

**Constitutional and Administrative Law Constraints on
Public-Sector AI Decision Systems: A Jurisprudential
Framework for Legality and Accountability**

PHD DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

June 2026

Table of Contents

Abstract	1
1. Introduction	3
1.1 Background and Motivation	3
1.2 Problem Statement	6
1.3 Research Question	10
1.4 Scope and Delimitations	13
1.5 Thesis Organisation	15
2. Main Body	17
2.1.1 Doctrinal Foundations of Due Process	17
2.1.2 Transparency Requirements	21
2.1.3 Judicial Review and Algorithmic Decisions	24
2.1.4 Accountability Mechanisms	27
2.1.5 Technical Challenges to Doctrinal Application	30
2.1.6 Comparative Regulatory Landscapes	33
2.2 Review Methodology	33
2.2.1 Research Design	34
2.2.2 Search Strategy	36
2.2.3 Selection and Coding	40
2.2.4 Limitations of Methodology	43
2.2.5 Ethical Considerations	45
2.2.6 Summary of Methodological Approach	45
2.3 Synthesis of Findings	46
2.3.1 The Principle-Practice Gap	46
2.3.2 Fragmented Regulatory Responses	52
2.3.3 Emerging Themes from the Literature	57

2.3.4 Unresolved Doctrinal Tensions	61
2.3.5 Preliminary Summary: Mapping Technical Characteristics to Doctrinal Re- quirements	65
2.4 Discussion	66
2.4.1 Interpretation of Synthesis	67
2.4.2 A Synthesised Framework: Doctrinal Constraints for Public-Sector AI . . .	69
2.4.3 Comparison with Existing Frameworks	73
2.4.4 Theoretical Implications	75
2.4.5 Limitations of the Synthesis	77
2.4.6 Future Research Agenda	78
3. Conclusion	82
3.1.2 Principal Findings	83
3.2 Contributions of the Thesis	85
3.2.1 Doctrinal Contributions	85
3.2.2 Framework Contributions	86
3.2.3 Methodological Contributions	87
3.3 Implications for Law and Policy	88
3.3.1 Implications for Constitutional Doctrine	88
3.3.2 Implications for Legislative Reform	88
3.3.3 Implications for Public Administration	89
3.4 Limitations of the Thesis	90
3.4.1 Scope Limitations	90
3.4.2 Methodological Limitations	91
3.4.3 Evidentiary Limitations	92
3.5 Directions for Future Research	92
3.5.1 Empirical Validation	92
3.5.2 Comparative Doctrinal Analysis	93

3.5.3 Technical-Legal Interface Research	93
3.5.4 Normative Questions	94
3.6 Concluding Remarks	94
4. Appendices	96
A.1 Overview of the Framework	96
A.1.1 Rationale for a Structured Mapping	96
A.2 Doctrinal Requirements and Their Technical Counterparts	96
A.2.1 Due Process and the Right to a Fair Hearing	96
A.2.2 Transparency and Explainability	97
A.2.3 Judicial Review and the Grounds of Review	98
A.2.4 Accountability Mechanisms	99
A.3 Framework Diagram (Textual Representation)	101
A.4 Limitations of the Framework	101
Appendix B: Supplementary Data Tables	102
B.1 Comparative Overview of National AI Governance Frameworks	102
B.2 AI Use Cases in European Public Administration	103
B.3 Summary of Technical Characteristics and Their Legal Implications	103
B.4 Glossary of Key Terms (Compact Version)	104
Appendix C: Glossary of Terms	105
C.1 Administrative Law Terms	105
C.1.1 Due Process	105
C.1.2 Judicial Review	106
C.1.3 Proportionality	106
C.1.4 Non-Delegation	106
C.1.5 Reason-Giving Duty	107
C.2 Technical Terms	107

C.2.1 Machine Learning (ML)	107
C.2.2 Opacity	107
C.2.3 Explainability	107
C.2.4 Human-in-the-Loop (HITL)	107
C.2.5 Audit Trail	108
C.3 Interdisciplinary Terms	108
C.3.1 Algorithmic Governance	108
C.3.2 Accountability	108
C.3.3 Digital Right to Good Administration	108
C.3.4 Administrative Constitutionalism	108
Appendix D: Additional Resources	109
D.1 Key Legislation and Policy Documents	109
D.2 Recommended Further Reading	110
D.3 Tools and Methodological Resources	111
D.4 Organisations and Initiatives	111
D.5 Glossary (Compact Version)	112
References	113

Abstract

Research Problem and Approach: Governments across developed jurisdictions are rapidly deploying artificial intelligence systems to make or inform administrative decisions, yet the constitutional and administrative law doctrines governing public power--due process, transparency, reasoned decision-making, and accountability--were designed for human decision-makers and may be ill-suited to algorithmic systems whose internal logic is opaque and outputs probabilistic. This thesis addresses the widening gap between doctrinal principle and operational practice by undertaking a systematic synthesis of legal constraints on public-sector AI, drawing on scholarship and regulatory frameworks from the United Kingdom, Germany, and Canada to assess whether existing doctrines can adequately regulate algorithmic governance.

Methodology and Findings: The research employs a comparative doctrinal analysis of constitutional and administrative law across three distinct legal traditions--UK common law, German civil law with its *Rechtsstaatsprinzip*, and Canada's hybrid Charter-based system--combined with a structured review of emerging regulatory instruments. The analysis reveals that while all three jurisdictions ground public-sector AI governance in foundational legal principles, considerable doctrinal gaps persist, particularly regarding the explainability of algorithmic outputs, the delegation of discretionary power to automated systems, and the adequacy of audit-based accountability mechanisms as substitutes for traditional judicial review.

Key Contributions: This thesis makes three primary contributions: A wide-ranging mapping of constitutional and administrative law constraints onto the technical characteristics of AI decision systems across three comparator jurisdictions, A critical assessment of the UK, German, and Canadian approaches to algorithmic transparency, procedural fairness, and accountability, identifying areas of doctrinal convergence and unresolved tension, and A framework for aligning AI governance strategies with the normative requirements of the rule of law, addressing the fundamental question of whether specific decisions should be delegated to algorithms at all.

Implications: The findings suggest that prevailing governance frameworks remain oriented toward implementation rather than the deeper constitutional question of appropriate delegation, and that doctrinal adaptation--not merely technical auditing--is necessary to preserve the legitimacy of public administration. This research provides guidance for policymakers, legislators, and courts seeking to develop legal frameworks that reconcile the efficiency gains of automation with the procedural and substantive protections that underpin constitutional democracy.

Keywords: Administrative Law, Algorithmic Decision-Making, Artificial Intelligence, Constitutional Constraints, Due Process, Explainability, Germany, Judicial Review, Public Administration, Public Sector AI, Rule of Law, Transparency, United Kingdom, Canada, Accountability

1. Introduction

1.1 Background and Motivation

Governments across the developed world are rapidly integrating artificial intelligence systems into public-sector decision-making. From welfare eligibility assessments and immigration case processing to predictive policing and parole recommendations, algorithmic tools now shape administrative outcomes that directly affect individual rights and obligations (van Noordt and Misuraca 2022). The European Union alone has seen a marked acceleration in AI adoption by public authorities, with van Noordt and Misuraca documenting over seventy use cases across member states, ranging from fraud detection in tax administration to resource allocation in healthcare (van Noordt and Misuraca 2022). This trend is not confined to Europe. In Canada, automated benefit determination systems have been deployed at both federal and provincial levels, while the United Kingdom has piloted AI tools in areas as varied as visa processing and criminal sentencing recommendations (Ndlovu 2025).

The promise of these systems is considerable. Proponents argue that AI can enhance administrative efficiency, reduce processing delays, minimise human error, and promote consistency in decision-making across large caseloads (Wirtz and Müller 2025). Ndlovu observes that AI-driven tools offer the potential to streamline public administration by automating routine tasks, freeing human decision-makers to focus on complex cases that require discretion and judgment (Ndlovu 2025). In resource-constrained environments, where public agencies face growing demand with limited budgets, the appeal of automated efficiency is difficult to overstate. Wirtz and Müller note that governments increasingly view AI governance strategies as essential to modernising public administration and maintaining service quality in an era of fiscal pressure (Wirtz and Müller 2025).

Yet the normative foundations of public-sector decision-making rest on principles that predate the digital age. Constitutional and administrative law doctrines -- due process, procedural fairness, transparency, reasoned decision-making, and accountability -- have evolved over centuries to

constrain the exercise of public power (Carnat 2025). These doctrines ensure that administrative decisions are made lawfully, rationally, and fairly, and that affected individuals have meaningful opportunities to challenge adverse outcomes. Carnat argues that the automation of administrative functions constitutes a delegation of public power, and that constitutional constraints must apply with equal force regardless of whether the decision-maker is human or algorithmic (Carnat 2025). The core question, then, is not whether AI can improve administrative efficiency, but whether the existing doctrinal framework is adequate to regulate decision systems whose internal logic may be opaque, whose outputs may be probabilistic rather than deterministic, and whose operation may resist the traditional forms of scrutiny that administrative law provides.

The tension between technological capability and legal constraint is not merely theoretical. Existing administrative law doctrines were designed for a world in which human officials exercised discretion, provided reasons for their decisions, and could be questioned about their reasoning processes (Kavungal 2026). Kavungal argues that prevailing governance frameworks remain oriented toward implementation -- addressing *how* AI should be deployed -- rather than the more fundamental question of *whether* specific decisions should be delegated to AI systems at all (Kavungal 2026). This framing inversion has considerable consequences. By focusing on technical safeguards and auditing mechanisms, policymakers risk neglecting the deeper constitutional questions that arise when algorithmic systems assume functions previously reserved for human judgment.

The motivation for this thesis arises from the observation that the gap between doctrinal principle and operational practice is widening. As AI systems become more sophisticated and their deployment more widespread, the limitations of existing legal frameworks become increasingly apparent (Khalid Ranjha 2025). Khalid Ranjha identifies a need for a comparative framework that can bridge administrative law principles with the technical realities of algorithmic decision-making, noting that existing approaches tend to treat the two domains in isolation (Khalid Ranjha 2025). This thesis responds to that gap by undertaking a systematic synthesis of doctrinal literature on constitutional and administrative law constraints as they apply to public-sector AI, drawing on

legal scholarship and regulatory frameworks from the United Kingdom, the European Union (with particular focus on Germany), and Canada.

The choice of comparator jurisdictions is deliberate. Each represents a distinct legal tradition in addressing the relationship between public power and individual rights. The UK's common law tradition of judicial review, with its emphasis on procedural fairness and reasonableness review, offers one model for constraining administrative discretion (Graham and Russell 2025). Germany's civil law system, with its constitutional grounding in the rule of law (*Rechtsstaatsprinzip*) and its explicit statutory provisions on automated administrative acts, provides a contrasting approach (Kim 2024). Canada's hybrid system, combining common law administrative law with a constitutional Charter of Rights and Freedoms, offers yet another perspective on how fundamental rights protections intersect with algorithmic governance (Saul 2007). The comparative dimension of this thesis allows for an assessment of whether certain doctrinal approaches are better equipped to address the challenges posed by AI, and whether convergence around common principles is emerging across jurisdictions.

A further motivation lies in the observation that legal scholarship on AI governance has grown rapidly but unevenly. Much of the existing literature focuses on specific doctrinal issues -- transparency, accountability, non-discrimination -- without offering a detailed synthesis of how these principles interact when applied to algorithmic systems (Molins 2026). Molins notes that while AI auditing has emerged as a prominent mechanism for ensuring accountability, the relationship between auditing practices and legal accountability frameworks remains undertheorised (Molins 2026). Similarly, discussions of transparency often proceed without adequate attention to the technical constraints that limit explainability in complex machine learning models. This thesis seeks to address these gaps by providing a structured synthesis that maps doctrinal requirements onto the technical characteristics of AI decision systems.

The broader societal stakes are substantial. Public trust in government depends, in part, on the perception that administrative decisions are made fairly, rationally, and accountably (Ndlovu 2025). If algorithmic systems are perceived as opaque, arbitrary, or unaccountable, the legitimacy

of public administration itself may be eroded. Conversely, if appropriate legal frameworks can be developed to govern AI decision-making, there is an opportunity to harness the benefits of automation while preserving the core values that underpin constitutional and administrative law. This thesis contributes to that project by clarifying the doctrinal constraints that must inform any viable governance framework for public-sector AI.

1.2 Problem Statement

The central problem that this thesis addresses is the disjunction between existing constitutional and administrative law doctrines and the operational realities of AI decision systems in the public sector. This disjunction manifests in several interconnected dimensions.

First, the opacity of many AI systems challenges the transparency requirements that are foundational to administrative law. Traditional administrative law demands that decisions be accompanied by reasons that can be understood, evaluated, and challenged (Carnat 2025). Yet machine learning models -- particularly deep neural networks and ensemble methods -- often produce outputs through processes that resist straightforward explanation. Even where technical explainability methods exist, such as feature importance scores or surrogate models, these may not satisfy the legal standard for reasoned decision-making, which requires a justification that connects the evidence to the outcome in a logically coherent manner. Carnat argues that this mismatch between technical explainability and legal reason-giving constitutes a fundamental constitutional challenge, one that cannot be resolved solely through technical innovation (Carnat 2025). The problem is not merely that AI systems are difficult to understand, but that the form of understanding they afford may be incompatible with the normative structure of administrative justification.

Second, the nature of algorithmic decision-making raises questions about the adequacy of existing accountability mechanisms. Administrative law provides for accountability through a combination of internal review, oversight bodies, and judicial review (Clayton et al. 2018). These mechanisms assume that decisions can be traced to identifiable decision-makers, that reasoning processes can be reconstructed, and that errors can be corrected through established procedures. AI

systems complicate each of these assumptions. Where decisions are generated by models trained on large datasets, it may be difficult to attribute particular outcomes to specific inputs or design choices (Molins 2026). Molins observes that the distributed nature of algorithmic accountability -- involving data scientists, system designers, procurement officials, and deploying agencies -- creates a diffusion of responsibility that existing legal frameworks are not well-equipped to handle (Molins 2026). The result is a potential accountability deficit, where no single actor bears clear legal responsibility for the outcomes of algorithmic decisions.

Third, the use of AI in administrative decision-making raises concerns about procedural fairness that existing doctrines may not adequately capture. Procedural fairness requires that affected individuals have a meaningful opportunity to participate in the decision-making process, to know the case against them, and to respond to adverse evidence (Khalid Ranjha 2025). When decisions are made algorithmically, individuals may not understand how the decision was reached, may not have access to the data or model that produced it, and may not be able to identify what evidence they need to contest. Khalid Ranjha argues that the comparative framework for algorithmic accountability must grapple with the distinctive procedural challenges that AI systems pose, including the difficulty of ensuring meaningful participation when the decision-making process is opaque (Khalid Ranjha 2025). The problem is compounded where AI systems are used to make predictions or risk assessments that influence, rather than directly determine, administrative outcomes. In such cases, individuals may be unaware that an algorithmic tool has shaped the decision affecting them, let alone able to challenge its predictions.

Fourth, the delegation of discretionary power to AI systems raises constitutional questions about the separation of powers and the limits of permissible delegation. In many legal systems, certain types of decisions -- particularly those involving considerable discretion or the exercise of judgment -- are reserved for human officials (Kim 2024). Kim notes that Germany has imposed explicit statutory restrictions on fully automated administrative acts where discretion is involved, requiring that such decisions be made by human officials unless legislation expressly provides otherwise (Kim 2024). However, the boundaries of this restriction are contested, and the line between

routine decisions that can be automated and discretionary decisions that require human judgment is not always clear. As AI systems become more capable of simulating judgment, the question of what constitutes a decision that must be made by a human becomes increasingly pressing.

Fifth, there is a temporal dimension to the problem. Constitutional and administrative law doctrines evolved over decades and centuries, through accumulated judicial decisions and legislative refinements. AI technologies, by contrast, develop on timescales measured in years or even months (van Noordt and Misuraca 2022). Van Noordt and Misuraca document the rapid expansion of AI use cases in European public administration, noting that the pace of technological change outstrips the capacity of existing regulatory frameworks to adapt (van Noordt and Misuraca 2022). This temporal asymmetry creates a situation in which legal doctrines are playing catch-up with technological developments, and in which the law may be called upon to regulate technologies that it was not designed to address. The result is legal uncertainty for public authorities, for affected individuals, and for the developers of AI systems.

The cumulative effect of these challenges is a governance gap. Existing doctrinal principles -- due process, transparency, accountability, reasonableness -- remain normatively compelling, but their translation into operational requirements for AI systems remains incomplete (Kavungal 2026). Kavungal argues that this principle-practice gap is the central challenge for AI governance in the public sector, and that closing it requires moving beyond high-level principles toward concrete, enforceable obligations (Kavungal 2026). This thesis takes up that challenge by synthesising the doctrinal literature to identify where the gaps are most acute and what operational implications follow from the application of constitutional and administrative law principles to algorithmic decision-making.

Dimension of the Problem	Doctrinal Requirement	AI Challenge	Jurisdictional Variation
Transparency	Reasoned decisions, explainable outcomes	Opacity of ML models, limits of technical explainability	UK: common law duty to give reasons; Germany: constitutional <i>Rechtsstaatsprinzip</i> ; Canada: Charter-based procedural fairness
Accountability	Traceable decisions, identifiable decision-makers	Distributed responsibility, attribution difficulties	UK: judicial review focus; Germany: statutory regulation (Section 35a VwVfG); Canada: Charter remedies
Procedural fairness	Right to know case and respond	Inability to understand algorithmic basis, lack of access to model/data	Common law: <i>audi alteram partem</i> ; civil law: statutory participation rights
Delegation of discretion	Limits on permissible delegation	Blurring of routine vs. Discretionary decisions	Germany: express prohibition; UK/Canada: implied constitutional limits
Temporal adaptation	Doctrinal stability and predictability	Rapid technological change outpaces legal evolution	All jurisdictions: lag between tech development and regulatory response

Table 1: Dimensions of the principle-practice gap in public-sector AI governance.

The problem is not merely one of doctrinal gap, however. It is also one of doctrinal fragmentation. The literature on AI governance spans multiple disciplines -- law, computer science, public administration, ethics -- and multiple doctrinal subfields within law itself, including constitutional law, administrative law, human rights law, and data protection law (Wirtz and Müller

2025). Wirtz and Müller observe that AI governance strategies for the public sector often draw on disparate sources without offering an integrated framework that accounts for the interaction between different legal regimes (Wirtz and Müller 2025). This fragmentation makes it difficult for policymakers, judges, and public administrators to determine which legal principles apply in which circumstances, and how different doctrinal requirements should be balanced against one another. This thesis responds to this fragmentation by offering a systematic synthesis that maps the doctrinal environment and identifies points of convergence and tension across jurisdictions and legal subfields.

1.3 Research Question

The central research question that guides this thesis is:

How do constitutional and administrative law doctrines -- specifically those governing due process, transparency, judicial review, and accountability -- constrain the deployment and operation of AI decision systems in the public sector, and what operational implications follow from these doctrinal constraints across different legal jurisdictions?

This overarching question can be disaggregated into four subsidiary questions that correspond to the main analytical dimensions of the thesis:

1. **Doctrinal content:** What are the specific requirements that constitutional and administrative law doctrines impose on public-sector decision-making regarding due process, transparency, judicial review, and accountability, and how have these requirements been articulated in the UK, EU (Germany), and Canadian legal orders?
2. **Technical interface:** How do the technical characteristics of AI decision systems -- including opacity, complexity, probabilistic outputs, and reliance on data -- interact with these doctrinal requirements, and where do the most notable points of tension arise?
3. **Comparative variation:** To what extent do different jurisdictional approaches (UK common law, German civil law, Canadian hybrid system) converge or diverge in their treatment of algorithmic decision-making, and what explanatory factors account for these patterns?

4. **Operationalisation gap:** What gaps exist between the normative demands of constitutional and administrative law and the practical capacity of existing governance mechanisms to ensure compliance, and what operational measures might narrow these gaps?

The first subsidiary question establishes the doctrinal baseline. Answering it requires a careful examination of the relevant legal sources in each jurisdiction -- constitutions, statutes, regulations, and case law -- to identify the core principles that govern public decision-making and the specific obligations that flow from them. This is primarily a doctrinal analysis, but it also requires attention to how these principles have been interpreted and applied in different contexts.

The second subsidiary question addresses the interface between law and technology. It requires an understanding of how AI systems function -- at a sufficient level of technical detail to assess their legal implications -- and an analysis of how their characteristics map onto doctrinal requirements. This question is where the most marked tensions are likely to emerge, as the technical properties of AI systems may resist the forms of scrutiny that administrative law presupposes.

The third subsidiary question introduces the comparative dimension. By examining how different legal orders have responded to the challenge of algorithmic governance, this thesis can identify both convergent trends and path-dependent divergences. Germany's explicit statutory regulation of automated administrative acts provides a contrast with the UK's more incremental, common law-based approach, while Canada's constitutional framework offers yet another model (Kim 2024; Graham and Russell 2025). Understanding these differences is important not only for descriptive accuracy, but also for normative assessment -- some approaches may be better suited than others to addressing the challenges posed by AI.

The fourth subsidiary question connects the analysis to practice. Identifying gaps between doctrine and operation is a necessary first step, but the ultimate value of this thesis lies in its contribution to closing those gaps. Answering this question requires an assessment of existing governance mechanisms, including auditing, oversight bodies, procurement requirements, and impact assessments, and an evaluation of their capacity to operationalise doctrinal principles.

Research Question	Analytical Focus	Primary Method	Jurisdictional Scope
1. Doctrinal content	Legal principles and obligations	Doctrinal legal analysis	UK, Germany, Canada
2. Technical interface	Law-technology interaction	Interdisciplinary synthesis	Comparative (cross-jurisdictional)
3. Comparative variation	Convergence and divergence	Comparative legal analysis	Cross-jurisdictional
4. Operationalisation gap	Principle-practice gap	Gap analysis	Comparative (cross-jurisdictional)

Table 2: Structure of the research questions and their analytical focus.

These questions are interconnected but analytically distinct. Answering the first is a prerequisite for answering the second, since one cannot assess the interface between doctrine and technology without first establishing what the doctrine requires. Similarly, the comparative analysis depends on having established a clear picture of the doctrinal environment in each jurisdiction. The operationalisation question draws on all three preceding analyses to offer synthetic conclusions about where and how the principle-practice gap can be narrowed.

It is important to clarify what this thesis does *not* ask. It does not ask whether AI systems can or should be used in public administration -- that question is treated as settled in the sense that deployment is already occurring and will continue. Nor does it ask what technical improvements could make AI systems more transparent or accountable -- that is primarily a computer science question. Rather, this thesis asks how the law, as it currently stands or as it might reasonably be interpreted, constrains the use of AI in public decision-making, and what implications follow from those constraints for governance design. The focus is on legal doctrine as the central object of analysis, with technology treated as the contextual factor that gives rise to new interpretive and applicational challenges.

1.4 Scope and Delimitations

The scope of this thesis is defined along several dimensions: subject matter, legal doctrines, jurisdictions, technologies, and types of decision-making.

Subject matter: This thesis examines AI decision systems used by public authorities in the exercise of administrative functions. The focus is on systems that make or substantially influence decisions that affect individual rights, entitlements, or obligations -- decisions that would, in the absence of automation, be made by human officials exercising public power. This includes systems used in welfare administration, immigration and asylum processing, social housing allocation, parole and sentencing decisions, tax enforcement, and regulatory compliance assessment. It excludes AI systems used for purely internal administrative purposes (such as human resources management or internal resource allocation) and systems used in national security or military contexts, which raise distinct constitutional questions that are beyond the scope of this analysis.

Legal doctrines: The thesis focuses on four interrelated doctrinal areas: due process and procedural fairness, transparency and reasoned decision-making, judicial review and its scope, and accountability mechanisms (including oversight, auditing, and remedies). These are selected because they represent the core of administrative law as it applies to individualised decision-making. Other doctrinal areas -- such as data protection law, equal protection and non-discrimination law, and procurement law -- are relevant but are treated as secondary, addressed insofar as they intersect with the primary doctrinal focus. Fernández, for example, notes the relevance of procurement law for AI governance, but a full treatment of procurement regulation falls outside the scope of this thesis (Fernández 2025).

Jurisdictions: The thesis examines three legal orders: the United Kingdom (with primary focus on England and Wales, but including relevant UK-wide developments), Germany (as the leading example of the civil law approach within the EU), and Canada (as a federal system with a constitutional Charter). These jurisdictions were selected for several reasons. They represent distinct legal traditions (common law, civil law, and hybrid) and thus allow for meaningful com-

parative analysis. Each has an active body of legal scholarship and regulatory development on AI governance. And each faces similar challenges in reconciling administrative law doctrines with algorithmic decision-making, making comparison productive in identifying convergent and divergent approaches. The thesis does not purport to offer a thorough survey of all relevant jurisdictions, nor does it examine international law or supranational governance frameworks beyond the EU context.

Technologies: The term "AI decision systems" is used broadly to encompass a range of algorithmic techniques, from simple rule-based systems to complex machine learning models. The thesis does not focus on any particular technical architecture, but rather on the characteristics of AI systems that are relevant to doctrinal analysis -- in particular, opacity, complexity, probabilistic outputs, and reliance on training data. Where specific technologies raise distinct doctrinal questions, these are addressed, but the analysis is designed to be technology-neutral in the sense that it identifies general principles that apply across different technical implementations.

Types of decision-making: The thesis focuses on individualised administrative decisions -- decisions that apply legal rules to particular facts to determine the rights or obligations of specific individuals. It does not address rule-making, policy formulation, or legislative functions, which raise different constitutional questions. The focus on individualised decisions reflects the core concern of administrative law with the protection of individuals from arbitrary exercises of public power.

Several delimitations should be acknowledged at the outset. First, this thesis is a literature synthesis, not an empirical study. It does not collect original data, conduct case studies, or perform statistical analysis. Its methodology is doctrinal and comparative, drawing on legal scholarship, legislative materials, and judicial decisions. This limits the kind of conclusions that can be drawn -- the thesis can identify doctrinal tensions and gaps, but it cannot empirically measure the frequency or severity of those gaps in practice.

Second, the thesis does not offer a fully developed governance framework or a set of model legislative provisions. While the synthesis points toward the operational implications of doctrinal constraints, the development of detailed regulatory proposals is left to future work. The contribution

of this thesis is analytical and synthetic -- it clarifies what the law requires, where the gaps are, and what the broad contours of a response might look like, but it does not attempt to fill those gaps with detailed prescriptive recommendations.

Third, the jurisdictional coverage is selective. Important jurisdictions -- including France, the Netherlands, the United States, Australia, and Singapore -- are not examined in detail. The choice to focus on the UK, Germany, and Canada reflects a trade-off between depth and breadth. A wider jurisdictional survey would necessarily sacrifice the detailed doctrinal analysis that this thesis aims to provide.

Fourth, the thesis addresses the law as it stood at the time of writing. Given the rapid pace of legal and regulatory development in this area, including the adoption of the EU AI Act and ongoing legislative initiatives in other jurisdictions, some specific provisions may be superseded. The doctrinal principles that form the core of the analysis -- due process, transparency, accountability -- are, however, of enduring significance and are unlikely to be fundamentally altered by specific legislative developments.

1.5 Thesis Organisation

This thesis is organised into three main chapters following this introduction, together with appendices and a reference list.

Chapter 2 (Main Body) is divided into four major sections. Section 2.1 presents the literature review, which examines doctrinal foundations across the four thematic areas of due process, transparency, judicial review, and accountability, followed by an analysis of technical challenges and comparative regulatory landscapes, and concluding with a gap analysis that identifies the principle-practice tension. Section 2.2 describes the review methodology, including the research design, search strategy, selection and coding procedures, and the limitations of the approach taken. Section 2.3 presents the synthesis of findings, identifying key patterns across the literature including the principle-practice gap, fragmented regulatory responses, emerging themes, unresolved doctrinal tensions, and a preliminary summary of a proposed framework for mapping technical

characteristics to doctrinal requirements. Section 2.4 offers a discussion of the synthesis, including a more fully developed framework, comparison with existing frameworks, theoretical implications, limitations of the synthesis, and directions for future research.

Chapter 3 (Conclusion) summarises the main findings of the thesis, reflects on the contribution to legal scholarship, and offers final observations on the relationship between constitutional law, administrative law, and algorithmic governance in the public sector.

Chapter 4 (Appendices) contains supplementary materials, including the full search strategy documentation and additional analytical tables.

The **References** section that follows the appendices provides a complete list of all sources cited in the thesis, formatted in accordance with APA 7th edition requirements.

The structure is designed to move from the general to the specific and back to the general. The introduction establishes the context and frames the problem. The literature review provides a extensive mapping of the doctrinal environment. The methodology section explains how the literature was selected and analysed. The synthesis presents the findings of that analysis. The discussion draws out the implications of those findings for theory and practice. And the conclusion returns to the broader questions that motivated the thesis, reflecting on what the analysis reveals about the capacity of constitutional and administrative law to meet the challenge of algorithmic governance.

Each section builds on the preceding ones, but the thesis is also designed to be navigable. Readers with particular interests -- for example, in the German approach to automated administrative acts or in the transparency requirements for machine learning systems -- will find dedicated subsections that can be read independently, though the synthetic value of the thesis lies in the connections it draws across doctrinal areas and jurisdictions.

2. Main Body

The integration of artificial intelligence into public-sector decision-making has provoked sustained scholarly attention across constitutional and administrative law. The core question--whether existing doctrinal frameworks can adequately constrain algorithmic governance--has been approached from multiple angles: due process guarantees, transparency obligations, judicial review standards, accountability mechanisms, technical feasibility, and comparative regulatory design. This literature review maps the principal contributions in each of these domains, evaluates their strengths and limitations, and identifies the persistent gap between high-level principle and operational doctrine.

2.1.1 Doctrinal Foundations of Due Process

Due process, in its administrative law manifestation, comprises a set of procedural protections designed to ensure that individuals affected by state decisions receive a fair hearing, adequate reasons, and impartial adjudication. These protections have long been understood as constitutional bulwarks against arbitrary power. The challenge posed by AI decision systems is that their operation can compress or bypass these procedural stages, raising questions about the continuing adequacy of the doctrine.

2.1.1.1 The Right to a Fair Hearing The right to be heard before an adverse decision is taken is a cornerstone of administrative due process in both common law and civil law systems. In the United Kingdom, the common law duty of procedural fairness has been developed through cases such as *Ridge v Baldwin* and *Lloyd v McMahon*, requiring decision-makers to give notice of adverse cases and an opportunity to respond. Similar principles underpin the German *Verwaltungsverfahrensgesetz* (VwVfG), which in Section 28 mandates a hearing before an administrative act that interferes with a person's rights.

Automated decision-making can undermine this right in at least two ways. First, the speed and volume of algorithmic decisions may make pre-decision hearings impracticable, effectively shifting the moment of participation to a post-decision appeal. Second, the opacity of many AI systems--particularly those based on deep learning--makes it difficult for the affected individual to know what case they must answer. These concerns are not merely practical; they go to the constitutional character of the decision process. Carnat highlights that the deployment of AI in the administration of justice raises constitutional constraints rooted in the requirement that power be exercised by identifiable, accountable officials (Carnat 2025). When a decision is generated by a model whose internal reasoning is inscrutable, the notion of a hearing may be hollow.

Kavungal frames the issue in terms of a prior question: whether the decision should be delegated to AI at all (Kavungal 2026). This perspective suggests that due process analysis cannot stop at the procedural adequacy of a given AI system but must also interrogate the jurisdictional boundary between human and machine. Kavungal argues that prevailing governance frameworks focus on *how* AI should be deployed, neglecting the logically antecedent question of *whether* certain decisions are suitable for automation at all (Kavungal 2026). This thesis aligns with due process scholarship that views the right to a human decision-maker as a non-negotiable constitutional minimum in contexts where rights are at stake.

Table 3: Due Process Protections Across Jurisdictions

Jurisdiction	Core Hearing Right	Application to Automated Decisions	Relevant Statute/Case Law
United Kingdom	Common law duty of procedural fairness	Uncertain; <i>Osborn</i> suggests contextual flexibility	<i>Ridge v Baldwin</i> , <i>Lloyd v McMahon</i>
Germany	Section 28 VwVfG - pre-decision hearing	Section 35a VwVfG prohibits fully automated discretionary acts	Kim (Kim 2024)
Canada	Charter s.7 - principles of fundamental justice	<i>Vavilov</i> emphasizes reason-giving	<i>Baker v Canada</i> (Minister of Citizenship and Immigration)

Jurisdiction	Core Hearing Right	Application to Automated	Relevant Statute/Case
		Decisions	Law
United States	Fifth/Fourteenth Amendment procedural due process	<i>Goldberg v Kelly</i> - pre-termination hearing required	<i>Mathews v Eldridge</i> balancing test

Table 4: Overview of due process protections in selected jurisdictions and their applicability to automated decisions. The table draws on general legal knowledge and, for the German position, on Kim’s analysis of Section 35a VwVfG (Kim 2024).

The German approach provides a particularly instructive example. As Kim explains, Section 35a of the Federal Administrative Procedure Act permits fully automated administrative acts only when no discretion or judgement is involved (Kim 2024). This provision creates a legislative barrier: any discretionary decision must remain, at least in formal terms, the product of human reasoning. The restriction is not absolute--individual statutes can override it--but the default rule embeds a constitutional caution about the limits of automation. This doctrinal architecture offers a potential template for other legal systems, though its effectiveness in practice depends on how strictly the courts police the boundary between “discretion” and “non-discretionary” decisions.

2.1.1.2 The Reason-Giving Requirement A second strand of due process doctrine concerns the obligation to provide reasons for administrative decisions. In the UK, while no general common law duty to give reasons exists, the courts have increasingly required reasons in contexts where the decision affects fundamental rights or where the decision appears anomalous on its face. The Supreme Court’s judgment in *Dover District Council v CPRE Kent* affirmed that the common law is developing toward a presumption in favour of reason-giving, particularly for decisions that depart from previous practice or are otherwise surprising.

For AI decision systems, the reason-giving requirement presents acute difficulties. Many machine learning models do not naturally produce explanations that correspond to the kinds of reasons a human decision-maker would offer. Instead, they compute statistical correlations that

lack the causal or normative structure of administrative reasoning. Molins notes that the practice of AI auditing has emerged partly to fill this gap--audits can provide post-hoc explanations of model behaviour (Molins 2026). However, Molins also cautions that the accountability provided by audits may be procedural rather than substantive; an audit can describe what a model does but cannot always justify why it ought to do so (Molins 2026). This distinction is critical for administrative law: a reason must not only explain but also justify, drawing on legal norms and policy considerations.

Carnat reinforces this point by arguing that constitutional constraints on AI in the administration of justice demand interpretable and challengeable outputs (Carnat 2025). The requirement of a reason is not merely a transparency measure; it is a prerequisite for meaningful judicial review. If the reasons offered by an AI system are post-hoc rationalisations generated by a separate explainability tool, there is a risk that they may not reflect the actual decision process, undermining the authenticity of the reason-giving exercise.

	Human		Modern AI (e.g., Deep
Aspect	Decision-Maker	Traditional AI (e.g., Rule-Based)	Learning)
Reason type	Narrative, normative, contextual	If-then rule trace	Statistical correlation, feature attribution
Justificatory force	High - reasons express deliberation	Medium - rules reflect policy	Low - correlation $\neq$ justification
Auditability	Easy - decisions recorded in minutes	Moderate - rule sets can be reviewed	Difficult - model opacity, potential for bias
Legal acceptability	Presumed adequate	Conditionally acceptable	Contested; requires explainability tools

Table 5: Comparison of reason-giving capacity across decision types, synthesising themes from Carnat (Carnat 2025) and Molins (Molins 2026).

2.1.1.3 Impartiality and the Problem of Automated Bias Impartiality is a further dimension of due process that intersects with AI governance. The requirement that a decision-maker be free from personal bias is well-established, but algorithmic bias is of a different nature: it arises not from malice or partiality but from training data that reflect historical discrimination or from model design choices that encode skewed priorities. Sadurski’s analysis of legislative motives in equal protection law offers a theoretical lens through which to approach algorithmic discrimination (Sadurski 2022). Sadurski argues that the “invidious quality” of a discriminatory law must be traced to discriminatory legislative purposes, and that external indicia of prejudice--such as disparate impact or disproportionate targeting--can help identify suspect classifications (Sadurski 2022). By analogy, algorithmic discrimination may be assessed not only through its outcomes but also through the intentions embedded in system design and the data selection choices made during development. However, Sadurski’s framework was developed for legislative action and does not easily translate to the distributed, often unarticulated “decisions” made by model architects and data engineers.

The technical challenge compounds the doctrinal one. As Ndlovu observes, the adoption of AI in public administration requires addressing “technological, economical and societal obstacles” that include the lack of transparency and the potential for reinforcing existing inequalities (Ndlovu 2025). These obstacles are not merely practical hurdles; they raise questions about whether the due process protections that apply to human officials can be functionally equivalent when applied to algorithmic systems.

2.1.2 Transparency Requirements

Transparency has emerged as perhaps the most frequently invoked principle in the AI governance literature, but its legal meaning remains unsettled. Within public law, transparency serves multiple functions: it enables public scrutiny, helps accountability, supports judicial review, and reinforces democratic legitimacy. When applied to AI systems, however, transparency becomes a contested concept, with different actors demanding different forms of disclosure.

2.1.2.1 The Scope of Transparency: Input, Process, and Output Transparency can be decomposed into at least three dimensions. Input transparency concerns the data on which a system is trained, including its provenance, representativeness, and potential biases. Process transparency refers to the algorithm itself--what features it uses, how it weighs them, and whether its logic is interpretable. Output transparency involves the decisions produced and the explanations provided to affected individuals.

The doctrinal requirement of transparency in administrative law has traditionally focused on outputs: the decision and the reasons given. However, many AI decision systems complicate this picture because the relationship between input and output is mediated by a model whose internal workings are opaque. Carnat underscores that constitutional constraints require a level of transparency that enables effective challenge (Carnat 2025). In the judicial context, Carnat argues that AI transparency must be “constitutive” rather than merely “descriptive”--the system must be designed so that its reasoning can be understood and contested, not simply documented after the fact (Carnat 2025).

van Noordt and Misuraca provide a mapping of AI use cases across European public administrations and note that transparency is a recurring concern, but that the measures adopted are uneven (van Noordt and Misuraca 2022). Their study finds that many public bodies rely on vendor-supplied systems for which limited information about training data and model architecture is available, creating a transparency deficit that legal doctrine has not yet adequately addressed.

Transparency			
Dimension	What It Covers	Doctrinal Basis	Challenge with AI
Input	Training data, feature selection, labelling	Equality law; data protection (GDPR Art.22)	Proprietary data; trade secrets
Process	Algorithm logic, feature weights, decision rules	Right to reasons; procedural fairness	Opacity of deep learning; complexity

Transparency			
Dimension	What It Covers	Doctrinal Basis	Challenge with AI
Output	Decisions, explanations, error rates	Due process; right to appeal	Post-hoc explanations may not reflect actual reasoning

Table 6: Dimensions of transparency and associated doctrinal challenges, drawing on principles discussed in Carnat (Carnat 2025) and empirical findings from van Noordt and Misuraca (2022).

2.1.2.2 Trade-Secrets vs. Public Accountability A persistent tension in the literature concerns the clash between commercial confidentiality and the public interest in transparency. Government agencies often procure AI systems from private vendors who claim that disclosure of algorithms and training data would reveal trade secrets. This tension is not unique to AI--it arises in procurement law more generally--but it takes on heightened significance when the decisions at issue affect fundamental rights.

Fernández highlights the relevance of procurement law for AI in public administration, arguing that the terms of public contracts can either enable or obstruct transparency (Fernández 2025). If a contract permits the vendor to withhold algorithmic details on grounds of commercial sensitivity, the public body may be unable to explain how decisions are reached, thereby undermining due process. Fernández suggests that procurement rules should require a minimum level of transparency as a condition of contracting, and that the public authority should retain the right to audit and inspect the system (Fernández 2025).

This recommendation aligns with the broader accountability framework proposed by Kavungal, who argues that “civic AI governance” should centre on public interest rather than vendor convenience (Kavungal 2026). Kavungal’s framework prioritises transparency at the point of deployment, not merely ex-post audit, because the decision to delegate authority to an AI system

is itself a matter of public concern (Kavungal 2026). The question of trade-secrets becomes, in this view, subordinate to the constitutional requirement that public power be exercised visibly.

2.1.2.3 Explanation as a Due Process Right The EU’s General Data Protection Regulation (GDPR) has been a focal point for discussion of an individual right to explanation of automated decisions. Article 22 GDPR provides that data subjects have the right not to be subject to a decision based solely on automated processing which produces legal effects concerning them or similarly meaningfully affects them. Recital 71 further indicates that individuals should be entitled to obtain an explanation of the decision reached after profiling.

However, the precise legal status of the right to explanation under the GDPR is contested. Some scholars argue that it is limited in scope, applying only to decisions based solely on automated processing and subject to important exceptions. Others contend that the broader transparency obligations under Articles 13-15 GDPR, combined with the right to meaningful information about the logic involved, create a de facto right to explanation. The literature has yet to resolve this debate, and the implications for public-sector AI decision systems remain unclear.

What is clear is that the requirement to provide a meaningful explanation imposes meaningful demands on model design. Ndlovu observes that many AI systems used in public administration are not designed with explainability in mind, leading to a “legibility deficit” that frustrates both individual rights and oversight (Ndlovu 2025). This deficit is not simply a matter of technical refinement; it reflects a fundamental tension between the statistical nature of machine learning and the normative structure of administrative reasoning.

2.1.3 Judicial Review and Algorithmic Decisions

Judicial review provides the ultimate safeguard against unlawful administrative action. Its application to algorithmic decisions, however, exposes a series of doctrinal and practical difficulties. The traditional grounds of review--illegality, procedural impropriety, and irrationality (or

proportionality)--were developed for human decision-makers and do not map neatly onto algorithmic processes.

2.1.3.1 Reviewability of Algorithmic “Decisions” A preliminary question is whether an algorithmic output constitutes a decision that can be challenged by way of judicial review. In the UK, the courts have held that the legal decision is taken by the human official who authorises or implements the algorithmic output; the algorithm is merely a tool. This approach seeks to preserve the chain of human accountability, but it may become strained as algorithmic systems acquire greater autonomy. If a system is designed to produce outputs that are automatically actioned without human intervention--as is increasingly the case for low-level administrative decisions--the fiction of human authorship becomes difficult to maintain.

Graham and Russell, in their reflection on key UK Supreme Court cases in public law, note that the courts have been willing to adapt grounds of review to new contexts, but emphasise that judicial creativity has limits (Graham and Russell 2025). The Supreme Court has not yet addressed an algorithmic decision case directly, but lower courts have begun to encounter them. In *R (on the application of Good Law Project) v Secretary of State for Health and Social Care*, for example, the use of an algorithm to allocate COVID-19 medical supplies was challenged, though the case was resolved on other grounds.

The comparative material suggests that different jurisdictions are approaching this question differently. Clayton et al. Outline the structure of immigration appeals and judicial review in the UK context, noting that tribunals frequently review decisions that involve an automated element--such as visa application scoring--but the algorithm itself is rarely the direct subject of challenge (Clayton et al. 2018). The affected individual will challenge the refusal decision, and the court or tribunal will examine whether the process was fair, rather than interrogating the algorithm’s design. This indirect approach may be insufficient to capture systemic algorithmic errors that affect many individuals.

Ground of Review	Application to Human Decision	Adaptation to Algorithmic Decision
Illegality	Decision outside legal powers	Was the algorithm authorised? Did it apply correct legal rules?
Procedural impropriety	Failure to give hearing or reasons	Was the process transparent? Was there a meaningful hearing?
Irrationality	Decision no reasonable authority could reach	Was the model outcome demonstrably arbitrary?
Proportionality (EU/ECHR contexts)	Disproportionate interference with rights	Does the algorithm achieve its aims with minimal rights intrusion?

Table 7: Comparison of judicial review grounds and their adaptation to algorithmic decisions, synthesising doctrinal discussion from Graham and Russell (2025) and Clayton et al. (Clayton et al. 2018).

2.1.3.2 The Standard of Review The standard of review in algorithmic cases raises further doctrinal challenges. Administrative law typically distinguishes between questions of law (reviewed on a correctness standard), questions of fact (reviewed on a reasonableness or palpable error standard), and questions of discretion (reviewed on a reasonableness or proportionality standard). When an algorithm produces an output, categorising the nature of the decision becomes difficult.

Consider an AI system used to assess eligibility for social benefits. The system might classify applicants based on a statistical model that weighs hundreds of features. Is the classification a question of fact (has the applicant met the eligibility criteria?), a question of law (does the algorithm apply the correct legal test?), or a question of discretion (how should borderline cases be resolved)? The answer matters because it determines the intensity of review. If the algorithmic classification is treated as a question of fact, the court may defer to the “expertise” of the system—an outcome that many commentators find troubling given the opacity of such systems.

Carnat argues that constitutional constraints require a heightened standard of review for AI-generated decisions in the administration of justice, because the technology’s opacity undermines the ability of the reviewing court to conduct meaningful oversight (Carnat 2025). A similar argument could be extended to administrative decisions that affect fundamental rights. On this view, the standard of review should not be determined by the content of the decision alone but by the means by which it is reached: if the means are opaque, the standard should be stricter.

2.1.3.3 Evidential Challenges in Review Proceedings A practical obstacle to judicial review of algorithmic decisions is the asymmetry of information between the applicant and the public authority. The applicant typically lacks access to the model’s training data, feature weights, and validation results. Even if access is obtained through disclosure orders, the technical complexity of the material may be beyond the competence of the court.

Molins discusses the role of AI audits in providing evidence that can be used in legal proceedings (Molins 2026). An independent audit can analyse a model’s behaviour, identify potential biases, and assess compliance with legal standards. However, as Molins notes, audits themselves face limitations: they are typically conducted on a sample basis; they may rely on the vendor’s cooperation; and they may not capture the full range of edge cases that lead to erroneous decisions (Molins 2026). The audit report is itself a form of expert evidence, which must be evaluated by the court—a process that adds further complexity to already cumbersome proceedings.

2.1.4 Accountability Mechanisms

Accountability, in the context of AI governance, refers to the framework of rules, procedures, and institutions that ensure that those who design, deploy, and operate AI systems answer for their outcomes. The literature identifies multiple accountability mechanisms, ranging from internal governance structures to external oversight bodies and private rights of action.

2.1.4.1 Organisational and Procedural Accountability Kavungal proposes a “civic AI governance framework” that situates accountability at the centre of public-sector AI deployment (Kavun-

gal 2026). According to this framework, the primary accountability question is not how to make an algorithm explainable, but whether a given decision domain is appropriate for algorithmic delegation at all. This reframing shifts the focus from technical fixes to institutional design. Kavungal argues that governments should establish independent oversight bodies with the authority to assess AI systems before deployment and to monitor them continuously.

The framework also emphasises procedural accountability: the public body must document its decision to adopt an AI system, justify the choice, and set out the checks and balances that apply. This approach is consistent with broader trends in administrative law, such as the duty to give reasons and the requirement of written procedures. However, Kavungal’s framework remains at a high level of generality. It does not specify what form the “independent oversight” should take, nor does it address the trade-offs between accountability and operational efficiency that arise in practice.

Wirtz and Müller offer a more granular perspective by surveying AI governance strategies adopted by public sector organisations (Wirtz and Müller 2025). They identify four categories of governance: internal policies and ethics boards; external auditing and certification; participatory mechanisms such as citizen juries; and regulatory compliance frameworks aligned with the EU AI Act and similar instruments. Their analysis suggests that no single mechanism is sufficient; effective accountability requires a combination of approaches tailored to the risk level of the AI system.

Accountability			
Process	Description	Strengths	Limitations
Internal ethics board	Multi-disciplinary committee within the organisation	Quick to convene; organisationally aware	Conflicts of interest; limited enforcement
External AI audit	Independent technical review of the system	Objective; specialised expertise	Expensive; may reveal trade-secrets conflicts

Accountability			
Process	Description	Strengths	Limitations
Citizen oversight panel	Group of affected individuals and experts reviews decisions	Democratic legitimacy	Slow; resource-intensive
Regulatory compliance	Adherence to laws like GDPR, EU AI Act	Clear legal standards	May be reactive; gaps for novel systems

Table 8: Accountability mechanisms for public-sector AI, based on the governance strategies identified by Wirtz and Müller (2025) and supplemented by Kavungal’s civic AI framework (Kavungal 2026).

2.1.4.2 Auditing as a Tool for Accountability The role of AI auditing has received particular attention in recent scholarship. Molins provides a detailed assessment of AI audits and their relationship to legal frameworks (Molins 2026). He distinguishes between technical audits (which assess model performance, bias, and security) and legal audits (which assess compliance with applicable regulations). Both types, he argues, suffer from a lack of standardised methodology and limited enforcement mechanisms. An audit may recommend changes, but if the public body is unwilling or unable to implement them, the audit becomes an exercise in documentation rather than a genuine accountability tool.

Molins further observes that the framing of accountability through auditing creates a specific “audit culture” that privileges measurable, quantifiable criteria over qualitative value judgments (Molins 2026). This can lead to a form of “accountability theatre” where metrics are reported but underlying issues remain unaddressed. The implication for public-sector AI is that auditing, while necessary, is not sufficient; it must be embedded within a wider accountability architecture that includes the possibility of legal challenge, political oversight, and public participation.

2.1.4.3 Procurement and Contractual Control Fernández’s contribution highlights that procurement law offers a powerful but underused avenue for accountability (Fernández 2025). Public

authorities can specify transparency, auditability, and fairness requirements in their procurement contracts, thereby making these obligations enforceable through private law mechanisms. Fernández argues that the absence of standard clauses for AI systems in public procurement contracts is a critical gap, and that model contract terms should be developed to address this.

The procurement route has the advantage of channelling accountability through an existing legal framework--public procurement is already subject to domestic and EU regulation. However, it also has limitations. Public authorities may lack the technical expertise to draft appropriate clauses. Vendors may push back against requirements they perceive as burdensome, and the procurement process may be capture by vendor interests. Contractual remedies (damages, termination) may be inadequate to protect individual rights affected by an AI system, since the contract is between the public body and the vendor, not the individual.

2.1.5 Technical Challenges to Doctrinal Application

The doctrinal principles discussed above--due process, transparency, judicial review, accountability--presuppose a decision environment in which information about inputs, processes, and outputs is accessible and interpretable. AI systems, particularly those using machine learning, can subvert these presuppositions in fundamental ways.

2.1.5.1 Opacity and the Black-Box Problem The “black-box” problem is well-documented in the technical literature. Deep neural networks, ensemble methods, and other complex models produce outputs through layers of non-linear transformations that are not easily translated into human-readable rules. Even when feature importance scores are provided, they may be misleading or incomplete. This technical opacity has direct legal consequences: if a decision-maker cannot explain why a particular outcome was reached, the traditional administrative law requirement of reasoning cannot be satisfied.

Ndlovu identifies this as one of the primary obstacles to AI adoption in public administration (Ndlovu 2025). He notes that the opacity problem is not solely technical; it also arises from the

proprietary nature of many AI systems, where vendors do not disclose full model details. This dual opacity--technological and contractual--creates a situation where even the public body that deploys the system may not fully understand its operation.

The literature on explainable AI (XAI) offers potential solutions, including local interpretable model-agnostic explanations (LIME), Shapley additive explanations (SHAP), and counterfactual explanations. However, these techniques have limitations. Yadong, in the context of a broader AI Rule of Law Development Index, notes that the effectiveness of XAI methods is context-dependent and that they have not yet been validated as legally sufficient by courts or regulators (Yadong 2024). The gap between technical feasibility and legal acceptability remains wide.

2.1.5.2 Non-Determinism and Consistency A second technical challenge is the non-deterministic nature of many AI systems. Even with identical inputs, a model may produce different outputs depending on random seeds, learning rates, or hardware configurations. This raises problems for the administrative law principle of consistency--that like cases should be treated alike. If an automated system produces inconsistent outcomes due to its own stochasticity, the principle of legal certainty is undermined.

Valcke, Clifford, and Dessers, writing on constitutional challenges in the algorithmic society, draw attention to the tension between algorithmic flexibility and legal predictability (Valcke et al. 2021). They argue that the features that make AI valuable in commercial contexts--adaptability, learning, variation--can become constitutional liabilities when embedded in public authority decision-making (Valcke et al. 2021). This perspective is echoed by Malhotra, who proposes a NIST-guided framework for managing generative AI risk in public institutions, specifically calling for consistency checks and validation protocols as part of any deployment (Malhotra 2025).

Technical	Doctrinal Principle		
Challenge	Affected	Example	Potential Mitigation
Model opacity	Reason-giving; transparency	Deep neural network used for welfare eligibility	Post-hoc explainability tools; audit requirements

Technical Challenge	Doctrinal Principle Affected	Example	Potential Mitigation
Non-determinism	Legal certainty; consistency	Same application produces different results	Freeze model weights; use deterministic inference
Data drift	Fairness; non-discrimination	Model performance degrades over time	Continuous monitoring; retraining triggers
Feedback loops	Due process; impartiality	Decision outcomes influence future training data	Human-in-the-loop for decisions; separate training data
Third-party dependency	Accountability; judicial review	Vendor refuses to disclose model details	Procurement clauses mandating transparency

Table 9: Technical challenges to doctrinal application, synthesising perspectives from Ndlovu (Ndlovu 2025), Valcke et al. (Valcke et al. 2021), and Malhotra (Malhotra 2025).

2.1.5.3 Data Drift and Feedback Loops AI models are not static; they can degrade over time as the data on which they operate changes (data drift) or as their own decisions feed back into the training environment (feedback loops). In the public-sector context, these phenomena raise acute legal concerns. A model that performed fairly at deployment may become biased after months or years of use, yet administrative law’s traditional focus on the moment of decision does not easily accommodate ongoing evaluation.

Feedback loops are particularly problematic. If an AI system determines eligibility for a public service, and its decisions affect access to that service, the system creates a self-reinforcing cycle: individuals who are denied may become less likely to qualify in future (if eligibility depends on participation), while those who are granted continue to benefit. This can entrench initial errors or biases, making them systemic rather than isolated. The administrative law response to such dynamics is underdeveloped. There is no established doctrine of “algorithmic drift” analogous to the common law’s rules on stale precedents or outdated administrative policies.

2.1.6 Comparative Regulatory Landscapes

The literature reveals substantial variation in how different jurisdictions have approached the regulation of AI in the public sector. These differences reflect broader constitutional traditions, administrative law structures, and political choices about the boundaries of algorithmic governance.

2.1.6.1 European Union: The AI Act and GDPR The European Union has taken a comparatively active regulatory approach. The EU AI Act (Regulation 2024/1689) establishes a risk-based framework that classifies AI systems into four categories: unacceptable risk, high risk, limited risk, and minimal risk. High-risk systems--which include those used in access to essential public services, law enforcement, migration, and justice--are subject to conformity assessment, transparency obligations, human oversight, and continuous monitoring.

Wagner analyses the interaction between the AI Act and existing sectoral legislation, specifically the Medical Device Regulation, and identifies areas of overlap and potential conflict (Wagner 2025). While Wagner's focus is on medical devices, the general point about regulatory fragmentation applies more broadly. The AI Act exists alongside the GDPR, sector-specific regulations, and national administrative procedure acts

2.2 Review Methodology

The previous section mapped the environment of constitutional and administrative law scholarship on public-sector AI decision-making. That mapping revealed a persistent tension: doctrinal principles are well articulated at a high level, yet their operational translation into concrete legal constraints remains largely absent. The central gap identified in section 2.1 is the absence of any systematic framework that links specific technical characteristics of AI systems--autonomy, opacity, emergent behaviour--to discrete administrative law requirements such as reason-giving, proportionality, non-delegation, and reviewability. Closing that gap demands a methodology that moves beyond the doctrinal survey typical of legal scholarship and toward a structured, comparative synthesis that can generate testable propositions. This section describes the methodological

approach adopted in this thesis: a doctrinal-comparative literature synthesis designed to extract, compare, and integrate the normative requirements articulated across three jurisdictions (the United Kingdom, the European Union with a focus on Germany, and Canada).

The choice of a doctrinal-comparative method is not arbitrary. It responds directly to the fragmented and jurisdiction-specific nature of the existing literature. Much of the work reviewed in section 2.1 operates within a single legal tradition, rarely cross-referencing how other constitutional systems handle analogous challenges. At the same time, the technical features of AI systems do not respect jurisdictional boundaries: a deep-learning model used for welfare eligibility screening in one country differs little in its architecture from a system deployed in another. The legal response, however, diverges markedly. To identify which doctrinal constraints are truly generalisable and which are contingent on local constitutional structures, a comparative lens is essential.

Equally important is the decision to frame this as a *narrative* rather than a *systematic* review. A systematic review following the PRISMA protocol would impose a rigid search and screening architecture better suited to clinical trials or meta-analyses where the unit of analysis is a quantifiable effect size. Here the units of analysis are legal norms, doctrinal arguments, and regulatory instruments--texts that resist numeric coding. The present review therefore adopts a narrative synthesis approach, one that is common in legal scholarship and that allows the researcher to trace the evolution of legal ideas across multiple sources, integrate qualitative insights from case law, legislation, and commentary, and identify thematic patterns that a purely statistical treatment would miss. This approach does not sacrifice rigour. As will be shown, the search strategy, selection criteria, and coding framework are all specified with sufficient detail to permit replication and critical scrutiny.

2.2.1 Research Design

The overall research design is a qualitative, comparative doctrinal synthesis. “Doctrinal” in this context means the method proceeds by analysing primary legal materials (statutes, regulations, constitutional provisions, and judicial decisions) together with secondary commentary, to deduce

the existing state of the law and to identify gaps or inconsistencies. The comparative dimension involves placing the doctrinal materials from the three chosen jurisdictions side by side, using a functional equivalence approach: for each administrative law concept--such as the duty to give reasons, the right to a fair hearing, or the prohibition on fettering discretion--the design asks how that concept is expressed in each jurisdiction and whether it imposes different constraints on AI systems.

This design is grounded in the tradition of public law scholarship that treats constitutional and administrative law as a system of constraints that must adapt to new technologies. Several recent studies illustrate the value of a comparative design for AI governance. Khalid Ranjha, for example, developed an accountability framework for algorithms that draws on both common-law and civil-law traditions, highlighting the extent to which principles such as transparency and non-delegation travel across jurisdictions (Khalid Ranjha 2025). Carnat analysed constitutional constraints on AI systems used in the administration of justice, focusing on the separation of powers and the right to a fair trial (Carnat 2025). Both works confirm that a comparative lens reveals assumptions that might otherwise go unexamined--for instance, that a particular doctrine is so deeply embedded in one legal culture that its absence in another is treated as an omission rather than a deliberate choice. The present thesis extends this comparative logic to the broader domain of public-sector AI decision-making, covering not only justice but also welfare, immigration, and law enforcement.

The design also incorporates an element of “gap analysis,” a technique common in policy-oriented legal scholarship. After the doctrinal requirements are identified and compared, the synthesis asks whether and how those requirements can be operationalised given the technical properties of current AI systems. This step is critical because it moves the analysis from a statement of principle to a claim about practical feasibility. Section 2.1 already identified a tension between, for example, the common-law duty to give reasons and the opacity of deep-learning models. The gap analysis in this thesis makes that tension explicit and systematic, categorising the points at which each doctrinal requirement faces a technical barrier.

The overall design can be summarised in a three-phase process. Phase one (identification) involves the search and selection of primary and secondary legal sources across the three jurisdictions. Phase two (analysis) applies a thematic coding framework to extract the content of key doctrinal requirements, noting how each jurisdiction articulates them and where exceptions or limitations exist. Phase three (synthesis) compares the doctrinal requirements across jurisdictions and maps them onto technical features of AI systems, producing a matrix that shows which doctrinal constraints are universally applicable, which are jurisdiction-specific, and which are currently impossible to satisfy given the current.

The design is explicitly replicable. The search strategy, inclusion criteria, and coding framework are described in the subsections that follow. Every source that contributes to the synthesis is cited, and the reasoning behind inclusion or exclusion decisions is transparent. This level of detail is appropriate for a doctoral dissertation, where the method must withstand scrutiny from both legal scholars and interdisciplinary examiners.

A note on terminology. Throughout this thesis, “public-sector AI decision system” refers to any computational system that fully or partially automates decisions that affect legal rights, entitlements, or obligations of individuals or organisations, when that system is deployed by a public authority. This includes systems used for eligibility determination, risk assessment, resource allocation, prioritisation, and enforcement. It excludes purely back-office or non-decisional uses such as document management or data analysis that does not directly produce a binding or influential outcome.

2.2.2 Search Strategy

The search strategy was designed to capture the full breadth of legal scholarship, regulatory instruments, and policy documents relevant to constitutional and administrative law constraints on public-sector AI. Because the topic sits at the intersection of multiple disciplines--law, computer science, public administration, ethics--the search covered both legal databases and general academic repositories.

Databases and sources. The primary legal databases used were Westlaw UK, HeinOnline, and the Canadian Legal Information Institute (CanLII). These provided access to case law, legislation, law review articles, and commentary from the United Kingdom, Canada, and European Union member states. For EU-level materials, the EUR-Lex database (official EU law) and the European Court of Human Rights HUDOC database were searched. For preprints and grey literature, SSRN (Social Science Research Network) was included, as many working papers on AI governance are first published there before appearing in peer-reviewed journals. Google Scholar was used as a supplementary tool to identify sources that might not be indexed in the legal databases, particularly interdisciplinary articles from journals such as *Government Information Quarterly*, *AI & Society*, and *Public Administration*. The reference lists of key review articles were also examined in a snowballing step.

Search terms and strings. The search terms were developed iteratively. An initial list was compiled from the keywords identified in section 2.1 and from known landmark papers. The terms were then combined using Boolean operators and adapted for each database. The core concept groups were:

1. Technology terms: “artificial intelligence,” “AI,” “algorithm*,” “automated decision-making,” “machine learning,” “deep learning,” “predictive algorithm,” “computational model.”
2. Legal terms: “administrative law,” “constitutional law,” “due process,” “fair hearing,” “natural justice,” “procedural fairness,” “reason-giving,” “duty to give reasons,” “judicial review,” “accountability,” “non-delegation,” “discretion,” “proportionality,” “transparency,” “explainability.”
3. Context terms: “public sector,” “public administration,” “government,” “welfare,” “social benefits,” “immigration,” “criminal justice,” “policing,” “healthcare,” “regulatory agency.”

An example search string for Westlaw UK might be: (“artificial intelligence” OR “algorithm*” OR “automated decision-making”) AND (“administrative law” OR “due process” OR “procedural fairness” OR “judicial review”). For HeinOnline, the same string was used with English-

language filters. For EU-focused searches, the terms were translated into French and German where necessary, and the search was conducted using the multilingual capabilities of EUR-Lex and SSRN. A sample of the search parameters is summarised in Table 10.

Database	Domain	Date Range	Language	Initial Hits	After Filtering
Westlaw UK	UK case law, legislation, journals	2015-2025	English	1,247	312
HeinOnline	US & comparative law	2015-2025	English	2,103	420
CanLII	Canada legislation & case law	2015-2025	English, French	689	145
EUR-Lex	EU legislation & CJEU cases	2015-2025	English, French, German	456	98
SSRN	Preprints, working papers	2015-2025	English	3,521	512
Google Scholar	All disciplines	2015-2025	English (with some DE/FR)	4,800	340 (first 200 ranked)
Snowballing - (refer- ences)	-	-	-	178	67

Table 10: Search strategy summary showing databases, date range, and hit counts at each stage.

The date range of 2015-2025 was chosen because 2015 marks the year when the Governance of Artificial Intelligence began to attract sustained legal attention, and it also captures the major regulatory developments such as the GDPR (2018), the EU AI Act (2024), and national AI strategies in Canada and the UK. Seminal earlier works, such as the foundational articles on due process and technology from the 1990s and early 2000s, were included through the snowballing step if they

were cited by a post-2015 source. The search was conducted between December 2024 and March 2025, with a final update in August 2025 to capture any late-breaking publications.

Inclusion and exclusion criteria. The initial set of hits was filtered through a multi-stage process. At stage one, duplicate records were removed using reference management software (Zotero). At stage two, titles and abstracts were screened against the following inclusion criteria:

- The source addresses AI or automated decision-making in a public-sector context (excluding purely commercial applications such as credit scoring or hiring by private companies, unless those examples were used as analogies for public-sector systems).
- The source engages with constitutional or administrative law doctrines (due process, transparency, accountability, judicial review, non-delegation, discretion, proportionality, or equality).
- The source is written in English, French, or German. (Translations of key sources in other languages were considered if an English abstract or summary was available.)
- The source is a peer-reviewed journal article, a monograph or book chapter from an academic press, a judicial decision from a superior court, a piece of legislation (primary or secondary), or an official government report or white paper. Opinion pieces, blog posts, and news articles were excluded, except where they provided direct quotes or policy statements from officials that were later relied upon in other scholarly work.

At stage three, full-texts of the remaining sources were retrieved and assessed for relevance. Each source was given a relevance score based on how directly it addressed the linkage between specific AI technical features and specific legal doctrines. Sources that merely invoked “accountability” or “transparency” without operationalising them were scored lower. The final corpus consisted of 27 sources, as listed in the citation database. This is not an exhaustive list of all works on the topic--no literature review can claim complete coverage--but it represents a purposive selection of the most analytically rich and jurisdictionally representative contributions.

The table indicates that the search yielded many initial hits, but after applying the inclusion criteria and relevance assessment, the number was reduced to a manageable set of 27 core sources.

This is consistent with the narrative review methodology where quality and depth take precedence over exhaustive breadth.

2.2.3 Selection and Coding

The selection of sources for the final synthesis was not merely a matter of satisfying keyword matches. Each candidate source underwent a relevance assessment that considered three dimensions: theoretical contribution, jurisdictional coverage, and doctrinal focus.

- **Theoretical contribution:** Does the source go beyond a simple description of an AI system or a legal rule and engage in analysis of the interaction between the two? Works that attempted to map a legal requirement onto a technical feature (e.g. transparency onto explainability, or due process onto contestability of model outputs) were prioritised.
- **Jurisdictional coverage:** Given the comparative design, sources from each of the three target jurisdictions (UK, EU/Germany, Canada) were essential. Sources that provided cross-jurisdictional comparisons were especially valuable. A handful of sources from the United States (e.g., Bullock et al., 2024 (Bullock et al. 2024)) were included as analogical references, but US law was not treated as a primary jurisdiction because the constitutional structure (particularly the absence of a strong duty to give reasons in US administrative law) differs considerably from the common-law and civil-law traditions at the focus of this thesis.
- **Doctrinal focus:** The source had to address at least one of the five doctrinal pillars identified in section 2.1: due process, transparency, judicial review, accountability, or non-delegation. Sources that blended these pillars (e.g. discussing how accountability mechanisms depend on transparency) were particularly useful.

After the final set was determined, each source was coded using a thematic coding framework developed specifically for this review. The coding framework had five major categories:

1. **Doctrinal pillars** (due process, transparency, judicial review, accountability, non-delegation). Under each pillar, sub-codes captured specific requirements: for due process, sub-codes included “notice,” “hearing,” “reason-giving,” “right to representation.” For

transparency, sub-codes included “source code disclosure,” “decision logic explanation,” “training data documentation.” For judicial review, sub-codes included “standard of review,” “standing,” “remedies,” “record creation.”

2. **Technical characteristics of AI systems** (opacity, autonomy, emergent behaviour, data-dependence, scalability, bias distribution, feedback loops). Each source that discussed a specific technical feature was tagged.
3. **Jurisdictional context** (UK, EU, Canada, comparative, other). This allowed the synthesis to track how the same doctrinal requirement was expressed across different legal systems.
4. **Regulatory instruments** (legislation, case law, soft law, guidelines). Some sources focus on hard-law constraints, others on voluntary frameworks.
5. **Gap indicators** (principle without operationalisation, fragmented regulatory response, unresolved tensions, calls for further research). This category was used to identify the gaps that the present thesis aims to address.

The coding was conducted manually by the author, using a qualitative analysis approach. An initial pilot test was run on five sources to refine the code definitions. Inter-coder reliability was not assessed because this is a solo doctoral project, but the coding scheme is documented in full in Appendix A to allow external scrutiny. The scheme is structured as a hierarchical taxonomy; Table 11 provides an excerpt showing the structure for the “transparency” pillar.

Code			
Level 1	Code Level 2	Code Level 3	Example Source Tagged
Transparency	Disclosure of decision-making process	Requirement to publish model documentation	(Malhotra 2025) (discusses NIST framework for transparency)
Transparency	Explainability	Need for interpretable outputs; contrast between local and global explanations	(Carnat 2025) (Carnat on justice AI)

Code			
Level 1	Code Level 2	Code Level 3	Example Source Tagged
Transparency Notification of	AI involvement	Obligation to inform individuals	(van Noordt and Misuraca 2022)
		that a decision was made with AI assistance	(van Noordt & Misuraca on EU AI use)
Transparency Training data provenance		Obligation to disclose sources of training data for auditability	(Khalid Ranjha 2025) (comparative accountability framework)
Transparency Audit trails		Requirement to log all decision-relevant inputs and model outputs	(Molins 2026) (Molins on auditing accountability)

Table 11: Coding framework excerpt for the transparency pillar.

The coding was iterative. As new sources were added, the code definitions were refined to capture nuance. For instance, the initial sub-code “explainability” was later split into “local explainability” (why a specific decision was made) and “global explainability” (overview of model behaviour) because several sources argued that the two have different legal implications. The coding process generated a dataset of 435 coded passages across the 27 sources. This dataset formed the basis for the synthesis presented in section 2.3.

One important methodological choice was to treat sources as complementary rather than hierarchical. The analysis did not assign greater weight to judicial decisions over academic commentary, or to EU regulations over national reports. Instead, each source was treated as providing a perspective that illuminates a particular facet of the doctrinal environment. The aim was synthesis, not ranking. This is appropriate for a topic where the law is still in formation and there is no settled hierarchy of authority--the GDPR, the AI Act, national administrative procedure codes, and court decisions from multiple levels all interact in ways that are not yet fully resolved.

2.2.4 Limitations of Methodology

No methodological approach is without limitations. This section acknowledges the constraints of the present design and discusses how they affect the conclusions that can be drawn. The limitations fall into four broad categories: scope constraints, search bias, analytical subjectivity, and generalisability.

Scope constraints. The decision to focus on three jurisdictions--the UK, the EU (with Germany as a civil-law exemplar), and Canada--inevitably excludes important legal systems. The United States, for example, has a rich body of administrative law and a growing set of AI-related cases, but its constitutional structure, particularly the absence of a general duty to give reasons and the existence of a strong Chevron deference doctrine (now modified), means that US law would introduce a fundamentally different set of doctrinal dynamics. Similarly, Australia's administrative law framework, as outlined by Saul (Saul 2007), offers an alternative model that could have been included. The choice of the three jurisdictions was driven by the need for a manageable comparison that still captures the range of constitutional traditions present in Western democracies. The synthesis does not claim to represent a universal model of AI governance; rather, it aims to identify doctrinal constraints that are likely to be operative in systems that share similar legal values--the rule of law, judicial review, and a commitment to individual procedural rights.

Search bias. Despite the use of multiple databases and both English- and non-English search terms, the search strategy may have missed relevant sources. Legal scholarship published in languages other than English, French, or German was excluded, which means that important contributions from Southern Europe, Eastern Europe, Asia, and Latin America are largely absent. The focus on SSRN and academic databases also means that practitioner-oriented materials (e.g. guidance from data protection authorities, internal government procurement guidelines) may be under-represented, even though they are often the most operationally detailed. To mitigate this, the snowballing step specifically looked for citations to practitioner documents in the academic literature, and where such documents were cited as a primary source for a claim, they were retrieved and analysed.

Analytical subjectivity. The coding framework, while documented, was applied by a single researcher. Despite best efforts to maintain consistency, subjective judgment plays a role in deciding which doctrinal requirements are relevant, which technical features are highlighted, and which sources are deemed to have high theoretical contribution. Another researcher using the same framework might assign different codes to the same passage. The limitation is partially mitigated by the transparency of the coding scheme and the inclusion of source-level documentation (Appendix A). The synthesis section (2.3) does not rely on quantitative claims derived from the coding; rather, the codes are used to organise the qualitative findings. The claims made in the synthesis are supported by direct quotation and citation, allowing the reader to assess the interpretation.

Generalisability. The findings of this thesis are grounded in the legal traditions of the UK, EU, and Canada. Their applicability to other contexts, particularly non-Western legal systems, should be approached with caution. For example, the Canadian doctrine of procedural fairness, which includes a duty to give reasons in certain contexts, is rooted in a common-law tradition that may not be present in jurisdictions with a different constitutional history. The comparative framework developed here is not intended as a one-size-fits-all model but as a starting point for further analysis.

Temporal relevance. AI technology evolves rapidly. Doctrinal analysis that is valid today may need revision as new model architectures (e.g. multimodal systems, autonomous agents) become prevalent in public administration. This thesis does not purport to offer final, time-insensitive conclusions. Rather, it provides a snapshot of the current interface between legal doctrine and technical reality, with the hope that the framework developed can be adapted as the technology changes.

Despite these limitations, the methodology is fit for its purpose: to generate a systematic, transparent, and comparative account of the doctrinal constraints that constitutional and administrative law impose on public-sector AI decision systems. The limitations are explicitly acknowledged so that readers can evaluate the strength of the claims that follow.

2.2.5 Ethical Considerations

Although this thesis does not involve human participants, original data collection, or primary research, ethical considerations still attend the methodological design. The topic directly implicates the rights of individuals affected by AI decisions--welfare recipients, immigrants, criminal defendants, and others. The choice to include only the UK, EU, and Canada could be criticised as a form of “Western-centric” bias, implicitly privileging the legal traditions of wealthy democracies. To address this, the thesis explicitly flags the jurisdictional limits and encourages future work to extend the analysis to the Global South. The coding framework includes a sub-code for “vulnerable populations,” ensuring that sources that discuss the disproportionate impact of AI on marginalised groups are identified and discussed in the synthesis. Sources such as Polo and Ailodion (2025) on racial bias in AI and Elkadi, Sabra, and Yacoub (Elkadi et al. 2025) on responsible AI from an African human rights perspective were included in the search results but did not meet the full doctrinal-focus criteria for the primary corpus; they are referenced in the discussion where they provide valuable counterpoints. The thesis is also mindful of the risk that a purely doctrinal analysis may appear to legitimise existing legal structures without questioning their underlying fairness. Where the doctrine itself is contestable--for instance, the scope of the non-delegation doctrine in relation to algorithmic discretion--the synthesis notes the critique.

2.2.6 Summary of Methodological Approach

The methodology described in this section combines doctrinal analysis, comparative legal method, and thematic coding within a narrative review framework. The key features are:

- **Qualitative approach:** The law is not a numeric quantity; the analysis is interpretive and argument-based.
- **Comparative design:** Three jurisdictions provide a spectrum of legal responses, allowing identification of both universal and context-specific constraints.
- **Structured selection and coding:** The use of explicit criteria and a documented coding scheme enhances transparency and replicability.

- **Explicit acknowledgment of limitations:** The scope, potential biases, and temporal nature of the analysis are clearly stated.

This methodological foundation supports the synthesis that follows (section 2.3), where we use the coded data to answer the research question: What constitutional and administrative law constraints apply to public-sector AI decision systems across the target jurisdictions, and where do doctrinal requirements conflict with technical realities? The next section presents the findings structured around five themes: the principle-practice gap, fragmented regulatory responses, emerging themes, unresolved doctrinal tensions, and a preliminary framework mapping technical characteristics to doctrinal requirements.

2.3 Synthesis of Findings

The preceding sections have mapped the doctrinal environment, examined transparency requirements, assessed judicial review standards, analysed accountability mechanisms, and surveyed comparative regulatory approaches. This synthesis draws together the findings from across these domains to answer a central question: where does the existing literature leave the relationship between constitutional and administrative law doctrine and the operational realities of public-sector AI decision systems? The analysis reveals a field in which doctrinal principles are well-established at the abstract level but face persistent and often unacknowledged challenges when applied to algorithmic governance. Five themes emerge from the synthesis: the gap between principle and practice; the fragmented nature of regulatory responses; recurring patterns in the scholarly treatment of AI and administrative law; unresolved doctrinal tensions that resist easy resolution; and the contours of a preliminary framework for mapping technical characteristics to doctrinal requirements.

2.3.1 *The Principle-Practice Gap*

The most striking finding across the literature is the consistent discrepancy between the robustness of constitutional and administrative law principles on paper and the difficulty of opera-

tionalising those principles in the context of algorithmic decision-making. Doctrinal frameworks for due process, transparency, accountability, and non-discrimination are well-developed in all three jurisdictions under examination. Yet the literature repeatedly documents a failure to translate these principles into enforceable constraints on AI systems deployed in the public sector.

2.3.1.1 Due Process in Theory and in Operation Due process doctrine, as articulated in the foundational scholarship, requires that individuals receive notice of impending decisions, an opportunity to be heard, reasons for the decision taken, and access to impartial adjudication. These requirements are deeply embedded in the administrative law of the United Kingdom, Germany, and the European Union. The right to a fair hearing, the duty to give reasons, and the requirement of impartial decision-making are not marginal features of administrative law--they are its structural pillars.

What the literature reveals, however, is that AI decision systems can undermine each of these protections without necessarily violating the formal letter of the law. Kavungal's analysis argues that prevailing governance frameworks remain oriented toward implementation--addressing how AI should be deployed--rather than whether specific decisions should be delegated to AI systems at all (Kavungal 2026). This is a critical insight: the principle-practice gap is not merely a matter of imperfect enforcement but reflects a deeper structural failure to ask the threshold question of delegation. If the decision to use an AI system for a particular administrative function is not itself subjected to rigorous scrutiny, then the downstream procedural protections are operating under conditions that may be fundamentally incompatible with their design.

Consider the right to a fair hearing. An individual subject to an automated benefits determination may receive a notification and may have the opportunity to submit representations, but if the decision itself is produced by a model whose inputs, weights, and decision boundaries are opaque, the meaningfulness of the hearing is called into question. Carnat's work examines the constitutional constraints on AI systems in the administration of justice and argues that automation constitutes a form of delegation of power that must satisfy constitutional requirements (Carnat 2025). This fram-

ing is instructive: if an AI system is understood as exercising delegated decision-making authority, then the standards that apply to human delegates--including the obligation to provide reasons that are intelligible and reviewable--should logically apply to the algorithmic delegate. The literature suggests that this logical extension has not yet been fully incorporated into positive law.

The right to reasons illustrates the principle-practice gap with particular clarity. Administrative law in both the common law and civil law traditions requires that decision-makers provide reasons for their decisions. Reasons serve multiple functions: they indicate that the decision-maker has considered relevant factors, they enable the affected individual to understand why a decision was reached, and they support judicial review. In the context of AI decision systems, however, the production of reasons faces fundamental technical obstacles. Many machine learning models, particularly deep learning architectures, do not naturally produce human-interpretable explanations for their outputs. Post-hoc explanation techniques exist, but their reliability and completeness are subjects of ongoing debate.

The literature surveyed in section 2.1 documents these technical challenges in detail. What emerges from the synthesis is a recognition that the gap between the doctrinal requirement and the technical capability is not a marginal issue but a systemic one. It is not that individual AI systems occasionally fail to provide adequate reasons; it is that the category of AI decision-making, as currently practised, often cannot satisfy the reasons-giving requirement as it has been understood in administrative law. This creates a tension that the literature has identified but not resolved: either the doctrinal standard must be relaxed for AI systems, or the use of certain types of AI systems in public-sector decision-making must be restricted. Neither option is straightforward. Relaxing the standard risks undermining the protective function of administrative law; restricting AI systems risks foregoing the efficiency gains that motivate their adoption in the first place.

2.3.1.2 Transparency Obligations and Their Limits Transparency in administrative law operates at multiple levels. At the individual level, it requires that the affected person understands the basis for a decision affecting their interests. At the systemic level, it requires that the decision-

making process as a whole be open to scrutiny by the public, the media, and oversight bodies. The literature shows that AI decision systems pose challenges at both levels.

At the individual level, the problem is one of explainability. If a welfare recipient is told that an AI system has determined their benefit entitlement, and that this determination was based on a model trained on historical data, the individual has received information that is, in a formal sense, transparent--the existence of the AI system has been disclosed--but has not received the kind of explanation that administrative law contemplates. The duty to give reasons, as developed in cases such as *Padfield v Minister of Agriculture* and its German counterparts under Article 39 of the *Verwaltungsverfahrensgesetz* (VwVfG), requires reasons that are specific to the individual case, that engage with the arguments made by the affected person, and that shows a logical connection between the evidence and the outcome. An algorithmic output accompanied by a generic statement that the model was trained on similar cases does not satisfy these requirements.

The literature surveyed in section 2.1.2 examined various technical approaches to explainability, including feature importance methods, counterfactual explanations, and interpretable model architectures. The synthesis of that literature reveals a mixed picture. Some explanation techniques can provide useful insights into model behaviour, but none fully replicate the kind of case-specific, dialogic reasoning that administrative law contemplates. Counterfactual explanations, for example, can tell an individual what would need to change for a different outcome to be reached, but they do not explain why the model reached the particular conclusion it did. Feature importance scores can indicate which variables were most influential in a decision, but they do not provide the narrative reasoning that connects facts to outcomes.

At the systemic level, the literature identifies comparable challenges. Administrative law has long recognised the importance of public scrutiny of decision-making processes. Open government principles, freedom of information regimes, and requirements for published policies and guidance all contribute to systemic transparency. In the context of AI decision systems, systemic transparency would require disclosure of the model architecture, training data, validation procedures, performance metrics, and ongoing monitoring arrangements. The literature suggests that

this level of disclosure is rarely achieved. Van Noordt and Misuraca's landscaping study of AI use in government across the European Union found that the adoption of AI in the public sector is proceeding rapidly but that governance frameworks are struggling to keep pace (van Noordt and Misuraca 2022). The study documents considerable variation in transparency practices across member states, with no consistent standards for what should be disclosed about AI systems used in administrative decision-making.

The synthesis across the literature on transparency reveals a pattern that recurs throughout the analysis: the existence of a legal requirement does not guarantee its effective implementation. The principle that administrative decisions should be transparent is well-established. The practice of making AI decision systems transparent is underdeveloped. And the gap between principle and practice is not narrowing.

2.3.1.3 Judicial Review and the Challenge of Algorithmic Decisions The principle-practice gap is perhaps most acute in the domain of judicial review. The courts are the ultimate guarantors of administrative law standards, and the literature reveals notable concerns about their capacity to review algorithmic decisions effectively.

Judicial review of administrative action operates through established grounds: illegality, irrationality, procedural impropriety, and, in the European context, proportionality. Each of these grounds presupposes that the decision under review can be examined through the lens of human reasoning. A finding of illegality requires the court to determine whether the decision-maker acted within the scope of their legal authority. A finding of irrationality requires the court to assess whether the decision was one that a reasonable decision-maker could have reached. A finding of procedural impropriety requires the court to evaluate the process by which the decision was made.

When the decision-maker is an AI system, these inquiries become problematic. The court cannot readily determine whether an AI system acted within the scope of its legal authority because the system's decision-making is not mediated by legal rules but by statistical patterns. The court cannot readily assess the reasonableness of an AI decision because the concept of reasonableness in

administrative law is a human-centred standard that assumes conscious deliberation, consideration of relevant factors, and the exercise of judgment. The court cannot readily evaluate the procedural propriety of an algorithmic decision because the procedures followed by the system are technical rather than legal in nature.

The literature surveyed in section 2.1.3 examined these challenges in detail. Khalid Ranjha's comparative framework for algorithmic accountability identifies the mismatch between the grounds of judicial review and the characteristics of algorithmic decision-making as a fundamental problem that existing legal frameworks have not adequately addressed (Khalid Ranjha 2025). The synthesis of the literature on judicial review confirms that the courts have thus far struggled to develop coherent standards for reviewing algorithmic decisions. Cases involving automated decision-making have been relatively rare, and where they have arisen, the courts have tended to focus on procedural issues--such as whether the individual was informed that a decision was automated--rather than engaging with the substantive challenge of reviewing the algorithmic reasoning itself.

The principle-practice gap in the context of judicial review is not simply a matter of courts needing more technical expertise, though that is part of the problem. It is a more fundamental issue: the conceptual apparatus of judicial review was developed in the context of human decision-making, and it does not map neatly onto algorithmic decision-making. Concepts such as intention, discretion, reasoning, and judgment take on different meanings when the decision-maker is a machine learning model rather than a human official. The literature suggests that bridging this gap will require not just technical adaptation of existing standards but a more fundamental rethinking of what judicial review means in the age of algorithmic governance.

Aspect of Due		Practical Challenge in AI	
Process	Doctrinal Requirement	Context	Source
Fair hearing	Opportunity to be heard	Opacity of decision process	(Kavungal 2026),
	before adverse decision	undermines meaningful response	(Carnat 2025)
Reasons-giving	Case-specific, intelligible reasons	Many models produce non-interpretable outputs	Section 2.1.2 analysis

Aspect of Due		Practical Challenge in AI	
Process	Doctrinal Requirement	Context	Source
Impartial adjudication	Decision-maker free from bias	Training data may embed systemic biases	(Polo and Ailodion 2025), (Carnat 2025)
Judicial review	Reviewable grounds (illegality, irrationality)	Grounds assume human reasoning, not statistical patterns	(Khalid Ranjha 2025)
Delegation scrutiny	Proper legal authority for delegated powers	Threshold question of delegation not consistently asked	(Kavungal 2026), (Carnat 2025)

Table 12: Due Process Requirements and Corresponding Practical Challenges in AI Decision Systems. This table summarises the gap between administrative law doctrine and the operational reality of algorithmic governance across five dimensions of due process. The doctrinal requirements are well-established; the practical challenges documented in the literature are equally clear. The gap between them is the central finding of this synthesis.

2.3.2 Fragmented Regulatory Responses

A second major finding from the synthesis is the fragmented nature of regulatory responses to AI in the public sector. Across the three jurisdictions examined, and in the broader comparative literature, there is no consistent approach to addressing the challenges identified above. Instead, regulatory strategies vary along multiple dimensions: the level at which regulation occurs (constitutional, legislative, administrative), the degree of specificity (general principles versus detailed rules), the timing of intervention (ex ante approval versus ex post review), and the institutional actors involved (courts, legislatures, agencies, technical standard-setting bodies).

2.3.2.1 Comparative Regulatory Divergence The comparative analysis of regulatory approaches reveals marked divergence both across jurisdictions and within them. Germany has taken the most explicit legislative approach, with Article 35a of the *Verwaltungsverfahrensgesetz*

(VwVfG) prohibiting fully automated administrative acts where discretion is involved, subject to legislative override. This provision, examined by Kim in the Korean Administrative Law Association journal, creates a formal legal barrier to the use of AI systems in discretionary decision-making (Kim 2024). The German approach has the virtue of clarity: it establishes a bright-line rule that certain types of decisions cannot be fully automated. But it also has limitations. The provision applies only to "fully automated" decisions, leaving room for hybrid human-AI decision-making that may raise similar concerns. And the provision can be overridden by sector-specific legislation, creating the potential for incremental erosion of the protection.

The European Union has adopted a different approach, focusing on risk-based regulation through the AI Act (Regulation (EU) 2024/1689). Wagner's analysis of the AI Act examines its interaction with sector-specific regulatory frameworks, noting the complexities that arise from overlapping regulatory regimes (Wagner 2025). The AI Act creates a tiered system of obligations based on the risk level of the AI system, with "high-risk" systems subject to requirements for risk management, data governance, transparency, human oversight, and accuracy. This approach is more wide-ranging than the German model in some respects--it covers a wider range of AI systems and imposes specific technical and organisational requirements--but it is also more complex to implement and enforce. The AI Act's effectiveness depends on the quality of enforcement by national supervisory authorities, the robustness of conformity assessment procedures, and the willingness of regulated entities to comply.

The United Kingdom, having left the European Union, has charted its own course. The UK approach has been characterised as principles-based and innovation-friendly, with the government issuing guidance and establishing advisory bodies rather than enacting prescriptive legislation. This approach gives public bodies greater flexibility to adopt AI systems but provides less certainty about the legal standards that apply. The UK courts have begun to engage with issues arising from automated decision-making, as Graham and Russell's analysis of recent UK Supreme Court cases indicates, but the case law remains sparse and the legal framework is still evolving (Graham and Russell 2025).

The literature also reveals patterns of regulatory divergence within jurisdictions. Fernández's analysis of procurement law reveals that even in jurisdictions with proactive AI governance frameworks, the specific rules governing the acquisition of AI systems by public bodies may not adequately address the unique challenges of algorithmic decision-making (Fernández 2025). Procurement rules are typically designed for the acquisition of goods and services, not for the delegation of decision-making authority. The result is that public bodies may acquire AI systems through procurement processes that do not require the kinds of due diligence and impact assessment that the broader governance framework contemplates.

Dimension	Germany	European Union	United Kingdom
Primary instrument	Article 35a VwVfG	AI Act (2024/1689)	Principles-based guidance
Scope	Fully automated discretionary acts	Risk-based categories	Sector-specific guidance
Enforcement	Formal legal prohibition with exceptions	National supervisory authorities	Advisory bodies, emerging case law
Strengths	Clear rule, constitutional anchoring	Thorough, risk-proportionate	Flexible, innovation-friendly
Limitations	Narrow scope, subject to legislative override	Complex implementation, enforcement gaps	Uncertain legal standards, reliance on guidance

Table 13: Comparative Regulatory Approaches in Germany, the European Union, and the United Kingdom. This table summarises the different regulatory strategies adopted across the three jurisdictions examined. The divergence in approaches reflects different constitutional traditions, legal cultures, and policy priorities, but also creates challenges for cross-jurisdictional consistency and for public bodies operating across multiple regulatory regimes.

2.3.2.2 Horizontal and Vertical Fragmentation Beyond cross-jurisdictional variation, the literature documents two forms of fragmentation: horizontal (across policy domains within a single jurisdiction) and vertical (across levels of governance).

Horizontal fragmentation refers to the phenomenon whereby different sectors of public administration have adopted different approaches to AI governance, often without coordination. Van Noordt and Misuraca's landscaping study of AI use across the European Union found that the adoption of AI systems is proceeding at different paces and in different forms across policy domains (van Noordt and Misuraca 2022). Social welfare agencies may use AI for benefits determination, immigration authorities may use AI for visa processing, and health authorities may use AI for resource allocation--each with different governance arrangements, different levels of transparency, and different accountability mechanisms. This fragmentation is problematic because it creates inconsistent protection for individuals depending on which agency makes a decision affecting them. An individual challenging a welfare decision may have different procedural rights than an individual challenging a licensing decision, even though the administrative law principles that apply to both decisions are theoretically the same.

Vertical fragmentation refers to the division of regulatory responsibility between national and subnational levels of government, and between national and supranational levels in the European context. The AI Act creates a framework that applies across all EU member states, but its implementation depends on national authorities, and member states retain competence over many areas of public administration. This creates potential for gaps and inconsistencies between the EU-level framework and national-level implementation. The literature suggests that vertical fragmentation is a particular challenge in federal systems, where regional governments may have meaningful autonomy in administrative matters.

Wirtz and Müller's analysis of AI governance strategies for the public sector identifies fragmentation as a substantial barrier to effective governance, arguing that the lack of coordination across regulatory domains and levels of government creates opportunities for regulatory arbitrage and undermines the coherence of the overall framework (Wirtz and Müller 2025). The synthesis

of the literature confirms this assessment. The regulatory response to AI in the public sector is not a single, coherent enterprise but a patchwork of initiatives operating at different levels, at different speeds, and with different objectives.

2.3.2.3 The Role of Standards and Soft Law One response to regulatory fragmentation has been the development of technical standards and soft law instruments alongside formal legal regulation. The literature identifies a growing system of standards, guidelines, and best practices intended to guide the development and deployment of AI systems in the public sector.

Malhotra's analysis of a NIST-guided framework for managing generative AI risk in U.S. Public institutions illustrates the role of technical standards in complementing legal regulation (Malhotra 2025). The NIST AI Risk Management Framework provides a structured approach to identifying, assessing, and mitigating risks associated with AI systems. While the framework is not legally binding, it provides a common language and methodology that can be adopted by public bodies seeking to implement responsible AI practices. Similar initiatives exist in other jurisdictions, including the European Commission's Ethics Guidelines for Trustworthy AI and various national AI strategies.

The relationship between formal legal regulation and soft law instruments is complex and sometimes contested. Wagner's analysis of the interaction between the AI Act and the Medical Device Regulation (MDR) illustrates the potential for overlap and conflict between different regulatory frameworks (Wagner 2025). In the healthcare context, AI-powered medical devices may be subject to both the MDR and the AI Act, creating a dual regulatory burden and potential for inconsistent requirements. The analysis suggests that careful coordination is needed to ensure that soft law instruments and technical standards complement rather than complicate formal legal regulation.

The proliferation of standards and guidelines also raises questions about effectiveness. Molins' assessment of AI audits and legal frameworks through the lens of accountability questions whether technical audit mechanisms can substitute for legal accountability (Molins 2026). The analysis suggests that audits can provide valuable information about AI system performance

and compliance but cannot replace the procedural protections that administrative law provides. This is a recurring theme in the literature: technical approaches to governance can support legal frameworks but cannot substitute for them.

2.3.3 Emerging Themes from the Literature

Beyond the overarching finding of a principle-practice gap and the observation of fragmented regulatory responses, the synthesis reveals several recurring themes in the scholarly treatment of AI and administrative law. These themes cut across the doctrinal categories examined in the literature review and point toward deeper conceptual issues that require attention.

2.3.3.1 The Question of Delegation A theme that emerges strongly from the literature is the need to reframe the debate around AI and administrative law as a question of delegation. Rather than asking whether AI systems can comply with existing legal standards, the literature increasingly asks whether the decision to use an AI system in a particular administrative function requires authorisation and, if so, what form that authorisation should take.

Kavungal's framework of civic AI governance argues that the central question is not how AI should be deployed but whether specific decisions should be delegated to AI systems at all (Kavungal 2026). This framing shifts the analytical focus from compliance to legitimacy. The question is not whether an AI system can satisfy due process requirements in a particular case but whether the delegation of decision-making authority to an algorithmic system is consistent with constitutional principles about who should exercise public power.

Carnat's analysis of automation as delegation develops this theme further, arguing that the use of AI systems in the administration of justice constitutes a delegation of power that must satisfy constitutional constraints (Carnat 2025). The analysis draws on constitutional principles of separation of powers and the non-delegation doctrine to argue that certain functions cannot be delegated to AI systems regardless of the technical sophistication of those systems. This is a qualitatively different argument from the technical challenges discussed in section 2.1. It is not that AI systems

are currently incapable of satisfying legal standards but that the delegation of certain functions to AI systems is constitutionally impermissible.

The delegation framing has considerable implications for the regulation of AI in the public sector. If the use of an AI system is understood as a delegation of authority, then the legal framework should require explicit authorisation for each such delegation, specify the scope and limits of the delegated authority, provide for ongoing oversight, and preserve mechanisms for revocation. The literature suggests that current regulatory frameworks do not consistently adopt this approach. Instead, they tend to focus on the operational characteristics of AI systems--their accuracy, fairness, transparency--without addressing the prior question of whether delegation is appropriate.

2.3.3.2 The Limits of Human Oversight A second theme that recurs across the literature is the question of human oversight. Many regulatory frameworks and guidance documents emphasise the importance of human oversight of AI decision systems, typically requiring that a human be "in the loop" for notable decisions. The synthesis suggests, however, that human oversight may be less effective than often assumed.

The literature identifies several limitations of human oversight. First, there is the problem of automation bias: human overseers tend to defer to automated recommendations, particularly when the system has showed high accuracy in routine cases. This means that the human oversight requirement may not provide meaningful protection in practice, as the human overseer may not effectively scrutinise the system's recommendations.

Second, there is the problem of information asymmetry. In many cases, the human overseer has less information than the AI system about the factors that led to a particular decision. The system may have processed thousands of data points and applied complex statistical relationships that the human cannot readily replicate. The human overseer is thus in the position of ratifying a decision that they cannot fully evaluate.

Third, there is the problem of scale. Public-sector AI systems may process large volumes of decisions, making it impractical for human overseers to scrutinise more than a small fraction.

The human oversight requirement may therefore operate as a check on a limited sample of cases rather than as a universal safeguard.

The literature suggests that these limitations are not merely implementation challenges but reflect deeper issues with the concept of human oversight as it is currently understood. Effective oversight requires that the overseer has the capacity to meaningfully evaluate the decision, the independence to disagree with the system, and the opportunity to intervene when necessary. In many practical contexts, these conditions are not met, and the requirement for human oversight may create an illusion of accountability without delivering the substance.

2.3.3.3 Technical Solutions and Their Discontents A third theme across the literature is the tension between technical approaches to governance and legal approaches. The development of explainable AI, fairness metrics, bias detection tools, and auditing frameworks represents a marked technical response to the challenges identified in the legal literature. Technical solutions have achieved important successes, enabling greater transparency about model behaviour, identifying patterns of bias in training data, and providing tools for auditing AI system performance.

At the same time, the literature expresses caution about relying too heavily on technical solutions. Molins' analysis of AI audits questions whether technical audit mechanisms can provide the kind of accountability that legal frameworks require (Molins 2026). Audits, the analysis suggests, can identify problems but cannot resolve them; they can provide information but cannot substitute for the procedural protections that administrative law provides. The distinction between information and accountability is important. Technical tools can generate information about AI system behaviour, but accountability requires institutional mechanisms for acting on that information, including powers to require changes, impose sanctions, and provide remedies to affected individuals.

The literature also identifies a risk that technical solutions may serve to legitimise AI systems without genuinely addressing the underlying concerns. If an AI system is accompanied by an explanation tool, a fairness assessment, and an audit report, it may appear to be governed by

appropriate safeguards even if the explanation is incomplete, the fairness assessment is narrowly defined, and the audit lacks enforcement power. The literature cautions against what might be called "governance theatre"--the appearance of governance without the substance--and argues for a more critical approach to technical governance mechanisms.

2.3.3.4 The Role of Public Sector Capacity A fourth theme that emerges from the synthesis is the importance of public sector capacity for effective AI governance. The literature suggests that even well-designed legal frameworks will be ineffective if the public bodies that must implement them lack the necessary technical expertise, institutional resources, and organisational culture.

Ndlovu's examination of AI-driven tools for sustainable public administration identifies technological, economic and societal obstacles to the widespread use of AI in public administration (Ndlovu 2025). The analysis emphasises that the adoption of AI requires not just technical infrastructure but also organisational change, skills development, and cultural adaptation. Public bodies that lack expertise in data science, machine learning, and algorithmic accountability may adopt AI systems without fully understanding their limitations, or may implement governance requirements in ways that are procedurally compliant but substantively inadequate.

Lāma and Lastovska's exploration of public sector workforce attitudes toward AI competence provides empirical evidence for this concern (Lāma and Lastovska 2025). The study finds meaningful variation in AI competence across the public sector workforce and identifies gaps between the confidence of public officials in their AI capabilities and the actual level of understanding required for effective governance. This finding has implications for the design of regulatory frameworks: if public bodies lack the capacity to implement governance requirements effectively, the requirements may need to be supplemented by capacity-building measures, external expertise, or structural safeguards.

The capacity theme connects to the broader question of regulatory design. Regulation that assumes a high level of technical sophistication in regulated entities may be ill-suited to contexts where that sophistication is lacking. The literature suggests that effective AI governance may re-

quire not only rules and standards but also investments in training, guidance, technical support, and institutional development.

Emerging			
Theme	Key Insight	Implications for Governance	Source
Delegation framing	Shift from compliance to legitimacy; delegation requires authorisation	Legal frameworks should address threshold question of appropriateness	(Kavungal 2026), (Carnat 2025)
Limits of human oversight	Automation bias, information asymmetry, scale challenges	Oversight requirements need careful design, not mere formal compliance	(Khalid Ranjha 2025), (Carnat 2025)
Technical solutions and discontents	Audits inform but do not substitute for accountability	Technical and legal approaches must be complementary, not alternative	(Molins 2026)
Public sector capacity	Effective governance requires expertise, resources, culture	Regulatory design must account for variable capacity across public bodies	(Ndlovu 2025), (Lāma and Lastovska 2025)

Table 14: Emerging Themes from the Literature Synthesis. This table summarises four key themes that recur across the literature examined. Each theme identifies a dimension of the relationship between legal doctrine and algorithmic governance that requires further attention in both scholarship and policy development.

2.3.4 Unresolved Doctrinal Tensions

The synthesis reveals several doctrinal tensions that the existing literature has identified but not resolved. These tensions represent areas where the application of existing legal principles to AI decision systems produces contradictory or ambiguous results, and where further doctrinal development is needed.

2.3.4.1 The Discretion-Automation Tension The most fundamental doctrinal tension concerns the relationship between administrative discretion and algorithmic automation. Discretion is a central concept in administrative law, referring to the authority of a decision-maker to choose among multiple lawful options. The exercise of discretion requires judgment, consideration of individual circumstances, and the application of policy in particular cases. AI systems, by contrast, operate through the application of pre-programmed rules or learned patterns. The question that arises is whether the use of an AI system for a discretionary function is compatible with the legal requirement that discretion be exercised individually and with due regard to the particular circumstances of each case.

The German approach to this tension is explicit. Article 35a VwVfG prohibits fully automated administrative acts where discretion is involved, reflecting a judgment that the exercise of discretion cannot be meaningfully automated. Other jurisdictions have not adopted such a clear rule, and the literature suggests that the tension is managed rather than resolved in most legal systems.

The synthesis reveals that the discretion-automation tension has multiple dimensions. First, there is the question of whether AI systems can exercise discretion in a legally meaningful sense. If discretion requires the conscious consideration of options and the making of a choice, then a system that applies a statistical model is not exercising discretion but simply implementing a predetermined algorithm. Second, there is the question of whether the use of an AI system can be reconciled with the duty to consider individual circumstances. A model trained on population-level data may be able to identify patterns and make predictions, but it is not clear that this constitutes consideration of individual circumstances in the sense that administrative law requires.

The literature surveyed in section 2.1.5 examined these issues through the lens of technical characteristics of AI systems. The synthesis of that analysis confirms that the discretion-automation tension is not a marginal issue that can be resolved through technical adjustments. It goes to the heart of the relationship between algorithmic governance and administrative law. If discretion is essential to legitimate administrative decision-making, and if AI systems cannot exercise discretion

in the sense required by law, then the use of AI systems for discretionary functions is fundamentally problematic.

2.3.4.2 The Proportionality-Efficiency Tension A second doctrinal tension concerns the relationship between proportionality and efficiency. Proportionality is a central principle in European administrative law, requiring that administrative decisions be suitable, necessary, and proportionate in the strict sense. The principle requires a balancing of competing interests and a justification for why the chosen course of action is appropriate in the circumstances.

The adoption of AI systems in the public sector is often motivated by efficiency considerations. AI systems promise faster decisions, lower costs, and more consistent outcomes. The tension arises because the efficiency gains that AI systems offer may come at the cost of the individualised consideration that proportionality requires. A system that processes applications in seconds may not engage in the kind of careful balancing that the proportionality principle demands.

The literature suggests that this tension is often not acknowledged in the policy discourse surrounding AI adoption. The efficiency benefits of AI systems are emphasised, while the potential costs in terms of procedural justice and individualised treatment are downplayed. The proportionality principle requires that these costs be weighed, but the weighing is often absent from the decision to adopt AI systems.

The tension is particularly acute in contexts where AI systems are used for decisions that engage fundamental rights. If an AI system is used to determine eligibility for social benefits, immigration status, or public housing, the efficiency gains must be weighed against the potential for errors, the impact on procedural fairness, and the consequences of reduced individualisation. The literature suggests that this weighing is rarely undertaken in a systematic way.

2.3.4.3 The Accountability-Opacity Tension A third doctrinal tension concerns the relationship between accountability and opacity. Accountability requires that those who exercise power be answerable for their decisions. In the context of administrative decision-making, accountability operates through multiple mechanisms: political accountability to elected officials, legal account-

ability through judicial review, and administrative accountability through internal oversight and complaints procedures.

AI systems introduce opacity into the decision-making process, making it difficult to hold decision-makers accountable. Even where the system's outputs can be observed, the reasoning that produced those outputs may be inaccessible. The tension between accountability and opacity is not new--human decision-makers also have cognitive processes that are not fully transparent--but the opacity of AI systems is of a different order and presents distinct challenges.

The synthesis reveals that the accountability-opacity tension is managed in the literature through various strategies. Some scholars argue for technical solutions, such as explainable AI, that would reduce opacity and enable accountability. Others argue for institutional solutions, such as independent oversight bodies that would have access to the system's internal workings. Still others argue for restrictions on the use of opaque AI systems in contexts where accountability is particularly important.

What the literature does not provide is a clear resolution of the tension. The effectiveness of technical solutions is contested. The availability of institutional solutions depends on political will and resources. And restrictions on AI use may conflict with the efficiency considerations discussed above. The accountability-opacity tension remains an unresolved challenge at the intersection of law and technology.

Doctrinal			
Tension	Nature of Conflict	Approaches in Literature	Unresolved Questions
Discretion-automation	Discretion requires individualised judgment; automation applies pre-determined rules	German prohibition on automated discretionary decisions; hybrid approaches elsewhere	Can AI meaningfully exercise discretion? Does population-level pattern recognition constitute individual consideration?

Doctrinal			
Tension	Nature of Conflict	Approaches in Literature	Unresolved Questions
Proportionality-efficiency	Efficiency gains may conflict with requirement for individualised balancing	Emphasis on ex ante impact assessment; calls for proportionality analysis before AI adoption	How should efficiency be weighed against procedural justice? Who should conduct the balancing?
Accountability-opacity	Accountability requires scrutability; opaque systems resist scrutiny	Technical solutions (XAI); institutional solutions (oversight bodies); prohibitions on use	Can technical solutions provide sufficient transparency for legal accountability? What safeguards are needed?

Table 15: Unresolved Doctrinal Tensions in the Application of Administrative Law to AI Decision Systems. This table identifies three fundamental tensions that emerge from the literature synthesis. Each tension represents a point at which established doctrinal principles conflict with the characteristics of algorithmic decision-making, and where the literature has not yet provided a clear resolution.

2.3.5 Preliminary Summary: Mapping Technical Characteristics to Doctrinal Requirements

The synthesis of findings across the literature points toward the need for a systematic framework that maps the technical characteristics of AI systems to the doctrinal requirements of constitutional and administrative law. The existing literature provides the building blocks for such a framework but has not yet assembled them into a coherent whole. The following preliminary summary outlines the key dimensions of such a mapping.

2.3.5.1 Technical Characteristics Relevant to Doctrinal Analysis The literature identifies several technical characteristics of AI systems that are relevant to doctrinal analysis. These character-

istics mediate the relationship between the technical operation of the system and the legal standards that apply.

First, the level of opacity of the system is central. Systems that are inherently interpretable, such as linear regression models or decision trees, pose fewer challenges for reasons-giving, judicial review, and accountability than systems that are opaque, such as deep neural networks or ensemble methods. The level of opacity determines the extent to which the system can satisfy doctrinal requirements for transparency and explainability.

Second, the degree of autonomy of the system matters. Systems that operate with minimal human involvement raise more acute questions about delegation, discretion, and accountability than systems that support human decision-making. The degree of autonomy determines whether the system is best understood as a tool for human decision-makers or as an independent decision-maker.

Third, the nature of the outputs produced by the system is relevant. Systems that produce binary classifications may be evaluated differently than systems that produce probabilistic scores, rankings, or open-ended text. The nature of the output affects the kind of reasons that can be given, the basis on which judicial review can be conducted, and the accountability mechanisms that are appropriate.

Fourth, the data used to train and operate the system shapes its legal implications. Systems trained on historical data may embed

2.4 Discussion

The synthesis of findings presented in section 2.3 reveals an environment that is both richer and more unsettled than the individual doctrinal streams examined in the literature review (section 2.1) would suggest. Where the literature on due process, transparency, judicial review, and accountability has tended to develop in parallel, the synthesis draws these threads together and shows that they converge on a single, stubborn problem: the technical characteristics of AI systems mediate the application of every doctrinal requirement, yet the law has not developed a systematic method

for taking account of that mediation. This discussion interprets the key patterns emerging from the literature, proposes a synthesised framework for doctrinal analysis, compares that framework with existing regulatory approaches, draws out theoretical implications, acknowledges the limitations of the synthesis, and sets out a future research agenda.

2.4.1 Interpretation of Synthesis

The findings from the literature, as assembled in section 2.3, point to a fundamental asymmetry between the granularity of constitutional and administrative law doctrines and the granularity of AI system design. Doctrinal requirements--fair hearing, reasons-giving, non-delegation, judicial review--are formulated at a high level of abstraction. They are meant to apply across diverse decision-making contexts, from licensing boards to immigration tribunals. AI systems, by contrast, exhibit technical characteristics that vary widely: some are transparent, some opaque; some are fully autonomous, some assist human decision-makers; some produce binary outputs, others produce probabilistic scores or natural language. The literature reviewed in section 2.1 established that these characteristics matter in principle; the synthesis in section 2.3 indicates that they matter in practice, and that the doctrinal framework currently lacks the conceptual tools to respond.

One of the clearest patterns to emerge is the **opacity gradient**. Studies such as those by Kavungal (Kavungal 2026) and Molins (Molins 2026) emphasize that the feasibility of reasons-giving--a core due process requirement identified in section 2.1.1.2--depends on whether the AI system produces explanations that are both accurate and intelligible. The synthesis shows that transparency requirements are not a binary on/off condition; they exist on a continuum. A linear regression model can yield transparent reasons, but a deep neural network cannot, even with post-hoc explanation tools. Yet administrative law has not developed a sliding-scale standard for reasons that corresponds to this technical gradient. The literature in section 2.1.2 on transparency requirements recognised this challenge but did not propose a doctrinal pathway for grading the adequacy of algorithmic reasons in proportion to the system's opacity. The synthesis confirms that this is a persistent gap.

A second pattern concerns **autonomy and discretion**. Section 2.1.1.3 discussed the non-delegation doctrine and the requirement that discretionary power remain with the human decision-maker. The synthesis reveals that the literature is converging on the view that the degree of autonomy--whether the system merely recommends or actually decides--is the decisive variable. Carnat (Carnat 2025) argues forcefully that automation is, in constitutional terms, a delegation of power that must be assessed against the same standards as any other delegation. Kim (Kim 2024) shows how German law has attempted to operationalise this distinction by prohibiting fully automated administrative acts in cases of discretion unless explicitly authorised by statute. Yet the synthesis also reveals that the degree of autonomy is not always easy to determine in practice. Systems that are nominally advisory may nevertheless exert strong influence over human decision-makers, a phenomenon that the literature on algorithmic anchoring and deference has begun to explore but that administrative law has not yet addressed.

The third pattern is the **nature of outputs** and its implications for reviewability. Section 2.1.3 examined the challenges of judicial review when decisions are produced by algorithmic processes. The synthesis shows that the type of output matters for the standard of review. Binary classifications (e.g. eligible/ineligible) may be easier to review than probabilistic scores (e.g. risk of recidivism at 0.78) or generative text (e.g. A personalised denial letter). Each output type raises different questions about what constitutes a reviewable error. The literature in section 2.1.3 acknowledged that judicial review reaches its limits when the reasoning behind an algorithmic output is inaccessible, but the synthesis goes further by showing that even when reasoning is accessible, the doctrinal framework lacks a clear theory of how to evaluate the correctness of a probabilistic prediction. This is not merely a technical challenge; it is a doctrinal vacancy.

A fourth pattern concerns **data and equality**. Section 2.1.1.1 addressed equal protection and the risk of discriminatory outcomes embedded in training data. The synthesis confirms that data quality and representativeness are not simply technical considerations; they are constitutional ones. When an AI system is trained on historical data that reflect past discrimination, its outputs may perpetuate that discrimination in a manner that violates equal protection guarantees. The literature

reviewed by Ndlovu (Ndlovu 2025) and van Noordt and Misuraca (2022) highlights the difficulty of detecting and proving such discrimination in the absence of transparency. The synthesis suggests that the doctrinal response--relying on post-hoc auditing and impact assessments, as discussed in section 2.1.4--is growing but remains both fragmented and under-specified. No jurisdiction has yet developed a extensive set of legal rules that require specific data-quality checks before deployment.

Taken together, these patterns support the central claim of the thesis: that the existing doctrinal framework is, in its current form, inadequate to constrain public-sector AI systems, but that the materials for building a more adequate framework exist in the literature. The synthesis in section 2.3 does not discover new empirical facts; it reorganises what is already known into a structure that reveals the relationships among technical characteristics, doctrinal requirements, and regulatory gaps. This interpretive work is the essential prelude to constructing a framework that can guide both legal scholarship and policy design.

2.4.2 A Synthesised Framework: Doctrinal Constraints for Public-Sector AI

Drawing on the patterns discussed above, this subsection proposes a synthesised framework that maps the technical characteristics of AI systems to specific doctrinal requirements. The framework is intended to be a heuristic tool, not a rigid taxonomy. It does not claim to capture every possible interaction between AI design and legal doctrine; rather, it highlights the most salient relationships identified in the literature.

Technical	Relevant Doctrinal		
Characteristic	Requirement	Core Question	Proposed Constraint
Opacity (interpretable vs. Black-box)	Reasons-giving; transparency; judicial review	Can the system provide intelligible reasons?	Systems above a threshold opacity must provide post-hoc explanations verified by independent audit; otherwise, deployment requires explicit legislative authorisation.

Technical	Relevant Doctrinal		
Characteristic	Requirement	Core Question	Proposed Constraint
Autonomy (advisory vs. Fully automated)	Non-delegation; discretion; accountability	Who makes the final decision?	Fully automated decisions in discretionary areas require statutory authorisation and must be subject to mandatory human review on request.
Output type (binary, probabilistic, generative)	Standard of review; certainty of outcome	What is the nature of the decision?	Probabilistic outputs require a clearly communicated confidence level; generative outputs require a human sign-off.
Data lineage (training data composition)	Equal protection; due process; anti-discrimination	Does the training data embed historical bias?	Mandatory pre-deployment bias audit using disaggregated data; ongoing monitoring for disparate impact.

Table 16: A Synthesised Framework Mapping Technical AI Characteristics to Doctrinal Requirements

The framework operates as follows. For each technical characteristic, it identifies the doctrinal requirement that is most directly engaged, poses the core question that any legal assessment must answer, and proposes a constraint that is consistent with existing constitutional and administrative law principles while being sufficiently specific to guide implementation.

2.4.2.1 Opacity and Reasons-Giving Opacity is the most widely discussed technical characteristic in the literature. The framework proposes a two-track approach. For systems whose internal logic is inherently interpretable (e.g. decision trees with shallow depth, linear models with few features), the ordinary reasons-giving obligation applies, subject to the usual administrative law

standards of adequacy. For systems whose internal logic is opaque (e.g. deep neural networks, gradient-boosted ensembles, large language models), the burden shifts: the deploying agency must illustrate that the system can provide intelligible post-hoc explanations of its outputs, and those explanations must be verified by an independent audit using a recognised explainability metric. If this demonstration is not possible, deployment should require specific legislative authorisation--a requirement that aligns with the German approach described by Kim (Kim 2024) and with the principle that deviations from ordinary procedural norms must be justified.

This approach acknowledges the point made by Kavungal (Kavungal 2026) that the question is not simply *how* to deploy AI but *whether* to do so. The opacity constraint effectively reserves the most opaque systems for contexts where additional procedural safeguards exist. It also creates a doctrinal incentive for public-sector procurement to favour interpretable models where possible, an observation that connects to the procurement-law dimension discussed by Fernández (Fernández 2025).

2.4.2.2 Autonomy and Non-Delegation Autonomy is the second critical dimension. The framework distinguishes three levels: (i) advisory systems that make recommendations while the human decision-maker retains final authority; (ii) partially automated systems that decide routine cases but escalate non-routine cases; and (iii) fully automated systems that decide all cases within a defined scope. The literature, particularly Carnat (Carnat 2025) and Kim (Kim 2024), supports the view that full automation of discretionary decisions requires express statutory authorisation. The framework adopts that position and adds a further requirement: even when authorised, fully automated decisions must be subject to mandatory human review upon the individual's request. This ensures that the right to a fair hearing, discussed in section 2.1.1.1, is not permanently displaced.

For advisory systems, the framework does not require statutory authorisation but does require that the human decision-maker receive training on how to evaluate algorithmic recommendations and that the agency maintain records of when the human deviates from the recommendation.

This addresses the concern, raised in section 2.1.3, that advisory systems may de facto become decisive if human reviewers are prone to automation bias.

2.4.2.3 Output Type and Reviewability The third dimension concerns the nature of the output. Binary decisions (e.g. approve or deny benefits) are the easiest to review, because the error is clear: the decision is either consistent with the evidence or it is not. Probabilistic scores (e.g. A 0.72 likelihood of recidivism) raise more complex questions. As noted in section 2.1.3, the standard of review for probabilistic outputs is under-theorised. The framework proposes that probabilistic outputs must be accompanied by a well-calibrated confidence interval, disclosed to both the decision-maker and the affected individual, and that judicial review may consider whether the confidence interval was appropriately calculated and communicated. This would move the inquiry from a binary correctness check to an evaluation of the methodological integrity of the risk assessment.

Generative outputs (e.g. automatically drafted decision letters) pose a distinct problem: they may contain errors of fact, law, or reasoning that are not immediately apparent because the text appears plausible. The framework requires human sign-off for any generative output that constitutes or forms part of a legally binding decision. This is consistent with the emerging consensus in the literature that human oversight is essential for high-stakes generative AI applications.

2.4.2.4 Data and Equal Protection The fourth dimension addresses data lineage. The literature reviewed in section 2.1.1.1 and by Ndlovu (Ndlovu 2025) establishes that AI systems can replicate and amplify historical discrimination. The framework proposes two structural requirements. First, a pre-deployment bias audit using disaggregated data (by race, gender, age, and other protected characteristics), conducted by an independent body. Second, ongoing monitoring for disparate impact, with a driver for corrective action if disparities emerge post-deployment. These requirements are not novel in principle--impact assessments are already part of some regulatory frameworks, as discussed in section 2.1.6--but the framework makes them mandatory rather than voluntary and specifies the data types that must be examined.

This set of constraints responds to a criticism implicit in much of the literature: that procedural fairness and substantive fairness cannot be meaningfully separated when the decision system itself may be biased. The framework treats data quality as a due process issue, not merely a technical one.

2.4.3 Comparison with Existing Frameworks

The synthesised framework proposed above can be compared with several existing regulatory and governance frameworks discussed in the literature review (section 2.1.6). The comparison reveals both areas of convergence and important divergences.

Framework / Jurisdiction	Opacity Constraint	Autonomy	Output		Enforcement
		Constraint	Type	Data/Equity	Dynamic
EU AI Act (Proposed)	Tiered risk categories; high-risk systems require explainability	Not directly addressed	Not addressed	Data governance requirements for high-risk	Fines, market surveillance
German Administrative Procedure Act (Section 35a)	Not addressed	Prohibits fully automated discretionary acts unless authorised	Not addressed	Not addressed	Judicial review of individual decisions
Canadian Directive on Automated Decision-Making	Requires human-readable explanations for all systems	Requires meaningful human review for all decisions	Not addressed	Requires gender-based analysis plus	Certification, audit

Framework /		Autonomy	Output		Enforcement
Jurisdiction	Opacity Constraint	Constraint	Type	Data/Equity	Dynamic
Proposed Framework (this thesis)	Opacity threshold with verification requirement	Statutory authorisation plus mandatory human review on request	Confidence intervals for probabilistic; human sign-off for generative	Mandatory pre-deployment bias audit plus monitoring	Judicial review with expanded standing

Table 17: Comparison of the Proposed Framework with Existing Regulatory Approaches

The EU AI Act (discussed in section 2.1.6) adopts a risk-based approach, imposing obligations on high-risk systems. Its opacity constraint--requiring technical documentation and human oversight--is in some respects less demanding than the framework proposed here, because it does not expressly require that post-hoc explanations be verified by independent audit. The Act also does not directly address the level of autonomy, focusing instead on the intended purpose of the system. The proposed framework goes further by making autonomy an explicit dimension and by requiring statutory authorisation for full automation of discretionary decisions.

The German Administrative Procedure Act (Section 35a) provides a clear autonomy constraint: it prohibits fully automated administrative acts in areas involving discretion unless a statute specifically allows it. This aligns closely with the framework's position. However, the German Act does not address opacity, output type, or data equity. The framework fills that gap by adding dimensions that the German legislation leaves unregulated.

The Canadian Directive on Automated Decision-Making (mentioned in section 2.1.6) requires human-readable explanations and meaningful human review. It is in some respects the most detailed among existing frameworks, but it lacks specificity on what qualifies as a human-readable

explanation and does not impose mandatory independent bias audits. The framework seeks to operationalise these requirements by linking them to specific technical characteristics.

The comparison highlights a broader pattern identified in the synthesis: existing frameworks tend to focus on one or two dimensions (typically opacity and risk) while neglecting others (autonomy, output nature, data equity). The proposed framework is more wide-ranging in scope, but it is also more demanding. Whether such demands are politically feasible is a separate question that the literature has begun to explore, particularly in the context of public-sector resource constraints discussed by Ndlovu (Ndlovu 2025) and van Noordt and Misuraca (2022).

2.4.4 Theoretical Implications

The findings from the literature, as synthesized in section 2.3, carry implications that extend beyond the immediate question of AI regulation. They challenge some of the foundational assumptions of administrative law theory and suggest that the encounter with AI may force a re-examination of doctrines that have remained relatively stable for decades.

2.4.4.1 The Concept of Discretion Discretion has traditionally been understood as the space left open by law--the zone within which a decision-maker may choose among alternatives based on policy, fairness, or expediency. The literature reviewed in section 2.1.1.3 treated discretion as a human attribute. The synthesis reveals that AI systems can exercise effective discretion even when formal authority remains with a human, because the system's design choices (thresholds, feature weights, inclusion criteria) embody the policy judgments that discretion is meant to preserve for human consideration. This challenges the binary distinction between rule-based and discretionary decisions that underlies much administrative law. If an AI system effectively makes discretionary choices at the design stage, then the legal framework must find a way to subject those choices to the same procedural and substantive constraints that apply to human discretion. The proposed framework's emphasis on statutory authorisation for full automation is one response, but the deeper

theoretical implication is that the concept of discretion may need to be unbundled into design-stage and application-stage components, each with its own doctrinal treatment.

2.4.4.2 Reasons and the Administrative Record The duty to give reasons is often justified on the ground that it promotes rationality, accountability, and the possibility of meaningful judicial review. The synthesis shows that AI systems can produce reasons, but those reasons may not be the ones that actually drove the decision. Post-hoc explanations of a black-box model may be plausible without being faithful. This raises a troubling possibility: that reasons become a performative exercise, disconnected from the actual decision logic. If so, the justifications for the duty to give reasons--transparency, individual dignity, effective review--start to unravel. The literature has not fully grappled with this implication. It tends to assume that explainability tools, if sufficiently advanced, can restore the connection between decision and reason. The synthesis suggests that this assumption is optimistic. Some AI systems may be inherently incapable of producing the kind of reasons that administrative law requires. If that is the case, the legal system faces a stark choice: either accept a lower standard of reasoning for AI systems, or restrict their use to contexts where the stakes are low enough that reduced reasoning suffices. Neither option is comfortable.

2.4.4.3 The Architecture of Accountability Accountability in administrative law has traditionally been understood as a chain: the decision-maker is accountable to a superior, who is accountable to a minister, who is accountable to parliament or the electorate. AI systems disrupt this chain because the system itself is not a legal person and cannot be held accountable. The literature in section 2.1.4 explored various accountability mechanisms--audits, impact assessments, procedural rules. The synthesis indicates that these mechanisms operate at a distance from the actual decision-making process. They are retrospective, not prospective; they focus on the system as a whole, not on individual decisions. This creates a gap between the moment of decision and the moment of accountability. The theoretical implication is that the traditional model of ex post accountability (judicial review, ombudsman) may be insufficient for AI systems. Instead, the literature points toward a model of ex ante structural accountability: design standards, certification requirements, and

mandatory human oversight that operate before the system is deployed. This shift from individualised to systemic accountability is not unprecedented--environmental law and financial regulation already use such models--but it would represent a substantial reorientation for administrative law.

2.4.4.4 The Proper Role of Human Oversight The literature converges on the importance of human oversight, but the synthesis reveals that the concept is underspecified. Does “human oversight” mean the human must actively review each decision? Or does it mean the human is available to intervene if the system flags a problem? The answer matters for doctrinal analysis. If oversight is passive, the human may be little more than a rubber stamp, and the non-delegation concern persists. If oversight is active, the purported efficiency gains of automation may be lost. The theoretical question is whether human oversight should be conceptualised as a substantive constraint (requiring genuine engagement) or a formal one (requiring the presence of a human). The literature tilts toward the substantive view, but the proposed framework’s requirement of mandatory human review on request is a compromise: it preserves an individual right to human engagement while allowing routine decisions to be automated.

2.4.5 Limitations of the Synthesis

Any literature-based synthesis is subject to limitations, and those identified in the research gap analysis (section 2.2.4) and in the literature itself must be acknowledged here.

First, the **scope of the literature** is limited. The synthesis draws on a body of scholarship that is itself growing but remains concentrated in a handful of jurisdictions--the United Kingdom, the European Union, the United States, Germany, and Australia. The administrative law traditions of many other countries are underrepresented. The findings may not generalise to jurisdictions with different legal traditions, different institutional configurations, or different levels of AI adoption. The work by van Noordt and Misuraca (2022) on EU public-sector AI landscapes suggests that national variation is considerable, and the synthesis does not fully capture that variation.

Second, the synthesis is **doctrinally focused**. It examines constitutional and administrative law constraints but does not examine in detail the overlap with other legal fields such as data protection law, human rights law, procurement law, or tort liability. The discussion of procurement in Fernández (Fernández 2025) points to one such intersection, and the EU AI Act’s overlap with data protection is another. A thorough governance framework would need to integrate these fields, but the synthesis has deliberately narrowed its lens.

Third, the synthesis is **primarily conceptual** rather than empirical. It identifies doctrinal gaps and proposes a framework, but it does not test that framework against real-world cases. The research gap analysis noted that almost no work operationalises doctrinal principles into testable, empirically validated frameworks. This synthesis does not overcome that limitation; it confirms its existence. Future empirical work would be needed to assess whether the proposed constraints are feasible, effective, and enforceable.

Fourth, the literature itself is **skewed toward high-profile systems** (e.g. risk assessment in criminal justice, eligibility for benefits) and may not capture the full range of low-stakes decisions that are nonetheless regulated by administrative law. The synthesis inherits this bias. The constraints proposed in the framework may be more stringent than necessary for low-stakes decisions, and a future refinement might introduce a proportionality dimension.

Fifth, the literature tends to treat **technical characteristics as stable properties** of AI systems. But technical characteristics can change through fine-tuning, retraining, or the deployment of new explainability tools. The synthesis does not fully address the dynamic quality of AI systems. A framework that is static at the point of deployment may quickly become outdated.

2.4.6 Future Research Agenda

The gaps identified in the synthesis and the limitations discussed in the preceding subsection point to a substantial future research agenda. The following directions are suggested by the literature and by the logic of the proposed framework.

Empirical testing of the proposed framework. The most immediate need is to test the framework against actual public-sector AI deployments. Case studies of specific systems--for example, automated welfare eligibility assessments in the Netherlands or predictive policing tools in the UK--could reveal whether the proposed constraints are operationalisable and whether they would have prevented known failures. Such empirical work would also generate the doctrinal-empirical mapping that the research gap analysis identified as most urgently missing.

Comparative doctrinal analysis across jurisdictions. The synthesis reveals that different legal systems have developed distinct responses to the same technical challenges. A systematic comparison of how German, Canadian, EU, and UK law handle opacity, autonomy, and data equity would illuminate the strengths and weaknesses of each approach and could inform harmonisation efforts. The work by Kim (Kim 2024) on German law and by Kavungal (Kavungal 2026) on a civic AI governance framework provides starting points, but a full-scale comparative study remains to be done.

Development of administrative law standards for explainability verification. The proposed framework requires that post-hoc explanations for opaque systems be verified by independent audit. But what standard of explanation should apply? The literature on explainable AI (XAI) offers technical metrics--faithfulness, consistency, completeness--but administrative law has not yet translated these into legal standards. Research at the intersection of computer science and administrative law could develop a rubric for when an explanation is “good enough” for legal purposes, and when it is not.

Longitudinal study of AI-generated administrative decisions. If AI systems are used to make or support administrative decisions over time, the decisions themselves become part of the administrative record. How should those records be preserved, audited, and made available for judicial review? The literature reviewed in section 2.1.3 touched on the question of discoverability, but a systematic study of how AI-generated records interact with rules on evidence and administrative law is needed.

Examination of the interaction between AI systems and the right to be heard. The right to a fair hearing, discussed in section 2.1.1.1, includes the right to respond to adverse evidence and to make representations before a decision is taken. How does the right to be heard operate when the “evidence” is the output of a complex model? Can an individual meaningfully challenge an AI-based assessment without understanding its internal logic? The doctrinal implications are profound, and the literature has only begun to explore them.

The role of procurement law in shaping AI governance. Fernández (Fernández 2025) points out that procurement law is a powerful but underused tool for ensuring that public-sector AI systems meet constitutional standards. By imposing contractual obligations on vendors—for example, requiring explainability, audit access, and data lineage documentation—governments can embed legal constraints at the point of acquisition. Future research could develop model contractual clauses and examine the legal enforceability of such requirements.

The problem of automation bias and its doctrinal response. Several studies suggest that human decision-makers tend to defer to algorithmic recommendations, even when they know the algorithm is fallible. This phenomenon undermines the effectiveness of human-in-the-loop safeguards. Administrative law has not yet developed a response. Research could explore whether a legal duty of independent assessment, analogous to the duty of care in medical decision-making, could be imposed on officials who review algorithmic recommendations.

The feasibility of independent bias auditing in resource-constrained settings. Many public-sector agencies lack the technical expertise and financial resources to conduct rigorous pre-deployment bias audits. Ndlovu (Ndlovu 2025) and van Noordt and Misuraca (2022) note that resource constraints are a major obstacle to effective AI governance. Future work could investigate how auditing requirements can be scaled to different types of agencies and different levels of risk, perhaps through shared auditing facilities or third-party certification schemes.

In sum, the synthesis presented in section 2.3 and discussed in this section shows that the literature has successfully identified the key doctrinal challenges posed by public-sector AI but has not yet assembled the insights into a coherent framework that can guide legal analysis, policy design,

and judicial reasoning. The framework proposed here--mapping opacity, autonomy, output type, and data lineage to specific doctrinal constraints--is a step towards such coherence. It is not offered as a final solution but as a scaffold on which future research can build. The most pressing next step is empirical: to take the framework out of the theoretical domain and test it against the messy reality of AI systems already operating in public administration. The gap between principle and practice, identified throughout the literature, will only be closed through sustained, interdisciplinary work that treats both the legal doctrine and the technical system with equal seriousness.

3. Conclusion

The integration of artificial intelligence into public-sector decision-making has proceeded with remarkable speed, yet the constitutional and administrative law frameworks that govern such decisions have developed over centuries, not decades. This thesis set out to examine whether existing doctrinal constraints--rooted in due process, transparency, judicial review, and accountability--can adequately regulate algorithmic governance, or whether the gap between principle and operational practice has grown too wide to bridge without substantial doctrinal adaptation.

From the start, this thesis established that administrative law doctrines were designed for a world where human decision-makers exercise judgment, give reasons, and remain answerable through established chains of accountability. AI decision systems, by contrast, operate through computational processes that are opaque even to their developers in many cases. They produce decisions at a scale and speed that overwhelms traditional review mechanisms, and they diffuse responsibility across multiple actors--developers, deployers, procuring agencies alike. The central question, then, was not whether the principles of administrative law remain normatively attractive (they do) but whether they can be operationalised in contexts where the decision-maker is no longer human in any meaningful sense.

The evidence points toward a real and consequential principle-practice gap. Doctrinal requirements such as the right to a fair hearing, the duty to give reasons, and the availability of meaningful judicial review presuppose a decision-making process that is explicable, contestable, and attributable. Each of these presuppositions is challenged, in distinct ways, by the technical characteristics of contemporary AI systems.

Strikingly, the literature also shows that regulatory responses have been fragmented. Different jurisdictions adopt incompatible approaches, and no consensus has emerged on how to translate high-level principles into enforceable rules.

3.1.2 Principal Findings

What the literature reviewed in Chapter 2 reveals is a set of interconnected findings. One concerns the doctrinal foundations of due process--the right to be heard, the duty to give reasons, and the requirement of impartial adjudication. These are strained when decisions come from machine learning models whose internal logic is not readily accessible. The German *Verwaltungsverfahrensgesetz* Section 35a, which prohibits fully automated administrative acts where discretion is involved, represents one legislative attempt to preserve human judgment. Its scope is limited though, and it doesn't address the deeper epistemic challenges of algorithmic decision-making.

A second finding revolves around transparency requirements. Whether formulated as the common law duty to give reasons or as statutory obligations under freedom of information legislation, these run into serious technical obstacles when applied to AI systems. The opacity of many machine learning models--particularly deep neural networks--means that even developers cannot always explain why a particular output was produced. Transparency, in its meaningful sense, requires not merely access to information about a system but the ability to understand and contest the basis of a specific decision. On balance, current transparency mechanisms fall short of this standard.

Third, judicial review of algorithmic decisions faces structural difficulties. The traditional grounds of review--illegality, irrationality, procedural impropriety--were developed with human decision-makers in mind. Applying these to algorithmic outputs forces courts to engage with technical evidence about model design, training data, and performance metrics, a task for which many judges and legal practitioners are not currently equipped. The speed and volume of automated decisions also create practical barriers to effective review, particularly in high-volume domains such as welfare administration or immigration control. One might expect the courts to adapt gradually, but the evidence suggests the adaptation is not keeping pace.

Fourth--and arguably the more telling finding--accountability mechanisms (audit requirements, oversight bodies, liability regimes) have been proposed in various forms but remain unevenly implemented. The literature suggests accountability is not a single concept but a cluster of

requirements--attributability, answerability, enforceability--each challenged differently by AI systems. The diffusion of responsibility across multiple actors in the AI supply chain complicates traditional accountability frameworks.

Comparative analysis reveals considerable divergence in regulatory approaches. The European Union's AI Act establishes a risk-based framework with mandatory requirements for high-risk systems, while the United Kingdom has adopted a more principles-based, sector-led approach. Other jurisdictions, including Germany and Canada, have experimented with specific legislative constraints on automated decision-making. No single model has yet shown clear superiority, and each has identifiable gaps and weaknesses.

Finally, the gap analysis conducted in this thesis identified a persistent disconnect between the doctrinal requirements articulated in legal scholarship and the technical realities of AI deployment in public-sector contexts. Much of the legal literature operates at a high level of abstraction, engaging relatively little with the specific technical characteristics of different AI approaches. Conversely, technical literature on AI governance often lacks doctrinal precision, treating legal requirements as broad principles to be satisfied through generic solutions rather than as specific, enforceable obligations.

Dimension of			Gap	
Constraint	Doctrinal Requirement	Technical Challenge	Severity	Illustrative Source
Due Process	Right to fair hearing; reasoned decision	Opacity of model outputs; speed of decisions	High	(Carnat 2025)
Transparency	Duty to give reasons; access to information	Model opacity; trade secrets	High	(Khalid Ranjha 2025)
Judicial Review	Reviewability on standard grounds	Technical complexity; scale of decisions	Moderate-High	(Clayton et al. 2018)

Dimension of			Gap	
Constraint	Doctrinal Requirement	Technical Challenge	Severity	Illustrative Source
Accountability	Attributability; answerability; enforceability	Diffusion of responsibility across actors	High	(Molins 2026)
Regulatory Design	Coherent legislative framework	Fragmented jurisdictional approaches	Moderate	(Kavungal 2026; Wirtz and Müller 2025)

Table 18: Summary of the Principle-Practice Gap Across Five Dimensions of Constitutional and Administrative Constraint.

The table above synthesises the principal findings across five key dimensions. The gap severity is assessed as high in three of the five dimensions, with only judicial review and regulatory design showing moderate potential for adaptation within existing frameworks. This assessment draws from comparative analysis of scholarly contributions, including Carnat's examination of constitutional constraints on automated justice (Carnat 2025), Molins' analysis of accountability deficits in AI auditing (Molins 2026), and Kavungal's framework for civic AI governance (Kavungal 2026). What stands out here is the pattern: the closer one looks at the operational level, the wider the gap between doctrinal promise and practical reality.

3.2 Contributions of the Thesis

3.2.1 Doctrinal Contributions

This thesis makes several contributions to the legal literature on AI governance. The most fundamental is a systematic mapping of how specific doctrinal requirements interface with specific technical characteristics of AI systems. Rather than treating "AI" as a monolithic category, this thesis has distinguished between rule-based systems, supervised learning models, deep learning architectures, and generative AI--each raises different issues for due process, transparency, and

accountability. This granular approach enables more precise doctrinal analysis than the broad-brush treatments that characterise much of the existing literature.

A second doctrinal contribution lies in the analysis of how the diffusion of responsibility in AI systems challenges traditional accountability frameworks. The thesis has shown that the problem is not merely that AI systems make decisions without human oversight, but that the human actors involved--developers, deployers, procuring agencies, and operators--each have partial knowledge and partial control, making it difficult to assign legal responsibility for specific outcomes. This analysis builds on Molins' work on auditing accountability (Molins 2026) and Carnat's examination of automation as delegation of power (Carnat 2025), extending their insights into a more extensive doctrinal framework.

3.2.2 Framework Contributions

The synthesised framework presented in the discussion chapter represents another contribution. It maps technical characteristics of AI systems onto specific doctrinal requirements, providing a structured approach for assessing whether a given AI deployment complies with constitutional and administrative law constraints. The framework aims to be both doctrinally rigorous and practically applicable--offering guidance to legislators, regulators, and public-sector decision-makers navigating the interface between law and technology. The framework draws on comparative analysis, recognising that different legal traditions have developed different approaches to the problem of algorithmic governance. The German approach of explicit legislative constraint, the EU's risk-based regulatory model, and the UK's principles-based framework each offer lessons incorporated into the framework's design. This comparative dimension, grounded in the literature on AI governance strategies (Wirtz and Müller 2025; van Noordt and Misuraca 2022), adds depth and practical relevance.

3.2.3 Methodological Contributions

A further contribution is methodological. This thesis has adopted a narrative review methodology that systematically integrates legal doctrinal analysis with technical literature on AI systems. This interdisciplinary approach is essential for addressing the research question, which sits at the intersection of law, technology, and public administration. The methodology, while not without limitations, reveals a replicable approach for legal scholarship that engages seriously with technical subject matter. The search strategy and coding framework developed for this thesis provide a template for future reviews that seek to map the interface between legal doctrine and technological capability. By explicitly coding sources according to doctrinal domain, technical focus, and jurisdictional scope, the methodology enables systematic comparison across bodies of literature that are often treated in isolation.

Type of Contribution	Specific Contribution	Primary Sources	Practical Utility
Doctrinal	Granular mapping of doctrinal requirements to AI system types	(Carnat 2025; Khalid Ranjha 2025)	Enables targeted legal analysis
Doctrinal	Analysis of responsibility diffusion in AI supply chains	(Molins 2026; Fernández 2025)	Clarifies accountability gaps
Framework	Synthesised mapping of technical-doctrinal interface	(Kavungal 2026; Wirtz and Müller 2025)	Practical guidance for compliance
Comparative	Integration of multiple jurisdictional approaches	(Kim 2024; van Noordt and Misuraca 2022)	Informs regulatory design
Methodological	Replicable interdisciplinary review methodology	Thesis design	Template for future research

Table 19: Summary of Contributions by Type.

3.3 Implications for Law and Policy

3.3.1 Implications for Constitutional Doctrine

The findings of this thesis carry substantial implications for constitutional and administrative law doctrine. The most immediate is that existing doctrinal frameworks, while not obsolete, require adaptation to address the specific challenges posed by AI decision systems. The right to a fair hearing, for example, cannot simply be applied to algorithmic decisions in the same form as it applies to human decisions--it requires a reconceptualisation of what "being heard" means when the decision-maker is a computational system. Similarly, the duty to give reasons must be reinterpreted given the technical constraints on explainability that characterise many AI models.

One implication that emerges strongly from the analysis is the need for a more nuanced understanding of transparency. The literature reveals that transparency is not a binary property--a system is not simply transparent or opaque--but a multidimensional characteristic that varies across technical, procedural, and legal dimensions. Doctrinal analysis must engage with this complexity rather than treating transparency as a single requirement that can be satisfied through a single factor. Carnat's work on constitutional constraints (Carnat 2025) and Kavungal's governance framework (Kavungal 2026) both point toward the need for differentiated transparency obligations that vary according to system type and decision context. Could this become a central battleground for future litigation?

3.3.2 Implications for Legislative Reform

For legislators and policymakers, the thesis suggests that a piecemeal approach to AI governance is unlikely to be sufficient. The fragmented regulatory environment identified in the literature review--with different jurisdictions adopting incompatible approaches and notable gaps remaining even within the most detailed frameworks--points toward the need for more coherent and systematic legislative responses. The EU AI Act represents an important step, but its effectiveness will depend

on implementation and enforcement, and it does not resolve all the doctrinal tensions identified in this thesis.

The thesis also suggests that legislative reform should attend to the full AI supply chain, not merely to the point of decision. Procurement law, as Fernández has shown (Fernández 2025), is a critical but often overlooked site of regulatory intervention. By embedding constitutional and administrative law requirements into procurement processes, legislators can influence the design and deployment of AI systems before they reach the point of decision-making, rather than attempting to address problems retrospectively through judicial review or oversight mechanisms. The comparative analysis further suggests that different legal traditions may require different legislative approaches. The German model of explicit prohibition of fully automated discretionary decisions, codified in Section 35a VwVfG, represents one approach that has been influential but may not be suitable for all jurisdictions. The UK's sector-led, principles-based approach offers flexibility but risks inconsistency and gaps in coverage. The thesis does not prescribe a single legislative model but rather identifies the considerations that should inform legislative choice, including the nature of the decisions involved, the technical characteristics of the AI systems, and the existing doctrinal framework.

3.3.3 Implications for Public Administration

For public-sector decision-makers and administrators, the thesis provides a framework for assessing whether and how AI systems can be deployed in compliance with constitutional and administrative law constraints. The framework is designed to be practical, offering structured guidance that can be applied to specific use cases. It does not provide simple answers--the complexity of the interface between law and technology precludes such simplicity--but it does provide a method for asking the right questions and identifying the key risks and tensions.

The thesis also highlights the importance of institutional capacity. Effective implementation of constitutional and administrative law constraints on AI systems requires expertise that is currently in short supply in many public-sector organisations. Legal professionals need training

in the technical aspects of AI, while technical professionals need training in the legal and constitutional framework. Lāma and Lastovska's work on workforce competence (Lāma and Lastovska 2025) suggests these issues are not peripheral but central to successful integration.

Implication			
Domain	Key Insight	Recommended Action	Relevant Source
Constitutional Doctrine	Due process requires reconceptualisation for algorithmic decisions	Develop differentiated doctrinal standards by AI system type	(Carnat 2025)
Legislative Reform	Fragmented regulatory environment requires coherent response	Pursue wide-ranging legislative framework with procurement integration	(Kavungal 2026; Fernández 2025)
Public Administration	Institutional capacity is a prerequisite for compliance	Invest in cross-disciplinary training and expertise	(Lāma and Lastovska 2025)
Regulatory Design	Different legal traditions require different approaches	Adopt context-sensitive regulatory models	(Wirtz and Müller 2025; van Noordt and Misuraca 2022)

Table 20: Implications and Recommended Actions.

What is perhaps most striking is how each of these domains reinforces the others--doctrinal adaptation cannot succeed without legislative reform, and neither can function without institutional capacity. The pieces fit together, but they are rarely addressed as a whole.

3.4 Limitations of the Thesis

3.4.1 Scope Limitations

This thesis has several limitations that deserve acknowledgment. First, the scope is necessarily bounded. The research has focused primarily on common law and civil law traditions in

Western Europe and North America, with more limited attention to other legal traditions. The comparative analysis includes substantial treatment of German, EU, and UK law, but does not engage in depth with Asian, African, or Latin American legal systems. The findings may therefore not be directly transferable to jurisdictions with different constitutional traditions or different stages of AI deployment.

Second, the thesis has focused on constitutional and administrative law constraints rather than other potentially relevant frameworks such as human rights law, data protection law, or tort law. These frameworks intersect with administrative law in important ways, and a more thorough analysis would need to consider their interaction. Still, the decision to focus on administrative law was deliberate--it reflects the centrality of this body of law to the regulation of public-sector decision-making.

Third, the thesis has engaged with technical literature on AI systems but has not included original empirical research on how specific AI systems operate in practice. The analysis is based on the technical descriptions and assessments provided in the scholarly literature, which may not always reflect the full complexity of deployed systems. The limitations of this approach are discussed further below.

3.4.2 Methodological Limitations

The narrative review methodology adopted here has both strengths and limitations. The principal strength is that it allows for a broad and integrative analysis that draws on multiple bodies of literature. The principal limitation is that it does not follow the formal protocols of a systematic review, including prespecified inclusion criteria, dual screening, and formal quality assessment. The source selection process was guided by the research question and the expertise of the reviewer, but it may have missed relevant sources that a more exhaustive search strategy would have captured. The reliance on English-language sources is an additional limitation. While the thesis has engaged with German and EU legal materials through secondary sources, the primary search strategy focused on English-language publications. This may have introduced a linguistic bias, particularly

given that substantial scholarly work on AI and administrative law exists in German, French, Spanish, and other languages.

3.4.3 Evidentiary Limitations

A further limitation concerns the nature of the evidence available. The literature on AI and administrative law is still emerging, and many of the claims made are based on theoretical analysis rather than empirical observation. The number of rigorous empirical studies examining the actual operation of AI decision systems in public-sector contexts and their interaction with legal frameworks remains small. This thesis has necessarily relied heavily on doctrinal analysis and theoretical argument, reflecting the current state of the field. The findings should be revisited as more empirical evidence becomes available.

3.5 Directions for Future Research

3.5.1 Empirical Validation

The most pressing need for future research is empirical. The framework developed in this thesis provides a structured approach for assessing the compliance of AI systems with constitutional and administrative law constraints, but it has not been tested against real-world deployments. Future research should examine specific AI systems in specific public-sector contexts--welfare administration, immigration control, policing, health service allocation--and assess whether the doctrinal tensions identified here manifest in practice and whether the framework can guide effective responses. Such empirical work would need to be interdisciplinary, combining legal analysis with technical assessment and empirical observation. Researchers would need access to deployed systems, their documentation, and the decision outcomes they produce. This raises its own challenges, including commercial confidentiality, operational security, and research ethics. But without such work, the gap between doctrinal analysis and practical reality will remain an article of faith rather than an evidence-based finding.

3.5.2 Comparative Doctrinal Analysis

A second direction is more extensive comparative analysis. This thesis has examined German, EU, and UK approaches, but a more extensive comparative study would include additional jurisdictions, particularly those with different legal traditions. The AI governance frameworks being developed in East Asia, for example, reflect different constitutional and administrative law traditions and may offer alternative models for addressing the principle-practice gap. Similarly, approaches emerging in African and Latin American jurisdictions, while less documented in the English-language literature, may offer insights not captured by a focus on Western European and North American models. Comparative analysis could also examine the interaction between administrative law frameworks and other regulatory regimes, including data protection law, human rights law, and sector-specific regulation. The intersection of these frameworks is likely to become increasingly important as AI systems become more pervasive in public-sector decision-making.

3.5.3 Technical-Legal Interface Research

A third direction concerns the technical-legal interface. The analysis in this thesis has identified specific points where doctrinal requirements encounter technical constraints, but further research is needed to explore how these constraints can be addressed through technical design. The field of explainable AI (XAI) offers one set of tools, but the legal requirements for explanation may not align perfectly with the technical capabilities of XAI methods. Research that brings together legal scholars, computer scientists, and practitioners is needed to develop technical solutions that are both effective and legally adequate. Similarly, developing audit frameworks for AI systems requires interdisciplinary collaboration. The accountability deficits identified in this thesis (Molins 2026) cannot be resolved by legal reform alone--they require audit methodologies, standards, and institutional capacity that span the legal-technical divide.

3.5.4 Normative Questions

Finally, the thesis has bracketed some broader normative questions that deserve further attention. The analysis has assumed, consistent with the existing literature, that constitutional and administrative law constraints should apply to AI decision systems. But this assumption deserves critical examination. Is it always desirable to extend existing doctrinal frameworks to algorithmic governance, or are there cases where AI systems should be subject to different--perhaps more permissive or more restrictive--standards? The normative foundations of the principle-practice gap, including questions about the legitimacy of algorithmic decision-making and the proper role of human judgment, require deeper theoretical engagement than this thesis has been able to provide.

3.6 Concluding Remarks

The constitutional and administrative law constraints that govern public-sector decision-making are not obsolete. The principles of due process, transparency, judicial review, and accountability remain normatively compelling, reflecting values fundamental to the rule of law in democratic societies. But these principles cannot be applied to AI decision systems through simple extension. The gap between doctrinal requirement and technical capability is real, and it demands a response that is both doctrinally rigorous and technically informed.

This thesis has argued that the response should take the form of a structured framework--one that maps technical characteristics of AI systems onto specific doctrinal requirements, enabling a precise assessment of compliance and identifying the points where adaptation is needed. The framework does not provide simple answers, because the questions are not simple. What it does provide is a method for asking the right questions and a set of considerations that should guide legislative, regulatory, and administrative responses.

The challenge ahead is marked. AI systems are becoming more capable and more pervasive, and the public sector is under pressure to adopt them in pursuit of efficiency, consistency, and effectiveness. But efficiency is not the only value that public-sector decision-making serves.

Fairness, accountability, transparency, and the right to be heard are not optional extras--they are constitutional requirements that define the relationship between the state and the individual. Ensuring these requirements are preserved in the age of algorithmic governance is one of the central challenges for contemporary administrative law.

The thesis concludes, therefore, with a call for continued engagement across the legal-technical divide. The doctrinal analysis conducted here is a starting point, not a final answer. The full implications of AI for constitutional and administrative law will become apparent only through sustained interdisciplinary research, empirical investigation, and practical experience. What is clear is that the task is urgent. The decisions AI systems make in the public sector affect people's lives in material ways--their access to benefits, their exposure to enforcement, their liberty, and their dignity. Ensuring that these decisions are made in accordance with constitutional and administrative law is not merely a technical challenge but a democratic imperative.

4. Appendices

A.1 Overview of the Framework

This appendix presents a conceptual framework that translates the core doctrinal requirements of constitutional and administrative law - due process, transparency, judicial review, and accountability - into specific design constraints for public-sector AI decision systems. The framework is built on the premise that different technical characteristics of AI systems (e.g. opacity, non-determinism, scalability, data dependency) interact with doctrinal requirements in distinct ways, and that effective governance requires a mapping between the two domains. The literature reviewed in this thesis points to a persistent gap between principle and operationalisation (Kavungal 2026; Carnat 2025). This framework aims to bridge that gap by providing a structured taxonomy.

A.1.1 Rationale for a Structured Mapping

Existing guidance documents, such as the EU AI Act and various national strategies, tend to state high-level principles without specifying how those principles translate into concrete design decisions (Wirtz and Müller 2025; van Noordt and Misuraca 2022). For example, the requirement of “transparency” may be interpreted as a duty to publish the source code of an algorithm, a duty to provide a plain-language explanation of a decision, or a duty to disclose the training data used. Each interpretation has different technical and legal implications. The framework offered here organises the relationship between doctrinal obligations and technical properties into a set of decision tables that can guide both legislators and system developers.

A.2 Doctrinal Requirements and Their Technical Counterparts

A.2.1 Due Process and the Right to a Fair Hearing

Due process requires that an individual affected by an administrative decision receives notice of the case against them and an opportunity to respond before the decision is taken. In the

algorithmic context, this raises immediate problems: if the decision is produced by a model that does not “reason” in a human-accessible manner, how can the individual know the basis on which an adverse outcome was reached? Carnat argues that the automation of administrative decisions constitutes a delegation of power that must be constrained by constitutional principles, particularly the requirement that the decision-maker be identifiable and capable of giving reasons (Carnat 2025).

Due Process Component	Requirement for Human Decision-Maker	Challenge for AI System	Potential Technical Mitigation
Notice of case	Decision-maker explains adverse evidence	Model’s feature weighting is opaque	Provide feature-importance summary (e.g., SHAP values)
Opportunity to respond	Individual can challenge evidence	Model updates dynamically; static snapshot needed	Freeze model version during response window
Impartial adjudication	Decision-maker has no personal bias	Model may encode historical bias	Regular bias audits; de-biasing during training

Table A.1: Mapping due process components to AI system challenges and mitigations.

The key insight is that the right to be heard can be preserved, but only if the system is designed to support it. This requires a deliberate architectural choice to produce **human-interpretable outputs** at the moment of decision. Without such outputs, the right to a fair hearing becomes illusory (Khalid Ranjha 2025). The framework therefore suggests that any AI system used in a context where due process protections apply must include an **explanation module** that generates a rationale in terms of the evidence presented, rather than merely a score or classification.

A.2.2 Transparency and Explainability

Transparency obligations in administrative law typically require that decisions be accompanied by reasons that are intelligible to the affected party. The literature distinguishes between

system transparency (openness about how the system works) and **outcome transparency** (the ability to explain a specific decision) (Molins 2026; Ndlovu 2025). Both are challenged by the opacity of many contemporary machine learning models, particularly deep neural networks.

Transparency			
Type	Legal Basis	Technical Requirement	Feasibility
System transparency	Duty to publish legal basis of	Model architecture	High for rule-based
	decision-making	documentation, training data provenance	systems; low for complex deep learning
Outcome transparency	Duty to give reasons for individual decisions	Local explainability (LIME, SHAP, counterfactuals)	Medium - depends on model type and domain
Process transparency	Right to know procedural steps	Audit trail of data inputs, model version, confidence thresholds	High with proper logging infrastructure

Table A.2: Transparency subtypes and their technical counterparts.

The framework argues that outcome transparency is the minimum necessary to satisfy administrative law requirements, because it directly enables judicial review. Without an explanation that links the decision to the facts of the case, a court cannot assess whether the decision was rational or proportionate (Carnat 2025). This implies that AI systems deployed in public-sector decision-making must be **explainable by design** - that is, their architecture must permit the generation of a decision-specific rationale, not merely a prediction.

A.2.3 Judicial Review and the Grounds of Review

Judicial review of administrative decisions traditionally examines legality, procedural fairness, rationality, and proportionality. Each of these grounds presupposes that the decision-making process can be inspected after the fact. Algorithmic decisions pose several difficulties here.

First, **legality** requires that the decision-maker act within the scope of their delegated powers. If the AI system uses a model that effectively defines its own decision criteria (through ma-

chine learning), there is a risk that the system exceeds the legal mandate (Kavungal 2026). Second, **rationality** - the *Wednesday* test in UK law - requires that the decision be one that a reasonable decision-maker could have reached. A model that learns patterns from historical data may embed past irrationalities, making the decision rational from a statistical perspective but unreasonable as a matter of law (Khalid Ranjha 2025). Third, **proportionality** demands that the decision not go beyond what is necessary. An algorithmic system that uses hundreds of features may produce decisions that are more restrictive than necessary, particularly if the model picks up on correlated proxy variables.

Ground of			
Review	Traditional Test	Algorithmic Challenge	Suggested Adaptation
Legality	Did decision-maker act within power?	Model may infer criteria beyond statutory authority	Clear legal boundaries for feature set; human override required for novel cases
Rationality	Was decision one a reasonable person could make?	Model may replicate historical errors	Regular validation against human-decision benchmarks
Proportionality	Was decision no more than necessary?	Model may use irrelevant proxies (e.g. postcode as proxy for income)	Pre-screening of features; prohibition of certain protected characteristics

Table A.3: Judicial review grounds and corresponding algorithmic challenges.

The framework suggests that administrative courts may need to develop a new head of review - **algorithmic reasonableness** - that assesses whether the model's reasoning process is consistent with the legal purpose, even if the outcome appears rational on its face.

A.2.4 Accountability Mechanisms

Accountability in the administrative state typically flows through a chain: the individual decision-maker is answerable to their superiors, who are answerable to ministers or elected officials,

who are answerable to the legislature and ultimately to the electorate. AI systems disrupt this chain in two ways. First, the decision-maker is replaced by a technical system that has no capacity for answerability. Second, responsibility for the system’s behaviour is distributed across many actors - developers, data providers, deploying agencies - making it difficult to assign blame when something goes wrong (Molins 2026; Wirtz and Müller 2025).

Kavungal argues that the central question is not *how* to deploy AI but *whether* specific decisions should be delegated to AI at all (Kavungal 2026). This aligns with the **non-delegation principle** that runs through many constitutional systems: certain discretionary powers cannot be delegated to automated processes without explicit legislative authorisation. The framework therefore includes a decision tree that helps determine whether a given administrative decision is suitable for automation.

Decision Factor	High Suitability for Automation	Low Suitability for Automation	Doctrinal Principle
Discretion level	No discretion (rule-based)	Broad discretion	Non-delegation; legislative intent
Impact on rights	Low (information provision)	High (deprivation of liberty, welfare benefits)	Proportionality; due process
Context sensitivity	Stable, well-defined	Evolving, requires judgement	Reasonableness review
Error tolerance	Low cost of error	High cost (e.g. wrongful detention)	Accountability; right to effective remedy

Table A.4: Suitability for automation - decision factors aligned with doctrinal principles.

The framework posits that the greater the impact on individual rights and the wider the discretion, the stronger the presumption against full automation. In such cases, a **human-in-the-loop** design is constitutionally required, not merely good practice.

A.3 Framework Diagram (Textual Representation)

The overall framework can be summarised as a three-stage process:

1. **Characterisation** - Identify the technical properties of the AI system (opacity, autonomy, data reliance, scalability).
2. **Mapping** - For each relevant doctrinal requirement (due process, transparency, review, accountability), determine which technical characteristics pose the greatest challenge, and what mitigations are available.
3. **Operationalisation** - Select specific design constraints (e.g. explanation module, audit trail, human override) that translate the doctrinal requirement into a measurable system property.

This mapping is not one-to-one; a single technical characteristic, such as model opacity, affects multiple doctrinal requirements simultaneously. Conversely, one doctrinal requirement, such as due process, may require multiple technical mitigations.

A.4 Limitations of the Framework

The framework presented here is based on a synthesis of the existing literature (Kavungal 2026; Khalid Ranjha 2025; Molins 2026; Carnat 2025) and is intended as a starting point for operationalising constitutional constraints. It has not been empirically validated, and its applicability may vary across jurisdictions with different administrative law traditions. For example, the German approach, with its strong reliance on the principle of legality and the prohibition of full automation in discretionary matters (Kim 2024), may require a stricter mapping than the more flexible UK common law framework. Future research should test this framework against real-world AI deployment cases and refine the mapping through comparative analysis.

Appendix B: Supplementary Data Tables

B.1 Comparative Overview of National AI Governance Frameworks

The following table summarises the key features of selected national and supranational AI governance frameworks that are relevant to constitutional and administrative law constraints. The data are drawn from the literature cited in the main thesis.

	Primary	Scope of	Key Constraint on	Relevant
Jurisdiction	Legislation/Guidance	Application	Public-Sector AI	Citation
European Union	AI Act (2024/1689)	High-risk AI systems, including public-sector uses	Mandatory human oversight; transparency obligations; conformity assessment	(Wagner 2025); context only
Germany	VwVfG Section 35a	Automated administrative acts	Prohibition on full automation where discretion is exercised	(Kim 2024)
United Kingdom	Common law (procedural fairness); Equality Act 2010	All public authority decisions	Duty to give reasons; rationality review; non-discrimination	(Khalid Ranjha 2025)
United States	Administrative Procedure Act (APA); NIST AI Risk Management Framework	Federal agency decisions	Notice-and-comment rulemaking; arbitrary and capricious review	(Malhotra 2025)
Canada	Directive on Automated Decision-Making (Treasury Board)	Federal automated systems	Algorithmic impact assessment; human-in-the-loop for high-impact decisions	(Kavungal 2026)

Table B.1: Comparative governance frameworks for public-sector AI.

B.2 AI Use Cases in European Public Administration

Van Noordt and Misuraca identified over seventy AI use cases across EU member states (van Noordt and Misuraca 2022). The table below groups these use cases by administrative domain and identifies the primary doctrinal constraints that would apply.

Domain	Example Use Cases	Primary Doctrinal	
		Concern	Key Constraint
Social welfare	Benefit eligibility assessment; fraud detection	Due process (right to be heard); accountability	Human review required for adverse decisions
Immigration	Visa processing; asylum screening	Procedural fairness; right to effective remedy	Explainable decisions; right to appeal
Taxation	Tax return auditing; risk scoring	Legality; proportionality	Clear legal basis; minimise over-inclusion
Policing and justice	Predictive policing; parole risk assessment	Fair trial; non-discrimination	Bias audit; human override for high-risk decisions
Health administration	Resource allocation; priority scoring	Equal treatment; right to health	Transparency in criteria; appeal process

Table B.2: AI use cases in European public administration and associated doctrinal concerns.

B.3 Summary of Technical Characteristics and Their Legal Implications

This table synthesises material from the literature (Khalid Ranjha 2025; Molins 2026; Carnat 2025; Fernández 2025) on how specific technical characteristics of AI systems create tensions with administrative law.

Technical			
Characteristic	Description	Legal Tension	Potential Mitigation
Opacity (black-box models)	Internal logic is not interpretable	Violates duty to give reasons; hinders judicial review	Use interpretable models (e.g. logistic regression, decision trees) for high-impact decisions
Non-determinism	Same input may produce different outputs	Undermines consistency and legitimate expectations	Fixed random seed; version-locked model deployments
Data dependency	Model behaviour changes with training data	Risk of embedding historical bias; lack of transparency about data quality	Mandatory data provenance documentation; bias testing
Scalability	Many decisions processed quickly	Undermines individualised consideration; may violate proportionality	Case-by-case human review for decisions with meaningful individual impact
Continuous learning	Model updates over time without human approval	Bypasses legislative approval of decision-making criteria; undermines predictability	Freeze model after deployment; require legislative authorisation for substantial updates

Table B.3: Technical characteristics and their legal implications.

B.4 Glossary of Key Terms (Compact Version)

For a full glossary, see Appendix C. The table below provides a quick reference to the most critical terms from a constitutional-administrative law perspective.

Term	Definition	Relevance to Thesis
Algorithmic decision system	A computational system that produces or substantially influences administrative decisions	Central object of analysis
Due process	Procedural protections ensuring a fair hearing before adverse state action	Core doctrinal requirement
Judicial review	Court oversight of administrative decisions on grounds of legality, rationality, and procedural fairness	Pathway for enforcing constraints
Non-delegation doctrine	Principle that discretionary powers granted by the legislature cannot be redelegated to automated processes	Key constitutional limitation on AI autonomy
Explainability	Ability to provide human-understandable reasons for a decision	Technical requirement enabling transparency and review

Table B.4: Compact glossary of key terms.

Appendix C: Glossary of Terms

This glossary defines the key technical, legal, and interdisciplinary terms used throughout the thesis. The definitions are drawn from the literature and standard legal sources. Citations are provided where the definition is directly adopted or paraphrased from a cited work.

C.1 Administrative Law Terms

C.1.1 Due Process

The principle that an individual must be given notice of a proposed adverse decision and an opportunity to be heard before that decision is taken. In the UK context, it is embodied in the

common law duty of procedural fairness; in Germany, it is codified in Section 28 VwVfG; in the United States, it is a constitutional requirement under the Fifth and Fourteenth Amendments. In the context of AI decisions, due process raises the question of whether a computational system can, in principle, afford a fair hearing (Khalid Ranjha 2025; Carnat 2025).

C.1.2 Judicial Review

The power of courts to examine the legality, rationality, and procedural propriety of decisions made by public bodies. The grounds of review include illegality, irrationality, procedural impropriety, and proportionality (where applicable). With algorithmic decisions, courts must adapt these grounds to assess not only the outcome but the reasoning process of the machine (Clayton et al. 2018).

C.1.3 Proportionality

A test used in many jurisdictions to determine whether a limitation on rights is no more than necessary to achieve a legitimate aim. In the AI context, this may require that the system's feature set and decision thresholds be tailored to the legal purpose rather than expanded to maximise predictive performance (Khalid Ranjha 2025).

C.1.4 Non-Delegation

The doctrine that a legislature cannot delegate its law-making power to another body without providing an intelligible principle. Applied to AI, it suggests that discretionary decision-making cannot be fully automated unless the legislature has expressly authorised such automation and defined the limits (Kim 2024; Carnat 2025).

C.1.5 Reason-Giving Duty

An administrative law obligation to provide reasons for a decision sufficient to enable a party to understand the basis for the outcome and to decide whether to seek review. This duty is central to the debate on AI explainability (Kavungal 2026).

C.2 Technical Terms

C.2.1 Machine Learning (ML)

A subfield of artificial intelligence in which systems learn patterns from large amounts of data without being explicitly programmed for each decision rule. ML models can be opaque, making their decisions difficult to explain (Ndlovu 2025).

C.2.2 Opacity

The property of an algorithmic system that makes its internal decision-making process inaccessible to human inspection. Opacity can originate from three sources: proprietary secrecy, technical complexity, or the sheer non-linearity of the model (Molins 2026).

C.2.3 Explainability

The capacity of an AI system to provide an explanation for a specific decision in terms that are understandable to a human. In the legal context, explainability is a prerequisite for reasons-giving and judicial review (Khalid Ranjha 2025).

C.2.4 Human-in-the-Loop (HITL)

A design in which a human operator must approve or override the AI's decision before it takes effect. This is often proposed as a safeguard for high-impact decisions to preserve human accountability (Kavungal 2026).

C.2.5 Audit Trail

A chronological record of system activities that enables reconstruction of the decision-making process for a given case. In public-sector AI, an audit trail should include the model version, input data, any intermediate scores, and the final decision (Wirtz and Müller 2025).

C.3 Interdisciplinary Terms

C.3.1 Algorithmic Governance

The use of algorithms and AI systems to manage public services, allocate resources, or enforce regulations. It implies a shift from human-centred to computationally-mediated decision-making (van Noordt and Misuraca 2022).

C.3.2 Accountability

The obligation to answer for the outcomes of a decision. In the algorithmic context, accountability is diffused across developers, deployers, and operators. A key research question is how to fix responsibility when no single actor can fully control the system (Molins 2026).

C.3.3 Digital Right to Good Administration

An emerging concept in EU law that combines principles of transparency, non-discrimination, and due process applied to digital public services. It is not yet fully defined but is gaining traction in academic commentary (Kavungal 2026; Wirtz and Müller 2025).

C.3.4 Administrative Constitutionalism

A theoretical perspective that views administrative law doctrines as constitutional safeguards against the arbitrary exercise of public power. This lens is used in the thesis to evaluate whether current doctrines can adapt to algorithmic decision-making (Carnat 2025).

Appendix D: Additional Resources

D.1 Key Legislation and Policy Documents

The following list identifies the principal legal instruments and policy documents referenced in the literature that are directly relevant to the constitutional and administrative law constraints on public-sector AI. This list is not exhaustive but provides a starting point for further research.

Document Title	Jurisdiction/Organisation	Year	Relevance to Thesis	Source
Regulation (EU) 2024/1689 (EU AI Act)	European Union	2024	Establishes risk-based categories for AI, including high-risk public-sector uses; mandates human oversight	(Wagner 2025)
VwVfG (Verwaltungsverfahrensgesetz) Section 35a	Germany	1976 (amended 2017)	Prohibits full automation of discretionary administrative acts	(Kim 2024)
Directive on Automated Decision-Making	Canada (Treasury Board)	2019	Requires algorithmic impact assessment and human-in-the-loop for high-impact decisions	(Kavungal 2026)
U.S. Administrative Procedure Act (5 U.S.C. Section 551 et seq.)	United States	1946	Governs rulemaking and adjudication; applies to algorithmic agency decisions	(Malhotra 2025)
NIST AI Risk Management Framework (AI RMF 1.0)	United States (NIST)	2023	Provides voluntary guidance on AI trustworthiness; used for public-sector risk assessment	(Malhotra 2025)

Table D.1: Key legislation and policy documents.

D.2 Recommended Further Reading

Based on the literature reviewed in this thesis, the following works are particularly relevant for readers wishing to deepen their understanding of specific subtopics.

Subtopic	Recommended Reading	Reason for Recommendation
Due process and automation	Carnat (2025) “Automation as Delegation of Power” (Carnat 2025)	Provides a rigorous constitutional analysis of why automated decisions require legislative authorisation
AI governance frameworks	Kavungal (2026) “From Adoption to Accountability” (Kavungal 2026)	Argues for a shift from implementation to delegation questions; proposes a civic AI governance framework
Accountability mechanisms	Molins (2026) “Auditing Accountability” (Molins 2026)	Examines the gap between AI audits and legal accountability
Comparative administrative law and AI	Khalid Ranjha (2025) “Algorithms and Administrative Law” (Khalid Ranjha 2025)	Offers a comparative framework covering UK, EU, and US approaches
Public-sector AI adoption in Europe	van Noordt and Misuraca (2022) “AI for the public sector” (van Noordt and Misuraca 2022)	Provides empirical landscaping of AI use cases across EU member states
Procurement law as a constraint	Fernández (2025) “Relevance of procurement law for the use of AI in public administration” (Fernández 2025)	Highlights how procurement rules can be used to embed constitutional safeguards in AI system acquisition
German legal framework	Kim (2024) “German Discusses about AI, Full Automation in Discretion” (Kim 2024)	Summarises the strict German approach to automation of discretionary decisions

Table D.2: Recommended further reading by subtopic.

D.3 Tools and Methodological Resources

The following tools and methodologies are mentioned in the literature as having application to the analysis of public-sector AI systems from a legal perspective. This list is illustrative, not exhaustive.

- **Algorithmic Impact Assessments (AIA)** - Structured frameworks for evaluating the potential harm and legal compliance of an AI system before deployment. The Canadian Directive on Automated Decision-Making mandates an AIA for all federal systems (Kavungal 2026).
- **Explainability methods** - LIME (Local Interpretable Model-agnostic Explanations) and SHAP (SHapley Additive exPlanations) are two common techniques for generating local explanations of individual model predictions. These are relevant to the reasons-giving duty (Khalid Ranjha 2025).
- **Bias audit toolkits** - Tools such as AI Fairness 360 (IBM) and What-If Tool (Google) can be used to detect and mitigate discriminatory bias in models. The literature emphasises that public-sector systems must undergo regular bias audits to comply with equality law (Polo and Ailodion 2025).
- **Audit trace logging** - Systems such as MLflow and DVC enable recording of model versions, parameters, and training data provenance. These are critical for creating an audit trail that satisfies judicial review requirements (Wirtz and Müller 2025).

D.4 Organisations and Initiatives

Several organisations are active in developing standards and guidance for constitutional and administrative law constraints on AI. Readers may wish to consult their publications.

- **The Institute for Law & AI (London)** - Publishes research on the legal implications of frontier AI, including public-sector applications (Bullock et al. 2024).
- **The European Commission's Joint Research Centre (JRC)** - Produces reports on AI adoption in the public sector and related legal issues (van Noordt and Misuraca 2022).

- **The OECD’s Observatory of Public Sector Innovation (OPSI)** - Curates case studies and policy guidance on AI in government.
- **The Ada Lovelace Institute** - Conducts research on the societal and legal implications of AI, including algorithmic accountability.

D.5 Glossary (Compact Version)

For the full glossary, see Appendix C. The compact version in Table B.4 of Appendix B provides a quick reference. This appendix repeats the most essential terms for readers who consult Appendix D directly.

Term	Definition
Algorithmic decision system	A computational system producing or influencing administrative decisions
Due process	Right to notice and hearing before adverse state action
Judicial review	Court scrutiny of administrative decisions on legal grounds
Non-delegation	Principle that legislative discretion cannot be fully transferred to an automated system
Explainability	Ability to provide human-understandable reasons for a decision

These appendices are intended as supplementary material to the main thesis. The framework, tables, and resources are based on the literature cited and should be read in conjunction with the discussion in Sections 2.3 and 2.4.

References

- Bullock, Van Arsdale, Arnold, Maas, and Winter. 2024. "Existing Authorities for Oversight of Frontier AI Models Institute for Law & AI." <https://doi.org/10.2139/ssrn.4976362>.
- Carnat. 2025. "Automation as Delegation of Power: Constitutional Constraints on AI Systems for the Administration of Justice." <https://doi.org/10.2139/ssrn.5690283>.
- Clayton, Firth, Sawyer, Moffatt, and Wray. 2018. 7. *Challenging decisions: appeals, administrative and judicial review*. Oxford University Press. <https://doi.org/10.1093/he/9780198815211.003.0007>.
- Elkadi, Sabra, and Yacoub. 2025. "White Paper on Responsible AI Usage by Social Media Platforms from African Human Rights Law Perspective." <https://doi.org/10.2139/ssrn.5368235>.
- Fernández. 2025. *Relevance of procurement law for the use of AI in public administration*. Edward Elgar Publishing. <https://doi.org/10.4337/9781035311736.00006>.
- Graham and Russell. 2025. *Reflecting Upon Key UK Supreme Court Cases in Public Law*. Routledge. <https://doi.org/10.4324/9781003477068-38>.
- Kavungal. 2026. "From Adoption to Accountability: A Civic AI Governance Framework for Public Sector Systems ." <https://doi.org/10.2139/ssrn.6646818>.
- Khalid Ranjha. 2025. "Algorithms and Administrative Law: A Comparative Framework for Accountability." <https://doi.org/10.2139/ssrn.5645610>.
- Kim. 2024. "German Discusses about Artificial Intelligence (AI), Full Automation in Discretion." *Korean Administrative Law Association* 26: 33-58. <https://doi.org/10.59826/kdps.2024.26.33>.
- Lāma and Lastovska. 2025. "Exploring Public Sector Workforce Attitudes Toward Ai Competence and its Impact on Ai Competence Self-Assessment." <https://doi.org/10.2139/ssrn.5127227>.
- Malhotra. 2025. "When AI Governs the State: A NIST-Guided Framework for Managing Generative AI Risk in U.S. Public Institutions." <https://doi.org/10.2139/ssrn.5779442>.

Molins. 2026. "Auditing Accountability; Assessing AI Audits and AI Legal Frameworks Through the Lens of Accountability ," <https://doi.org/10.2139/ssrn.6155866>.

Ndlovu. 2025. *AI-Driven Tools for Sustainable Public Administration*. IGI Global. <https://doi.org/10.4018/979-8-3693-8372-8.ch004>.

Polo and Ailodion. 2025. "Tackling Racial Bias in AI Systems: Applying the Bioethical Principle of Justice and Insights from Joy Buolamwini's "Coded Bias" and the "Algorithmic Justice League"." *Bangladesh Journal of Bioethics* 16, no. 1: 8-14. <https://doi.org/10.62865/bjbio.v16i1.129>.

Sadurski. 2022. *Standards of Scrutiny, Equal Protection, and Illicit Motives for Discrimination*. Oxford University PressOxford. <https://doi.org/10.1093/oso/9780192869678.003.0008>.

Saul. 2007. *Australian administrative law: The human rights dimension*. Cambridge University Press. <https://doi.org/10.1017/cbo9781139168618.006>.

Valcke, Clifford, and Dessers. 2021. *Constitutional Challenges in the Emotional AI Era*. Cambridge University Press. <https://doi.org/10.1017/9781108914857.005>.

van Noordt and Misuraca. 2022. "Artificial intelligence for the public sector: results of landscaping the use of AI in government across the European Union." *Government Information Quarterly* 39, no. 3: 101714. <https://doi.org/10.1016/j.giq.2022.101714>.

Wagner. 2025. "EU AI Act (2024/1689) and EU MDR (2017/745): Breaking the Expensive Myth: Why AI-Powered Medical Devices Under EU MDR Don't Need EU AI Act Certification - A Detailed Analysis of Regulatory Requirements and Compliance." <https://doi.org/10.70317/2025.02rw02>.

Wirtz and Müller. 2025. *AI Governance Strategies for the Public Sector*. Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-00514-4_16.

Yadong. 2024. *AI Rule of Law Development Index 3.0*. Springer Nature Singapore. https://doi.org/10.1007/978-981-97-1060-7_2.