About the IET

The IET is a trusted adviser of independent, impartial, evidence-based engineering and technology expertise. We are a registered charity and one of the world's leading professional societies for the engineering and technology community with over 158,000 members worldwide in 148 countries. Our strength is in working collaboratively with government, industry and academia to engineer solutions for our greatest societal challenges. We believe that professional guidance, especially in highly technological areas, is critical to good policy making. For further details on the evidence submitted, please contact policy@theiet.org.

Executive Summary

The IET welcomes the opportunity to provide input on the AI Growth Lab, and we support the overall proposal. The production of the AI Growth Lab is an ideal opportunity to demonstrate best practice of AI development, considering net zero and ensuring that the infrastructure to run the lab is sustainable.

AI is influencing how businesses, industries and technologies operate now and in the future. Digital technologies can help bolster resilience across a range of sectors by providing a greater insight into scenario modelling and provide insightful data analytics to help inform decision making. As with any other technology, it can also pose its own risks, and it is important to be aware of vulnerabilities that digital technologies may introduce to a system of systems. AI solutions must be aligned, encompassed, and influenced by software engineering, software architecture, management, governance, technology operations, and service delivery/service management. By showing how AI can be used successfully, it will encourage other sectors to adopt AI. The AI Growth Lab has the potential to achieve this level of robustness before deploying to a real-life scenario.

For the AI Growth Lab to enable testing of genuine, end-to-end applications, it must function like a collaborative research lab. The AI Growth Lab should actively facilitate multi-party interactions, bringing the innovators, their required supply chain partners (e.g., a hospital, a logistics firm, and an AI startup), and the regulator into the same process and testing environment.

It is important for the AI Growth Lab to have access to representative data of key sectors such as healthcare, the built environment, energy, transport, etc. In future, it should aim to be compatible with existing digital twins within these sectors/systems. Having appropriate digital twin models enables effective testing of AI behaviours, for example, discovering if the AI model drives undesirable outcomes as well as desirable ones.

Recommendations

- **A centralised model with oversight** – The IET recommends that the AI Growth Lab should be designed as a centralised model, in the Cabinet Office's portfolio, so there is more coordination across government. Oversight should consist of public scrutiny and reporting, with Parliamentary scrutiny when modifying or disapplying regulation within the AI Growth Lab.
- **High quality and real time data** – The IET recommends designing the AI Growth Lab so that it can access representative data.

- **Regulation** – AI safety and the assessment of risk must go beyond the physical, to look at financial, societal, reputational and risks to mental health, amongst other harms.
- **Data governance** – There should be firm rules on which data can and cannot be used to train AI systems.
- **Ensuring competency** – Competency is pivotal to the successful uptake of AI. Cross-sector collaboration, and capable workforces are key to wide-spread usage of AI. If government shows to be successfully using AI, then other sectors will be encouraged to suitably adopt AI.

Modelling and oversight of the AI Growth Lab

The IET recommends that the AI Growth Lab should be designed as a centralised model, so there is more coordination across government. The Cabinet Office would be best placed to have this centralisation due to its work cutting across government. Oversight should consist of public scrutiny and reporting, with Parliamentary scrutiny when modifying or disapplying regulation within the AI Growth Lab.

There is currently a number of different AI sandboxes being developed across government that are siloed, in addition to the recently announced AI incubator. The proposal, however, does not cover how the AI Growth Lab would interact with these other sandboxes. A centralised model will allow for coordination of work, so that it is not duplicated, and it will also allow for overlap between sectors. For example, an NHS sandbox would need to consider transport legislation. It is important then to consider how these sandboxes interact with each other and what the long-term unintended consequences might be.

There should be life cycle assessments on the pilots, considering the enablers for long-term success, what will happen to the projects after the pilot period and what success will look like. A centralised model would allow greater oversight. Having a coordinated approach across sandboxes (and the associated pilots) at different phases of their implementation, will allow for life cycle assessments and to consider what should happen after the pilot period.

High quality and real time data

The IET recommends designing the AI Growth Lab so that it can access representative data. The AI Growth Lab is a good way to establish high quality and use of data. There is concern over the lack of understanding and information surrounding AI, the data used to train the models and ethical considerations. This has given rise to a discussion around the ownership of data needed to train these algorithms, as well as the impact of bias and fundamental data quality in the information they produce. It is imperative that regulators co-operate globally, so there is consistency, clarity and sharing of best practice.

The availability of data is a concern as the proposal does not consider how the AI Growth Lab would interact with open sources of data. There is a need to be able to share date from sandboxes that already exist and are working well (anonymised/pseudonymised). A mechanism on how to unblock this needs to be considered. The National Data Library is well placed to achieve this as its aim is to have open data.

Data Governance

There needs to be greater clarity on the legal use of data in the research and development of AI systems. Many developers do not understand their legal obligations in terms of permission to use data, or taking decisions based on that data. This lack of clarity hinders research and innovation. While there is much openly accessible data that could be used to ‘train’ and test algorithms in their early stages, it is unclear the extent to which such data is legally allowed to be used, especially when the data is not accompanied by explicit terms and conditions.

Greater legal clarity on these issues would give researchers and companies more confidence to develop AI systems, supporting the UK's pro-innovation approach.

The IET recommends that there should be firm rules on which data can and cannot be used to train AI systems.

A major concern surrounding the implementation of AI is the lack of information and trust in it. IET research has found that 29% of engineering employers surveyed had concerns on the lack of information around AI (Source: The IET, Artificial Intelligence behind 3 times more daily tasks than we think). When developing AI, it is imperative that there are strong data foundations, competency and full transparency, as this will enhance the national public trust and uptake of AI.

Regulation

It is critical that the appropriate legal and regulatory structures are in place to allow AI's safe development and use but also do not stifle innovation. It needs greater transparency around the training and operation of AI systems. Although the onus is on developers to prove that the product is fit for purpose and has no unintended consequences, further guidelines and standards around how this should be reported would support a regulatory environment that is pro-innovation and provides safeguards against harm. As AI models have very large compute and storage demands, there will be a need for new data centres.

The IET recommends that AI safety and the assessment of risk must go beyond the physical, to look at financial, societal, reputational and risks to mental health, amongst other harms.

There is concern over the lack of broader understanding and information surrounding AI, the data used to train the models and ethical considerations. AI has created a discussion around the ownership of data needed to train these algorithms, as well as the impact of bias and fundamental data quality in the information they produce. As AI spans every sector, it is imperative that regulation is coordinated, so there is consistency and clarity.

It is necessary to ensure AI is used safely and to help prevent incidents from occurring, and it is fundamental to maintaining public trust, which underpins the economic and social benefits AI can bring. A sandbox provides an opportunity to achieve this, however, there needs to be clear governance frameworks set and this sandbox is a great opportunity to spearhead this.

This highlights a clear gap in the ecosystem that needs to be addressed with regards to supporting best practice for development and adoption, ensuring that knowledge is not lost along the pipeline. A clear example of this is in products such as AI, where product development knowledge is not passed onto product users and maintainers in a robust way. The Responsible Handover of AI report highlights this example specifically as a break in the pipeline for information/knowledge handover (Source: Sense about Science, The Responsible Handover of AI).

Whilst a regulatory sandbox provides some value it is extremely limited without the regulators issuing recognised good practice. For example, when considering the safety assurance of autonomous systems and AI based information systems, these types of systems are subject to the Health and Safety at Work Act 1974 and other relevant health and safety regulations. Therefore, there is a duty to reduce the risk of harm to as low as reasonably practicable. This is normally achieved by demonstrating the application of the good practice that is recognised by the HSE or by following an approved code of practice.

Currently the HSE has not recognised good practice, nor has it approved a code of practice or issued operational guidance. So, the HSE would not be able to define what good practice looks like or provide a benchmark against which developers can demonstrate that they have

reduces the risk of harm to as low as reasonably practicable for autonomous systems or AI based information systems. This can hinder growth within some AI applications and the HSE should be more active in defining what is required by developers to demonstrate that they have discharged their health and safety duties. The true value of a regulatory sandbox will not be realised in industry until the HSE have stated what is good practice.

Ensuring competency

To enhance competency, there needs to be further investment into bridging the skills gap. The current lack of skills in AI is not only a safety concern but is hindering productivity and the ability to deliver contracts. As among employers that expect AI to be important for them, 50% say they don't have the necessary skills, 32% of employers reported an AI skills gap at technician level and 46% say that senior management do not understand AI (Source: 2022 Skills for a digital future survey – summary report, IET). Therefore, proper training and skills means safer AI. This will be aided by having tools and techniques available to AI developers that can help them prove they are safe and fit for purpose to regulators.

Extending the AI Growth Lab to other technologies

Quantum

There are opportunities within quantum technologies where a growth lab would be beneficial. The UK already has a strong background in systems engineering and is a global leader in developing quantum technologies (ranked 2nd for number of quantum companies). In the future, the UK can stand out globally in the future by harnessing its strength in bringing both quantum and systems engineering together. There needs to be greater quantum literacy within industry – particularly at C-suite level.

There is an opportunity for the UK to take a leadership position in global regulation and standards for quantum technology. Government and regulators should expand on the 2024 Regulation of Quantum Technology Applications report by the Regulatory Horizons Council to focus on the systems element of regulation in this field. A quantum growth lab has the potential to move from research to deployment of quantum technologies as well as finding early routes to market. The opportunity to test before deployment will be particularly useful due to its non-local effects.

Digital Twins

Digital Twins are widely used across different sectors to test decisions before we make them and understand how different actions might have real life implications. However, they are often not used to their full potential, particularly as they are excellent models for systems that have complex requirements such as the energy sector. The availability and quality of data is important as poor data quality or insufficient data can lead to inaccurate simulations and therefore misguided insights and suboptimal decisions. If the AI Growth Lab is to integrate with a digital twin in future it will need to be able to work with real representative data. It will also need to be cross-checked over time, for example, verifying that it is still a faithful representation of the real-world energy system.