European Review of Digital Administration & Law - Erdal
2026, Volume 7, Issue 1, pp. 5-8
ISSN 2724-5969 – ISBN 979-12-218-2846-7– DOI 10.53136/9791221828467 01

Artificial Intelligence, Liability and Public Administration

Valentina Giomi

The use of artificial intelligence by public authorities does not merely introduce new technologies into established administrative processes. It affects the manner in which public decisions are prepared, adopted and justified, while also transforming the conditions under which liability for those decisions may be established.

The traditional architecture of State liability has generally been constructed around identifiable conduct, human fault, intelligible causal sequences and a relatively clear distinction between the person who decides and the instrument through which the decision is implemented. Artificial intelligence unsettles each of these assumptions.

An administrative outcome produced with the support of an AI system may depend on a complex combination of factors: the decision to automate a particular activity, the procurement of the technological solution, the design of the model, the selection and quality of the data, the formulation of operational parameters, the interaction between different information systems, the organisation of human oversight and the manner in which the final output is interpreted or implemented. Harm may therefore arise without being readily attributable to a single act or actor. Responsibility appears dispersed across an institutional and technological chain involving public authorities, private providers, programmers, data suppliers, system operators and officials entrusted with the final decision.

This fragmentation does not, however, entail the disappearance of responsibility. Public authorities remain bound by legality, impartiality, transparency, reason-giving and effective judicial protection. The administration cannot evade these obligations by presenting artificial intelligence as an autonomous source of decision-making. Nor can responsibility be transferred entirely to a technological provider or dissolved within the complexity of the system's lifecycle. The adoption of AI may alter the conditions under which responsibility is identified and proved, but it cannot render public power less accountable.

The contributions gathered in this monographic section address this tension from different legal and comparative perspectives. A first group of questions concerns the functions of liability itself. In the context of algorithmic administration, liability continues to perform its classical compensatory function: individuals must be able to obtain redress where automated or AI-assisted processes cause unlawful harm, discrimination, loss of access to public services or violations of fundamental rights. Yet compensation alone may prove insufficient where harmful effects are systemic, difficult to detect or capable of affecting large categories of persons.

Liability therefore increasingly assumes a broader preventive role. It may influence the organisational and technological choices made by public authorities before harm occurs, encouraging the adoption of adequate safeguards, effective monitoring mechanisms and systems designed to minimise foreseeable risks. From this perspective, liability is no longer merely an *ex post* reaction to a completed administrative failure. It becomes one of the mechanisms through which the development and deployment of public-sector AI are governed.

This preventive dimension is particularly important because algorithmic harm may be cumulative, diffuse and difficult to reverse. A defective administrative decision can ordinarily be annulled or corrected. An AI system operating at scale, by contrast, may reproduce the same error across thousands of cases, consolidate discriminatory patterns or progressively alter administrative practice before its defects are identified. The law of liability must therefore be connected with anticipatory instruments, including impact assessments, audit duties, documentation requirements, internal controls and judicial remedies capable of addressing systemic risks rather than merely compensating individual losses.

A second major issue concerns the continuing adequacy of traditional requirements of fault, causation and proof. In

many legal systems, the unlawful use of an algorithm or an unlawful decision based upon it can still be examined through established doctrines of State liability. Artificial intelligence does not necessarily require the construction of an entirely autonomous liability regime. The main difficulties often arise not from the absence of applicable legal rules, but from the claimant's inability to reconstruct how a harmful outcome was produced.

The opacity of AI systems may make it difficult to establish whether the error originated in the data, in the model, in the system's design, in its configuration, in its interaction with other systems or in the conduct of the officials who relied on its output. Linear causation becomes harder to demonstrate when the result emerges from the combined operation of multiple technological and organisational components. Similarly, fault may be difficult to identify where the internal functioning of the system is inaccessible to the affected individual and perhaps not fully understood even by the administration that deploys it.

These evidentiary asymmetries raise the question whether traditional rules concerning the burden of proof should be adjusted. Presumptions of fault, disclosure obligations, facilitated access to technical documentation and forms of strict or risk-based liability may be necessary in certain contexts. Such mechanisms would not amount to treating every harmful output as automatically unlawful. They would instead respond to the structural imbalance between the public authority that selects and controls the system and the individual who suffers its consequences without having access to the information needed to challenge it.

The possible transition from fault-based liability to forms of liability grounded in technological risk is especially significant in fields such as taxation, social security and the allocation of public benefits. In these areas, AI systems may process large volumes of personal and financial data, generate risk profiles and influence decisions with immediate effects on individual legal positions. Where the State benefits from the efficiency, speed and predictive capacity of these systems, it is questionable whether the risks arising from their opacity or malfunction should be borne primarily by the affected person.

The problem of attribution also extends to the internal liability of public officials. The use of artificial intelligence may appear to reduce the sphere of individual decision-making, yet it simultaneously creates new duties relating to procurement, data management, system configuration, verification, supervision and intervention. Administrative and accounting liability must therefore confront a paradox. Officials may be held responsible for failing to control technological systems whose functioning they are not technically equipped to understand, while the fear of liability may discourage the adoption of innovative tools or induce excessive reliance on preventive opinions and formal legality controls.

The solution cannot consist simply in transferring full responsibility to the individual official situated at the end of the process. Effective responsibility presupposes adequate organisational conditions. Public authorities must possess the internal expertise necessary to evaluate the systems they acquire and use. They must define clear competences, establish procedures for escalation and correction, preserve relevant documentation and ensure that those entrusted with oversight have both the authority and the practical ability to intervene. Where these conditions are absent, the failure is not merely individual. It is organisational.

This leads to one of the central themes of the section: the relationship between State liability and the human-in-the-loop. The preservation of human involvement is frequently presented as the principal safeguard against the risks of automated decision-making. Yet the mere presence of a person within the decisional chain does not guarantee meaningful control. Human intervention may be purely formal where the official lacks the technical knowledge, information, time or institutional power required to question the system's output.

There is also a risk that the human-in-the-loop becomes a mechanism for attributing legal responsibility to an individual without ensuring genuine decisional autonomy. Automation bias may lead officials to accept machine-generated recommendations because they appear objective or scientifically grounded. Conversely, selective reliance on AI may allow the human decision-maker to invoke the system when its output supports a preferred conclusion and disregard it when it

does not. In both cases, the formal attribution of the decision to a human being may conceal rather than clarify the actual distribution of power within the process.

Meaningful human oversight therefore requires more than the possibility of approving or rejecting an output. It requires sufficient understanding of the system's purpose, limitations and conditions of use; access to the information needed to assess its reliability; the capacity to identify anomalies; and the authority to override, suspend or discontinue its operation. The most appropriate model may be that of an artificial co-pilot: a system capable of supporting administrative analysis without displacing the irreducible responsibility of the public decision-maker. Such a model, however, depends on sustained investment in public-sector expertise and on organisational arrangements designed to prevent both uncritical deference and fictitious control.

Responsibility is also inseparable from explanation and justification. Technological explainability and legal justification must not be confused. An account of the variables that influenced a model, the statistical correlations it identified or the computational path through which it produced an output does not necessarily explain why the resulting administrative decision is legally justified. Public authorities must still identify the applicable legal rule, establish the relevant facts and explain how those facts have been legally qualified.

This justificatory process is inherently contextual. Administrative law does not operate through the mechanical application of abstract propositions to neutral data. Facts acquire legal significance within a particular normative and institutional setting. The duty to give reasons therefore requires situational awareness and an intelligible synthesis of fact and law. AI systems may assist this process, but they must not erase the contextual elements on which lawful administrative reasoning depends.

The distinction becomes particularly important in relation to large language models and other generative systems. Their ability to produce fluent and apparently persuasive explanations may create the appearance of legal justification without ensuring that the reasons provided correspond to the actual basis of the decision. A generated explanation may be plausible yet inaccurate, generic or disconnected from the system that produced the substantive outcome. The duty to give reasons cannot be satisfied by linguistic coherence alone. Reasons must be truthful, case-specific, reviewable and connected to the real decisional process.

The same concerns extend beyond administration in the strict sense to algorithmic justice. Courts, prosecutors, law-enforcement authorities and legal professionals increasingly rely on technological tools for research, case management, document processing, risk assessment and other activities connected with the administration of justice. In this field, civil liability is only one component of a broader accountability framework. Judicial review, professional ethics, institutional governance, internal compliance mechanisms and soft-law standards may play an equally important role, particularly where judicial and prosecutorial immunities limit ordinary compensatory remedies.

The comparative analyses collected in this section reveal that legal systems are responding to these challenges through different combinations of continuity and innovation. Some rely primarily on the adaptation of traditional doctrines of administrative liability. Others place greater emphasis on presumptions, strict liability, organisational duties or risk-based approaches. Yet a common tendency emerges: responsibility is progressively being connected to the entire lifecycle of the AI system rather than exclusively to the final administrative act.

The decision to automate, the selection of the provider, the design of the system, the governance of data, the allocation of institutional competences, the training of officials, the provision of reasons, the availability of review and the correction of errors all become legally relevant. Liability, oversight and organisational governance can therefore no longer be treated as separate domains. They form part of a single architecture of public accountability.

Artificial intelligence may improve the efficiency, consistency and analytical capacity of public administration. It may also make errors more difficult to detect, causal relationships harder to reconstruct and responsibility easier to disperse. The central challenge for public law is to ensure that technological complexity does not become

legal irresponsibility.

The decisive question is not simply whether a human signature remains at the end of the procedure. It is whether public institutions retain effective control over the technological systems through which they exercise their powers. Such control requires identifiable responsibilities, adequate expertise, transparent organisational choices, legally sufficient reasons and remedies capable of addressing both individual and systemic harm. Artificial intelligence may transform the modalities of administrative action, but it cannot alter the fundamental principle that public power must remain answerable to those whose rights and interests it affects.