RAND EUROPE

RESEARCH BRIEFING

AI ecosystems in the Netherlands and the UK

Compute for AI

Shyam Krishna, Lara Marie Otto, Adam Urwick,
Scott Warnier, Ishita Puri, Manuel Hernandez Gonzalez

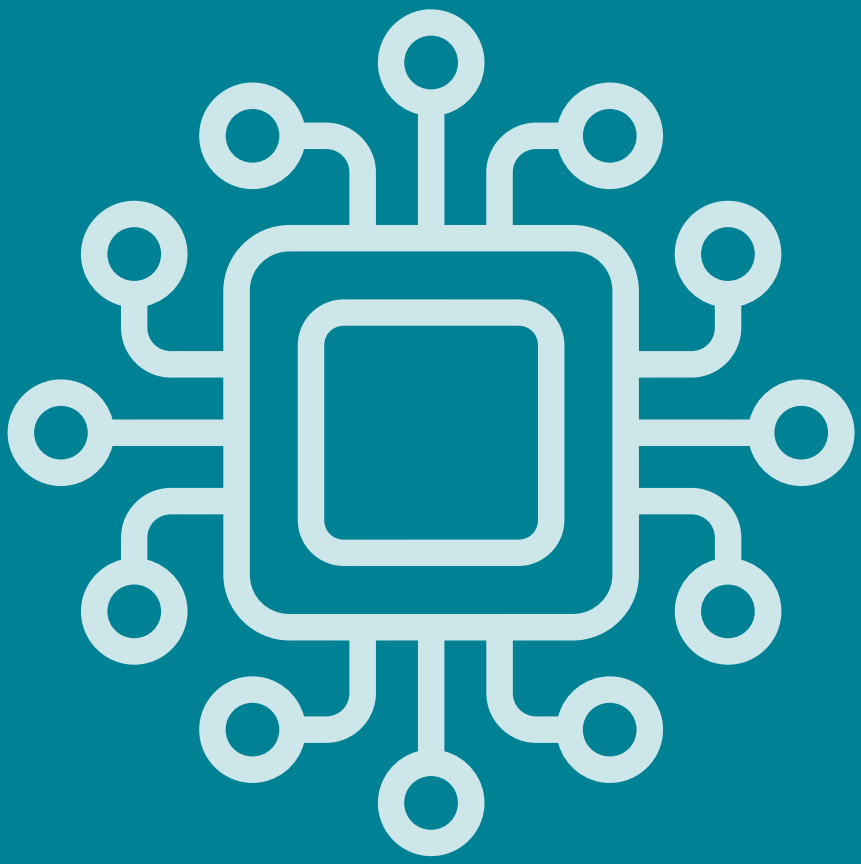

This research is sponsored by the UK Foreign, Commonwealth & Development Office (FCDO)

1. Purpose and context

This briefing highlights key findings on opportunities for UK–Netherlands collaboration with compute for artificial intelligence (AI) for the life sciences and health technologies, drawing on RAND Europe's assessment for the Foreign, Commonwealth & Development Office (FCDO) of AI ecosystems in both countries. The report was commissioned by FCDO in 2025 to provide an evidence base on the strengths and challenges in both countries' national AI ecosystems, as well as to explore areas for complementary action and synergy in light of the 2025 UK–Netherlands Innovation Partnership.

2. Overview

The UK and the Netherlands are both strong AI ecosystems with complementary strengths. The UK offers scale, global connectivity and research depth, with a tech sector worth around £1 trillion and AI revenues growing by 68% in 2024.[1] Public institutions such as the Department for Science, Innovation and Technology (DSIT), UK Research and Innovation (UKRI) and the Alan Turing Institute anchor the ecosystem, while start-ups can benefit from greater investment depth and a more advanced hyperscale and data centre capacity compared to the Netherlands. The Netherlands has a smaller but highly networked ecosystem supported by strong public-private coordination and strategic importance in AI value chains through semiconductor and enabling technology strengths.

Both countries share similar strategic priorities around innovation-led growth, responsible regulation, compute and data resilience, and skills development, while also facing similar constraints in talent and infrastructure. Existing cooperation through Horizon Europe, the G7 and the OECD AI Principles, alongside the 2025 UK–Netherlands Innovation Partnership, mean there is already a coherent foundation for bilateral cooperation between the two countries.

3. Sectoral overview and strategic positioning

This briefing is part of a wider analysis organised around three priority sectors for UK–Netherlands AI ecosystems: (1) **compute for AI**, (2) **life sciences and health technologies** and (3) **agrifood and agritech**. These sectors were selected because they capture different but complementary forms of bilateral opportunity and can demonstrate a strong industrial foundation and policy prioritisation in both countries. Compute functions as a foundational enabler of wider AI capability and resilience; life sciences and health technologies represent an area of shared strategic strength, with strong research, innovation and policy alignment in both countries; and agrifood and agritech represents a domain of complementary advantage, combining Dutch applied leadership with UK strengths in R&D and digital innovation.

---

1 Department for Science, Innovation and Technology & UK Research and Innovation (2025). 'UK Compute Roadmap'. Department for Science, Innovation and Technology (2025). 'Artificial Intelligence sector study 2024'.



Compute for AI

Foundational enabler

- Underpins AI capability across sectors
- Combines UK compute scale with Dutch connectivity and sustainability strengths
- Relevant for sovereignty, interoperability, and infrastructure resilience



Life sciences & health technologies

Shared strategic strength

- Strong in both countries across research, innovation, and policy priorities
- Existing collaboration base in medical research and translational health AI
- Offers scope for deeper coordination, joint validation, and clinical adoption



Agrifood & agritech

Complementary strength

- Dutch leadership in applied agrifood innovation
- UK strengths in upstream R&D, sensing and digital tools
- Offers scope to combine different capabilities in joint pilots and testbeds

4. Compute for AI overview

Compute capacity is a strategic enabler of AI leadership, economic resilience and national security. Both the UK and the Netherlands treat compute as critical infrastructure underpinning productivity, innovation and sovereign AI capability. Policy choices on compute now shape who can train, deploy and govern advanced AI systems in the next decade.

The UK's approach prioritises national scale and speed. This involves direct public investment in AI-dedicated supercomputers, accelerated data centre planning and strong partnerships with global hyperscalers. The Netherlands pursues strategic relevance through integration, embedding national capability within EU-level compute, cloud and research infrastructures.

The distinction is not ambition but governance. UK policy emphasises domestic control and rapid capacity expansion; Dutch policy emphasises interoperability, coordination and shared European assets.

In the UK, policy has unlocked large-scale public and private investment. Government has committed multi-billion-pound funding to expand AI compute twentyfold by 2030, established central coordinating bodies (including a sovereign AI unit) and attracted major hyperscaler and data centre investments. This has reinforced the UK's position as a leading European hub for AI-optimised compute infrastructure while stimulating innovation in hardware design, cooling and modular data centre systems. However, rapid expansion has also exposed fragmentation, oversubscription of public compute, grid bottlenecks and skills shortages.

The Netherlands has focused on coordinated provision rather than scale. Public funding through AINed, SURF and the AI Coalition for the Netherlands (AIC4NL) supports shared access to compute, strong integration of high-performance computing (HPC) and cloud, and participation in EuroHPC.

Dutch strengths lie in world-class digital connectivity, energy-efficient infrastructure and system integration. Structural constraints – limited grid capacity, land-use restrictions, funding gaps and weaker growth-stage finance – cap near-term expansion and increase reliance on EU-level assets. **Both countries face acute shortages in AI, compute and energy-system skills.**

Both countries are investing to support AI model training and deployment but follow different paths. The UK is building frontier-scale national systems alongside commercial cloud capacity, positioning itself as a location for training and running large models. The Netherlands relies more heavily on EU-scale supercomputing and federated access, accepting that frontier compute will be shared rather than nationally owned.

Despite these differences, both remain heavily dependent on US hyperscalers, which control the majority of available AI compute. Shared constraints are acute: grid congestion, long connection timelines, rising energy costs, planning restrictions and environmental pressures. Data centre location choices are increasingly constrained by energy availability rather than demand. These factors risk slowing AI deployment, increasing costs, and deepening dependence on foreign infrastructure.

The UK and the Netherlands face common strategic challenges including insufficient compute at scale, energy and grid constraints, skills shortages and reliance on non-European providers. These shared pressures create clear opportunities for bilateral cooperation. **Priority areas include** aligning standards for sustainable data centres; coordinating access to high-end compute; joint approaches to talent development; and exploring cross-border, energy-linked compute solutions, particularly in the North Sea. Cooperation would strengthen resilience, reduce duplication and reinforce Europe's strategic position in AI.

Key findings for compute for the AI sector



4.1. Policy insights

UK and Dutch compute policies are aligned in intent but diverge in execution. In the UK, the National AI Strategy, Modern Industrial Strategy, and Digital and Technologies Sector Plan highlight data and compute resources as foundational to AI leadership and critical enablers of productivity nationwide. The UK Compute Roadmap commits £2bn by 2030 to a 'modern public compute ecosystem' largely geared towards AI. Overall, the UK treats compute as a sovereign capability and growth lever, deploying significant public capital, central coordination and planning reform to accelerate capacity. The Netherlands embeds compute within EU frameworks, prioritising interoperability, shared infrastructure and strategic relevance over national ownership.

EU legislation such as the Cloud and AI Development Act (CADA) is bolstering the diffusion of AI and the adoption of cloud across the Netherlands through Europe-wide efforts. At the same time, national efforts such as the Strategic Action Plan for AI highlights the importance of compute and data resources alongside skills and research for developing the Dutch AI economy.

Both approaches face constraints. In the UK, fragmentation across public, academic and commercial compute, oversubscription of national resources, grid delays and skills shortages risk limiting returns on investment. In the Netherlands, constrained funding, land-use restrictions, grid saturation and limited access to frontierscale compute restrict scalability. Both ecosystems remain exposed to dependence on US hyperscalers.

These differences matter because they are complementary. UK scale and capital mobilisation combine well with Dutch strengths in coordination, connectivity and EU integration. A more deliberate alignment – on standards, access models, sustainability and talent – could deliver greater collective resilience than either approach alone.



4.2. Industry insights

Private investment is driving compute expansion in both countries, but at different scales. The UK has attracted exceptionally large hyperscaler and data centre commitments, reinforcing its role as a European AI compute hub and supporting a growing ecosystem of specialist suppliers. The Netherlands continues to attract investment from major players, enabled by connectivity and sustainability credentials, but expansion is constrained by zoning and grid limits. **Industry faces shared barriers** including high and volatile energy costs, grid access delays, land scarcity, skills shortages and lock-in to proprietary cloud ecosystems. These factors raise costs, slow deployment and limit the emergence of competitive European alternatives.

Shared barriers also mean that future competitiveness is increasingly dependent on the interlinkage of infrastructure development and factors such as permitting, energy prices, grid access and connectivity. Additionally, cross-border AI development is constrained not only by compute access but also by data sensitivity and sovereignty concerns, which might limit how firms collaborate.

Bilateral cooperation could improve industry outcomes by enabling pooled procurement, shared standards for green data centres, joint testing environments and coordinated engagement with hyperscalers. This would strengthen negotiating leverage, reduce duplication and support a more resilient European compute supply chain. Over the longer term it could also support more coordinated approaches to sustainable compute infrastructure.



What industry can do

- **Explore collaboration on privacy-preserving model development, where cross-border AI development is limited by data regulatory or sovereignty concerns.** In these cases, federated learning is a promising route for joint work without requiring raw data to be shared across jurisdictions.

- **Use bilateral engagement to better understand complementary infrastructure environments.** Firms can benefit from learning how UK scale and Dutch coordination create different advantages in deployment, connectivity and infrastructure use.



What further support could provide

- **Shared work on sustainable compute and energy integration** is a promising opportunity, but would require stronger policy and institutional support to move beyond dialogue into coordinated action. At the level of policy, a bilateral compute treaty could facilitate a coordinated approach to tackling challenges to sovereign compute such as energy access, data access, sustainability, and the spatial distribution of data centre capacity.
- **More structured coordination on standards, access models and resilience** could benefit industry, although these are primarily areas for future bilateral strategy and policy development. One option would be a bilateral compute strategy that aligns approaches to sovereign compute, infrastructure resilience and sustainable expansion.



4.3. Academic and research landscape

The UK research landscape is highly concentrated. Elite universities command significant compute resources and attract major investment, but oversubscription and fragmentation limit access for smaller institutions. This risks slowing experimentation and widening gaps between research leaders and the wider system.

The Netherlands operates a more integrated model through SURF, providing broad access to HPC and cloud resources and strong links to EU-level compute. However, national capacity is insufficient for frontierscale model training and chronic funding constraints limit expansion. Talent retention is also a challenge.

Across both countries, academia faces common barriers including limited access to cutting-edge GPUs, rising compute costs, skills shortages in research support roles and regulatory friction around data sharing. As mentioned above, barriers to cross-border collaboration include not only infrastructure but also the inability to work on sensitive datasets across institutional or country boundaries. This underlines the usefulness of privacy-preserving approaches for joint model development and training.

Closer UK–Netherlands cooperation could expand effective access to compute, support joint training and mobility schemes, and improve utilisation of European supercomputing assets. This would enhance research competitiveness, support innovation diffusion and strengthen the UK–Netherlands' contribution to the European AI research base.

What researchers can do now

- **Explore privacy-preserving approaches to joint model training and development where cross-border collaboration is constrained by data sensitivity.** Federated learning is one pragmatic route for collaboration without requiring data sharing.
- **Use existing European infrastructures and relationships where possible.** The UK's renewed connection to EuroHPC improves the basis for researchers to deepen collaboration with Dutch and wider European compute actors.

What further support could provide

- **Shared training, fellowships and researcher mobility schemes** are identified as a significant opportunity but would require policy and institutional backing to be implemented at scale. Now that the UK has rejoined the EuroHPC ecosystem, joint upskilling programmes could be leveraged to tackle skill gaps and promote the exchange of expertise across supercompute clusters. Similarly, joint participation in instruments such as Marie Skłodowska-Curie fellowships and Erasmus+ Knowledge Alliances could permit development of cross-border educational programmes in advanced computing.
- **Closer alignment around compute access and the use of shared infrastructure** would also require institutional coordination beyond what individual researchers can establish alone.

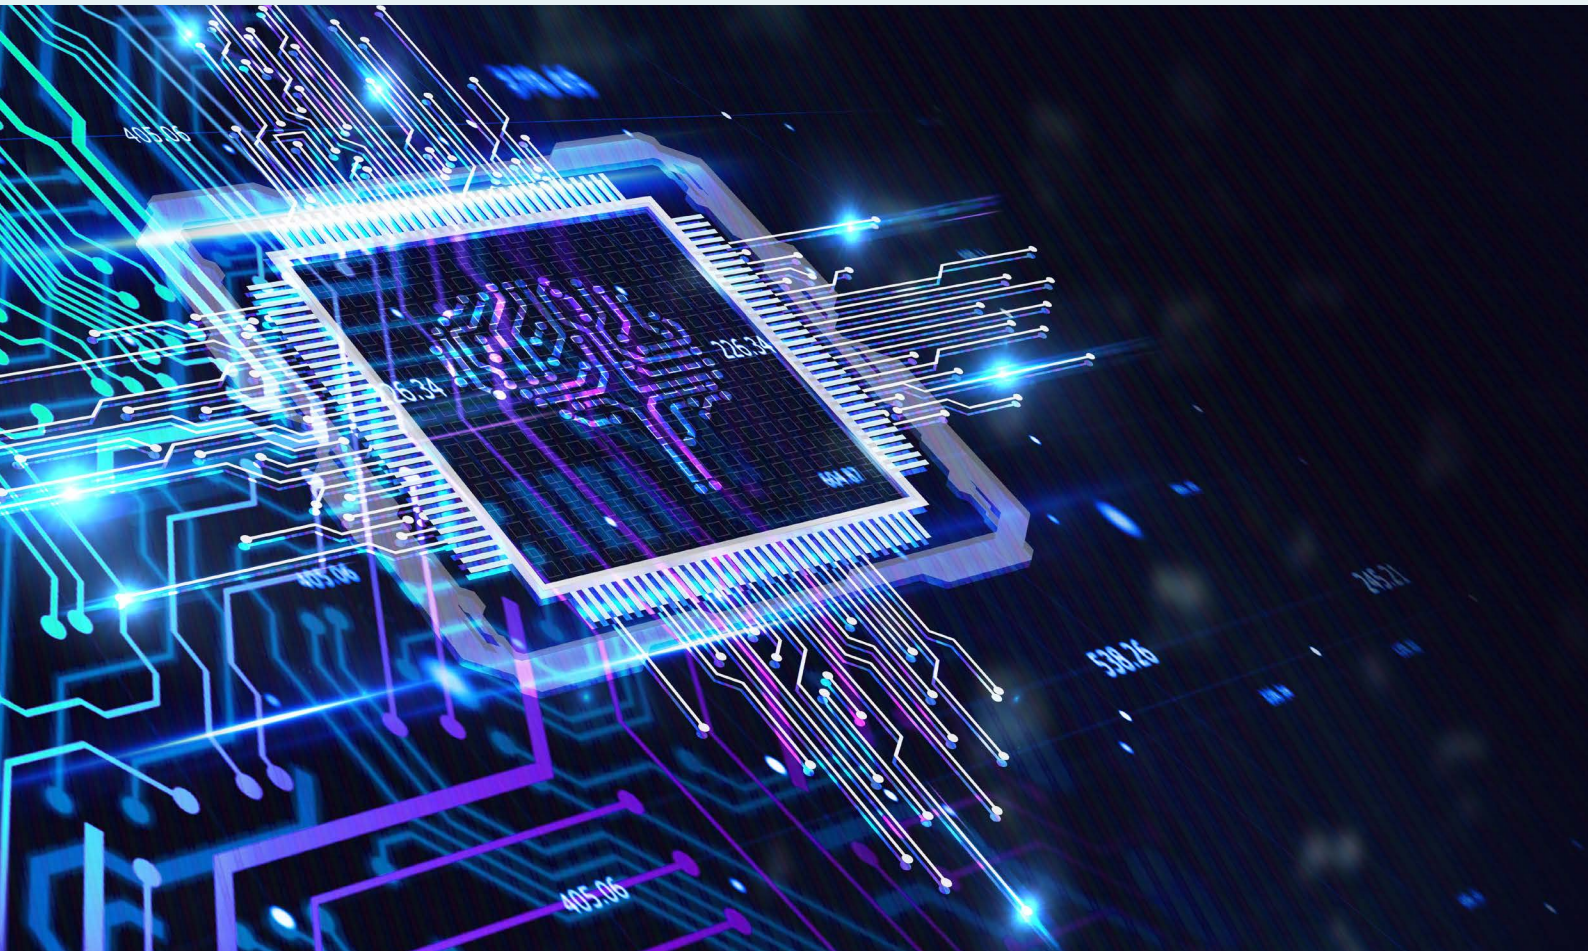

5. Opportunities for bilateral collaboration in compute for AI

The opportunities below are primarily areas where policy and institutional support could strengthen the conditions for collaboration between researchers, firms and innovation actors in both countries. Some, such as exchanges, networking and joint training, are also directly relevant to external stakeholders in the near term. Others, particularly those related to compute access, infrastructure coordination and standards, are more likely to depend on coordinated policy action and longer-term institutional alignment.

Table 1. Compute for AI opportunity matrix

	High impact	Low impact
High feasibility	**Federated learning** to enable secure, privacy-preserving collaboration on sensitive data, supporting joint AI model development without breaching sovereignty or GDPR constraints. **Bilateral compute strategy** to align national approaches to compute ownership, accountability and infrastructure resilience, creating a shared foundation for sovereign AI.	**Researcher exchange programmes** to build long-term talent pipelines through joint training, fellowships and HPC upskilling initiatives under EuroHPC and Erasmus+.
Low feasibility	**A North Sea compute treaty** offers high strategic potential but would be complex to implement due to legal, technical and jurisdictional challenges in maritime planning and compute-energy integration.	

This brief describes work done by RAND Europe in *Mapping AI ecosystems in the Netherlands and the United Kingdom* by Shyam Krishna, Lara Marie Otto, Adam Urwick, Scott Warnier, Ishita Puri, Manuel Hernandez Gonzalez, RR-A4654-1, 2026 (available at www.rand.org/t/RRA4654-1). To view this summary online, visit www.rand.org/t/RBA4654-1. RAND Europe is a not-for-profit research organisation that helps to improve policy and decision making through research and analysis. RAND Europe's publications do not necessarily reflect the opinions of its research clients and sponsors.
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights: This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorised posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.