

numero 1/2026

ISSN 3035-1839

DOI: 10.69099/RCBI-2026-1-05-M4L

costituzionalismo britannico e irlandese

Evoluzione costituzionale e transizioni politiche

From Competence to Capacity: Governing AI Territoriality in the UK

Andrea Gatti

Research fellow in Public Comparative Law
Università degli Studi di Padova

FROM COMPETENCE TO CAPACITY: GOVERNING AI TERRITORIALITY IN THE UK**

by ANDREA GATTI*

ABSTRACT (ITA): Il saggio propone una lettura costituzional-comparata dell'*AI Opportunities Action Plan* britannico (2025), sostenendo che la sua novità è soprattutto infrastrutturale, più che legislativa. La “portata nazionale” non viene perseguita tramite una riallocazione formale di competenze nell’assetto devoluto, ma attraverso capacità amministrativa, standard comuni, componenti riusabili e meccanismi centrali di diffusione dall’esperimento alla scala. La sequenza “Scan → Pilot → Scale” funziona come metodo di governo: lo *Scan* orienta priorità e missioni; il *Pilot* impone valutazione e documentazione per soluzioni trasferibili; lo *Scale* supera i silos tramite piattaforme condivise, interoperabilità e *procurement mission-oriented*. In questo modello lo Stato tende a operare come piattaforma e il *procurement* assume una funzione quasi-regolatoria. La standardizzazione e uno “stack” comune possono produrre *path dependency* e una centralizzazione tecnicamente mediata, riducendo la discrezionalità locale sul come adottare l’IA anche quando resta la scelta formale sul se adottarla. Poiché il baricentro si sposta su infrastrutture e filiere contrattuali, diventano centrali: visibilità e contestabilità delle scelte progettuali e tracciabilità *by design*, e chiara allocazione di responsabilità.

ABSTRACT (ENG): This article offers a constitutional-comparative reading of the UK *AI Opportunities Action Plan* (2025), arguing that its novelty is primarily infrastructural rather than legislative. “National level reach” is pursued less through reallocating competences within the UK’s devolved settlement than through capacity-building, common standards, reusable components, and centrally supported diffusion from pilots to scale. The Plan’s “Scan → Pilot → Scale” sequence operates as a governing method: Scan shapes upstream priorities and missions; Pilot produces transferable solutions through documentation and evaluation; Scale breaks organisational silos via shared platforms, interoperability, and mission-oriented procurement. The article shows how, in this model, the state increasingly acts as a platform and procurement functions as a quasi-regulatory tool. Standardisation and a shared tech stack can generate path dependency and a technically mediated centralisation, narrowing local discretion over how adoption occurs even where devolved bodies retain formal choice over whether to adopt. Because governance shifts into infrastructure and supply chains, the key constitutional requirements become visibility and contestability of design choices, traceability and auditability by design, and clear allocations of responsibility.

PAROLE CHIAVE: Intelligenza artificiale, Pluralismo amministrativo, Decentralizzazione, Sussidiarietà digitale.

KEYWORDS: AI, administrative Pluralism, Growth zones, Devolution, Digital subsidiarity.

SOMMARIO: 1. Introduction: a consistent constitutional way to deal with AI use; 2. The methodology: a centralised function with diffuse operational reach; 2.1 Legal implications of centralised AI-coordination; 3. The impact on devolution; 4. Growth zones; 5. Conclusion.

1. Introduction: a consistent constitutional way to deal with AI use.

Artificial intelligence has moved from experimental promise to an infrastructure question for the modern state: how to convert dispersed pilots into reliable, lawful, and measurable improvements in productivity and public services is an increasingly urgent challenge. In the United Kingdom, this challenge is sharpened by a territorially differentiated administration, uneven digital capacity, and a regulatory approach that has preferred principles and sectoral implementation over a single, comprehensive AI statute. Against that background, government faced a familiar problem in a new technological coordination. The 2025 AI Opportunities Action Plan responds to these pressures by setting out a programme intended to make AI adoption scalable and coherent across the UK, so that innovation does

* Research fellow in Public Comparative Law, Università degli Studi di Padova (Italy).

** Contributo sottoposto a *double blind review*.

not entrench a patchwork of local solutions and does not undermine territorial cohesion. The Action Plan frames the “territorial” operability of AI — one of the Plan’s three pillars, alongside “investing in the foundations of AI”, and “Industrial Sovereignty” — less through a formal allocation of powers than through capacity-building and standard-setting arrangements designed to generate effects at a pan-UK scale. Its organising core is the “Scan → Pilot → Scale” sequence, in which the scale phase is defined as the dismantling of administrative silos and the attainment of “national level reach”, supported by central resources (funding, tools, and scaling services) and by an explicit commitment to interoperability, code re-use, and — where appropriate — open sourcing².

This approach builds on a regulatory baseline that predates the Plan and largely remains intact, such as public-law constraints (in particular legality, rationality and fairness of the procedure), statutory duties in fields such as data protection³ and equality⁴, and the ordinary accountability routes (*i.e.* judicial review, audit, and parliamentary scrutiny).

More recently, the Government’s “pro-innovation” policy framework for AI regulation — centred on cross-sector principles to be operationalised by existing regulators rather than through a single horizontal AI statute — has reinforced a preference for sectoral application and coordinated guidance. In the UK, this principles-based, regulator-led model has become a particularly salient channel for steering the field.

At a sub-regulatory level, this orientation is reflected in the report “A pro-innovation approach to AI regulation: government response” (DSIT, 6 Feb 2024), in which the Department for Science, Innovation and Technology confirms the principles-based approach, relying on existing regulators and cross-government coordination rather than a single comprehensive AI Act; and the Algorithmic Transparency Recording Standard (ATRS) v2.1, which standardises public disclosure about algorithmic tools used in the public sector, together with the policy on mandatory scope and exemptions (ATRS Mandatory scope and exemptions policy, 17 December 2024)⁵. A further territorial strand concerns AI Growth Zones (discussed below, at § 4).

The Plan’s novelty is administrative rather than normative: it seeks to convert general expectations into repeatable pathways, shared artefacts, centrally supported mechanisms of diffusion, and building cross-government instruments that are formally non-statutory, while institutionalising a more centrally supported pathway to scaling. What remains constant is the UK’s familiar constitutional technique of achieving uniformity through coordination rather than by reopening formal competence boundaries. What changes is the scale and explicitness with which that technique is applied to AI adoption.

2. The methodology: a centralised function with diffuse operational reach.

² See UK GOVERNMENT (DEPARTMENT OF SCIENCE, INNOVATION AND TECHNOLOGY - DSIT), *AI Opportunities Action Plan* (January 2025) (from now on Action Plan), Section 2.2, Recommendations 41-42.

³ See Data Protection Act 2018 and UK GDPR (art. 22).

⁴ Equality Act 2010, s. 149 (Public Sector Equality Duty).

⁵ In partic. Sect.s 1 and 2.

The “Scan → Pilot → Scale” sequence should be read not as a generic project-management cycle, but as a very method of territorial governance because it reflects an approach of governing across multiple public bodies and administrative layers. In particular, the Scan phase somehow goes beyond cataloguing existing capabilities or identifying potential applications of AI. It rests on a central strategic function tasked with spotting systemic challenges — like recurring inefficiencies or capacity constraints — and determining how AI might help address them. This includes monitoring the technology landscape, collaborating with suppliers and start-ups, and assigning leadership to key public service missions. Importantly, it also plays a role in shaping institutional demand by setting the agenda for which technologies are prioritised and what kinds of problems AI should be applied to. This ability to influence upstream choices — even before individual departments make their own decisions — has constitutional significance. It determines what becomes feasible for local and sector-specific actors, effectively limiting the scope of action available to them. By setting the technological direction early, the centre not only guides procurement and development decisions (basically defining the options available to local administrations), but also narrows the range of tools others can reasonably choose from. This model is not entirely new, it builds on well-established mechanisms of digital governance in the UK, where central co-ordination has long been exercised through standards⁶, spending controls⁷, and the provision of shared capabilities⁸. What is different, now, is that these tools are being extended to the AI domain, with the intention of embedding AI across the public sector as a routine part of service delivery. In the Pilot phase, the focus shifts to quick experimentation, access to datasets and models, and more agile procurement routes⁹. The idea is that pilot projects should not remain isolated examples but should be designed from the outset with re-use and scalability in mind. To make this possible, the Plan requires pilots to be evaluated and documented in a way

⁶ GOVERNMENT DIGITAL SERVICE (UK GOVERNMENT), *Service Standard*, points 3, 10, 12 and 13 (joined-up delivery across channels; publication of performance data; openness of new source code; contribution to open standards and common components), read together with Government Digital Service - *The Technology Code of Practice*, points 3-4 and 8-9 (openness/open source; open standards; “share, re-use and collaborate”; “integrate and adapt technology”), and the *Technology Code of Practice* statement that it is “a cross-government agreed standard used for the Cabinet Office spend control process”.

⁷ GOVERNMENT DIGITAL SERVICE, *Digital and technology spend control, 2018*, in partic. the section “What the digital and technology spend control is for and how it works”, and the paragraph explaining that assurance functions assess how proposed activity complies with digital and technology standards used for assurance.

⁸ CABINET OFFICE (UK GOVERNMENT), *Government Functional Standard GovS 005: Digital 005*, sections 4.2.1 “Cross-government digital strategy” and 4.3 “Assurance”. Section 4.2.1 expressly frames a cross-government strategy as setting expectations for “synergy, efficiency, and interoperability (...) within and among organisations”, while section 4.3 requires a defined approach to assurance integrated with broader assurance frameworks and aligned to cross-government strategy and centrally provided requirements.

⁹ According to A. SANCHEZ-GRAELLS, *Responsibly Buying Artificial Intelligence: A ‘Regulatory Hallucination’*, in *Current Legal Problems*, Vol. 77, Issue 1, 2024, pp. 81-126, in particular p. 100: nowadays «Procurement is usually carried out by a unit or team within the public entity seeking to adopt the technology, which places it within a hierarchical organisation and thus with limited (functional) independence. As a result of the process of procurement centralisation, or in the context of collaborative procurement, it is possible for procurement to be carried out by a separate organisation. However, in that case, the organisational relationship will still be highly dependent on the interests of the adopting public entity». The Plan reinforces this coordination between center and local administration.

that makes their outcomes understandable and transferable across different institutions¹⁰. This leads to a broader proposal that AI solutions should be reusable and, where appropriate, open source. Two more aims follow. First, the Plan helps to translate the UK's "pro-innovation" regulatory philosophy into practical, testable actions. Second, it reinforces commitments to transparency and good information governance, particularly with the new requirement to publish details of algorithmic tools under the Algorithmic Transparency Recording Standard, which treats documentation, evaluation, and reproducibility not as optional good practice but as conditions for legitimate scaling¹¹.

Finally, the scale phase is explicitly characterised as the breaking of organisational boundaries and as a necessary condition for "meaningful" impact on productivity and citizen experience¹². Hence the envisaged "scaling service", the promotion of "mission-focused" national tenders for adoption "across decentralised systems", and, above all, the horizon of a scalable technological platform. The result is a territorial model in which unity is not (only) a juridical-formal fact, but an outcome of administrative coordination. The centre's action — embodied in the Plan by the role of the Department for Science, Innovation and Technology as the "digital centre of government" — operates by creating enabling conditions (mostly funding and procurement standards) that allow different layers of administration to converge on shared solutions without repeatedly rebuilding them from scratch¹³. Stated so, the prospect of a scalable AI tech stack capable of supporting "tens or hundreds of millions of citizen interactions across the UK" signals a move towards a common administrative infrastructure functioning as a vector of functional homogenisation in public action. Territorial unity is pursued through standards and platforms more than through constitutional norms in the strict sense, with a corresponding "technical constitutionalisation" of inter-administrative coordination¹⁴. This idea — often developed in recent scholarship under the headings of infrastructure governance or the normativity of architecture — provides a useful lens for the Plan. Infrastructure is not a neutral medium: it is an environment that structures possibilities and constraints for public action. If citizens interact with

¹⁰ The key move here is to treat the "pilot" phase not as a local exception, but as a governed precondition for re-use. The Plan's insistence on evaluability and publishability is meant to make pilots legible to other public bodies and, therefore, transferable. Sections 20(2)(b) of the Procurement Act 2023 (establishing the competitive flexible procedure) and 20(3) (requiring the procedure to be proportionate to the nature, complexity, and cost of the contract) help explain why the Plan can credibly favour multi-stage procurement routes.

¹¹ DSIT, *A pro-innovation approach to AI regulation: government response* (6 February 2024), Section 5 ("A regulatory framework to keep pace"); paragraph 10 of the Section lets out the five cross-sectoral principles: "Safety, security and robustness", "Appropriate transparency and explainability", "Fairness", "Accountability and governance" and "Contestability and redress". See also GOVERNMENT DIGITAL SERVICE, *Digital and technology spend control*, cit., which is similarly oriented towards eliminating duplication and strengthening assurance for "novel/complex/high-risk" projects, thereby enabling cross-institutional comparability of decisions and outcomes.

¹² Section 2.2. of the Plan stresses that scaling is essential for "meaningful" impact on productivity and citizen experience. This appears as the pinpoint that supports the idea behind the plan: from "project delivery" to territorial governance, because scale is framed as a deliberate break with organisational boundaries, rather than as mere growth in volume.

¹³ *Action Plan*, Section 2.2 (introductory paragraph) read with Recommendation 46 (the Digital Centre of Government to identify "quick wins" to enable scan-pilot-scale). These are the most direct pinpoints for attributing the coordinating role to DSIT.

¹⁴ L. LESSIG, *Code and Other Laws of Cyberspace*, Basic Books, 1999.

the state through automated interfaces, APIs, assistants, triage systems, and decision-support tools, then the technical configuration of those instruments contributes, in practice, to defining the scope of traditional principles such as equality, reasonableness, transparency, and reviewability. It is therefore unsurprising that the Plan couples the envisaged “stack” with requirements of interoperability and code modularity, and proposes that digital infrastructure be used to create opportunities for innovators by exposing data and functions via APIs and by standardising documentary interactions¹⁵.

A broader shift is discernible here. The state is not merely an “end-user” of AI; it increasingly becomes a platform that orders the ecosystem by setting integration standards and, indirectly, compliance standards. In this sense, procurement acquires a role that administrative law scholarship has increasingly described as regulatory: public contracts do not simply purchase services; they shape markets and techniques, establish quality requirements, specify evaluation methods, and include adaptive clauses — such as “contemplation clauses” — that compel continuous updating. Yet the same dynamic exposes a further difficulty. The drive towards standardisation, necessary to secure economies of scale and minimum uniformity, can generate path dependency and a form of technically mediated centralisation. When an infrastructural programme crystallises into a common tech stack, what is presented as facilitation may operate as a constraint, because early design choices tend to stabilise and make differentiation costly.

It is, at this point, that British territoriality, marked by devolved administrative pluralism and by uneven local capacities, encounters a point of tension. Functional unity is pursued through infrastructure that can narrow the organisational discretion of peripheral bodies: they may still decide whether to adopt a tool, yet how they can adopt it is largely pre-set by central design choices (data fields, work-flows, risk thresholds, vendor terms). In effect, local autonomy shifts from making substantive policy and organisational decisions to merely configuring a centrally defined system”.

2.1. Legal implications of centralised AI-coordination.

The current state of the field confirms that this tension has become central, particularly in debates that distinguish regulatory models from implementation models. Comparative discussion has often contrasted the European approach, precautionary and risk-based, with the British approach, more flexible and reliant on coordination among regulators and on soft-law instruments¹⁶. At the same time, it has been observed that this dichotomy becomes less sharp when one examines the concrete mechanics of adoption, because even a “light” model can become highly consequential through technical standards, procurement, and policy guidance¹⁷.

¹⁵ M. HILDEBRANDT, *Algorithmic regulation and the rule of law*, in *Philosophical Transactions of the Royal Society*, Volume 376, Issue 2128, 2018.

¹⁶ M. MOURBY, *The ‘DPIA+’: Aligning data protection with UK equality law*, in *Computer Law and Security Review*, Volume 59, November 2025, pp. 1-12.

¹⁷ S. CAREY, *Regulating Uncertainty: Governing General-Purpose AI Models and Systemic Risk*, in *European Journal of Risk Regulation*, 2025, pp. 1-17; R. ASIF, *Equality and Inclusivity through Accountability: Regulating Automated Decision-Making in the Public Sector*, in *Queen Mary Law Journal*, Vol. 6, Issue 1, 2025, pp. 96-108.

Recent scholarship has noted, for instance, that the British approach tends to entrust equilibrium to market dynamics and to competition among authorities, whereas the European approach privileges stratification by risk; it has also suggested that the more plausible long-term response may lie in a multi-level system combining distributed competences with cross-cutting controls. The Action Plan sits within this trajectory. It does not replace a general regulatory framework; rather, it attempts to make adoption coherent through centralised capacity, shared infrastructure, and coordination mechanisms, with the stated ambition of making AI core to mission delivery and to public services.

Once governance moves into standards, legitimacy and control are particularly important phases: the risk is not limited to the well-known “black box” of automated decision-making, it also concerns a form of “institutional opacity”, in which choices most consequential for rights and substantive equality are not made solely in traditional sites of political deliberation but are embedded in standards, architectures, and contractual supply chains¹⁸.

Accordingly, the constitutional adequacy of “national level reach” will depend less on formal competence design than on the visibility and contestability of infrastructural choices which will translate in, at least, two requirements: a) assuring traceability and auditability by design (*i.e.* applying the rule-of-law principles) and b) constructing intelligible documentation and accountability allocations of resources. On this latter point, the Plan shows awareness when it discusses compute infrastructure and the growth of domestic capacity, emphasising the need to avoid excessive reliance on single providers and to preserve a plurality of technical options, and when, on the application side, it insists on modularity and open sourcing, where possible¹⁹. In my opinion, two further consequences follow.

First, this infrastructural turn gives operational bite to the earlier, principles-based regulatory stance: interoperability, modularity, documentation, and auditability become procurement and architecture requirements rather than aspirations.

Second, it sharpens the familiar constitutional difficulty of opacity. I mean, the risk is not limited to black-box model behaviour but it includes institutional opacity. In fact, where decisions with significant implications for equality, reasonableness, transparency, and reviewability are embedded in standards and contractual supply chains, the pre-existing legal framework supplies constraints²⁰, but the Plan’s innovations “raise the stakes” by amplifying the effects of early design choices. In that setting, what remains unchanged is the formal distribution of competences; what is newly contested is the locus of practical control; in this respects the Plan does not answer to a number of questions: who sets the data formats, technical baselines, contractual conditions, and audit arrangements that determine whether local variation is feasible or prohibitively costly? But this is in the very nature of every Plan, which is normally by itself generic. A final, strictly constitutional question remains: if AI territoriality is built through standards and platforms, then the

¹⁸ SUPREME COURT (LORD SALES), *AI and Public Law: Automated Decision-Making in Government (Keynote Lecture)*, 5 November 2025.

¹⁹ *Action Plan*, Section 2.2, Recommendations 41-42.

²⁰ More generally, the statutory duties under the Data Protection Act 2018 and the Equality Act 2010.

democratic quality of territorial governance will increasingly depend on three factors. The transparency and contestability of those standards, the participation of peripheral administrations in defining them, and the ability of oversight bodies to scrutinise what happens within the infrastructure. Without these conditions, the promise of “national level reach” risks hardening into only an apparent uniformity. Such uniformity would neither correct territorial inequalities nor leave meaningful space for pluralism; instead, it would reproduce disparities in technical form and shift the centre of gravity from the formal distribution of competences to the more elusive domain of system architecture.

3. The impact on devolution.

As regards devolution, the Action Plan does not address the issue in terms of constitutional architecture, nor does it dwell on the formal boundary between devolved and reserved competences. The territorial dimension emerges, rather, through an administrative diagnosis: the adoption of AI in the public sector is depicted as uneven and fragmented, often confined to isolated experiments²¹. This underpins the decision to rely on central instruments capable of supporting the transition to scale, that is, to solutions that are replicable and usable across different institutional settings. The provision for mission-oriented national procurement procedures — explicitly designed to accelerate adoption even within de-centralised systems — and the location of overall direction in the “digital centre of government” indicate a model of innovation governance that assigns an ordering function to procurement, technical capability, interoperability standards, and evaluation criteria²². This design accords with a British tradition of administrative pragmatism: unity is not pursued through major constitutional re-engineering, but through operational coordination aimed at reducing systemic heterogeneity, standardising minimum requirements, and making adoption sustainable at a national scale. From a regulatory standpoint, this approach sits alongside the existing “pro-innovation” governance model for AI: both seek coherence chiefly through some administrative instruments (*i.e.* procurement templates, shared standards and central capability) and their practical harmonisation rather than through a new AI statute, or a sectoral normative regulation. By doing so, it relies on legal tools already available to the centre, whose effects can be territorially expansive without formally reallocating devolved competences²³. The state of the debate on AI governance in the United Kingdom helps to clarify this approach. In its more recent phase, the British orientation has privileged flexible instruments, general principles, and sector-specific applications, placing greater emphasis on the capacity of sectoral regulators and on inter-institutional coordination than on the construction of a rigid horizontal framework. The literature has, however, highlighted the limits of an ecosystem

²¹ *Action Plan*, Section 2.1.

²² *Action Plan*, Recommendation n. 40 (“Mission-focused national AI tenders... across de-centralised systems”) and the Plan’s description of DSIT as “the digital centre of government” in Section 2.1 support the claim that procurement and central digital coordination operate as the Plan’s principal ordering devices, including explicitly within decentralised settings.

²³ *Action Plan*, §§ 2.2. See also Government Digital Service, *Government Functional Standard GovS 005: Digital* (v2.1, updated 15 July 2024), and GOV.UK *Service Standard* (Service Manual), both foreseeing cross-government “controls by standard” without changing the devolution of competences.

overly dependent on guidance and soft law²⁴: the proliferation of standards that are not always coherent, variations in implementation across regulated domains, and difficulties of accountability and transparency when the centre of gravity shifts towards technical choices and supply chains. Against this background, the Action Plan can be read as an attempt to “create a system” without abandoning the flexible posture, namely strengthening the diffusion of AI by setting common conditions and supporting administrations in moving from pilots to scale. This agenda interacts directly with established procurement law and practice: the Plan’s reliance on mission-oriented tenders and reusable technical components presupposes that contractual specifications, evaluation criteria, and performance requirements can function as quasi-regulatory mechanisms²⁵. Equally, the emphasis on interoperability and common architectures engages existing digital and data-governance obligations — in particular information governance and data-protection compliance pursuant to Article 22 UK GDPR and the Data Protection Act 2018 —, since “scaling” across administrations increases the salience of lawful data sharing and accountability for automated or AI-assisted processes.

At the same time, the Plan acknowledges that coordination cannot dispense with forms of territorial cooperation, especially where interventions affect material factors that are locally salient. The reference to securing local “buy-in” and to aligning AI Growth Zones with territorial needs (Recommendation n.º4 of the Plan) may indicate that the legitimacy and practical effectiveness of public action depend on collaboration with sub-state and local actors. Yet this does not translate into a formalized (namely, normative) intergovernmental design: what emerges instead is a working method based on technical support and administrative capacity, whereby the centre enables and steers, while territorial levels render implementation workable. It is here that the theme of administrative-constitutional pluralism, long developed in the reflections associated with the Devolution Club, acquires particular salience: devolution is not questioned at the formal level, but may be indirectly constrained when the instruments of scale determine, in practice, the range of available options. From this perspective, the decisive question is not only where competence is located, but who controls the infrastructural premises of administrative action: data formats, contractual terms, audit and accountability arrangements, and the technical requirements that enable — or preclude — local adaptation. The resulting risk is a functional recentralisation, difficult to perceive as such because it does not proceed through constitutional reform but through operational choices and standardisation. Conversely, it would be reductive to treat this dynamic as purely pathological. In a context of extensive public services and rising expectations of

²⁴ A.A. GIKAY, *Risks, innovation, and adaptability in the UK’s incrementalism versus the European Union’s comprehensive Artificial Intelligence regulation*, in *International Journal of Law and Information Technology*, Vol. 32, Issue 1, 2024. See also H. ROBERTS *et al.*, *Artificial intelligence regulation in the United Kingdom: A path to good governance?*, in *Internet Policy Review*, Vol. 12, Issue 2, 2023. Both the Authors characterise the UK model as flexible and stress that this approach, nevertheless, would rise concerns about fragmentation, inconsistency and accountability.

²⁵ At the same time, this use of procurement remains compatible with the prevailing obligations of transparency, equal treatment, and value for money in public purchasing, as reflected in section 12 of the Procurement Act 2023 and in the Public Contracts Regulations 2015 (spec. Regulation n.º 18). See also A. SANCHEZ-GRAELLS, *Responsibly Buying Artificial Intelligence*, cit., which also states that UK aim to place «AI and its desired competitive advantage as the major plank for its post-Brexit economic strategy» (82-83).

minimum uniformity in provision, central coordination may constitute a condition for substantive equality and reliability. The tension that the Plan implicitly brings into view is therefore structural: reconciling the drive towards scale and service coherence with the preservation of genuine margins for territorial adaptation and with forms of transparent responsibility, so as to avoid standardisation becoming a mode of territorial governance lacking the ordinary democratic and legal safeguards. Legally, this tension is sharpened by the differentiated territorial settlement: procurement and digital-service initiatives designed for “UK-wide” scaling must, in practice, pass through distinct administrative systems and legal regimes in the devolved nations, particularly where implementation intersects with devolved public services and local government. Moreover, if standardisation is pursued through central frameworks and contractual architectures, the key question becomes whether existing accountability mechanisms — judicial review, audit institutions, and administrative complaint pathways — can effectively scrutinise decisions embedded in technical standards and procurement choices, rather than only the downstream decisions taken by individual devolved or local bodies.

4. Growth zones.

The third main area of the “territorial” operability of AI pillar is the issue of “growth zones” (Recommendation n.° 4 of the Plan). This priority is treated as a place-based policy for the infrastructuring of compute, in which the siting of data centres is tied to two enabling conditions regarded as decisive: the acceleration of planning procedures and access to “clean” energy within timelines compatible with the investment horizon of private capital.

The underlying premise is straightforward, yet rich in implications: computational capacity is not merely an immaterial factor of production, but a heavy, energy-intensive, territorially embedded infrastructure, requiring land, connectivity, permits, and the availability and stability of the electricity network. The Plan therefore approaches compute as a new component of national infrastructure policy, comparable, by strategic importance, to networks, ports, or large industrial plants. In doing so, it makes the feasibility of national-scale ambitions for AI adoption, in public services as well as in the private sector, depend in significant part upon that infrastructural build-out²⁶. In continuity with the document’s overall approach, territoriality is here constructed through capacity-building arrangements and coordination instruments rather than through abstract legal propositions: the “zone” is a site in which administrative conditions (planning), technical conditions (connection and network capacity), and financial conditions (investor attractiveness) are concentrated, with the aim of making more predictable and faster what, absent intervention, would remain exposed to delay and uncertainty. From the standpoint of the existing regulatory environment, this move operates primarily through the planning and energy consent frameworks already in force: the Plan does not, as such, create a new AI-specific legal regime for land-use or grid access, but seeks to re-orient decision-making through guidance, prioritisation,

²⁶ On the topic, see J. MONSTADT, K. SALTZMAN, *How Data Centers Have Come to Matter: Governing the Spatial and Environmental Footprint of the ‘Digital Gateway to Europe’*, in *International Journal of Urban and Regional Research*, volume 49(4), 757-778.

and administrative coordination within those frameworks.

In the most recent debate, and now also in comparative legal-institutional scholarship, the view has consolidated that AI is best read through the “data–compute” pairing: a system dependent upon material resources and infrastructural constraints, not reducible to software or abstract “models”. Attention to large-scale digital infrastructure, in particular data centres and cloud, has consequently become a central theme of technological sovereignty and industrial policy scientific debate²⁷. The Plan situates itself within this trend, but does so with a specifically British inflection: rather than proposing a single, centralised, wholly public programme, it sets up a device aimed at steering and accelerating private investment by removing administrative bottlenecks and, as far as possible, rendering project conditions reliable. The growth zone, as formulated, thus functions as a mechanism for “valorising” territory: certain places are selected because they already possess energy capacity or other favourable conditions, and a package of facilitations is concentrated there with the expectation of producing an agglomeration effect. This choice sits at the intersection of infrastructure policy and administrative law: it presupposes that speed and predictability can be raised not by displacing legal constraints, but by reorganising institutional capacity, sequencing consents, and aligning the incentives that the current regulatory apparatus already permits.

The most salient element, with regard to the growth of less advantaged areas, is the characterisation of AI Growth Zones as an instrument of “local rejuvenation”, capable of channelling investment towards areas with existing energy capacity, with express reference to post-industrial towns and coastal areas of Scotland. Here the notion of territoriality shifts register: it is not only the efficient location of data centres, but an attempt to connect the AI strategy to a British tradition of territorial regeneration and reconversion, long familiar with investment-attraction policies implemented through incentives, simplification, and special “zones”. On this reading, the Plan suggests that the emerging compute economy might perform, for some places, a role analogous to that played in other historical cycles by industrial or logistical infrastructure: bringing jobs, supply chains, services, and fiscal returns, thereby contributing to a reduction of territorial disparities. This argument remodulate the British politico-administrative context by redeploying a familiar legal vocabulary in relation to a sector perceived as strategic and future-oriented, and seeking to link technological innovation with territorial cohesion. Yet the regulatory architecture matters: regeneration-through-compute can only be credible insofar as the incentives and accelerations proposed are consistent with the prevailing rules on planning and assess local benefits. Local institutions retain recognisable roles within their respective statutory.

Precisely for that reason, however, the choice of growth zones warrants a more stringent critical assessment than the Plan’s programmatic language might suggest. The literature on place-based policies and technological clusters has long insisted on one point: infrastructural concentration yields high benefits when it generates local externalities, that is, when investment translates into an ecosystem of skills, firms, services, research, training, and public demand. In the absence of such conditions, the risk is that of “enclaves”: infrastructures that occupy land and

²⁷ *Ex multis*, M.T. TAKCI *et al.*, *Data centres as a source of flexibility for power systems*, in *Energy Report*, Volume 13, June 2025, 3661-3671.

absorb energy but produce circumscribed economic effects, with weakly rooted supply chains and limited local spillovers. In the case of data centres this risk is especially concrete: after the construction phase, jobs may be relatively few; effects on local SMEs are not automatic; and value added may be captured elsewhere, particularly if advanced functions of design, management, and high-end services remain concentrated in other territories. It follows that “local rejuvenation”, if understood robustly, cannot be entrusted to the mere physical presence of compute. It requires a policy of territorial integration that complements infrastructural development with measures on training, attraction of complementary firms, local procurement, digital public services, and, more generally, the institutional capacity of the territory to govern the investment. In legal-administrative terms, this points to the limits of relying on accelerated consenting alone: the distributive and developmental claims attached to zones must be translated into enforceable governance arrangements if they are to withstand scrutiny within the existing regulatory order.

From this perspective one can see why, in the Plan, the insistence on the involvement of “local regions” should not be read as a stylistic detail, but as a marker of a structural problem: without territorial cooperation the growth zone risks remaining a top-down project, perceived as extractive or, in any event, misaligned with local needs. The Plan suggests a cooperative trajectory, but it does not fully specify which institutional instruments should ensure that cooperation is substantive rather than merely consultative, nor how the conflict is to be composed between national objectives and local interests (such as land use, environmental impact, network infrastructure, opportunity costs, distribution of benefits)²⁸. In a territorially composite polity such conflicts are unavoidable, and their handling is precisely the test case for the administrative-constitutional pluralism so often discussed in reflection on devolution. This is also where present regulatory structures become decisive: planning competence and territorial development powers are not uniform across the United Kingdom, and any zone-based acceleration must be reconciled with distinct consent procedures, accountability channels, and — where applicable — separate environmental and energy governance arrangements²⁹.

On this point, growth zones bring into focus a feature already visible in the Plan’s other nuclei: the tendency to pursue unity and “scale” not through formal redefinitions of competence, but through technical-infrastructure and financial levers. In the case of zones, this occurs through a package combining planning and energy infrastructure with incentives and administrative coordination. This mode of action is readily intelligible within the Plan’s logic: remove brakes, create predictability, reduce time, and make investment “bankable”. Yet an approach of this kind produces indirect constitutional consequences. In particular, simplified planning and any strengthening of central intervention capacity may — if not carefully calibrated reduce the space within which territorial discretion is exercised, not because it is formally withdrawn, but because the urgency of the

²⁸ See HOUSE OF COMMONS, *Data centres: planning policy, sustainability, and resilience* (research briefing), 3 November 2025, in part. 28 ss., online on <https://researchbriefings.files.parliament.uk/documents/CBP-10315/CBP-10315.pdf>.

²⁹ *Ivi*, 54 and 68.

national objective pushes towards standardisation and acceleration. The United Kingdom's territorial asymmetry increases the sensitivity of the point: what can be pursued in England through certain instruments (including more assertive central administrative involvement) intersects, in Scotland, Wales, and Northern Ireland, with different regimes and distinct administrative and political cultures. Hence intergovernmental cooperation is not an accessory, but a condition of legitimacy and effectiveness. Normatively, the Plan's ambitions therefore sit alongside, rather than above, the existing constitutional settlement: the legal durability of zone policies depends on their capacity to operate through, and be justified within, the differentiated territorial regimes of planning and economic governance.

On growth zones, following the Action Plan, the UK Government published an "AI Growth Zones Policy Paper"³⁰ which launches AI Growth Zones to speed the deployment of UK AI data centres via faster planning, priority grid connections and targeted power discounts, supported by a delivery unit and an investor pipeline. The central question is whether growth zones constitute a device capable of connecting the national compute strategy with territorial plurality, without reducing devolution to a merely executive role and without shifting onto territories the costs (energy, environmental, land-use) in exchange for uncertain benefits. Recent discussion on the governance of AI and digital infrastructure has shown that the social and institutional acceptability of large projects depends on the transparency of choices and the credibility of compensation and redistribution mechanisms. It is not enough to claim that investment "will create jobs": it is necessary to show how benefits will be retained locally, how skills and supply chains will be strengthened, what assurances will be offered regarding grid impacts and energy costs, and how the territory's capacity to orient development in line with its priorities will be preserved. Absent these elements, the redistributive objective risks being perceived as rhetorical, while the infrastructuring of compute may be read as a further concentration of economic and decision-making power. Here the interaction with current regulatory disciplines becomes explicit: where financial inducements, preferential treatment, or accelerated access are employed, they must be structured so as to remain compatible with the prevailing procurement and subsidy-control constraints, as well as with planning and environmental requirements that condition the permissibility of speed.

Finally, growth zones can also be understood as part of a broader transformation in the manner in which the state "makes territory" in the digital age. If, as the Action Plan maintains elsewhere, AI adoption requires common standards, scalable platforms, and procurement procedures capable of diffusing solutions at national scale, then territory is governed not only through boundaries and competences, but through access to infrastructure and the capacity to insert oneself into implementation chains. In this sense, the growth zone is the point at which the material dimension of AI (energy, land, networks) meets the administrative dimension (planning, coordination, market rules) and the constitutional dimension (territorial pluralism, democratic responsibility, the balance between centre and periphery). The challenge for the United Kingdom is to avoid reducing the place-based policy of compute to a bundle of technical facilitations lacking an adequate

³⁰ DSIT, *Delivering AI Growth Zones*, 13 November 2025, online on <https://www.gov.uk/government/publications/delivering-ai-growth-zones/delivering-ai-growth-zones>.

institutional anchoring. If, instead, zones are accompanied by instruments capable of making territorial cooperation effective, guaranteeing verifiable local spillovers, and preserving margins of adaptation across devolved settings, they may become a significant laboratory: not only for expanding compute capacity, but also for a form of AI governance compatible with the United Kingdom's complex territoriality and its administrative-constitutional pluralism. In this perspective, the Plan's zone strategy functions less as an autonomous "AI policy" and more as an interface between AI objectives and established bodies of law whose constraints and procedures will largely determine whether the promised territorial dividends materialise.

5. Conclusion.

The 2025 AI Opportunities Action Plan exemplifies a familiar feature of British constitutional practice: pursuing coherence in public administration less through formal reallocations of competences than through capacity, coordination, and common operating frameworks. In this sense, the Plan shifts the centre of gravity from competence to capacity: it treats "national level reach" as an infrastructural achievement, to be delivered via shared technical standards, reusable components, mission-oriented procurement, and centrally supported pathways from pilots to scale. The constitutional novelty lies not in rewriting the devolution settlement, but in how a centrally steered "stack" and procurement architecture can reallocate practical control by defining data formats, workflows, performance metrics, and risk thresholds that shape what devolved and local bodies can realistically do. This makes legitimacy depend less on the formal location of powers than on the visibility, contestability, and accountability of infrastructural choices. If standardisation and scaling are to support territorial cohesion without collapsing pluralism into mere configuration, the Plan's approach requires rule-of-law conditions to be built in upstream: traceability and auditability by design, intelligible documentation and responsibility allocation across supply chains, and meaningful participation of devolved administrations and local government in setting shared standards. Similar considerations apply to Growth Zones: accelerated planning and energy coordination may be compatible with territorial regeneration only if accompanied by governance arrangements that secure verifiable local spillovers, preserve statutory roles in devolved regimes, and make distributive trade-offs transparent. Read this way, the Action Plan is best understood as an attempt to construct infrastructure-led territorial governance for AI: a model that can deliver functional unity, but only if its technical and contractual architecture is treated as a constitutional object — in particular open to scrutiny and democratic oversight — rather than as a merely managerial means to scale.