18] L. Sundberg and J. Holmström. Democratizing artificial intelligence: How no-code ai can leverage machine learning operations. *Business Horizons*, 66(6):777–788, 2023.

[19] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.

[20] A. Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.

[21] Welsh Government. Chatgpt learns welsh, 2024. Accessed: 2024-11-24.

[22] W. Zhao, J. P. Queralta, and T. Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pages 737–744. IEEE, 2020.