Royal Academy of Engineering

Vision for an AI-enabled clean energy system

Response to the Department for Energy Security and Net Zero's call for evidence

About the Royal Academy of Engineering

The Royal Academy of Engineering creates and leads a community of outstanding experts and innovators to engineer better lives. As a charity and a Fellowship, we deliver public benefit from excellence in engineering and technology and convene leading businesspeople, entrepreneurs, innovators and academics from every part of the profession. As a National Academy, we provide leadership for engineering and technology, and independent, expert advice to policymakers in the UK and beyond.

Our work is enabled by funding from the Department for Business, Innovation, Science, and Trade, corporate and university partners, charitable trusts and foundations, and individual donors.

Evidence base

This response draws on evidence from the Academy's Fellowship and Awardees with relevant expertise. The response also draws on existing Academy and National Engineering Policy Centre work where directly relevant. This includes Smart Systems for Clean Power, and the Academy's response to the AI Growth Lab call for evidence.

Summary

AI can add near-term value by improving defined engineering and operational decisions. Greater autonomy is not a goal in itself. Applications should be judged by measurable effects on system and consumer costs, connection times, curtailment, interruptions, asset life, emissions and resilience.

The response makes the following key points:

- The UK has strong energy system and AI expertise, a complex and increasingly decentralised electricity system and innovative companies and institutions for testing and deployment. Its competitive advantage will depend on combining these strengths and turning them into safe, secure and reliable deployment at scale in the UK energy system.
- Data access, data quality and technical integration must b. Common definitions, machine-readable context, open interfaces and proportionate portability are needed to avoid rebuilding bespoke integrations for each application.
- Effective digitalisation of the electricity system will be critical to support data access and quality and to underpin the application of AI. This requires Government to urgently establish a permanent digitalisation coordination function, to establish strong strategic oversight of digitalisation and to create dedicated oversight for each aspect of digitalisation.
- Government should focus on shared enabling infrastructure where the system benefits are clear but, at present, no single organisation has sufficient incentive or authority to provide the necessary capability. Priorities include system data and interoperability, forecasting and observability, storage and demand-side flexibility and shared testing and validation environments.
- DESNZ and Ofgem, working with Innovate UK and other relevant public funders, should give successful applications a defined route from pilot to business-as-usual. Funding, procurement and price-control arrangements should reflect different regulatory timetables, align costs with system benefits and cover not just only model development

but integration, cyber security and safety assurance, monitoring, maintenance, replacement and replication. Cases in which implementation costs and system benefits fall on different organisations should be addressed.
- Deployment should build on existing energy, IT and operational technology expertise, bringing it together with AI, cyber security, safety assurance and human factors. Applications should define which decisions remain with people, which actions may be automated and when authority must return to a human.
- Faster or machine-speed coordination is not an objective in itself. Market arrangements should use timely data and distributed flexibility where this creates value, while distinguishing market decisions from direct physical control and preserving clear authority, operator visibility and system security.
- Effective assurance must address the system architecture and the full operational lifecycle, not only the accuracy of an individual model.
- Autonomy should be introduced progressively and remain reversible. Decision support and shadow operation should precede automation, with stronger evidence and independently enforced limits required as authority increases.

Response

1. Where and how can AI add the most value in the energy system?

AI is most likely to add near-term value to the energy system where it improves the quality, speed or consistency of a defined engineering or operational decision, producing a measurable improvement for the system or consumers.

Forecasting, planning and connections

AI can improve accuracy and consistency of forecasts of renewable generation, electricity demand, available flexibility and network conditions in a way that can support both short-term system planning and longer-term decisions about capacity, reinforcement and investment. AI can help forecasters use larger and more varied datasets and compare a wider range of scenarios.

AI can also support the analysis of how much capacity is available on electricity networks and the connection studies for new generation, storage and demand. It can do so by supporting the analysis of multiple network capacity and connection scenarios in parallel, reducing the time needed to compare likely outcomes. This could help network operators identify where spare capacity may exist, and where additional grid capacity or physical network upgrades are needed, avoid unnecessary investment and manage connection applications quicker. Where these outputs inform safety-critical or high-value decisions, professional engineering verification should remain in place.

System operation, flexibility and resilience

AI can help operators understand and manage an increasingly complex system and make better-informed operational decisions as renewable generation, storage, flexible demand and distributed assets grow. Applications include electricity system balancing, constraint and curtailment management, outage planning, congestion management and coordination between NESO, transmission and distribution system operators.

AI could also support resilience by improving forecasts of demand or weather-dependent generation, and by helping operators assess available responses. This includes coordinating

storage and demand management when supply is constrained. In the near term, these tools should improve situational awareness and present options to operators, rather than replacing accountable human decision-making in safety-critical control.

Asset management and predictive maintenance

Combining inspection, condition, maintenance, failure and environmental data can help identify emerging faults, prioritise repairs, select the right intervention at the right time and extend the useful life of infrastructure. The Academy's Smart Systems for Clean Power work cited an estimate that predictive-maintenance approaches could produce savings of around £5.5 billion by 2030, although these benefits depend on sufficient digitalisation of infrastructure and usable data.

Consumer and distributed flexibility

AI can help coordinate electric vehicle charging, heat pumps, batteries and other flexible household or commercial assets. Automation could enable demand increase when electricity is abundant, potentially reducing wind curtailment costs, and reduce when the system is strained, limiting the need for dispatchable power which is often expensive. This would lower system costs and potentially reduce consumer bills.

This depends on good visibility of the assets on the system and reliable communication between individual assets, service providers, networks and control rooms. Consumers must retain choice, and the use of household and smart-meter data must be supported by clear consumer consent, privacy and security arrangements.

Data centres also present a flexibility opportunity, ramping up or down energy use in response to signals from the grid. As the energy consumption associated with AI fine-tuning and inference is highly dynamic – and heavily influenced by user behaviour, AI enables data centres to shift where and when workloads are processed – and thereby reduce strain on generation and power infrastructure during peak demand (when ramping down) and minimise curtailment (when ramping up). This will likely require incentives for data centre operators to participate in such initiatives and alterations to customer service level agreements. (See NEPC work on sustainable AI for more).

Engineering productivity and system understanding

AI can support engineers by cleaning and reconciling disparate datasets, accelerating design and connection studies, and assisting with scenario analysis. Whole-system modelling can help planners explore interactions between electricity, heat, transport and industrial energy use and assess trade-offs between cost, resilience, energy security and decarbonisation.

These tools should augment engineering judgement. Their assumptions, limitations and uncertainty must remain visible, particularly where an output influences an infrastructure investment or operational decision.

AI use in a digitalised energy system

Digitalisation, as described by NESO, is the 'act of using digital technologies to change the way a system operates, with the goal of increasing efficiency, supporting decision making and unlocking new value and opportunities'. While this definition and others focus on technologies and their uptake, it is important to stress that the wider process of digitalisation extends beyond technology to involve people and processes. This includes their decision-making and behaviours within the energy system – across organisations, communities and individual households.

Digitalisation is not an end-goal in itself. It is a critical enabler for the clean power transition in which the system will become increasingly decentralised and complex. Without rapid progress on digitalisation – involving a clear and coordinated plan, backed by clear leadership – the decarbonisation of the GB electricity system will become significantly more difficult and costly.

The French electricity system illustrates the importance of the underlying digital infrastructure. France uses AI-enabled local forecasting at approximately 2,300 substations, alongside remote-controlled switches used to monitor and respond to local generation and demand. France's high smart-meter coverage and communications infrastructure also support remote fault detection and the management of connected assets.

France has a more consolidated distribution-network structure, a different approach to smart-meter rollout and a different generation mix but the example provided shows that, for AI to provide operational value in the UK, operators need timely visibility of supply and demand, reliable data flows between assets, network operators and authorised service providers, and the authority and infrastructure to act on the resulting information.

2. What are the UK's main strengths in applying AI in the energy system and where may the UK have a competitive advantage over other countries?

The UK's potential advantage lies in combining energy system expertise and AI research with a rapidly changing electricity system and institutions capable of testing and deploying new technologies.

The UK combines strong energy-system and AI expertise

The UK has established capabilities in power-system modelling, optimisation and control, sensing, meteorology, engineering and nuclear technology. It also has internationally recognised AI research, universities, startups, energy companies and specialist engineering and consultancy expertise. The opportunity lies in bringing these capabilities together around operational energy system problems.

A further strength is in hybrid approaches that combine machine learning with established physical models. These approaches can retain known engineering constraints and provide an engineering basis for interpreting model outputs. The UK also has capability in digital twins and synthetic environments, which can support the testing of future assets and network configurations where operational data are limited or does not yet exist. Within a defined validation envelope, these environments can be used to compare scenarios, including rare faults and unusual operating conditions, before deployment.

The scale of planned investment in networks and low-carbon assets creates an opportunity to build data standards, connectivity and interoperability into new infrastructure from the outset. The transformation also offers a demanding environment in which to develop and test applications. Increasing renewable generation, distributed energy resources, electrification and the growing need for flexibility create practical forecasting, planning, coordination and resilience challenges. Tools developed successfully in this environment could be relevant to other countries managing similar transitions.

The UK has institutions and facilities that can support deployment

The UK has several potential routes through which research can be tested and applied, including NESO, network operators, Ofgem's Regulatory Laboratory and Technical Sandbox, Ofgem's Strategic Innovation Fund, the Energy Systems Catapult, universities and established energy innovation programmes.

The Strategic Innovation Fund, operated by Ofgem with Innovate UK, has already supported work using AI and machine learning for electricity network asset management. Ofgem describes the Fund as a route for trialling or launching new products, services and business models.

Evidence received also highlighted the Power Networks Demonstration Centre at the University of Strathclyde. Its representative high- and low-voltage networks, supervisory control systems, real-time simulation and hardware-in-the-loop facilities allow technologies to be tested under realistic conditions, including abnormal events and rare faults, before they are introduced into live infrastructure.

There is already practical activity on which to build. Evidence submitted to the Academy described UK Power Networks prioritising around 150 potential AI use cases and moving selected applications into live operation; National Grid Electricity Distribution developing predictive maintenance and network visibility tools alongside an MLOps platform; and NESO's Volta programme applying AI to forecasting, optimisation and decision support. These examples illustrate an emerging deployment base, although they do not by themselves establish a national competitive advantage.

The strongest opportunities are in whole-system applications and assurance

The UK could develop a comparative advantage in AI applications for forecasting, system planning, network operation, flexibility and the coordination of distributed assets. The strongest opportunity is not a single application but the connections across organisational boundaries. The UK's experience of operating a complex and increasingly decentralised electricity system provides a useful setting in which to develop whole-system tools that account for interactions between markets, networks, generation, storage and consumers.

Another potential advantage lies in assurance and standards for AI used in critical infrastructure. The UK combines engineering expertise, technical standards capability and facilities for realistic testing. If these capabilities are brought together, the UK could develop methods for validating and assuring AI-enabled systems that are useful internationally.

The advantage will not arise automatically

The main risk is fragmentation and lack of robust digitalisation and data sharing infrastructure that would underpin the deployment of AI applications. Isolated pilots, inconsistent data, proprietary interfaces and slow procurement can prevent strong research and demonstrations from becoming repeatable products and services. Evidence also suggested that relevant sectors and institutions do not yet work together consistently enough.

Competitive advantage should be judged by whether the UK can deploy useful tools in its own energy system, demonstrate measurable benefits and translate the resulting products, services and assurance expertise into wider markets. Where an established solution can meet the need safely and on acceptable terms, procurement need not favour a domestic supplier. However, for high-consequence energy functions, buyers should assess dependence on foreign-owned or concentrated technology across models, cloud, hardware, software and maintenance, and require appropriate continuity, substitution and switching arrangements. Commercial advantage should remain secondary to reducing costs, supporting decarbonisation and maintaining energy security.

3. Do the six challenges reflect the main barriers to safe and effective deployment and which is most urgent to address?

The six themes in the call broadly reflect the main barriers. However, they are not independent or sequential. A project can have access to data but still fail because that data lacks context, cannot be integrated with operational systems or cannot be used within a clear assurance and accountability framework.

Data, governance and assurance are the immediate priorities

Contributors to the Academy's response differed on which challenge was most urgent. Some identified access to data, particularly in contexts where data is sensitive, fragmented, held in non-digital formats or unsuitable for the intended use. Others identified regulation and governance, or risk, security and trust, as the most urgent challenge because energy is safety-critical and organisations need confidence that AI-enabled decisions will be reliable, resilient and appropriately overseen.

The most immediate practical constraints are gaps in effective digitalisation and data-sharing infrastructure, while governance and assurance are equally necessary for higher-consequence uses. Data access is of little value if the data is inaccurate, incomplete or lacks context. Data must be representative, of known quality and provenance, and usable outside the system that created it. An application must also have a defined route into operational systems, with clear accountability, validation and liability. These are connected conditions for deployment, not competing priorities.

Capability, people and culture will be the most persistent challenge

Although data and governance require immediate action, capability, people and culture may be more difficult to resolve over the longer term. Successful deployment requires people who understand AI alongside power engineering, operations, safety, cyber security and human factors. Senior decision-makers also need sufficient AI literacy to judge when AI is useful, which form of AI is appropriate and when the risks outweigh the expected benefits.

This is not only a question of recruiting more AI specialists. Energy organisation needs the capacity to evaluate systems, supervise their use and challenge their outputs. AI developers also need a clear understanding of the physical, operational and regulatory constraints of the energy system. Organisational readiness and confidence among engineers, planners and operators will determine whether technically promising tools can be adopted safely.

The framework should make architecture, interoperability and actuation explicit

The framework would be clearer if it treated access to data and technical integration as a single architecture and interoperability problem. Scaling beyond one asset or site requires common definitions, machine-readable context, reusable information models and open interfaces, so that each deployment does not require a new bespoke integration.

The framework should also distinguish between applications that advise a human, influence prices or instructions sent to market participants, and directly control a physical asset. These forms of actuation carry different implications for integration, accountability and safety and should not be assessed as if they were equivalent.

Different forms and uses of AI should not be treated as presenting the same risks

The nature of each barrier depends on the technology, application and operating context. Supervised machine learning used to improve a demand forecast does not raise the same assurance questions as a system controlling critical infrastructure. The framework should avoid treating AI as a single category or assuming that frontier or generative AI is required for every use-case.

Deployment should favour techniques that are appropriate to the problem and can be tested and assured to the level required. In some cases, an established statistical or machine learning method may deliver most of the available benefit with fewer risks than a frontier system. The framework must also support a decision not to use AI where the expected benefits do not justify the additional risk.

4. For the most significant barriers identified what specific constraints are limiting deployment?

There are four connected constraints: operational data often lack the context and interoperability needed for reuse; integration with heterogeneous and ageing systems is expensive; assurance and accountability remain unclear; IT, operational technology and operations are typically organised in separate silos and organisations lack cross-disciplinary capability. These constraints reinforce one another and should not be addressed in isolation.

Data quality, access and interoperability

Relevant data is distributed across network operators, suppliers, aggregators, asset owners, manufacturers and consumers. This data uses different definitions, formats and update cycles, varies substantially in quality and granularity and may be commercially or operationally sensitive.

Distribution network operators do not always have sufficiently timely information about load profiles or the number and location of low-carbon assets such as electric vehicles and heat pumps. This limits their ability to assess network headroom, run the network efficiently and plan capacity.

Gaps in smart meter coverage and functionality, including the projected 7 million meters affected by the 2G and 3G switch-off, incomplete registration of low-carbon assets and the absence of sufficiently usable consent arrangements also constrain applications involving domestic demand and flexibility. At the end of June 2026, 75% of homes had a smart electricity meter, 71% had one operating in smart mode. A quarter still had a non-smart electricity meter. This reduces the number of households that can reliably participate in data-enabled flexibility services.

Access charges and controls around half-hourly meter data may be justified by privacy and security considerations, but can also restrict legitimate innovation.

Access alone does not make data usable by AI. A system needs machine-readable context, including what a data point represents, the physical asset and operating state to which it relates, its units and provenance and its relationship to other data. This is where the NESO Data Sharing Infrastructure will be crucial. Common identifiers, definitions and reusable information models are needed if that context is to carry across assets and organisations. This is especially important where historical records are incomplete or not representative: in the nuclear sector, for example, some records exist only on paper and data from older reactors may not be relevant to new designs.

Improving access doesn't require aggregating all energy system data in a single central platform. Federated approaches can allow authorised applications to use data held by different organisations, while clear governance, privacy and cyber security arrangements define who can access which data and for what purpose. This reduces unnecessary concentration, although shared interfaces and dependencies must themselves be treated as part of the system's security architecture too.

Technical integration, architecture and modes of actuation

Energy infrastructure includes long-lived operational technology that may be proprietary and subject to strict reliability and cyber security requirements. A pilot on one asset can rely on bespoke integration, but replication across many assets, sites and suppliers makes the underlying architecture and the cost of each further integration decisive. Open interfaces, portability and the ability to replace components without rebuilding the wider system are therefore resilience requirements as well as commercial considerations.

A further constraint is the lack of clarity about how an AI output is converted into action. Advice to an operator, a price or instruction sent to a market participant, and direct control of an asset have different technical and safety implications. Decision rights, interfaces and operating limits therefore need to be defined before integration and assurance can be completed.

Regulatory clarity, assurance and cyber security

Appropriate regulation is essential. The constraint is uncertainty about the evidence required to demonstrate safety, how changes to adaptive models will be treated and how liability will be allocated when an AI system influences a decision.

There are also too few representative environments in which applications can be tested before entering live operation. Normal operating data may not reveal how a system will respond to rare faults, unusual conditions, malicious inputs or failures elsewhere in the system. Confidence can be further weakened by the risk of attacks on data, models, communications or control systems. If operators are not confident that control rooms and data flows are adequately protected, they will be reluctant to deploy AI even where the potential benefits are clear.

Digital twins of existing assets and synthetic simulation environments for future configurations could help teams compare many scenarios before an application enters live operation. They should be used to test normal operation, rare faults, unusual conditions, malicious inputs and failures elsewhere in the system, especially where live data are sparse or the future configuration has not yet been built. Cyber security must be designed in from the outset because attacks on data, models, communications or control systems could undermine confidence in otherwise valuable applications.

The Academy's response to the Department for Business, Innovation, Science and Trade's AI Growth Lab call for evidence identified another risk where assurance processes designed for fixed systems can require extensive reassessment following relatively minor software changes. Energy sector change-control arrangements should distinguish between an update that remains within an approved operating domain and a material change to an application's function or risk.

Skills and organisational capability

Energy organisations already have substantial power system, IT and operational expertise. The constraint is that these capabilities are not always brought together with AI, cyber security, safety assurance and human factors in teams able to design, assure and operate an integrated service. Responsibility for data and deployment is also often divided across IT, operational technology and operational teams, with different suppliers and incentives. This makes it harder to design, assure and operate an integrated service.

Capability development should extend the existing skills base rather than assume organisations are starting from zero. AI developers need a clear understanding of the physical, operational and regulatory constraints of the energy system. Energy organisations need sufficient expertise to determine which technique is appropriate, whether AI is needed at all

and what evidence is required before deployment. Contributors also identified limited AI literacy among senior decision-makers and a lack of trusted independent guidance and knowledge exchange between the energy and AI communities.

Deployment teams should bring together power engineers, operators, AI specialists, cyber security and assurance professionals, human factors expertise and procurement staff from the outset. Training and independent guidance will also be needed so that senior decision-makers and staff supervising AI can judge its limits and challenge its outputs.

5. Where are current market procurement or regulatory incentives failing to support AI deployment and why are promising applications struggling to move from pilot to full scale?

A recurring problem is that current incentives support experimentation more readily than adoption. A pilot may demonstrate that a model can perform a task, but this does not establish who will integrate it into operational systems, maintain and assure it, pay for its continued use or remain accountable for its decisions.

Funding and procurement often stop at the pilot stage

Innovation funding commonly supports model development and proof-of-concept pilots, but not the less visible work needed for sustained deployment: data preparation, interfaces, validation, operational integration, staff training, monitoring, maintenance and supplier support. These requirements should be designed and costed from the outset. A demonstration on one asset or site may therefore succeed without showing that the approach can be replicated across different assets, systems or organisations. Programmes may compound this problem by emphasising frontier AI when a simpler, application-specific approach would be easier to deploy and assure.

DESNZ and Ofgem should align innovation funding, procurement expectations and price-control arrangements, working with Innovate UK and other relevant public funders, so that successful applications have a defined route into business-as-usual. Transmission and distribution price controls operate on different timetables, so each deployment plan should identify when regulated organisations can recover costs and what interim route is available if a project becomes ready between determinations. Funding and procurement should test replication and interoperability as well as technical performance and should provide a mechanism for recovering reasonable ongoing costs where measurable benefits have been demonstrated.

The organisation paying for an application may not receive all the benefits

Advanced decision-support tools can require one organisation to meet the costs of preparing data, integrating systems and completing assurance work, while the benefits are distributed across the wider energy system. This weakens the business case for individual organisations to invest, even where an application would produce an overall system benefit.

DESNZ and Ofgem should develop consistent methods for identifying where costs fall and where benefits accrue, then use price controls, procurement or market incentives to align them where necessary. Incentives should reward successful deployment and measurable outcomes, rather than the number of pilots undertaken.

Pilots are not always connected to a clear adoption decision

A technically successful pilot may still fail to progress if the intended user, evidence requirements and route to regulatory approval or procurement have not been agreed.

Uncertainty about what evidence regulators will accept, how updates to an AI system will be treated and who is responsible for the resulting decisions makes it difficult for organisations to plan for deployment.

The Academy's response to the AI Growth Lab call for evidence recommended that pilots begin with a roadmap beyond the sandbox. Applied to the energy system, this means agreeing the intended user and accountable owner, the outcomes against which success will be measured, the evidence required and the approval, procurement or price-control decision that will follow if the pilot succeeds. Evidence should be reusable within a clearly defined operating context, although a successful pilot should not be treated as proof that an application is safe in every setting.

Procurement can create long-term dependence on individual suppliers

Contracts for AI and digital services may not give energy organisations adequate access to operational data, model history, interfaces or the information needed to transfer a service to another provider. Proprietary technology stacks can therefore create a resilience risk as well as a commercial dependency. For high-consequence energy systems, procurement should require a continuity case showing how the service will continue safely if a model, cloud service, connectivity route, hardware component or supplier becomes unavailable, or if its terms change. The assessment should cover dependency concentration, time to recovery and the ability to replace a component or supplier without rebuilding the wider architecture.

Contracts should support export of data and necessary model artefacts, access to maintenance and skills, and workable exit arrangements. The Academy's response to the Smart Data multi-sector call for evidence recommended clear requirements for data access, interoperability, provenance, reuse, retention and handover, including machine-readable export and application programming interfaces where appropriate. These principles should be applied proportionately to energy-sector AI procurement, alongside requirements for data and model portability.

Market incentives should reward measurable system value

Question 7 considers market design in detail. For deployment purposes, the main incentive issue is that more accurate forecasting or optimisation has limited value if market and operational arrangements cannot use the information in time. Any change should therefore be tied to demonstrated reductions in cost, congestion or curtailment and should preserve system security, transparency and accountability.

6. Which areas of the energy system should be priorities for government action to unlock near term AI deployment?

Government should focus on enabling infrastructure and capabilities that offer a clear system benefit but that no single organisation has sufficient incentive or authority to provide. In those cases, DESNZ and Ofgem should identify who will pay, who will benefit and which policy, regulatory or market mechanism will align the two. They should coordinate and enable deployment rather than prescribe particular AI technologies or suppliers.

Accelerate digitalisation and build data sharing infrastructure

At present, there is no designated body who has overall responsibility and ownership of architectural governance and coordination of a nationwide digitalisation strategy. Our ongoing sector engagement suggests that roles and responsibilities are confused and disconnected between multiple actors; there is confusion over who has ownership of digitalisation as a policy area and there is no coherent oversight framework. The result of this

lack of ownership is fragmented efforts, duplicated activity, technical divergence across the system, slow progress towards agile and flexible energy systems and confusion for industry participants about where to direct conversations and efforts. This lack of clarity and direction is impeding progress and risks the delivery of systems that fail to integrate.

Strong strategic oversight will be required to ensure standardisation across markets and to maximise efficiency.

We welcome the greater clarity set out in the recent DESNZ Energy Digitalisation Framework but remain concerned that it states the process of setting up the permanent coordination function will take 2-3 years, thus potentially taking until 2029. This slow pace is concerning, as is the lack of committed timelines beyond the end of the calendar year. It is important that timelines are more ambitious, set out clearly, and not permitted to drift.

On data, the first priority is timely, reliable and usable information about the system. Government should accelerate delivery against the Energy Digitalisation Framework, with measurable targets for data quality and access. Specific priorities should include accelerated delivery of the NESO Data Sharing Infrastructure, effective system-wide implementation of Market-wide Half-Hourly Settlement, creating a central asset register for low-carbon technologies such as electric vehicles and heat pumps, and faster smart-meter rollout, including proactive replacement of meters affected by the 2G and 3G switch-off.

Data access should be accompanied by a common data architecture: shared identifiers and definitions, machine-readable context and provenance, and open interfaces through which authorised users and applications can discover and use information. Procurement should require proportionate data and model portability. Clear governance, privacy and cyber security arrangements, including federated approaches where appropriate, are needed to support trusted sharing without relying unnecessarily on a single central platform. These features should be folded into the emerging Data Sharing Infrastructure being developed by NESO.

Improve forecasting and visibility of assets and system conditions

Government should support forecasting and visibility of assets and system conditions where better information can improve planning and operation across organisational boundaries. Priorities include renewable generation and demand forecasting, network headroom and connection assessment, constraint and curtailment management, and improvements in asset management decisions.

Support should be tied to measurable outcomes such as improved forecast accuracy, shorter connection times, lower constraint or operating costs, fewer interruptions and reduced curtailment. These applications do not necessarily require frontier AI; mature statistical, machine-learning or engineering methods should be used where they are better suited to the task.

Enable distributed storage and demand-side flexibility

Storage, electric vehicles, heat pumps and flexible loads can provide system value, but coordination depends on effective digitalisation, reliable communications, common interfaces and market arrangements that allow flexibility to be used. Government should urgently establish a permanent digitalisation function, as envisaged in the recently published Energy Digitalisation Framework, prioritise the enabling infrastructure and rules needed to aggregate and coordinate these assets while preserving consumer choice, privacy and cyber security.

Adoption may scale unevenly, so government-supported testing should examine different scenarios, including mixed participation by households and businesses, different degrees of reliance on aggregators or supplier-provided agents, and many assets responding to the same

signal. This would help identify when local network visibility, safeguards or market changes are needed.

Provide shared testing and validation environments

Where possible and appropriate, government should connect Ofgem's Technical Sandbox and Regulatory Laboratory with engineering facilities such as the Power Networks Demonstration Centre. This would allow regulatory requirements and technical performance to be examined together and would give organisations access to realistic conditions that they cannot reproduce individually.

Shared environments should test normal operation, degraded conditions, rare faults and adversarial scenarios and should generate evidence that can be reused within a defined operating context. Where possible, programmes should also share validated results, unsuccessful trials and known failure modes so that organisations do not repeatedly fund the same demonstration work.

Build capability to deploy and supervise AI

These priorities depend on people who can combine power engineering and operations with AI, safety, cyber security and human factors. Government should work with energy organisations and professional bodies to build the capability needed to deploy and supervise AI safely. This should build on existing power engineering, operational and IT expertise and include:

- role-specific guidance and training for senior decision-makers, operators, engineers, procurement teams and AI specialists on how to select, assure, supervise and challenge AI in energy settings
- cross-disciplinary deployment teams and knowledge exchange that bring AI developers together with power engineers, operators, cyber security and safety assurance specialists to work on operational problems.

7. Can current market structures support coordination at machine speed or do existing design features constrain system performance and what changes could shape a market that supports this?

Current market structures can support some AI-assisted forecasting, optimisation and trading, but the evidence does not support the assumption that the whole energy system should be coordinated at machine speed. It is important to distinguish market coordination from automatic physical control. Parts of the electricity system already respond at machine speed through established protection and control mechanisms. These mechanisms are not necessarily improved by replacing transparent controllers with AI. For safety-critical functions, well-defined interfaces and independently assured operating limits remain more important than increasing the speed of optimisation.

In relation to market processes, contributors noted that many existing arrangements were designed for a slower moving and less decentralised energy system. The evidence identified settlement periods, fixed product definitions, gate closure and limited locational signals at distribution level as features that may restrict the value of faster forecasting, optimisation and local flexibility. At the same time, AI is already being used in short-term energy trading, demonstrating that existing structures can accommodate some faster and more adaptive participation. Future market design could enable decisions to be made closer to real time where this produces value and where system operators retain sufficient visibility and control. It

should not be assumed that every market or planning process needs to operate at the same speed.

Changes should focus on making market arrangements more capable of using timely information and distributed flexibility, including machine-readable rules and interfaces, more granular signals where justified, standardised flexibility products and clearer coordination between NESO and distribution system operators. Market design must also distinguish between AI providing decision support, influencing behaviour through a market signal and being authorised to control an asset. In each case, the responsible actor, the permitted action and the applicable limits should be clear. Market objectives should also reflect resilience, security of supply and decarbonisation alongside cost.

Faster coordination can create new risks if many assets respond to the same information in the same way. For example, large numbers of batteries or electric vehicle chargers responding simultaneously to one price signal could create a new local constraint or increase volatility. Measures such as ramp limits, diversity of response and local constraint signals may be needed to prevent individually rational actions from producing a poor system-wide outcome. Proposed changes should therefore be tested through simulation and controlled pilots before being introduced into live markets. The objective should be better coordination within clear market and physical boundaries, rather than machine-speed operation for its own sake.

8. What does effective assurance look like at system level and how can accountability be maintained where outcomes reflect cumulative system dynamics rather than discrete decisions?

Effective assurance should focus on the architecture within which models and automated agents operate, not only on the performance of each model. Components can behave as designed and still create an unsafe system-level outcome through their interaction. Assurance should therefore cover data, operational technology and communications, dependencies on suppliers or shared infrastructure, and the people and organisations responsible for operation. The evidence required should increase with the consequences of failure and the authority delegated to the system.

At a minimum, system-level assurance should establish six matters: what the system is authorised to do; the physical and operational limits it cannot exceed; the identity and provenance of the systems, data and instructions involved; whether an audit trail allows decisions and actions to be reconstructed after an incident, whether automated authority can be reduced or withdrawn and human control safely restored when conditions deteriorate, and; how assurance will be maintained as models, data and interactions change.

Explainability and accountability should be treated separately. It may not always be possible or useful to explain every internal inference, but it must be possible to establish which system or agent produced a decision, how that decision was translated into action, what information was used, what authority and constraints applied, and which organisation remains accountable. A deterministic method should remain the default for a high-consequence function where it can perform the task reliably and transparently.

Assurance must continue after deployment. Testing should cover normal, degraded, emergency and adversarial conditions, including correlated responses across multiple components, using simulation and, where proportionate, hardware-in-the-loop testing and physical testbeds. Where AI agents interact, testing should also examine behaviour under different goals, information, model capabilities and combinations of participants. Operators should monitor performance, model drift, changes in agent behaviour, control major updates,

investigate incidents and near misses, and retain tested recovery, rollback and non-AI fallback arrangements.

9. What are the most material risks arising from more autonomous and integrated AI deployment, how can these best be managed and how do we stage autonomy safely?

The most material risks are those that can propagate beyond the component where they begin. For example, large numbers of electric vehicle chargers, batteries or flexible loads using the same model or responding to the same price or network signal could create a new local peak or constraint. Common dependence on an AI foundation model, cloud platform, communications service or data source could cause several services to fail together. If an AI agent is authorised to control charging or storage, a cyber compromise could move from a digital service into coordinated physical action.

These include cascading failure, correlated responses by systems that share similar models, data or objectives, and common-mode failure where several services depend on the same foundation model, cloud provider, software platform or data source. A cyber compromise could then move quickly from a digital service into physical action across generation, storage or demand.

Harm can also arise when no component has failed. Several systems may act rationally against narrow local objectives and together create positive feedback, new constraints or greater volatility elsewhere in the network.

A further risk is that operators lose visibility of how automated decisions are made and translated into physical action, alongside the practical experience needed to intervene. They could remain formally accountable without having the information or time needed to understand what the system is doing or why. This is particularly serious where recommendations lead directly to physical action. For high-consequence functions, transparent and verified control may remain more appropriate than AI-based real-time control unless the latter can be shown to be at least as reliable.

Risk should be managed through the system architecture. Each AI-enabled component should have a defined identity, interface and authority and should remain within physical and operational limits enforced independently of the AI. Cyber security should be treated as part of system safety, and the monitoring, recovery and fallback arrangements described in Question 8 should be in place before authority is delegated.

Autonomy should be introduced through clear stages: monitoring, observing and reporting; recommending an action; acting with human approval; acting within bounded authority; coordinating autonomously within a defined domain; and, only where the evidence supports it, wider system autonomy subject to deterministic constraints. Each stage should require evidence from the preceding stage and a fresh decision about whether additional authority produces a proportionate system benefit.

Shadow operation is a useful step between recommendation and control. The AI system can run alongside the live process without controlling assets, allowing its performance, interactions and failure modes to be assessed under normal, unusual and adversarial conditions. Bounded automation should follow only where the permitted actions and operating limits are clear, independently enforced and supported by continuous monitoring, rollback and a tested non-AI fallback.

Authority should also be reversible. A resilient architecture must be able to reduce or withdraw authority when operating conditions deteriorate, a cyber event occurs or data quality falls. Coordination between multiple autonomous systems should face a substantially

higher evidential threshold, and there should be no presumption that every application will progress to full autonomy.

Human oversight must be specified in operational terms. Human oversight is only meaningful if operators can see the relevant system state, understand what the AI is permitted to do, have enough time and competence to assess its behaviour, and can intervene. Where action occurs too quickly for human review, safety must depend on predefined limits, automated checks and escalation points. Accountability should remain with identifiable organisations responsible for design, integration, deployment and operation; references to emergent behaviour should not obscure those responsibilities or prevent investigation after an incident.

10. Which issues highlighted in the transformative future are most relevant to your organisation and what steps are you already taking to realise the opportunities?

The Academy is not an energy system operator. The issues most relevant to our role are those requiring coordination across the engineering system. We are already addressing those through the Clean Power Advisory, including its work on smart systems and digitalisation; our Practical AI programme, which shares lessons from real-world adoption and convenes problem owners, engineers, innovators and policymakers to address deployment barriers; our wider work on responsible AI and data, and; wider support for research and innovation.

These activities are informed by the Academy's Fellowship, awardees, Enterprise Hub members and partners within the National Engineering Policy Centre, providing expertise across energy, AI, cyber security, safety, assurance and innovation. The Academy can support government and the sector by convening the right technical and operational expertise around specific energy system challenges, testing proposals against engineering practice, identifying what prevents promising applications from scaling, and helping to develop practical approaches for policy delivery. As data interoperability, legacy integration, procurement, assurance and supplier dependence recur across critical infrastructure, the Academy can also bring together regulators, operators, procurers and technical experts from different sectors to share practice, validated results, failures and near misses.