Cambridge Forum on AI: Law and Governance (2026), **2**, e15, 1–17
doi:10.1017/cfl.2026.10063

Cambridge Forum

RESEARCH ARTICLE

Limiting artificial intelligence (with the Rule of Law)

Paul Burgess

Faculty of Law, Monash University, Melbourne, Australia
Email: paul.burgess@monash.edu

(Received 30 September 2025 UTC; revised 17 May 2026 UTC; accepted 3 July 2026 UTC)

Abstract

Artificial intelligence (AI) is increasingly used in public administration, yet its opacity raises concerns about arbitrary exercises of state power. This article argues that alongside regulatory frameworks, the Rule of Law supplies an independent basis for limiting state deployment of AI. Drawing on widely accepted ideas associated with the Rule of Law – predictability (decisions must guide conduct ex ante) and procedural clarity (rationales must be knowable and contestable ex post) – the article shows how these standards constrain government use of AI. It examines why black box systems frustrate legal guidance and contestability, and why explanations are often the minimally sufficient remedy. A welfare-benefits hypothetical illustrates the stakes of opacity versus explainability. The article concludes that Rule of Law limits offer a distinct and operationalizable basis for determining when state deployment of AI contravenes fundamental commitments, demonstrating that not all government uses of AI should be permitted.

Keywords: artificial intelligence; Rule of Law; predictability; procedural clarity; provider/developer obligations; user/deployer obligations

1. Introduction

Artificial intelligence (AI) is demonstrably capable of shaping humankind's future (Aliu, 2025; OECD, 2025). There are, however, few voices that promote the completely unregulated use or development of AI (Veale & Edwards, 2018). The rationale for limitation may include ethical, sociological, economic or practical reasons. This paper argues that public-sector use of AI is legitimate only where it satisfies commonly accepted and identifiable ideas associated with the Rule of Law standard: decisions must be predictable (capable of guiding conduct ex ante) and procedurally clear (rules and rationales are knowable and contestable ex post). In doing this, an additional avenue of limitation – over and above regulatory approaches that have been adopted in some jurisdictions – is proposed. While the limitations that stem from the concept of the Rule of Law are not as extensive or as fine-grained as some regulatory approaches, they do facilitate a broad-based approach to ensuring that the exercise of public power using AI systems is subject to some limitations (even where specific regulations do not exist).

For this paper, "AI" is used to refer to AI systems in the sense set out in the updated OECD definition: an AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments, and that different AI systems vary in their levels of autonomy and adaptiveness after deployment (OECD, 2025). It is acknowledged that definitions

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

of AI vary across regulatory frameworks (compare, e.g., the European Union's (EU's) AI Act [2024]); the OECD formulation is adopted here both for its breadth and because it is the formulation that has been widely adopted across the regulatory frameworks that this article engages with.

The paper contributes three things. First, it maps thin Rule of Law requirements, showing why black box systems frustrate legal guidance and contestability (Pasquale, 2016) and why explanations from AI systems are often the minimally sufficient remedy. Second, it reconciles Rule of Law constraints with emerging legal frameworks – clarifying ways in which the Rule-of-Law limitations that are identified fit within existing frameworks that include developers/users and private/state entities. Third, it operationalizes the Rule of Law limitations to arrive at a simple conclusion: not all government uses of AI should be allowed. The Rule-of-Law limitation that is developed can, however, provide a mechanism to identify clear-cut (and different) ways in which states' uses of AI to exercise power may be limited. Although AI developers and private actors also raise important concerns, this paper focuses primarily on the state as user of AI systems in the exercise of public power.

In this paper, the broad-based position that AI should be limited is accepted, and it is the way in which this limitation should be effected that is put in issue. The argument put forward in this paper is that the frequently used regulatory schemes are not the only, and may not necessarily be the most effective, way to achieve this end. This is obviously the case in jurisdictions where no specific regulations exist. Another way to achieve this limitation is proposed: through the Rule of Law. In short, states' continued purported adherence to the concept of the Rule of Law creates limitations on the way those states may use AI when exercising power.

A note on terminology is appropriate here. Throughout this paper, "limit" is used in a specific sense: as a constraint on the kinds of AI use that a state may legitimately engage in. "Regulation," by contrast, is used to refer to the broader administrative apparatus through which states (and sometimes private actors) seek to govern the development and deployment of AI – typified, for example, by the EU AI Act (2024). The two terms are related but not equivalent: not every regulation operates by way of a (substantive) limit, and not every limit is regulatory in form. The paper's core claim concerns the former – the Rule of Law as a substantive limit – and not the latter.

The primary claim of this paper – and its key contribution to the literature – is that, although the Rule of Law does not operate as a direct legal limitation on the use of AI by states in exercising power, the existence of the Rule of Law in states that profess to adhere to it generates an indirect limitation: where a technology contravenes the Rule of Law (and the state would not wish to do this) then a limitation, based on the Rule of Law, exists in relation to that technology. Putting this another way, if a state purports to adhere to the Rule of Law, it should not take actions that contravene the Rule of Law; and if certain uses of AI contravene the Rule of Law, then those uses should not be deployed by the state. This is the article's primary claim and is offered as its key – and, in this form, novel – contribution: rather than treating the Rule of Law as an external regulatory tool, the paper develops it as an internal, political-normative limit that arises from a state's own commitment to Rule-of-Law compliance.

The use of the Rule of Law in this way may not be the intuitive choice for imposing limitations on the use or application of a new technology. The concept is the subject of a substantial secondary literature exploring its origins and historical development. This includes engagement with classical and premodern sources (see, e.g., Krygier, 2009; King, 2018; Sempill, 2017; and the special issue of the *Hague Journal on the Rule of Law* [volume 9, 2017] on the historical origins of the concept as well as Burgess (2020) who identifies the assumption that the Rule of Law has existed for over 2000 years as one of four common assumptions about the Rule of Law). Against this backdrop, it may seem peculiar to seek to use the idea in the context of 21st-century technological advancements; yet, the concept's meaning has been changed, adapted and evolved over the intervening millennia so that it is still frequently touted as being the most important political ideal – and it has remained entrenched in political systems around the world for centuries (Sevel, 2024). While the centuries of differing ideas

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

raise their own complications, the Rule of Law has, over that period, been a normative force that seeks to constrain the arbitrary exercise of power (Bingham, 2010; King, 2019; Krygier, 2017). It is in this sense that it is possible to see the overlap with AI's use: the substantial efficiencies and benefits that AI may offer make it an attractive way in which to aid or facilitate the exercise of power. If AI is used in this way, there is a Rule of Law-relevant connection.

The Rule of Law offers a distinct mode of limitation: it operates indirectly, through a state's own commitment to the concept, even where the term is not defined in (or absent from) the state's constitutional documents. It is this indirect mode of limitation – rather than the substitution of one regulatory regime for another – that the paper develops in what follows. There may be pros and cons of a purely regulatory approach toward AI. Given the way in which the Rule of Law may operate, as described above, there could also be pros and cons to the *way* in which the Rule of Law may do this – especially where there is both little agreement as to the terms of the concept and, relatedly, where the concept is not defined or included in most states' constitutional documents.

The paper proceeds as follows. Section 2 sets out the need to limit AI and introduces a developer/user × public/private taxonomy of attempts to do so. Section 3 surveys current regulatory approaches and notes their limitations. Section 4 – the substantive core – develops the argument that the Rule of Law, through its requirements of predictability and procedural clarity, itself supplies a limit on state use of AI; an illustrative welfare-benefits hypothetical is used to demonstrate this before the limits on those limits are considered. It should be noted in advance that the Rule of Law-based limitation developed in this paper is not without significant complications: the precise location of the nuanced positions it implies is difficult to identify (explored in subsection 4.5), and the resulting limit, even when operationalized through procurement, is at best one component of any broader regulatory scheme. This aspect is explored in section 5. That section considers operationalization, with particular attention to procurement. The novel conclusion of the paper is that, in a state that professes adherence to the Rule of Law, the Rule of Law itself generates an internal political-normative limit on the state's use of AI: where a particular use of AI contravenes the Rule of Law, that use should not be deployed by the state; and the limit can be made operational through ordinary state mechanisms – most directly, procurement – without the Rule of Law needing to be either codified in constitutional documents or justiciable in form.

2. The need to limit AI

The prominence of attempts to limit AI (in some way) has increased in recent years. Some of the more prominent limitations have included the EU AI Act, which after extensive negotiation and phased implementation became law in mid-2024, and the United States' executive orders in recent years such as EO 13859 (2019) and EO 14110 (2023) (United States, 2019, 2023) (noting that EO 14110 was rescinded on 21 January 2025). However, attempts have been underway to limit AI – and other methods of automated decision-making – for far longer (Veale & Edwards, 2018). While the trend is not necessarily new, it demonstrates that there is not only an appetite for the limitation of AI but also an increasingly explicit desire on the part of states and publics for such limitations to be introduced (OECD, 2025; UNESCO, 2022) – as well as being periodically updated (e.g., through the OECD's revised definition and explanation of an AI system [OECD, 2025]).

Attempts to limit AI can broadly be grouped into two categories: measures directed at developers/providers and measures directed at users/deployers. The limits imposed on developers/providers – for example, those seen in the EU's AI Act – can be seen to emphasize transparency, nondiscrimination and accountability in the design and the deployment of systems. Limits on users of AI, whether state agencies or private actors, may impose duties concerning procurement, application and oversight of AI tools (Jacobs & Simon, 2022). For example, public-sector procurement guidance recommends integrating explainability, access to information and assurance provisions into AI tenders and contracts (UK Government, 2020). On this basis, these categories can be considered in terms

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

	State entity	**Private entity**
Developer / provider	State developers / state providers	Private developers / private providers
User / deployer	State users / state deployers	Private users / private deployers

Figure 1. Simplified matrix of users/deployers and developers/providers according to their status as either public/private entities.

of the form of entity that may operate as a developer/provider or as a user/deployer: state or private entities. These are illustrated in Figure 1.

It is worth pausing on a related question: why focus on AI specifically, rather than on complex technical systems more generally? Several factors justify the AI-specific focus adopted here. First, the qualitative scale and opacity of contemporary AI systems – particularly those based on machine learning – present a degree of inscrutability that is qualitatively different from that found in conventional algorithmic or rule-based systems (Burrell, 2016). Second, AI usage has become popular in many aspects of public-sector activity, in part because of the substantial economic benefits that states associate with its adoption: cost reductions, efficiency gains and the scaling of decisional capacity all incentivize rapid uptake (Aliu, 2025; OECD, 2025). Third, AI systems have policy salience: their deployment by states is the subject of dedicated international (OECD and UNESCO) and supranational (EU AI Act) regulatory activity, indicating that they have been identified as warranting distinct legal and policy attention. These three features – opacity, ubiquity (driven by economic incentives) and policy salience – together justify treating AI as a regulatory target in its own right, even though, as Lyria Bennett Moses (2007) has shown more generally, law's response to technological change always involves difficult questions about the proper unit of regulation.

The quadrants in Figure 1 consider developers as relating to those entities that create AI systems in comparison to those that are the end users of those systems; and the relative status of users or developers as being either state (public) or private entities. While this distinction may not specifically and explicitly be adopted elsewhere, facets can be seen in terms of both the focus of key regulatory regimes and in the way that they are discussed; for example, in exploring public authorities developing and acquiring systems (Mir, 2025); in the similarities and differences between public and private entities in their use (Hickok, 2024); or, public buyers purchasing privately developed systems (Sanchez-Graells, 2024). Later in this paper, an example will focus on the state-user through a hypothetical example of the use of AI in the deployment of welfare benefits.

For reasons that relate to the widest possible application of the concept of the Rule of Law as a limiting factor, this article is mainly focused on the potential limitations that operate in only one quadrant of the illustration in Figure 1: the state-user quadrant. However, as will be seen later in the

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

article, there are wider potential implications should the Rule of Law be construed in less-agnostic ways.

3. The limitation of AI (commonly)

The previous section outlined, in broad terms, some of the ways in which AI has been limited. This section provides a brief account of key regulatory frameworks that currently exist to note some of the high-level problems that exist – including the difficulties associated with the regulation of technologies that may not yet exist within a regulatory regime that is slow-moving, together with the problems that may stem from industry-led or industry-monitored schemes.

The risk-based approach adopted by the EU imposes a duty on both those providing the AI systems (developers) and those deploying the systems (the users) to ensure certain standards are met. For example, developers of systems categorized as high-risk must take certain steps to ensure, among other things, appropriate data and data governance (Article 10) and human oversight (Article 14) (EU AI Act, 2024). The EU AI Act differentiates the levels of risk based on whether an entity is a user (deployer) or a provider (developer) – with more obligations being placed on providers than users. The applicable standards mean that high-risk systems are the most regulated, limited-risk systems are less regulated and minimal-risk systems are largely unregulated. The majority of obligations are placed on providers (developers) of high-risk AI systems (see Chapter III, including Article 16 [providers] and Article 26 [deployers] [EU AI Act, 2024]). Of course, the EU AI Act is not the only potential form of regulation, and while alignment could be seen to be beneficial (Alex Engler, 2023), this is not what has happened (Wang et al., 2025). Not only do the approaches in the EU and the United States differ in various ways (Shahin, 2024), but they also pose challenges for one another (Cantero Gamito & Marsden, 2024).

Some problems have been noted in utilizing the EU AI Act in a wholesale form, particularly outside of the EU (Cantero Gamito & Marsden, 2024). These concerns can be extended to the potential of future-proofing regulatory regimes in a fast-moving field – even if the EU's approach is designed as a way to avoid overregulation (Ebers, 2025). These sorts of concerns can be seen in general terms that are not solely limited to the EU AI Act. As a further potential problem, attempts to create laws or regulations that are technology-neutral – in an effort to keep pace with changing technologies – may not result in ideal outcomes but, instead, it may be possible to design a legal system that treats new technologies fairly and is resistant to difficulties associated with technological change (Lyria Bennett Moses, 2007).

What can be seen is that the limitation of AI commonly takes the form of regulatory schemes focused on either the developer or the user, that the proliferation of different schemes can cause tension in their application and that it may be useful to have a legal system that treats all forms of technologies fairly (despite future forms being impossible to conceive). It is with these things in mind that consideration moves to an existing and accepted component of the legal system that, while not generally considered to be a useful mechanism for limiting AI, will be shown to be one way to address the various challenges that common regulatory mechanisms have faced.

4. How to limit AI with the Rule of Law

This section explains how the Rule of Law can provide a limit on the use of AI in the exercise of state power. In this respect, and as noted earlier, focus is placed on the category of state-users of AI; however, through considering other ways in which the Rule of Law may be (and often is) conceived, limitations associated with other categories will also be explored.

To do this, the section commences by briefly sketching what is meant by "the Rule of Law" – which necessitates a consideration of the various contested forms of the concept. This takes place in the following subsection. Those ideas then lead on to a brief discussion of the ways in which the

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

Rule of Law is commonly seen to provide a limit on the exercise of public power in general before an explanation is provided, in subsection 4.3, of how this translates to Rule-of-Law limitations on the use of AI (in the exercise of public power). In subsection 4.5, the relative limitations associated with using the Rule of Law as a limitation on the way that AI is used are briefly noted. These discussions then provide the basis for exploring, in section 5, the ways in which the use of the Rule of Law as a limitation on the exercise of public power using AI can be operationalized. For coherence with the paper's central claim, the analysis throughout this section uses, and explains the rationale for adopting, two Rule of Law ideas – *predictability* and *procedural clarity* – as the operative standards for assessing state use of AI.

4.1. The Rule of Law

In its most general and least controversial form, the Rule of Law can be seen as a concept that opposes the arbitrary application of power. In this form, it has been said that nobody is against the Rule of Law (Burgess, 2024b; Krygier, 2011). Once any attempt is made to move beyond this abstract idea, consensus or agreement is lost; the concept's meaning is contested – perhaps even essentially contested (Fallon, 1997; Waldron, 2002). A substantial literature attempts to move beyond this abstract notion; only the briefest sketch of those ideas can be included within this subsection. While a detailed consideration of the complex literature is not feasible, the discussion below, wherever possible, seeks to illuminate common themes. This is an approach that is also common in the literature (Burgess, 2017; Craig, 1997; Tamanaha, 2004).

In providing this sort of brief sketch of the concept of the Rule of Law, a common approach is to identify that not only is the Rule of Law an ancient concept – often first attributed to Aristotle – but also it is seen to have a rich history – through contributions from authors like Locke, Dicey, Hayek, Fuller and Raz; and, in consequence of these very different historical conceptions, the precise concept is seen to be contested (Burgess, 2020; Shklar, 1987). As noted earlier, there is also a frequently stated suggestion that the concept is essentially contested (Fallon, 1997; Waldron, 2002).

The contested nature of the concept, and the different conceptions that have historically proved to be popular, have seen various attempts to categorize the ideas. Prominently, this includes distinguishing conceptions into thin (formal) and thick (substantive) forms (Craig, 1997; Tamanaha, 2004); this would typically see the historical conceptions noted earlier as being thin constraints on the arbitrary exercise of power – for example, by requiring certain steps to be taken in the passage of laws; whereas other conceptions are framed as thick conceptions that seek to attain a particular outcome – for example, and most prominently, those provided by Bingham or the United Nations (Bingham, 2010; UN Security Council, 2004). What is, however, notable is that the thick conceptions do still require or rely on a thin or formal base (Burgess, 2017; Craig, 1997). So, while Bingham suggests that the Rule of Law requires steps to be taken that maximize human rights, the passage of laws that achieve this end must be achieved through means that reflect thin Rule of Law ideas – for example, laws that are prospective, transparent and certain (Bingham, 2010).

It is worth flagging an important distinction at this point. Contestation does not entail vagueness. Within any single conception of the Rule of Law, the relevant terms may be (and frequently are) defined with considerable precision; what is contested is the choice between conceptions or the accuracy of the desiderata within those conceptions as a representation of the idea/meaning of the Rule of Law (Burgess, 2017; Tamanaha, 2004). The notion of contestedness in relation to the concept of the Rule of Law does not imply vagueness in the conceptions. The argument that follows therefore proceeds by adopting one well-defined thin/formal conception and developing it on its own terms; it does not rely on, or appeal to, any vagueness in the concept of the Rule of Law itself.

It is for this reason that the most unobjectionable form of the Rule of Law – that goes beyond the abstract notion that opened this subsection – can be seen to be a thin/formal conception of the Rule of Law; this would, in broad terms, be acceptable to those that favor a thin/formal conception

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

and, additionally, would not be fundamentally objected to by those that favor a thick/substantive conception. Although they would, of course, wish the extent or application of the concept to go *beyond* those thin ideas.

Simple adoption of a thin idea of the Rule of Law does not, however, assist in taking us far beyond the most abstract notion of the concept; it does not provide any analytical precision with which we can move forward. With this in mind, and for brevity, an approach to the concept that seeks to identify points of commonality across desiderata of popular thin conceptions will be adopted (Burgess, 2017; Craig, 1997). This approach identifies two ideas that undergird many of the desiderata associated with the thin ideas: the ideas of predictability and procedural clarity. These two ideas will be used later, where additional analytical precision can illuminate discussions. They also provide the paper's working test for whether particular deployments of AI by the state are legitimate.

Before turning to the engagement of Rule-of-Law scholarship with new technologies (which is set out immediately below), it is worth observing the broader literature landscape. There is, on the one hand, an extensive survey-type literature on the Rule of Law – see, for example, Burgess (2017) and Tamanaha (2004) – and, on the other, a similarly extensive (and rapidly growing) literature on technology regulation (see, e.g., Marsden, 2011; Veale et al., 2023; Veale & Zuiderveen Borgesius, 2021). There is not, however, an extensive literature that sits specifically at the intersection of the Rule of Law and technology regulation. That intersection – the body of work focused on the implications of new technologies for the Rule of Law, and on the use of the Rule of Law as a constraint on technological deployment – is the discrete and comparatively under-developed strand that the next paragraph addresses, and it is precisely that gap that motivates the present paper.

The interaction between technology – and AI in particular – and the Rule of Law has rapidly become an increasingly active area of scholarship. Importantly, several of the scholars who shaped the modern conceptual debate about the Rule of Law have either extended their attention to its implications for new technologies or have provided the conceptual foundations on which that engagement has been built. In the former category are Brownsword (whose work on technological management and the Rule of Law is a key reference [Brownsword, 2016]), Burgess (Burgess, 2017, 2020, 2024a), Gowder (whose Rule-of-Law scholarship has been extended to digital, algorithmic and platform governance [Gowder, 2018, 2023]) and Zalnieriute et al. (2019). In the latter category sits Tamanaha, whose *On the Rule of Law* (Tamanaha, 2004) has been heavily drawn upon by authors writing on AI and whose more recent work continues to develop the conceptual architecture. Tamanaha himself does not principally write on AI and the Rule of Law, but his conceptual work is heavily relied upon by authors who do (see, e.g., Huq [2024] at p. 265; Greenstein [2022a, 2022b] [the latter at pp. 1–2]; and Burgess [2024a], most prominently in chapter 2).

With this in mind, it should be noted that not every conceptual scholar of the Rule of Law has chosen to engage deeply with the AI discussion; this does not, however, diminish the relative importance of their contributions, as their ideas have been found foundational by those who do so engage. Across this body of work, three broad concerns recur: the difficulty AI systems pose for predictability and procedural transparency in the exercise of public power; the related problem of algorithmic opacity (Burrell, 2016; Pasquale, 2016); and the question of where, and whether, meaningful human oversight is required for AI-assisted government decisions to comply with Rule-of-Law values. The careful conceptual work previously undertaken by these scholars on the Rule of Law itself provides an essential foundation for this newer technology-focused Rule of Law-related literature: without that prior conceptual mapping, the field's engagement with AI would lack the analytical scaffolding required to identify both the threats AI poses to the Rule of Law and the limits the Rule of Law may, in turn, impose on the use of AI.

Before moving on, a further point of distinction must also be considered in any attempt to identify an agreeable conception: whether the Rule of Law should be constrained to the exercise of public power (from the state to the individual) or, instead, should also extend to the exercise of private power (from an individual to an individual). The historical conceptions noted earlier, which in consequence

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

of seeking to limit the exercise of power through the imposition of conditions associated with the passage of laws, all reflect the exercise of public power. These can be contrasted with more contemporary ideas that seek to extend the prevention of arbitrary power to the private sphere (Kampourakis et al., 2023; King, 2018; Krygier, 2009). What can be seen is that while the historical view of the concept was limited to the public exercise of power, there is increasingly an acceptance that the Rule of Law's remit should be extended to private exercises of power. In seeking a point of agreement, it is, however, less controversial to say that the Rule of Law relates to the exercise of public power; those who would advocate for the application of the Rule of Law in the private sphere would accept its operation (as well) in the public sphere.

In consequence of the distinctions and the commonalities explored, albeit very briefly, above, the Rule of Law will principally be considered to be a concept that seeks to limit the arbitrary application of public power through the application of rules that are both predictable and procedurally clear. These distinctions will be considered in the following two subsections (subsections 4.2 and 4.3) in relation to, first, the Rule of Law as a limitation on public power and, second, in a more specific consideration of how this translates to the Rule of Law operating as a limitation on the state's use of AI in exercising its powers.

Both of these requirements – predictability and procedural clarity – are distilled from popular primary Rule of Law conceptions (Burgess, 2017). The elements were considered to represent points of commonality across those popular conceptions. They required that an individual be able to comprehend the nature, content and operation of the rules to which they are subject; and that the creation and application of any rules be in terms capable of being clear and obvious to all concerned. These two ideas operate as paired but distinguishable requirements: predictability requires that an individual be able to know in advance how a rule will be applied to them; procedural clarity requires that the rule-making and rule-applying process – including the reasons given for any decision – be capable of being clear and obvious to all concerned and, accordingly, that the process and the decision be open to interrogation and contestation. The implicit relationship signaled later in section 4.4 – where the welfare-benefits hypothetical illustrates how predictability operates at the ex ante stage and procedural clarity at the contemporaneous decision-making stage – is therefore here made explicit: the two ideas function as paired but distinguishable requirements on the use of AI in the exercise of state power.

4.2. The Rule of Law as a limit on public power

The definition of the Rule of Law that has been adopted – as a commonly accepted and (as much as possible) theoretically agnostic conception – is one that considers the Rule of Law to be a concept that seeks to prevent the arbitrary application of powers by limiting the exercise of public power through a need for predictability and procedural clarity. In this respect, it is readily apparent that the Rule of Law would be seen to operate as a limit on public power. This subsection seeks to add a small amount of detail to this point by illustrating the way in which the concept achieves this in practice.

The concept of the Rule of Law, while seen as a universal good (when put in the abstract terms noted earlier) (Krygier, 2011), is a political concept. States all around the world will profess to be states that operate in accordance with the Rule of Law, but few states include adherence to the Rule of Law in their constitutions or founding documents. For example, states like the United States and Australia would doubtless agree that they are bound by and adhere to the Rule of Law, but there is no mention of the concept in their constitutions. Even states that do make mention of the Rule of Law in their constitutions – like South Africa (Michelman, 2013; The Constitution of the Republic of South Africa, 1996, Chapter 1) – may not provide more than a broad statement relating to the need to adhere to the Rule of Law (as an undefined idea). The effect of the absence of the concept within states' constitutions means that the concept is invariably not the subject of precise definitions that

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

have a precedential impact; the concept is a political – not a legal – concept in its application; the consequence of a failure to adhere to the concept is political sanction and not legal penalty.

The fact that the Rule of Law is a political concept has an effect on the normative impact of the concept. It is worth saying a little more about how, exactly, a concept like the Rule of Law operates as a limitation. The argument advanced in this paper is that the limitation operates not through direct legal sanction (it is not justiciable, and breach does not trigger a court-ordered remedy) but through an indirect political-normative constraint. Where a state has publicly represented itself as Rule-of-Law-compliant – through constitutional statement, policy commitment or international undertaking – the political cost of acting in ways that contravene the Rule of Law generates a practical constraint on what that state may legitimately do. The mechanism, in other words, is one of political-normative coherence: a state that publicly commits to a concept is, by virtue of that commitment, subject to political accountability for departures from it (see, on related sensibility- and culture-based accounts of how the Rule of Law constrains, Burgess, 2024a; Krygier, 2011, 2017). This is the broad form of limitation that the remainder of the paper develops and assumes – not a justiciable prohibition, but a normative requirement that operates through the state's own self-representation and the political consequences of departing from it.

The normative force generated by the concept in this respect is neither merely rhetorical (a purely presentational concern about how the state appears) nor justiciable (a basis for court-ordered remedy). It is sufficient, however, to operate as a meaningful limitation on the kinds of AI systems a Rule-of-Law-committed state may legitimately deploy, because departures from the commitment have politically significant consequences – contested legitimacy, reduced public trust, accountability through political processes and reputational cost in international Rule-of-Law assessments. It is in this sense that the Rule of Law functions here, as a limit on state action.

The nonjusticiable status of the concept in most jurisdictions around the world, together with the concept's essentially contested nature, does not appear to diminish the potential impact or the perceived benefit of the concept. The concept remains, however, a concept that has some potential benefit. For present purposes, this political (rather than justiciable) character explains why, in practice, the constraints identified here operate through design, documentation and oversight obligations that structure the exercise of public power.

4.3. Rule of Law Limitations on AI

What has been established thus far is that the Rule of Law is a commonly deployed concept that – in its most abstract and unobjectionable sense – can be seen to operate as a check on the arbitrary exercise of state power. What is more, the concept remains a popular way to hold those in power to account notwithstanding its ancient roots and notwithstanding its nonjusticiable status. In this subsection, these traits of the concept are used to illustrate that the concept itself can also be seen to operate as a limitation on the way in which states exercise power using AI.

The interaction of the Rule of Law and either AI or ADM more broadly has recently been a popularly discussed topic (Aliu, 2025; Burgess, 2024a; Suksi, 2023; Zalnieriute et al., 2019). However, those works explore the interaction of the concept and the technology in ways that do not seek to identify the way in which, or the extent to which, the Rule of Law may limit states' use of AI or what this means. The claim being made here is not that the Rule of Law operates as a direct limitation on the use of AI by states in exercising power. Instead, the suggestion is that the existence of the Rule of Law, in states where the Rule of Law is held out as being a concept to which the state adheres, means that if a technology contravenes the Rule of Law (and the state would not wish to do this) then there exists a limitation – based on the Rule of Law – relating to that technology. This claim is the article's key contribution and, in this form, is offered as a novel contribution to the Rule-of-Law-and-AI literature. Putting this another way, if the state purports to adhere to the Rule of Law, then the state should not take actions that contravene the Rule of Law; and, if the state should not take actions that contravene

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

the Rule of Law, then the use of technologies (including AI) that contravene the Rule of Law should not be used by the state.

While this may seem to be intuitively attractive, it is pregnant with potential problems – not least of which is the essentially contested nature of the concept of the Rule of Law as well as the compounding problem as to what it may mean for a technology to contravene the Rule of Law. These problems are briefly considered before exploring what this may mean.

4.3.1. Problem 1: the essentially contested rule of law

As noted earlier, the Rule of Law is widely accepted to be a contested – perhaps an essentially contested – concept. This means that the entire argument that follows may be seen to hinge on the necessity of identifying a precise and agreed meaning of the Rule of Law for any aspect of contravention (explored below) to be of any impact or effect. The problem of its contested nature is not, however, intractable – at least in terms of providing some form of practical guidance for the operation of a limitation.

As noted earlier, there is a broadly agreed abstract idea of the Rule of Law – as it is a concept opposing the arbitrary application of power – but that provides little practical application potential; however, the more refined version of the concept – associated with the most commonly seen core to the concept as a concept opposing the state's arbitrary application of power through the identification of core ideas that undergird other popular conceptions by grouping these together as requiring either predictability or procedural clarity (Burgess, 2017; Tamanaha, 2004), as one example, affords a less abstract way to consider and identify what the Rule of Law requires. When this is applied to the present consideration, this means that the state should not deploy technologies like AI in circumstances where that technology may result in the application of power in ways that cannot be seen to be either predictable or procedurally clear.

4.3.2. Problem 2: technology contravening the rule of law

The second problem is how to ascertain when a technology may contravene the Rule of Law. If the use of AI in the exercise of public power fails to be predictable or procedurally clear, then its use contravenes the Rule of Law. These sorts of contraventions could be seen in terms of the way that decisions are made – as has been illustrated in terms of ADM (Araujo et al., 2020; Suksi, 2023; Zalnieriute et al., 2019) – as reflected in instances of unpredictability of the application of the existing rules. It has commonly been stated that the use of complex algorithms involved in the use of AI in decision-making takes place within a black box (Pasquale, 2016); the opacity of the way in which decisions are made results in both unpredictability as well as an absence of procedural clarity (Burgess, 2024a; Burrell, 2016). At this level of generality, it could therefore be suggested that *any* use of AI in the exercise of public power may contravene the Rule of Law. However, a more nuanced approach may afford a more usable way from which limitations may be derived. Before illustrating potential limitations regarding the consideration of the Rule of Law in this way, section 4.4 introduces an illustrative hypothetical regarding a welfare-benefit determination to demonstrate what is meant by this idea of a more nuanced approach.

4.4. Nuanced Rule-of-Law limitations: an illustrative hypothetical

Given the huge potential efficiency benefits that may result from the use of AI in the exercise of public power (Aliu, 2025), it seems likely that AI will be used and adopted in a wide variety of processes in which state power is exercised. To demonstrate a more nuanced approach, an illustrative hypothetical will be briefly explored. While it is hypothetical, it is not fanciful. The hypothetical is the use of AI in determining state benefits/welfare payments.

While the example used is hypothetical, there are a number of recent ADM instances that could be seen as indicating both a desire to deploy technologies and an indication of the potential scope of failures that can occur. These include Robodebt in Australia, the Post Office scandal in the UK, and the

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

Dutch childcare benefits algorithm scandal (Bouwmeester, 2025; Moynihan et al., 2025; Priergaard, 2025). These ADM systems can be distinguished from the use of AI exemplified here because they were not, strictly speaking, applications of contemporary AI systems. They were systems through which the, albeit flawed, processing could be tracked and mapped readily. By contrast, the hypothetical here relates to the deployment of an AI system that determines individuals' benefit payments based on having access to all relevant state-based data sources, including information provided by the individual for the application, previous benefit payment histories, and any relevant legislation, rules and regulations. In the hypothetical, it is also assumed (unlike with the real examples previously noted) that the systems are fit for purpose and that they produce accurate outcomes. By "accurate outcomes," what is meant is that the outcomes would mirror the outcomes that a highly experienced human reviewer would achieve within acceptable bounds of error.

While the outcome itself may be accurate, this is not, however, sufficient to satisfy the two common ideas of the Rule of Law that have been noted: predictability and procedural clarity. The *process* associated with arriving at that decision, not just an accurate outcome, is vitally important in satisfying these ideas. This reflects the thin/procedural idea of the Rule of Law that represents the common understanding. For thick conceptions of the Rule of Law, the outcome being right would be important – but, as noted earlier, those thick conceptions also require the thin aspects themselves to be satisfied. In other words, something more than mere accuracy is required to satisfy the Rule of Law – this represents the broadest idea of the Rule of Law as a limitation.

In a hypothetical like this one, the ideas of predictability and procedural clarity can be further unpacked into what are, respectively, a retrospective and a contemporary consideration. The idea of predictability in the Rule of Law can be seen to relate to ways in which the applier of state power may act based on the way that they have previously acted. This idea can be seen across Rule of Law accounts that are commonly seen to be determinative of the concept's content (Fuller, 1964, p. 79; Hayek, 2007, p. 74; Locke, 1988, pp. 350–353; Raz, 2009, pp. 214–215). In this respect, it requires that previous acts can be seen, interrogated and understood. The idea of procedural clarity relates to a contemporary decision being able to be seen, interrogated and understood. While similar, satisfaction of both could be seen as being vital to satisfy the Rule of Law across many popular conceptions (Aristotle, 2004, Book 1, ch 1; Dicey, 1979; Raz, 2009, pp. 214–217).

The practical effect illustrated by this hypothetical is that, to satisfy the Rule of Law, an AI system used to determine benefits must be able to be seen, interrogated and understood. If this cannot be achieved within a system, then the system's use would contravene the Rule of Law (where it is assumed that the state will adhere to the concept). In this sense, and under these assumptions, the state is limited in the forms of AI that it can use or deploy. The way in which this is operationalized will be explored in the section 5 after the shortfalls of this sort of limitation are noted.

4.5. Limits on the limits

While it has been stated previously, it is useful at this point – following the consideration of the hypothetical – to illustrate the ways in which the limitations that may be seen to follow from the Rule of Law may, themselves, be limited. These are what I will call limits on the limits. Principal among the limits on the limits is the difficulty that stems from the concept's contested nature. While, as has been shown across the previous subsections, it is possible to identify broad limitations, the identification of precise limitations relating to the concept of the Rule of Law in any given situation is difficult (if not impossible).

The hypothetical explored in the previous subsection can be used to illustrate this. While it may be clear that, to satisfy the Rule of Law, the AI systems used need to facilitate both predictability and procedural clarity, it is not easy to establish precisely where the nuanced positions may lie; in other words, a system that provides no rationale or explanation for a decision regarding benefit allowances – that simply determines, for example, "the benefit payable is $500" – clearly cannot satisfy either of

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

the Rule of Law ideas. However, systems that output more detailed decisions may be less obvious in their potential to satisfy the ideas. If a system provides a rationale illustrating how particular factual scenarios for an applicant have been matched to the relevant regulations, and how that translates to the derivation of the decision regarding the ultimate decision "the benefit payable is $500," that represents an improvement over the bare decision, but this improvement is not, of itself, necessarily sufficient to satisfy the Rule of Law ideas. This is because those ideas themselves are broad and open to interpretation (Tamanaha, 2004; Waldron, 2002); there would be different levels of compliance or otherwise based on the precise formulation of the Rule of Law adopted/preferred. However, even based on this consideration, it becomes apparent that there are limitations to the limitations.

It is also possible to consider the limit on the limits from a different perspective – one that relates to the sufficiency of the process of decision-making (and not the output, as noted immediately above). As has been stated earlier, complex algorithms represent a black box in terms of processing (Burrell, 2016; Pasquale, 2016); the relative explainability of the decisions represents a largely intractable problem in computer science (Rudin, 2019). While it is possible to infer the process that has taken place to arrive at an output – based on both the relative inputs and the system parameters – any conclusion about what has taken place remains an inference. It may be because of this opacity that a further limit on the limitation can be seen: can predictability and procedural clarity requirements *ever* be satisfied if we cannot precisely determine what process has been applied to arrive at a decision regarding the exercise of state power?

The immediate answer, as previously identified, is that the use of AI systems in this way is sufficient to contravene the Rule of Law in all circumstances; thus rendering *all* uses of AI in exercising state power contrary to the Rule of Law (Brownsword, 2016; Burgess, 2024a). An argument could, however, be made that this is too fast a conclusion. It relies on the basis that decisions that have previously been made by humans have not necessarily been seen as Rule of Law-objectionable in terms of predictability and procedural clarity by virtue of an accurate explanation being provided by the human for their decision. While we, as humans, may labor under the belief that this is possible, it may not always be possible to accurately and precisely articulate reasons for a decision (Wegner, 2003). Neuroscientists may question our ability, as humans, to precisely communicate, understand or comprehend our reasons for action (Soon et al., 2008). On this basis, if human decision-making cannot be perfectly communicated, but it is still seen as being able to satisfy basic ideas associated with the Rule of Law, the relative bar for AI decision-making to clear in order to avoid Rule of Law contravention (in consequence of its relative opacity) may not be as absolute; systems that do provide an explanation (that may on reflection appear to be objectively at least a feasible way to arrive at a conclusion) could be seen to reflect human decision-making processes.

The intention in this section has not been to identify precisely where the boundary regarding what is or is not considered to be sufficient in terms of Rule of Law contravention or compliance. Instead, the aim was to identify that there are limitations on the limitations. These reflect one potential difficulty that may exist when seeking to identify the ways in which Rule-of-Law limitations may be identified. After considering these, even briefly, it is clear that while there may be clear instances where common Rule of Law ideas may be contravened, the identification of clear compliance is not an obvious and distinct line of demarcation. In section 5, the ways in which this may impact the potential operationalization of the Rule of Law limitations will be considered. For present purposes, the baseline remains that deployments must enable guidance ex ante (predictability) and contestation ex post (procedural clarity); where that baseline cannot be met, the Rule of Law counsels non-deployment.

5. Operationalizing Rule-of-Law limits

The potential for the Rule of Law to operate as a limitation on the ways in which AI is used to exercise public power has been explored in the previous sections. In doing this, it has been noted that – although the Rule of Law is a contested concept – commonly accepted ideas, built upon a foundational

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

thin idea of the concept, may provide some basis for the limitation. This limitation can, most unobjectionably, be seen to operate in relation to the exercise of state power. With these things in mind, it is possible to conceive of a way in which the Rule of Law can represent a limitation on the way in which AI systems are deployed by those exercising the power of the state – and, further, that this represents at least one of four areas in which attempts can be seen to regulate the use of AI in the broader sense (see Figure 1). In this brief section, the aim is to consider what the relative difficulty associated with the challenge of translating abstract legal principles may mean for the operationalization of the Rule of Law as a limitation on the use of AI in the exercise of public power.

The first consideration to note is that the Rule of Law, in its least objectionable form, cannot hope to achieve a similar form of limitation to that which is apparent from the desire to regulate AI. The least objectionable form of the Rule of Law as a limitation cannot apply to all four areas identified in Figure 1. While some areas may argue that the Rule of Law could operate in other areas of the structure adopted – for example, in relation to the exercise of private power (Bhatt et al., 2021; Kampourakis et al., 2023) – those areas are, for the reasons provided above, less widely accepted, which means that consensus relating to the Rule of Law is lost. On this basis, at most, the idea of the Rule of Law as a limitation can only represent one component of any scheme of regulation/limitation. Furthermore, for the reasons noted below, even within the state-users section in which the least objectionable Rule of Law may operate, it is not suggested that the Rule of Law should be the only, nor is it necessarily the best, form of regulation/limitation; the argument here is merely that it is one form of limitation that does exist but has not previously been considered.

It is worth acknowledging directly a related tension that the literature has highlighted. The empirical reality is that public-sector AI is expanding – and that this expansion has, in many jurisdictions, been observed to weaken adherence to Rule-of-Law values in practice (Bouwmeester, 2025; Zalnieriute et al., 2019). The claim advanced here is, accordingly, a normative one: it concerns what states ought to do where they purport to be committed to the Rule of Law, rather than a descriptive claim about what they currently do. The point of operationalizing Rule-of-Law requirements – particularly through procurement, as discussed below – is precisely that doing so provides a tractable mechanism for translating that normative commitment into checks that bite even where political commitment to the Rule of Law is, in practice, attenuated.

The second and third operational considerations are intimately connected. The second is that there is an inherent difficulty associated with the identification of precise limitations from (contested) abstract principles; but there are at least some commonly identifiable requirements that must be adhered to for any AI system's use in the exercise of public power to be potentially compliant with the Rule of Law – like the need for a system to be both predictable and procedurally clear. In practical terms, this may be implemented at the procurement phase – requiring developers to commit to providing explanations sufficient to permit external scrutiny of system outputs and to maintain documentation that supports such scrutiny over time. These sorts of requirements can be seen in some jurisdictions' guidance documents which recommend integrating explainability, access to information and assurance provisions into AI tenders and contracts (UK Government, 2020; World Economic Forum, 2019). Yet, the third consideration is that the relative complications return as an attempt is made to identify what these things may mean beyond the most obvious instances.

Building on this, the procurement of AI systems by state actors offers a particularly useful operational locus for Rule-of-Law requirements. Procurement contracts and tender specifications are precisely the documents in which Rule-of-Law-derived requirements – explainability sufficient for ex ante guidance (predictability) and sufficient documentation for contemporaneous interrogation (procedural clarity) – can be made enforceable, well before a system reaches deployment. Coglianese (2024) has argued that procurement functions as a form of AI governance in its own right: pre-contractual specifications, vendor obligations and audit rights enable states to embed legal and constitutional standards into systems they themselves do not design (see also UK Government, 2020; Hickok, 2024; Sanchez-Graells, 2024; World Economic Forum, 2019). On the view advanced in this

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

paper, the Rule of Law supplies the substantive criteria – predictability and procedural clarity – that should be operationalized through such procurement instruments.

As the welfare-benefits hypothetical earlier in the paper indicates, even where outcomes are accurate, compliance turns on whether people can foresee how inputs are mapped to legal criteria (predictability) and whether contemporaneous reasons permit interrogation and contestation (procedural clarity). The operational effect of this is that while limitations can be identified, they can only – clearly and unobjectionably – translate to requirements that may already have been seen as self-evident; for example, that a system that only provides a simple output that does not have *any* form of explanation or rationale cannot be seen to comply with the Rule of Law. Beyond this, issues of adequacy arise that mean, to continue the example, if an explanation or rationale is provided, that does not necessarily mean that the system-provided response is Rule-of-Law-compliant – as deeper questions relating to what it may mean for the rationale to be considered adequate (perhaps in comparison to human-provided levels) may need to be answered.

This does not, however, relegate the idea of the Rule of Law as a limitation on the exercise of state power with AI systems to a merely academic point. The Rule of Law has never been a precisely definable concept; yet it has been, and remains, a concept of principal importance in states around the world (Sevel, 2024). In recent decades, the concept's definition, from its thin/procedural roots, has evolved (Burgess, 2024b; Tamanaha, 2004). It has been conceived, variously and for example, as a thicker concept or one that applies to private relationships (Bingham, 2010; Kampourakis et al., 2023). Unsurprisingly, historical conceptions of the concept – including those formulated more than a few years ago – did not necessarily conceive of the concept as a potential limiting factor for technologies like AI.

This examination presents the concept of the Rule of Law as an additional form of limitation where regulatory approaches already exist in certain jurisdictions and as a standalone form of limitation in jurisdictions where regulatory approaches may not exist. The broad-based form of the limitations that have been outlined is neither extensive nor fine-grained; however, they are readily identifiable (notwithstanding the Rule of Law's contested nature) and can be seen to represent a form of guidance for states regarding the sorts of actions that they can (and cannot) take in implementing AI in the exercise of power. On this basis, it does appear that the concept's evolution – at least in terms of the scope of its application – has continued into the age of AI; it can, despite its ancient pedigree, operate as a way to limit the use of AI.

Acknowledgements. The author made use of ChatGPT5 to provide proofreading and reference checks on portions of this article. The author is responsible for the academic content of the article, and it adheres to the journal's authorship policy.

Funding statement. The author did not receive any funding specific to this project.

Competing interests. The author has no competing interests.

References

Alex Engler. (2023, April 25). The EU and U.S. diverge on AI regulation: A transatlantic comparison and steps to alignment. *Brookings*. https://www.brookings.edu/articles/the-eu-and-us-diverge-on-ai-regulation-a-transatlantic-comparison-and-steps-to-alignment/

Aliu, A. (2025). *Artificial Intelligence and the Rule of Law: The Age of Legal Tech and Digital Governance in a Fractured Digital World*. Springer Nature Switzerland.

Araujo, T., Helberger, N., Kruikemeier, S., & De Vreese, C. H. (2020). In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & SocietyY*, *35*(3), 611–623. https://doi.org/10.1007/s00146-019-00931-w

Aristotle. (2004). *Rhetoric*. W. R. Roberts, Trans. Dover.

Bhatt, K., Lander, J., & Taekema, S. (2021). Introduction: The rule of law in transnational development projects – Private actors and public chokeholds. *International Journal of Law in Context*, *17*(1), 91–99. https://doi.org/10.1017/S1744552321000021

Bingham, T. (2010). *The Rule of Law*. Allen Lane.

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

Bouwmeester, M. (2025). Checks and balances under pressure in the welfare state: Disentangling the rule of law risks of automation and welfare conditionality. *European Journal of Social Security*, *27*(2), 98–120. https://doi.org/10.1177/13882627251330438

Brownsword, R. (2016). Technological management and the Rule of Law. *Law, Innovation and Technology*, *8*(1), 100–140. https://doi.org/10.1080/17579961.2016.1161891

Burgess, P. (2017). The rule of law: Beyond contestedness. *Jurisprudence*, *8*(3), 480–500. https://doi.org/10.1080/20403313.2017.1341203

Burgess, P. (2020). The rule of lore in the rule of law: Putting the problem of the rule of law in context. *Hague Journal on the Rule of Law*, *12*, 333–361. https://doi.org/10.1007/s40803-019-00096-0

Burgess, P. (2024a). *AI and the Rule of Law: The Necessary Evolution of a Concept*. Hart Publishing.

Burgess, P. (2024b). At present, the future of the rule of law is history. In Sevel Michael (Ed.), *Routledge Handbook of the Rule of Law*. Routledge.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, *3*(1), 2053951715622512. https://doi.org/10.1177/2053951715622512

Cantero Gamito, M., & Marsden, C. T. (2024). Artificial intelligence co-regulation? The role of standards in the EU AI Act. *International Journal of Law and Information Technology*, *32*, eaae011. https://doi.org/10.1093/ijlit/eaae011

Coglianese, C. (2024). Procurement and artificial intelligence in Regine Paul. In J. Cobbe & E. Carmel (Eds.), *Handbook on Public Policy and Artificial Intelligence* (pp. 235–248). Edward Elgar.

The Constitution of the Republic of South Africa (1996).

Craig, P. (1997). Formal and Substantive Conceptions of the Rule of Law an Analytical Framework. *Public Law, Autumn*, 467–487.

Dicey, A. V. (1979). *Introduction to the Study of the Law of the Constitution* (10th ed.). Palgrave Macmillan UK.

Ebers, M. (2025). Truly risk-based regulation of artificial intelligence how to implement the EU's AI Act. *European Journal of Risk Regulation*, *16*(2), 684–703. https://doi.org/10.1017/err.2024.78

Fallon, R. H., Jr. (1997). "The rule of law" as a concept in constitutional discourse. *Columbia Law Review*, *97*(1), 1–56. https://doi.org/10.2307/1123446

Fuller, L. L. (1964). *The Morality of Law*. Yale University Press.

Gowder, P. (2018). Transformative legal technology and the rule of law. *University of Toronto Law Journal*, *68*(Suppl. 1), 82–105. https://doi.org/10.3138/utlj.2017-0047

Gowder, P. (2023). *The Networked Leviathan: For Democratic Platforms*. Cambridge University Press.

Greenstein, S. (2022a, March 30). Artificial intelligence, human flourishing and the rule of law. *Verfassungsblog*. https://verfassungsblog.de/roa-human-flourishing/

Greenstein, S. (2022b). Preserving the rule of law in the era of artificial intelligence (AI). *Artificial Intelligence and Law*, *30*, 291–323. https://doi.org/10.1007/s10506-021-09294-4

Hague Journal on the Rule of Law. (2017). Special issue on the historical origins of the rule of law. *Hague Journal on the Rule of Law*, *9*, 1–2. https://doi.org/10.1007/s40803-017-0057-y

Hayek, F. A. V. (2007). *The Road to Serfdom* B. Caldwell, (Ed.). University of Chicago Press.

Hickok, M. (2024). Public procurement of artificial intelligence systems: New risks and future proofing. *AI & SOCIETY*, *39*(3), 1213–1227. https://doi.org/10.1007/s00146-022-01572-2

Huq, A. Z. (2024). Artificial intelligence and the rule of law. In Sevel Michael (Ed.), *Routledge Handbook of the Rule of Law*. Routledge.

Jacobs, M., & Simon, J. (2022). Assigning obligations in AI regulation: A discussion of two frameworks proposed by the European Commission. *Digital Society*, *1*(1), 6. https://doi.org/10.1007/s44206-022-00009-z

Kampourakis, I., Taekema, S., & Arcuri, A. (2023). Reappropriating the rule of law: Between constituting and limiting private power. *Jurisprudence*, *14*(1), 76–94. https://www.tandfonline.com/doi/abs/10.1080/20403313.2022.2119016

King, J. (2018, April 27). *The Social Dimension of the Rule of Law: I-CONnect Symposium on "Constitutional Boundaries."* http://www.iconnectblog.com/2018/04/i-connect-symposium-on-constitutional-boundaries-the-social-dimension-of-the-rule-of-law/

King, J. (2019). Martin Krygier and the tempering of power. *Hague Journal on the Rule of Law*, *11*(2), 363–370. https://doi.org/10.1007/s40803-019-00125-y

Krygier, M. (2009). The rule of law: Legality, teleology and sociology. In G. Palombella & N. Walker (Eds.), *Relocating the Rule of Law*. Hart 45-70.

Krygier, M. (2011). Four puzzles about the rule of law: Why, what, where? and who cares? In J. Flemming (Ed.), *Getting to the Rule of Law* (pp. 64–104). New York University Press.

Krygier, M. (2017). Tempering power. In M. Adams, E. H. Ballin & A. Meuwese (Eds.), *Constitutionalism and the Rule of Law: Bridging Idealism and Realism*. Cambridge University Press 34–59.

Locke, J. (1988). *Two Treatises of Government* (P. Laslett, Ed.). Cambridge University Press.

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

Lyria Bennett Moses. (2007). Recurring dilemmas: The law's race to keep up with technological change. *University of Illinois Journal of Law Technology and Policy*, *2007*(2), 239–285. https://illinoisjltp.com/journal-archive/article/recurring-dilemmas-the-laws-race-to-keep-up-with-technological-change

Marsden, C. T. (2011). *Internet Co-Regulation: European Law, Regulatory Governance and Legitimacy in Cyberspace*. Cambridge University Press.

Michelman, F. I. (2013). The rule of law, legality and the supremacy of the Constitution. In S. Woolman & M. Bishop (Eds.), Constitutional law of South Africa (2nd ed., pp. 11–44). Juta & Company (Pty) Ltd.

Mir, O. (2025). The AI act from the perspective of administrative law: Much ado about nothing? *European Journal of Risk Regulation*, *16*(1), 63–75. https://doi.org/10.1017/err.2024.54

Moynihan, D., Hybschmann, M., Gimborys, K., Loudin, S., & McClellan, W. (2025). Digital administrative burdens and clawback logics: The Australian Robodebt scheme. *Perspectives on Public Management and Governance*, gvaf011. https://doi.org/10.1093/ppmgov/gvaf011

OECD. (2025). *Recommendation of the Council on Artificial Intelligence* (No. 449). https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449

Pasquale, F. (2016). *The Black Box Society: The Secret algorithms That Control Money and Information* (Reprint ed.). Harvard University Press.

Priergaard, J. (2025). Reform after robodebt: Lessons from the Netherlands. *Australian Journal of Political Science*, *60*(2), 192–206. https://www.tandfonline.com/doi/abs/10.1080/10361146.2025.2549868

Raz, J. (2009). *The Authority of Law: Essays on Law and Morality* (2nd ed.). Oxford University Press.

Regulation (EU) 2024/1689 … (Artificial Intelligence Act), 2024/1689 1 (2024). https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, *1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Sanchez-Graells, A. (2024). Responsibly buying artificial intelligence: A 'regulatory hallucination. *Current Legal Problems*, *77*(1), 81–126. https://doi.org/10.1093/clp/cuae003

Sempill, J. A. (2017). The lions and the greatest part: The rule of law and the constitution of employer power. *Hague Journal on the Rule of Law*, *9*, 283–314. https://doi.org/10.1007/s40803-017-0058-x

Sevel, Michael (2024). Introduction. In Sevel Michael (Ed.), *Routledge Handbook of the Rule of Law*. Routledge.

Shahin, J. (2024). Dancing to the same tune? EU and US approaches to standards setting in the global digital sector. *Journal of European Integration*, *46*(7), 1111–1131. https://doi.org/10.1080/07036337.2024.2398430

Shklar, J. (1987). Political theory and the rule of law. In A. C. Hutchinson & P. Monahan (Eds.), *The Rule of law: Ideal or Ideology* (pp. 1–16). Carswell.

Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543–545. https://doi.org/10.1038/nn.2112

Suksi, M. (Ed.). (2023). *The Rule of Law and Automated Decision-Making: Exploring Fundamentals of Algorithmic Governance*. Springer.

Tamanaha, B. Z. (2004). *On the Rule of Law History, Politics, Theory*. Cambridge University Press.

UK Government. (2020). *Guidelines for AI procurement*. https://www.gov.uk/government/publications/guidelines-for-ai-procurement/guidelines-for-ai-procurement?

UN Security Council. (2004). *The Rule of Law and Transitional Justice in Conflict and Post-conflict Societies—Report of the Secretary-General* (No. UN doc S/2004/616). https://www.un.org/ruleoflaw/files/2004%20report.pdf

UNESCO. (2022). Recommendation on the Ethics of Artificial Intelligence— UNESCO Digital Library (No. 9988). https://unesdoc.unesco.org/ark:/48223/pf0000381137.

United States. (2019). *Maintaining American Leadership in Artificial Intelligence* (No. Exec Order No 13859). https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence?

United States. (2023). *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* (No. Exec Order No 14110). Federal Register. https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence

Veale, M., & Edwards, L. (2018). Clarity, surprises, and further questions in the Article 29 Working Party draft guidance on automated decision-making and profiling. *Computer Law & Security Review*, *34*(2), 398–404. https://doi.org/10.1016/j.clsr.2017.12.002

Veale, M., Matus, K., & Gorwa, R. (2023). AI and global governance: Modalities, rationales, tensions. *Annual Review Law and Social Science*, *19*, 255–275. https://doi.org/10.1146/annurev-lawsocsci-020223-040749

Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU artificial intelligence act. *Computer Law Review International*, *22*(4), 97–112. https://doi.org/10.9785/cri-2021-220402

Waldron, J. (2002). Is the rule of law an essentially contested concept (in Florida)? *Law and Philosophy*, *21*(2), 137–164. https://www.jstor.org/stable/3505128

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.

Wang, S., Zhang, Y., Xiao, Y., & Liang, Z. (2025). Artificial intelligence policy frameworks in China, the European Union and the United States: An analysis based on structure topic model. *Technological Forecasting and Social Change*, *212*, 123971. https://doi.org/10.1016/j.techfore.2025.123971

Wegner, D. M. (2003). *The Illusion of Conscious Will*. MIT Press.

World Economic Forum. (2019). *World Economic Forum Guidelines for AI Procurement*. https://www3.weforum.org/docs/WEF_Guidelines_for_AI_Procurement.pdf?

Zalnieriute, M., Moses, L. B., & Williams, G. (2019). The rule of law and automation of government decision-making. *The Modern Law Review*, *82*(3), 425–455. https://doi.org/10.1111/1468-2230.12412

Paul Burgess is an associate professor in the Faculty of Law at Monash University, where he is Cofounder and Lab Lead of the Technology and AI in Legal Education Research (TAILER) Lab and Deputy Director of the Digital Law Group. His research sits at the intersection of law and artificial intelligence, examining how AI is used in large-scale public decision-making and in the exercise of state power. His monograph, *AI and the Rule of Law: The Necessary Evolution of a Concept* (Hart, 2024), was the first book-length treatment of the topic, and his work on the Rule of Law and emerging technologies has appeared in leading international journals and been cited extra-judicially by Lord Sales of the UK Supreme Court. His methodology draws on jurisprudence, constitutional theory and political philosophy, alongside frequent collaborations with colleagues in computer science on the evaluation of legal AI systems. He is Lead Chief Investigator on an Australian Research Council Linkage Project. Before joining Monash, Paul was a lecturer at University College London.

Cite this article: Burgess, P. (2026). Limiting artificial intelligence (with the Rule of Law). *Cambridge Forum on AI: Law and Governance*, *2*, e15, 1–17. https://doi.org/10.1017/cfl.2026.10063

Downloaded from https://www.cambridge.org/core. 19 Nov 2026 at 06:46:23, subject to the Cambridge Core terms of use.