The UK government and artificial intelligence: embracing the vampire

Carl Rowlands

The government's uncritical embrace of AI is binding the UK to companies that are draining the economy and embedding technological dependence.

The UK government has increasingly committed itself to full-scale adoption of Artificial Intelligence (AI) across government services and planning, without serious caveat. This can be described as *positivist*, in both senses: it sees AI primarily as an economic driver of growth; and (as a complement to this) it takes a functional, 'data-driven' perspective on the human mind and its relation to assistive technology. An apparently downbeat geopolitical evaluation of the UK's competitive position has persuaded the government that a wholesale adoption of hyperscale, proprietary and corporate-provided AI is an absolute imperative, based mostly on co-operation with US tech companies - the latter for ostensible reasons of national security. This is accompanied by complete commitment to all the associated infrastructure, which is being pushed through without serious consideration of objections.[1]

DOI: 10.3898/SOUN:90-91.08.2025

This article looks at the various spheres impacted by this approach, and the associated objections to it, and contrasts these with the stated primary motivations for the Starmer administration's advocacy of AI. In particular, it will argue that the UK government's relationship with 'Artificial Intelligence' cannot be understood solely as relating to what is technically or scientifically appropriate: it is also intimately linked to public discourses around the value of creativity and learning; particular axes of capital, power and influence; and, in geopolitical terms, Britain's increasingly subservient relationship to its erstwhile ally, the United States of America.

Data protection and data predation

Whilst much of the activity of the largely unreformed UK Parliament can be regarded as anachronistic, looking briefly at separate items of legislation from different periods remains one of the best ways of understanding rapid change in this field. This makes it possible to map intense social and technological changes onto areas which are subject to legislation.

The differences between the Data Protection Act, 1984 and the Data (Use and Access) Act, 2025 - both in terms of practical mechanisms and underpinning values - are extremely instructive. Both Acts are the product of powerful, one-party governments, and, at least superficially, are aimed at providing a framework for a flow of information, with a view to establishing a general, aggregated European standard. But the similarities end there.

Under the provisions of the 1984 Act, organisations which held details on more than 50 individuals were required to register the database, and the different details being held, with the Data Protection Registrar. Both the sources of the data and the recipients of any processed data were specified as being part of the registration process. The Registrar had the power to place an explicit ban on the transfer of this data to other countries. The subjects of the database, the individuals listed in the database, could make a request to the Registrar for full details of their records. There were, however, a number of exemption clauses added, which affected the ability to request data for national security, financial services, social care and healthcare.

Many of these provisions remain in place, with the Information Commission providing oversight, but - in contrast to the stated aims of the 1984 Act - the

Data (Use and Access) Act, passed in June 2025, has the explicit aim of enabling innovation. The Act could, in theory, prevent holders of copyright from asserting their rights over generated materials which use the original works as a source, or as 'training material'; and, in the interests of innovation, data on individuals can be used for anything that could be described as 'scientific research' - whether undertaken by private or public institutions, and with no precise definition of what 'scientific research' could mean. Transfer of data overseas is practically encouraged.

It is here that we can see the difference. In 1984, concern was based on misuse of what were essentially digitised versions of existing records on paper, film and tape. In 2025, the primary concern is that the UK shouldn't be left behind in the dash to commodify this data, to use personal data for a whole range of commercial and scientific purposes. Critical to this change has been the growth of machine learning, generative assistive tools and Large Language Models (LLMs), all of which are described as AI, and all of which require huge amounts of data input for 'training'.[2]

These engines are usually based upon neural networks. They require feeding with human-generated and human-centred data in order to calculate pathways, probabilities and models of deviance. The sources which 'feed' and 'train' AI models need to be huge - both comprehensive and diverse. Even then, there is a considerable amount of human intervention required - most usually provided by raters and moderators, to ensure that users see something intelligible and maybe even accurate. Data is critical: masses of tabulated raw data input are required for LLMs, which are, essentially, a predictive technology. They are based upon tokens representing different elements of a huge corpus of information, and calculations of proximity, which can then be determined, on the basis of probability, as a 'continuation' of an existing text. That the UK government would prioritise companies' access to this huge expanse of data is a clear indication of how it is now responding to, and, one might say, obeying, the prerogative to feed AI with everything: the personal details, all aspects of the lives and even the artistic contributions of its citizens. This data already seems critical to the functioning of capitalism.

Palantir and the National Health Service

As a company, Palantir has been well-documented. Founded by Peter Thiel, a particularly right-wing Silicon Valley billionaire, it has succeeded in inserting itself into numerous US government contracts, basing itself initially within the defence

and security sectors. Given this as a starting point, Palantir's move into the health sector and its success in winning a multi-million-pound contract with the NHS in England is alarming. In fact, as we will address later, defence relationships remain absolutely critical as the starting point for many AI-related projects.

Palantir's NHS project is, seemingly, straightforward enough - delivery of a comprehensive system of patient records, following continued IT implementation failures. But, for those who have followed this otherwise dry IT-related saga, there is more than a little suspicion about what other uses Palantir intends for these patient records. The apparently deliberate vagueness in the 2025 Data (Use and Access) Act's definition of the 'scientific research' that can be conducted upon huge data sets, means that this project effectively involves the signing over of a huge amount of patient data to an unaccountable US company and third parties, with no scrutiny or accountability. The data - including the results of blood samples and reports of visits to doctors - could potentially be used in many different ways, including, for example, for the mapping of genetics, or incidences of mental illness. This could represent a UK data set that Palantir could, in theory, then 'repatriate' and combine with various other databases.

There has been intense lobbying behind the scenes for Palantir to be adopted across a range of government departments, including the police, the military, schools and job centres. Despite the company's deliberate stance of aloofness from the Starmer administration's plans to introduce a 'Digital ID' system, there is an ongoing effort to embed Palantir across public administration and to absorb as much data as possible from everywhere. It should be noted that it was the Conservative government led by Rishi Sunak that first introduced Palantir into public services. Paid advocates for Palantir include former Labour deputy leader, Lord Tom Watson; ex-Ministry of Defence strategy director, Polly Scully; and former MI6 head, Sir John Sawers. On his last visit to the US, Keir Starmer paid a visit to the Palantir office in Washington DC. Now itself a controversial and widely-discussed topic, this was a meeting seemingly organised by the disgraced Labour politician Peter Mandelson, currently under investigation for misconduct in public office. Mandelson at that time was the UK's ambassador to the United States - and, until early 2025 the head of Palantir's lobbying firm, General Counsel.[3]

There is, of course, a range of ethical questions that should inform use of medical data to train AI systems. Yet scientific research projects around the UK are

being pressured to sell or donate existing data to support the training of various AI models. The fact that this is generally happening without significant public scrutiny, or consistent evaluation of public benefit, reflects a failure of the public domain in tracking ongoing developments.

This is not to say that machine learning will not be useful in developing cures or therapies - it is possible, and perhaps even likely, that it will provide improvements in treating cancer and a range of illnesses. But there is a very important question about who will provide these therapies, and who will ultimately profit from a data set which is collectively provided and collated from public sources and public services.

AI for Life

Since coming to power in July 2024, scarcely a week has passed without the Labour government making another grand announcement about AI. The flagship policy is something called an AI Opportunities Action Plan, within which educational initiatives are included. Specifically, there are plans to introduce 'AI lesson assistants' and develop various AI tools for teachers, to 'speed up lesson planning and reduce workloads'.

The government has shown a degree of restraint in some regards. There is an 'edtech evidence board' to evaluate the impact of AI upon teaching and learning, and the implications for future policy. But, as with so many of the UK government's policy decisions, there are contradictions between this and other government or public initiatives which seem difficult to resolve. For example, the evidence-based approach which is proposed by Ofsted is contradicted by a 'memorandum of understanding' signed by the Department of Science, Innovation and Technology (DSIT) and OpenAI, the provider of ChatGPT, in July 2025. This followed similar agreements with Anthropic and Google, with promises that OpenAI will work with DSIT to find areas where 'advanced AI models' can be deployed. One of the uses cited is that homework could be marked and graded by machine.[4]

There is little doubt that ChatGPT and other LLMs will change education, regardless of what the UK government plans. But one possible end result is that asking students to perform tasks such as homework - routine writing and numerical tasks outside of controlled environments, especially tasks that can be

quickly resolved by predictive algorithms - will be rendered practically useless. It may also well be that the technology makes redundant existing forms of knowledge assessment.

Rather than considering the possibility that ChatGPT or Claude might make traditional concepts of 'homework' somewhat futile, the instinct of government is to lean into the idea that, regardless of whether the students' tasks have been already automated, the teachers' feedback could also be automated. This results in levels of technological absurdity - in essence, LLMs squawking ineffectively at each other - that would not look out of place in one of Douglas Adams's later novels. Given the degree of antipathy which is likely to result, it is likely that the 'tech revolutionaries' - including Bill Gates, a leading advocate for replacing teachers with AI machines - will find classrooms among the first places where their aspirations for an AI utopia will hit a brick wall.

However, in light of the suggestion - somewhat discounted by government sources - that Chat GPT could be provided to everybody in the UK free of charge, one suspects that Big Tech has particular plans for the education system. The talk is of a personalised AI assistant - an AI that is with you for your whole life, which is contoured and moulded to your own personality (and which, inevitably, would also mould your personality and responses). This may initially be marketed as a 'pedagogical assistant' of some sort, potentially in combination with Augmented Reality devices, such as the 'smart glasses' being advocated by Meta's unhappy CEO, Mark Zuckerberg, or pendants which act as monitoring and intervention devices. These would potentially also capture and analyse individual speech and other learning interactions. A very particular example exists in one elite Silicon Valley school that is being cited as a flagship project, whereby pupils spend two hours a day interacting with dynamic, AI-calibrated learning materials.

The aim, ultimately, would be to have an AI avatar in a classroom rather than a teacher. It is not hard to see that making AI essential to learning is one of the ultimate goals of Big Tech. There is already a concerted attempt to incorporate AI as a part of writing any kind of document in Microsoft Word, or email in Microsoft Outlook, something that many users compare to 'junk food' - a shoal of unverified, loosely relevant content generated by a Large Learning Model (LLM). Any consequent negative effects on the cognitive and social development of young people will, of course, be left to the care of a grossly underfunded and overstretched

public healthcare sector. In the final analysis, because so much of the education of children relies upon Microsoft Corporation, which already has huge power over various parts of the UK's economic and social activities, the truth is this: if Microsoft Corporation wish to roll out different AI functionalities to students of any age, they will do so - and this will accompany pressures on educators to demonstrate their usage of this functionality. Only a more fundamental, public, questioning of the role of technology in education, and the role of a monopoly provider, would begin to address this.

National security

The government's AI Opportunities Action Plan is mostly - albeit not exclusively - about extending the reach of large tech corporations across public services. It is therefore unlikely that calls for digital sovereignty - including the concept of the 'commons' underpinning ownership and usage of data - can ever be satisfied under the operational model outlined in the Plan. The nature of the cloud militates against a separation of jurisdiction - given that it is a massive stack of computing resources which can be dynamically reallocated, according to demand, and the way in which data centres operate. There are disturbing precedents in this regard: the United States government has openly preserved the ability to demand that US companies hand over data, even if the data centre is (supposedly or otherwise) located in another jurisdiction. Any assignment of data to US cloud providers effectively allows extraterritorial US government access, no matter where the data centre is located.[5]

Whilst this has wide implications for personal privacy and commercial confidentiality, perhaps the most striking ramifications in the UK will be seen in the military adoption of AI. Incorporating foreign companies into weapons systems, both strategic and tactical, further embeds third parties into the UK's defence infrastructure, and even into decision-making processes. It incorporates UK test data into data lakes that are available to any US-based company - for example, Palantir - which chooses to work with them.

At this point we need to consider the implications of the leading role that Israelis and Israeli companies have assumed in the cybersecurity and techno-military spheres. This role reflects the ongoing relationship between Google - and, especially the Microsoft Corporation - and the Israeli military, and the particular set of skills that have been nurtured over the years by units aligned with Mossad: a significant

number of startup companies have drawn from Mossad personnel.[6]

The use of the 'cloud' in itself presents particular questions around data sovereignty. But the use of data from live warfare, and the harvesting of data resulting in the potential targeting of civilians, opens up ethical dimensions which are completely unaddressed by the UK government's current AI strategy. One is forced to conclude that this omission is entirely deliberate, and that the consequent introduction of technical dependencies is entirely by design.[7]

'Full Spectrum Surveillance' - facial recognition and beyond

A small section in the effusive AI Opportunities Action Plan is dedicated to 'Automated Threat And Anomaly Detection'. What this actually means is not really spelled out, but the example cited is that it is 'used to 'clean up social media'. As so much of our lives is now conducted online - all the way from work meetings, through family life, interactions with friends, dating, consumption of information, banking and food shopping - the idea of using AI to address fraud and cyber-crime might be interesting. However this does not, apparently, reflect the government's plans.

On a very basic level, one of the first uses for machine learning, based upon neural networks, was to recognise patterns - for example to identify tanks in imagery. Another of the early uses for machine learning, documented by Kate Crawford, was based on thousands upon thousands of different human faces, as expressions and different identities could be inferred from huge sets of data.[8] There is therefore a crushing inevitability in the use of mobile Live Facial Recognition (LFR) systems in policing, reflecting expanding real-world databases. And this data can be at least partially augmented by tapping the mobile devices with which we now surround ourselves. The combination of reading information from a SIM, via a transmission mast which interrogates a mobile phone, together with an external camera device designed to extract facial biometrics, is probably enough to provide a tentative identity to the vast majority of figures in the street.

The early trials of LFR cameras by the Metropolitan Police, with static cameras located at Kings Cross, were considered a success, and the arrest of a few crime suspects was cited. It is probably the data collected at Kings Cross which determined that LFR was successful, and that it could be matched and linked to various other forms of personal data. And this combination effectively opens the door to a form of

'full spectrum surveillance', where, if necessary, information can be harvested easily from specific locations.

Observers have specifically noted that LFR tends to be used around areas and events with a large black participation - for example, Notting Hill Carnival. Another example would be political demonstrations, with pro-Palestinian activists already targeted by LFR systems.[9] Before recent demonstrations in support of the proscribed Palestine Action group, the Metropolitan Police even warned participants, that their presence, confirmed by 'full spectrum surveillance', could affect their ability to become teachers at a later stage in their lives.

Needless to say, there is a pressing need for a clear set of guidelines on how surveillance data should be used and deployed - something that even some of the police themselves are asking for. The lax and ill-defined regulatory environment provided by the Data (Usage and Access) Act, also means that there is a possibility that data resulting from facial recognition cameras could be passed onto third parties, and even potentially used commercially. The main risk, however, is undoubtedly that of 'doxing' - the publication of individual data containing personal details of participants. This could be used to flatten and suppress valid, legal forms of political dissent, whether left- or right-wing. An example of the possibilities for repression was provided by the Hungarian government, which recently threatened use of facial recognition technology to identify people attending the banned Budapest Pride demonstration, and the subsequent fining of participants. The technology in this case was supplied by Chinese companies, where facial recognition has become a critical component of asserting police authority.

Other questions are raised by the Ministry of Justice's plans for AI: trials are being conducted on prisoners to establish the predicted likelihood of repeat offences, especially of murder. In true Benthamite fashion, the monitoring of prisoners offers the possibility of feeding in data from just about every aspect of an individual's life, and interactions.

In the meantime, the Department for Work and Pensions is set to introduce machine learning to identify fraud, claw back overpayments, and manage benefits sanctions, in combination with the access to individuals' bank accounts enabled by the Public Authorities (Fraud Error and Recovery) Bill.[10] The implications of this initiative, given the growth in both in-work and out-of-work poverty in the UK since the 2008 financial crisis, could be profound; it is part of an increasing

tendency to 'digitally police' those in poverty, through processes which have limited accountability, and in which there is virtually no transparency on how algorithms are weighted and trained. It is possible to see the amalgamation of facial recognition, phone, tax and social security data being used for a variety of purposes - not least of which is political repression.

It should also be noted that the potential for workplace surveillance is also greatly enhanced by the use of AI; and, again, this relates to the ubiquity of Microsoft 365 as a package of utilities. Microsoft 365 can already provide management with a number of reports on all sorts of aspects relating to employees: the number of files accessed and created, and how many messages are being sent and received across different channels. AI, it is clear, will be able to provide scores for each employee on 'projected productivity', according to Key Performance Indicators (KPIs) provided by management. Ultimately, for many companies, the temptation to address behavioural aspects with tech-centred, 'data-driven' analysis is likely to prove too great a temptation to resist.

The Plan For Change

Yet we are not suggesting here that the main reason for the government's excitement over machine learning is related to its latent potential for repression: the motivation is, in the main, the government's hope that it can help revive the economy that it inherited in such poor shape. At least some of the economy's weakness seems to have a basis in demographics - an ageing population, with increasing levels of incapacity, and extremely uneven regional development. Artificial Intelligence is seen as a way of addressing some of these issues, and providing a boost to wider productivity and growth. Labour's fixation on 'AI' and its preferential treatment for corporations suggests a reluctance to adopt serious regulation, let alone rebalance economic power away from the big corporations. This contrasts with its enforcement of strict identity requirements for access to certain material on the internet. As is often the case, the government finds itself awkwardly poised between its desire to ingratiate itself with the powerful (laissez-faire for the rich) and its existential fear of being seen as in any way connected to permissive 1960s liberalism (authoritarian tendencies and surveillance for the rest of us).

To get a clear view of the UK government's thinking on this we can refer to a speech made by Peter Mandelson, delivered just prior to his fall from grace

in September 2025. This gives a fairly detailed description of the official view.[11] Despite a career shrouded in intrigue, it is fair to say that Mandelson is not an especially original political thinker: so he can safely be assumed to be speaking not just for Keir Starmer's weak administration, but also for the UK's 'Treasury' perspective more generally.

The starting point is that the UK, post-Brexit, finds itself troubled economically and politically. Donald Trump's anti-migrant policies, which result in a continual process of rounding-up and deportation of 'illegals', is, *sotto voce*, regarded as opening up the possibility of persuading US tech companies to come to the UK - to feel at 'home' in Britain, as Mandelson puts it, and in a place where companies can access labour from around the world more easily. In return, there is a steadfast commitment to open up the UK's public infrastructure as part of the 'Atlantic alliance', including relaxing regulations on nuclear power operations. According to Mandelson, it doesn't matter if US companies buy up UK tech startups - the two countries are, to all intents and purposes, indivisible, allies against the threat posed by China. To provide some kind of simple summary of Mandelson's speech: the absorption of the UK into the structures of US capital and defence is worth the costs, risks and diminution of sovereignty.

In technological terms, this means that an alternative approach - for example one based upon open-source models, as generally favoured by policy in China and by some in the EU - is not specifically mentioned by the AI Opportunities Action Plan, even though it is probably less demanding in environmental terms, and allows much more freedom for independent development. Open-source projects - which aim to optimise and reduce the processing loads of AI, and allow machine learning and generative content to be run from the ground up - are sometimes described as 'frugal AI'. Whilst some governments in Europe support the initial development of these autonomous approaches, the UK government's commitment to US Big Tech is unwavering, in return for some inward relocation of resources.

As part of this economic boost, the government has declared the construction of data centres to be in the national interest, thereby reducing the scope for planning objections. AI Growth Zones facilitate the construction of data centres as part of a renewed national infrastructure. Even a cursory analysis, however, would conclude that, though they do initially provide jobs for less prosperous areas, at least in their construction and configuration, once constructed they provide very

few jobs, largely in service engineering, maintenance and security. There may also be some lower-grade jobs created, involving the provision of prompts to cloud-based AI engines, or the rating of results. However - though the wider effect of AI-related automation on white-collar employment is still, even in 2026, rather sketchy - initial pointers from the tech sector itself, after waves of redundancies, point to more pessimistic outcomes.

Climate costs

Whilst data centres might not have a huge impact on employment, they would potentially make a huge impact on the environment, resulting in potentially enormous demands for additional electricity, and for water - in a country already stressed by climate variations and an uneven concentration of population. One unintended consequence of the government's AI Opportunities Action Plan could therefore be increases in the costs of water delivery and electricity prices, further impacting upon disposable incomes - and especially affecting those in lower income categories. It will also affect existing, and already somewhat overextended, commitments to Net Zero targets.

As well as affecting energy, the use of AI exacerbates previous tendencies towards generating e-waste, with the average lifespan of the Graphical Processing Units (GPUs) on which it relies reckoned to be two to three years - after which they effectively become useless and extremely difficult to recycle in any meaningful sense. In fact, the data centre itself - the different racks, routers and other components - has a life cycle of approximately ten years, after which much of the infrastructure requires replacement. The demand for GPUs hugely accelerates an already damaging supply chain, which demands raw materials from around the world, and generates large quantities of toxic and radioactive residues.

The 'silver bullet' intended to meet the growth in energy demands is the widespread rollout of nuclear Small Modular Reactors (SMRs), a technology which is still in the prototype stage. As of 2026 there are no working commercial SMRs in the world, and a number of studies have concluded that the cost efficiencies of 'cookie cutter' SMR construction have been overstated. The economic model underpinning the development of SMRs requires a huge number of the same type to be commissioned at the same time if they are to be economically viable, and replacing the 400 existing large nuclear reactors around the world would demand several

thousand SMR sites. Despite these questions, as well as reservations surrounding proliferation, waste and decommissioning, the construction of these atomic power stations is regarded as integral to the UK government's AI strategy.[12]

Financial and operational precarity

The environmental instability of AI and its extremely resource-heavy supply chain are mirrored by the economic precarity of the companies on which the government is relying. OpenAI, the provider of ChatGPT, has yet to show it is capable of generating large amounts of income, and it is even more uncertain that consumers and businesses would be happy to pay additional money for Google or Microsoft's AI services. Ultimately, the utility of something like Microsoft's Copilot is unlikely to be fundamental to businesses, or to add value in the same way that, for example, word processors originally did. This gives a big clue as to what is going on. In keeping with some previously observed characteristics of neoliberal economics, these hugely powerful companies are effectively turning to government - especially the US government, and now the UK government - to preserve what is an otherwise flimsy business model. Governments are huge, usually stable entities, and responsible for critical infrastructure worth billions of dollars.

As critics such as Ed Zitron and Cory Doctorow have noted, AI is a financial bubble: stock in chip manufacturers such as NVDI, along with anything AI-related, has become massively inflated, to an extent that belies the actual business case for such resource-intensive models.[13] It is beyond the scope of this article to predict how this apparent unprofitability will be managed over the next decade, but it is certainly true that the government's plans do not include AI partner insolvency as a potential occurrence.

It should also be noted that the finance and ownership model of technology companies has increasingly based itself on private equity: private equity has become the effective owner of multiple competitors in the same sphere. This is something that is especially relevant to companies which are themselves dependent on the big cloud infrastructures; but it also raises the question of contagion - which could arise if a collapse in, for example, AI-related tech companies then threatened the funding of other companies owned by the same private equity vehicle. If AI is not going to yield a quick financial return in the majority of cases, as most evidence currently suggests, the impacts on the wider tech sector could be magnified.

There is also the question of operational precarity. The computing capabilities of AI - upon which many essential services are likely to depend - are concentrated in huge data centres, representing large-scale potential points of failure - either physically, via natural or human-related disaster or attack, or through software bugs and cyberattacks. Given that, generally speaking, all IT projects have technical issues, delays and budget overspends, it would seem ridiculous to assume that the inaccuracies, hallucinations and biases of all currently existing AI outputs will be effectively resolved. One has to question whether this faith is a result of belief in the power of divination itself - part of a 'religion' of technological optimism.

Building motorways, again

We can advance a couple of explanations as to why this is happening, which can be partially deduced from Mandelson's address. But we should not overestimate politicians' abilities, or over-represent their opinions in the development of these policies. They are largely the result of a serious lobbying effort: as we have seen from the political complexion of the different lobbyists, the consensus on AI runs through different parties. In particular the continuities between the Sunak and Starmer administrations are notable. The language has been slightly altered, but much of the content remains. The basis for the current approach seems to be the Science and Technology Framework, drawn up by the civil service in March 2023.

The emphasis on AI reflects a degree of desperation, or at least an absence of alternative ideas. It is not clear, given the nature of the UK economy, that high levels of growth are achievable, or that it is possible to reduce welfare spending without increasing levels of hardship and poverty. This means that the scope for manoeuvre is limited in the short-term. Arguing that a combination of welfare cuts and an aggressive AI rollout guarantees more hardship without guaranteeing economic growth invites the response of: 'what, then?'. It is also true that the additional investment coming from the United States might add value in the near-term, whilst benefits from developing native capacity would only accrue, at best, in the medium-term.

The reasoning underpinning the idea that AI can provide the possibility of some kind of economic resurgence is probably based on two precedents from the 1990s. The first can be seen in the Major government's response to recession and unpopularity in the early 1990s, in the aftermath of sterling devaluation: it drew up plans for a huge road-building scheme. Whilst public opinion eventually turned

almost completely against the programme, largely thanks to a concerted roads protest movement, one suspects that people inside both the Sunak and Starmer governments perceive data centres as somehow being similar to roads. They are a visible result of intervention, they provide construction employment, they are somehow about expansion, and are thoroughfares which somehow drive traffic.

The same search for short-term, growth-at-all-costs economy boosters has driven the Starmer administration to approve a third runway at Heathrow Airport; and it was the same logic that drove the Blair and Brown administrations to embrace out-of-town shopping centres in the late 1990s and early 2000s (which came, ultimately, at the expense of smaller, more diverse retail options in the centres of towns, currently a major political issue). In reflecting on such unintended consequences, one has to be sceptical that SMRs will actually emerge as a cheap nuclear energy solution. But one can easily foresee a situation in which US companies increasingly determine the shape of the UK's energy grid.[14]

The second precedent is the mixed investment in what was then the 'information superhighway', beginning in the 1970s, which led initially to the development of the internet as we know it. Of course, this was a complex, gradual process, taking place all through the 1970s, 1980s and 1990s. The New Labour governments were convinced that we were seeing the beginning of a knowledge economy. But in spite of all the transformations these developments have brought about, they have proved incapable of delivering long-term economic security for the UK. In fact, the bedrock of this knowledge economy, the higher education sector, is itself now extremely imperilled.

One can see why hopes might be currently pinned on a similar 'AI revolution' - but only if one is prepared to suspend all disbelief.

Conclusion

In assessing the prospects for the Starmer government's flagship AI plans, one is struck by the vulnerabilities they reveal. The government is drawn from a managerial class which tends not to pay too much attention to either long-term strategy or the devils in the detail; and this means they find themselves signing up to vague plans, the implications of which across different areas of government have not been properly investigated. They draw inaccurate parallels to fixes that apparently worked thirty years ago, and in doing so hitch themselves to variants of Big Tech

which are unproven and precarious. They then fall into the role of unconscious host for expensive systems that in many cases simply do not work without all kinds of manual and technical intervention.

Perhaps the most telling vulnerability in these plans relates to geo-political realities. As we can see from Mandelson's revealing speech, the appeal is to the better nature of the Trump administration, disregarding its inherently reactionary, nationalistic nature in its appeal to an idea of a 'West' (something which now has little meaning outside its use as a 'white Christian' racist construct). Yet it is becoming increasingly clear that the US is simply a predator, even in relation to its 'allies', and is using its powerful multinational companies to farm Britain's various resources, initiatives and even cultural existence. This means that any 'British Plan' for economic growth under existing circumstances, where the US effectively dominates almost all information infrastructure, can only serve to bind the country further to the companies that are draining the country, leading the economy further into a labyrinth of dependency.[15] This dependency is upon a technology which its central actors proclaim as not being 'woke' - in other words, one that is likely to be ideologically weighted against a loaded definition of 'woke' and therefore towards far-right ideologies.[16]

At the time of writing (March 2026), the billions of promised US investment in the 'US-UK Technology Prosperity Deal' is still under discussion, following the 'failure' of the UK government to reduce the 'trade barriers' for US companies in the UK - namely, to reduce food standards and various regulatory conditions to the extent the Trump administration seeks. In the recently reopened discussions, the focus has shifted to co-operation on atomic energy. That the delay to the deal revolves around the interests of the USA's big agricultural interests, of course, gives a further geopolitical and realpolitik context for the redistributive power of AI in practice. It is a technology which stands, at least in the US context, to make more powerful, more wealthy, actors still more wealthy. For the UK to be subordinate is not enough for the tech moguls now; the UK is required to surrender much of its sovereignty, as represented, in this case, by data, but *in addition*, the ability to set standards over commodities.

To develop an alternative strategy to the 'realist' logic of the government (which could be regarded as defeatist, or even treacherous) requires a radical change of approach. Such an approach would need to be based on informed long-term planning, a broad consensus within society, and an acceptance of no immediate

fiscal rewards - at a time when the economy lacks a growth dynamic and when public finances are under extreme stress. One wonders, in fact, whether the choice is perceived by - for example - the Treasury, as binary: between, on the one hand, accepting a deepening subjugation to US business interests at the expense of future economic autonomy, as well as increasing levels of technical infringement on basic human rights and the displacement of intellectual labour; and, on the other, attempting to build a degree of independent technical capacity from limited resources while entering a phase of uncertainty, in a 'sub-Roman Britain' (a deeply unglamorous option for those with a post-imperial mindset), within which there will be a generalised slow but perceptible worsening of conditions and decreasing large-scale possibilities - a United Kingdom that is 'left out'.

However, there are risks and potential consequences in posing this as a binary question. Microsoft's use of AI to develop code within its Windows operating system has been described as 'software rot'; similarly, the more the UK relies upon AI to deliver basic functions, the more it will experience 'core infrastructure rot', as accountability diminishes. If US investment, assuming some of it actually occurs, leads to a further abdication of control over what's left of the UK's strained public service infrastructure, and if the delicate financial basis for AI also collapses, the UK risks not just a slow economic deterioration, but something much more dramatic and precipitous.

Even its most devout former advocates now find it almost impossible to defend the Starmer administration, which itself looks increasingly precarious. It is true that, post-Brexit, the UK itself is in a fragile strategic position, seemingly ill-equipped by its political system and economy to handle multiple, overlapping crises. On top of this, it also seems apparent that the Starmer administration is hard-wired into a technical and interpretative mindset which rejects any interrogation of technocratic goals, or even a democratic process regarding how, why and to what extent it is desirable to implement potentially deadly computing processes across society. This wiring is both physical, through the UK's very real conjunctions with the US military, and - as we have seen - psychological.

Carl Rowlands has worked as a trainer and documentation specialist in the cybersecurity sector for some years. He is also a journalist and lecturer, focusing on politics, society and technology.

Notes

1. UK Government, 'AI Opportunities Action Plan', 13 January 2025: https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.

2. P. Krugman, 'Talking with Paul Kedrodsky', 6 December 2025: https://paulkrugman.substack.com/p/talking-with-paul-kedrosky.

3. I. Overton and M. Colbert, *Byline Times*, 'Peter Mandelson's Fixing of Keir Starmer's Visit to Spytech Firm Palintir Raises Serious Questions', 25 April 2025: https://bylinetimes.com/2025/04/25/peter-mandelsons-fixing-of-keir-starmers-visit-to-spytech-firm-palantir-raises-serious-questions/.

4. H. Shearing, 'Teachers can use AI to save time on marking, new guidance says', BBC News, 10 June 2025: https://www.bbc.co.uk/news/articles/c1kvyj7dkp0o.

5. J. Meyer, 'Foreign Companies: Does the US Government Now Have Access To Your Overseas Data?', 5 December 2018: https://www.jdsupra.com/legalnews/foreign-companies-does-the-u-s-22621/.

6. Y. Abraham, 'Leaked documents expose deep ties between Israeli army and Microsoft', 23 January 2025: https://www.972mag.com/microsoft-azure-openai-israeli-army-cloud/

7. A. Walsh, 'Questions for Ireland as investigation reveals Microsoft data centre use by Israeli military', 7 August 2025: https://www.thejournal.ie/microsoft-data-storage-israel-military-6784411-Aug2025/.

8. K. Crawford, *Atlas of AI*, Yale University Press, 2022.

9. M. Borak, 'London police deploy facial recognition during Palestine and Israel protests', Biometric update.com, 15 January 2024: https://www.biometricupdate.com/202401/london-police-deploy-facial-recognition-during-palestine-and-israel-protests.

10. C. Smith, 'DWP's upcoming bank account checks to target claimants on eight benefits', 18 September 2025: https://www.manchestereveningnews.co.uk/news/uk-news/dwps-upcoming-bank-account-checks-32503035.

11. P. Mandelson, 'Britain and America in the age of Trump - And Beyond', Ditchley Foundation, , 6 September 2025: https://www.ditchley.com/sites/default/files/61st%20Annual%20Lecture%20delivered%20by%20Lord%20Mandelson%20%20-%20transcript.pdf.

12. S. Thomas, 'Prospects for SMRs in the UK', 1 February 2024: https://100percentrenewableuk.org/wp-content/uploads/Thomas-SMR-Webinar-3.pdf.

13. E. Zitron, 'The Hater's Guide To The AI Bubble', 21 July 2025; https://www.wheresyoured.at/the-haters-gui/.

14. E.A. Farhat and W. Mathis, 'Tech Pledges From Trump's UK Visit Will Test Strained Power Grid', 20 September 2025: https://www.msn.com/en-us/money/markets/tech-pledges-from-trump-s-uk-visit-will-test-strained-power-grid/ ar-AA1MWNLS.

15. T. Bradshaw and D. Shepperd, 'US tech giants pledge billions for UK AI

infrastructure during Trump visit', 16 September 2025: https://www.ft.com/content/5b760b6c-3e8b-4853-8aaa-754343b98858.

16. E. Maiberg, 'Anthropic Promises Trump Admin Its AI Is Not Woke', 21 October 2025: https://www.404media.co/anthropic-promises-trump-admin-its-ai-is-not-woke/.