Public Governance, Administration and Finances Law Review Vol. 11. No. 1. (2026) • 169–190.

DOI: https://doi.org/10.53116/pgaflr.8853

Governing the Perceptual Infrastructure

Public Administration in the Age of Algorithmic Perception

Haris Alibašić*

* Professor and Whitman Faculty Fellow, Department of Business Administration, University of West Florida, Pensacola, Florida, USA, e-mail: halibasic@uwf.edu

Submitted: 17 February 2026 | Accepted: 01 June 2026 | Published online: 29 June 2026

Abstract: Artificial intelligence systems increasingly shape not only what information citizens access but the interpretive frameworks through which they perceive public affairs, posing a governance challenge that public administration scholarship has yet to theorise. Through conceptual and normative analysis, the paper develops the construct of *perceptual infrastructure* – the cognitive and informational substrate of democratic deliberation – drawing on information theory, public administration and regulatory-governance scholarship, and tests it against the EU regulatory architecture. The analysis shows that the algorithmic construction of perceptual frameworks constitutes a distinct governance domain that the EU AI Act and the Digital Services Act do not reach, owing primarily to a regulatory omission rather than an implementation deficit: existing provisions take systems, use cases and identifiable harms as their object, not the cumulative, longitudinal construction of perception. An accountability framework is proposed – incorporating aggregate transparency, perceptual sovereignty as a citizen right, proportional responsibility and meta-perceptual literacy – with concrete implications for administrative accountability, regulatory capacity and democratic resilience.

Keywords: AI governance, algorithmic accountability, democratic governance, public administration, perceptual infrastructure, EU AI Act

1. Introduction

Democratic governance depends upon citizens sharing sufficient common ground to perceive, deliberate about, and respond to public problems collectively. The common ground is not merely a matter of shared values or civic commitment; it requires sufficient overlap in the informational frameworks through which citizens perceive political events, policy proposals and social reality. When algorithmic systems systematically construct

Corresponding author: Haris Alibašić

divergent perceptual frameworks for different segments of the population, they undermine the epistemic foundations upon which democratic self-governance rests.

The paper addresses a governance challenge that public administration scholarship has yet to adequately theorise: artificial intelligence systems exercise power not merely by controlling what information citizens access but by shaping the interpretive frameworks through which citizens perceive public affairs. The process can be termed the *algorithmic construction of perceptual frameworks* – the systematic, large-scale and typically opaque process through which AI systems build the perceptual conditions under which citizens encounter and process information relevant to democratic governance. Taken together, these frameworks constitute what may be termed the *perceptual infrastructure of democracy* – the cognitive and informational substrate upon which deliberation, accountability, and collective decision-making depend.

The distinction between content and perceptual infrastructure matters profoundly for public administration. Consider two citizens encountering identical news coverage of a public health policy. Both see the same data: infection rates, hospitalisation figures and economic projections. Yet one perceives a responsible government response to a genuine crisis, while the other perceives authoritarian overreach justified by manipulated statistics. The factual content is identical; the information each citizen perceives is fundamentally different because years of algorithmically curated exposure to different sources, arguments, and interpretive patterns have constructed divergent perceptual frameworks. The problem is not that citizens disagree about shared facts; it is that they perceive different facts from identical data. The condition characterises the post-factual governance environment, in which the erosion of shared epistemic foundations creates administrative and policy crises that conventional informational remedies cannot resolve (Alibašić, 2024).

For public administration specifically, the stakes are threefold and concrete. First, *administrative accountability:* when public servants rely on AI decision-support tools, the perceptual frameworks those tools construct shape discretionary judgment in ways that current accountability mechanisms – designed to audit decisions, not perception – cannot easily reach. Second, *regulatory capacity:* agencies tasked with overseeing algorithmic systems require analytical competencies for assessing cumulative framework effects that few presently possess. Third, *public value and democratic resilience:* the capacity of administrative systems to coordinate collective responses to crises depends on a citizenry able to perceive shared evidence, a precondition that framework fragmentation directly erodes. These three concerns – accountability, capacity and resilience – organise the practice-facing analysis in Section 6, to which the recommendations return.

Existing governance responses focus primarily on content moderation and algorithmic transparency (Diakopoulos, 2015; Gillespie, 2018). While essential, these approaches address the surface manifestations of a deeper structural challenge. The EU AI Act [Regulation (EU) 2024/1689], despite representing the most comprehensive regulatory effort to date, classifies recommendation systems and algorithmic curation as limited-risk or minimal-risk applications, underestimating the cumulative effects on democratic governance of sustained algorithmic shaping of citizens' perceptual frameworks (Yadav, 2025). The Digital Services Act (DSA), while imposing algorithmic transparency

and risk assessment obligations on very large online platforms (VLOPs), similarly addresses individual content decisions and systemic risks without recognising the gradual construction of perceptual frameworks as a distinct governance domain.

The paper makes three contributions. First, it develops a theoretical account of how algorithmic systems reshape the perceptual infrastructure of democracy, grounded in the principle that information emerges through the interaction of perception and interpretive frameworks. Second, it analyses the implications for democratic governance, public sector accountability and administrative decision-making, with attention to the regulatory limitations of current instruments, including the EU AI Act and the DSA. Third, it proposes an accountability framework tailored to the challenge of governing algorithmic influence on democratic perception, extending existing regulatory approaches to address the gap.

The paper proceeds as follows. Section 1.1 states the methodological approach and research question. Section 2 establishes the theoretical foundations. Section 3 develops a taxonomy of algorithmically constructed frameworks. Section 4 analyses the EU regulatory architecture's limitations regarding perceptual infrastructure. Section 5 proposes an accountability framework. Section 6 examines implications for public administration practice, with attention to contexts of democratic fragility. Section 7 acknowledges limitations and identifies future research directions. And Section 8 concludes.

1.1. Approach and research question

The paper is a work of conceptual and normative analysis. Rather than testing hypotheses against new empirical data, it develops a conceptual apparatus for an emerging governance problem and derives normative implications for regulatory and administrative practice. The guiding research question is: *How does the algorithmic construction of perceptual frameworks challenge existing democratic-governance and public-administration arrangements, and what accountability mechanisms would address the resulting governance gap?* Three sub-questions structure the analysis: 1. What conceptual account best captures the phenomenon (Section 2)? 2. Where do current EU regulatory instruments fail to reach it (Section 4)? 3. What governance response follows (Sections 5–6)?

The analysis proceeds in three steps. First, *conceptual construction:* foundational work in information theory (Shannon, 1948; Bateson, 1972; Floridi, 2011), surveillance capitalism and algorithmic regulation scholarship (Zuboff, 2019; Yeung, 2018), and democratic theory (Habermas, 1996; Landemore, 2013; Dryzek, 2000) is synthesised to develop the construct of *perceptual infrastructure.* Literature was selected purposively for conceptual relevance to the relationship between information, perception and democratic deliberation, prioritising foundational and high-citation works that define the relevant constructs rather than aiming at exhaustive coverage of any single field. Second, *regulatory gap analysis:* the EU AI Act and the Digital Services Act are read against the construct, asking which risks each provision recognises and which it excludes (Section 4). Third, *normative derivation:* an accountability framework is inferred from the gap identified, distributing responsibility across the actors with the capacity to shape outcomes (Section 5).

The limitations inherent in this approach are noted in Section 7: conceptual-normative work generates frameworks and hypotheses rather than validated findings, and the constructs proposed here await empirical operationalisation.

2. Theoretical foundations

2.1. Information as perception given a framework

The analysis is grounded in a specific definition: information is perception given a framework. The formulation draws on and extends foundational work in information theory while resolving tensions that have limited governance applications.

Shannon's (1948) mathematical theory of communication provided quantitative rigour but explicitly excluded semantic meaning, focusing instead on the probability of signal transmission. Shannon's framework operates powerfully within engineering contexts but cannot address the question of what transmitted signals mean to human perceivers: a question central to democratic governance, where the meaning citizens derive from public data determines their political behaviour. Bateson's (1972) relational definition – information as "a difference that makes a difference" – restored meaning but left unaddressed the question of what determines which differences matter for different perceivers. A rustling in the grass constitutes information for a predator but not for a geologist examining rock formations; the same economic indicators constitute different information for a fiscal conservative and a Keynesian economist, not because they apply different interpretations to shared data but because their frameworks render different patterns salient.

Floridi (2011) offered the most comprehensive recent philosophical integration, defining semantic information as well-formed, meaningful and truthful data. However, the truth requirement creates practical difficulties for governance: disinformation functions as information in democratic systems precisely because it is perceived as meaningful within particular frameworks. Citizens act on false beliefs – voting, protesting, withdrawing civic participation, supporting authoritarian measures – because those beliefs emerged as genuine information through their perceptual frameworks. A governance framework that defines false beliefs as "not information" cannot adequately theorise the mechanisms through which AI systems shape democratic behaviour.

The definition adopted here resolves these tensions by centring the perceiver. Information does not exist independently of perceivers; it emerges through the interaction between a perceiver and data within interpretive frameworks. The same data generates different information for different perceivers because they bring different frameworks to bear. It is not merely a matter of different *interpretations* of shared information. The interpretation model assumes that information exists prior to and independently of perceivers, who then apply different lenses to shared content. The framework developed here locates the phenomenon at a more fundamental level: different perceptual frameworks produce different perceptions, yielding genuinely different information from identical data.

For public administration, the distinction has concrete implications. If the problem is a biased interpretation of shared information, the governance response involves promoting media literacy and critical thinking: interventions targeting the interpretive process. If the problem is divergent information emergence from shared data – citizens literally perceiving different realities from the same evidence – these interventions are insufficient because they presuppose shared informational ground that does not exist. The governance response must address how perceptual frameworks are constructed in the first place. As the post-factual governance literature demonstrates, when the perceptual infrastructure itself is fragmented, interventions targeting content – fact-checking, source labelling, content removal – address symptoms while leaving the underlying structural condition unexamined (Alibašić, 2024; Alibašić & Rose, 2019).

2.2. Algorithmic shaping of perceptual infrastructure

Three principles follow from the definition, each with direct governance implications.

First, *perceptual frameworks reside in the perceiver.* Although physical and social contexts are commonly spoken of as external features, the interpretive frameworks enabling perception are features of the perceiving entity: cognitive architecture, training, experience and habitual patterns of attention. External circumstances, including algorithmically curated information environments, shape what frameworks develop over time, but the resulting perceptual capacities are internalised. A citizen whose information environment has been algorithmically curated for years does not simply receive biased content that can be corrected by exposure to alternative sources – that citizen has developed perceptual frameworks that determine what patterns appear salient, what connections seem meaningful, and what conclusions appear obvious even when encountering new, uncurated information. The internalisation is precisely what makes algorithmic framework construction so consequential for governance: the effects persist beyond the immediate algorithmic environment.

Second, *controlling framework construction controls information emergence.* Whoever shapes perceptual frameworks determines what information citizens derive from public data, without the visible coercion of censorship or the obvious distortion of propaganda. It is a form of power that operates below conscious awareness: citizens experience their perceptions as direct apprehensions of reality, unaware that their frameworks have been systematically shaped by algorithmic exposure patterns. The power is more subtle than propaganda because it does not require false content – it shapes the perceptual conditions under which true content is received and processed. Two citizens watching the same legislative debate perceive different information not because one has been deceived but because their frameworks render different elements salient.

Third, *algorithmic framework construction operates at unprecedented scale and precision.* While institutions, media organisations, educational systems and political parties have always shaped perceptual frameworks, AI systems construct them for billions of users simultaneously, personalised to individual behavioural profiles, optimised for commercial metrics that may diverge from democratic governance needs, and refined through

continuous feedback loops that accelerate framework development. The personalisation dimension is particularly significant: unlike broadcast media, which constructed relatively uniform frameworks for large audiences, algorithmic systems can construct millions of individualised frameworks, creating a fragmentation of perceptual infrastructure that has no historical precedent.

Zuboff (2019) identified the economic logic driving the process: surveillance capitalism operates through the extraction of behavioural surplus – data generated by user engagement – for prediction and behavioural modification. The framework developed here specifies the mechanism through which the modification operates: the perceptual frameworks within which certain behaviours become more probable are themselves constructed through algorithmic curation. The behavioural modification that surveillance capitalism pursues is not merely nudging within a fixed perceptual landscape; it reshapes the landscape itself. When platforms derive revenue from advertising markets that reward engagement, the economic incentives structurally favour framework construction that maximises emotional activation over epistemic accuracy – a misalignment with democratic governance needs that is not a design flaw but a feature of the underlying economic model.

Yeung (2018) similarly demonstrated that algorithmic regulation constitutes a distinct system of social ordering, operating through design-based techniques that embed value-laden decisions in computational architectures users cannot easily perceive or resist. Her taxonomy of algorithmic regulation – distinguishing reactive from pre-emptive systems, and classifying systems by their configuration across standard-setting, monitoring and behavioural-change functions – provides analytical infrastructure for understanding how platforms exercise regulatory power through perceptual framework construction. The shaping of perceptual infrastructure represents the deepest layer of the regulatory power: it operates not at the level of specific behavioural nudges but at the level of the perceptual conditions enabling all subsequent information processing.

2.3. Implications for democratic theory

Democratic governance theories from deliberative democracy (Habermas, 1996) to epistemic democracy (Landemore, 2013) presuppose that citizens can share sufficient common ground to engage in productive deliberation about collective problems. The presupposition is not merely normative, as it identifies an epistemic condition: citizens must be able to perceive sufficiently compatible information from shared data to enable meaningful collective decision-making.

Habermas's (1996) account of communicative rationality assumes interlocutors capable of engaging with shared validity claims: claims to truth, rightness and sincerity that can be redeemed through argumentative discourse. When algorithmic systems construct divergent perceptual frameworks, the shared validity claims presupposed by communicative rationality may not exist: citizens operating within different frameworks may not be making competing claims about shared phenomena but perceiving different phenomena entirely. Landemore's (2013) argument for the epistemic value of cognitive diversity – the claim that diverse groups make better collective decisions than

homogeneous expert bodies – similarly assumes diverse perspectives operating on shared informational terrain. Cognitive diversity is epistemically productive when diverse perceivers bring different analytical lenses to common data; it becomes epistemically destructive when perceivers inhabit frameworks so divergent that they cannot recognise the data others perceive.

Dryzek's (2000) discursive democracy tradition, which emphasises the role of competing discourses in shaping democratic outcomes, provides additional theoretical resources. He argues that democratic legitimacy depends on the contestation of discourses in the public sphere. However, discursive contestation presupposes that competing discourses address recognisably shared problems; a condition threatened when algorithmic systems construct frameworks in which different citizen groups perceive fundamentally different problem landscapes.

Three mechanisms through which algorithmic shaping of perceptual frameworks threatens democratic governance warrant particular attention.

First, *perceptual fragmentation:* when algorithmic systems construct divergent frameworks for different population segments, citizens may perceive such different information from shared data that productive dialogue becomes impossible. It is distinct from political polarisation as conventionally understood – the problem is not that citizens hold different opinions about shared facts but that they perceive different facts from identical data. The distinction matters for governance: polarisation of opinion can be addressed through deliberative mechanisms that bring opposing perspectives into dialogue, while perceptual fragmentation requires addressing the underlying framework divergence before dialogue can be productive.

Second, *engagement–truth misalignment:* platforms optimised for engagement construct perceptual frameworks where emotionally arousing content appears more informative than accurate but complex analysis. Levy's (2021) field experiment demonstrated that Facebook's algorithm substantially increased exposure to partisan news sources. On YouTube, by contrast, Hosseinmardi et al. (2021) found little evidence that the recommendation algorithm drives attention to radical content, with consumption concentrated among users who actively seek it – a reminder that the degrading mechanism operates through engagement optimisation and emotional salience rather than through demonstrated large-scale algorithmic steering. These studies provide evidence of mechanisms that systematically degrade the epistemic quality of democratic discourse by constructing frameworks in which emotional intensity functions as a proxy for informational importance.

Third, *adversarial exploitation:* state and non-state actors can deliberately construct incompatible perceptual frameworks to fragment democratic publics and prevent collective action (Alibašić, 2024; Alibašić, 2019). Effective adversarial campaigns exploit the primacy of framework construction: rather than directly introducing false claims, they first construct perceptual conditions ensuring targets will perceive desired information from partially true data. A campaign establishing systematic distrust toward democratic institutions creates perceptual conditions under which routine governance operations are perceived as evidence of conspiracy or authoritarian intent.

The empirical literature on aggregate effects presents a contested picture. Some studies find modest or limited effects of algorithmic curation, with consumption driven more by user choice than algorithmic steering (Bakshy et al., 2015; Dubois & Blank, 2018; Hosseinmardi et al., 2021; Guess et al., 2023). Others document significant impacts in specific contexts (Levy, 2021). The framework adopted here navigates the terrain by distinguishing the *mechanism* of framework construction from its *aggregate effects*. Even if current systems produce modest measurable polarisation at the population level, the demonstrated capacity for perceptual framework construction at scale creates structural risks to democratic governance that warrant proactive regulatory attention, consistent with precautionary approaches in regulatory governance (Yeung, 2018). The question for governance is not solely whether harm has been demonstrated at aggregate levels but whether the mechanism for producing harm at scale exists and remains ungoverned.

3. Taxonomy of algorithmically shaped frameworks

Understanding governance challenges requires distinguishing how different types of algorithmic framework construction operate and interact within democratic systems.

The taxonomy is organised along a single primary axis – the *intentionality and origin of framework construction* – running from frameworks that emerge as unintended byproducts of optimisation (curated) through deliberately engineered frameworks (adversarial) to frameworks arising within organisational AI deployment (institutional), with real-world frameworks typically combining these (hybrid/contested). Secondary dimensions – the actor responsible, the operative mechanism, the governing objective and the resulting risk type – are summarised in Table 1. The categories are treated as analytically distinct but empirically interacting; Section 3.2 addresses the curated–adversarial interaction as a structural feature rather than an exception.

Table 1
A taxonomy of algorithmically constructed perceptual frameworks

Framework type	Primary actor	Mechanism	Objective	Typical example	Dominant risk	Governance tool
Curated	Platforms	Engagement-optimised selection, sequencing, emphasis	Commercial (attention, ad revenue)	Social media feed; AI-synthesised search	Perceptual fragmentation; engagement–truth misalignment	Aggregate transparency; perceptual-sovereignty controls (Sections 5.1–5.2)

Framework type	Primary actor	Mechanism	Objective	Typical example	Dominant risk	Governance tool
Adversarial	State / non-state operators	Deliberate seeding of distrust-building frameworks; exploitation of partial truths	Political / strategic manipulation	Influence operation establishing institutional distrust	Self-reinforcing manipulation; collective-action paralysis	Prohibition of covert construction; circuit breakers; auditing (Section 5.3)
Institutional	Public organisations	Decision-support tools reshaping professional perception	Administrative efficiency / decision support	Recidivism risk scoring; welfare-eligibility tools	Institutional framework capture; displaced professional judgment	Framework impact assessments; hybrid intelligence design (Section 6.1)
Hybrid / contested	Multiple, interacting	Combination of the above across one perceiver	None unified; emergent	Contested zoning dispute; polarised public health response	Divergent factual perception; deliberative breakdown	Sequenced intervention: address framework divergence before deliberation (Sections 3.4, 6.3)

Source: Compiled by the author.

3.1. Curated frameworks

Recommendation systems, search engines, and social media feeds algorithmically construct curated frameworks through continuous exposure to selected content optimised for engagement, time on platform, or advertising revenue. These represent a historically novel phenomenon: the automated construction of perceptual frameworks at scale, personalised to individual behavioural profiles, and refined through continuous feedback loops.

Facebook's news feed algorithm, serving approximately two billion users daily (Statista, 2024), exemplifies the process. The algorithm shapes not just individual content decisions but the overarching perceptual framework through which users encounter public affairs, operating through several interrelated mechanisms. *Source selection* builds familiarity and perceived credibility for frequently displayed sources, creating anchoring effects that shape how all subsequent information is processed. Over time, regularly encountered sources become default interpretive authorities – not because users have consciously chosen them but because algorithmic repetition has constructed familiarity frameworks. *Topic emphasis* makes certain policy issues appear more salient and important while rendering others perceptually invisible; citizens whose feeds emphasise immigration may perceive it as the defining issue of their era, while those whose feeds emphasise

economic indicators perceive a fundamentally different political landscape. *Narrative framing* juxtaposes content in ways that construct implicit causal narratives – consistently placing immigration stories alongside crime reports builds associative frameworks linking these topics without any individual item containing a false claim. *Emotional conditioning* prioritises emotionally arousing content over nuanced analysis, constructing frameworks where emotional intensity serves as a proxy for informational importance, gradually eroding the capacity to engage with complex policy analysis that lacks emotional activation.

For democratic governance, the critical issue is that traditional content moderation addresses individual posts while leaving the larger framework unexamined. A citizen might encounter no individual item qualifying as "misinformation" yet develop systematically distorted perceptions of public affairs because the algorithm has constructed a perceptual framework rendering certain patterns salient while making others invisible.

The shift toward AI-powered search with large language models intensifies these dynamics considerably. Traditional search engines presented multiple sources in ranked order, allowing users to compare perspectives and construct their own interpretive frameworks. AI-synthesised search responses consolidate multiple sources into unified answers – effectively constructing a single framework for the user. The AI determines which sources inform the synthesis (often opaque to users), chooses how to frame topics and which aspects to emphasise, and adopts particular perspectives while projecting neutrality and comprehensiveness. Citizens may perceive AI-synthesised responses as objective summaries when they reflect specific frameworks embedded in training data, ranking algorithms and synthesis processes. The governance challenge is acute: as AI-powered search displaces traditional search for an increasing share of information seeking, the capacity for individualised framework construction through independent source evaluation diminishes.

3.2. Adversarial frameworks

Adversarial frameworks are deliberately constructed to manipulate perception for political, commercial, or strategic objectives. Unlike curated frameworks arising as side effects of engagement optimisation, adversarial frameworks are intentionally designed to produce specific perceptual outcomes in target populations.

The sophistication of contemporary adversarial operations lies in their targeting of perceptual infrastructure rather than individual beliefs. Rather than directly introducing false claims – which can be fact-checked and debunked – effective campaigns first construct frameworks ensuring targets will perceive desired information from partially true data (Alibašić, 2024). A campaign establishing systematic distrust toward mainstream media, scientific institutions, and democratic processes creates perceptual conditions under which routine government operations are perceived as evidence of conspiracy, ordinary journalistic practices are perceived as coordinated manipulation, and public health measures are perceived as authoritarian control mechanisms. Critically, once these frameworks are established, they are self-reinforcing: information that would challenge

the framework is perceived through it, and thus appears as further confirmation rather than disconfirmation.

The intersection of adversarial framework construction with algorithmic curation creates compound threats. When adversarial actors seed content designed to construct particular frameworks, algorithmic systems optimised for engagement may amplify the content – not through any deliberate collaboration but because content designed to provoke strong emotional responses performs well within engagement-optimised systems. The adversarial operator constructs the initial framework; the platform's algorithm reinforces and extends it through its normal operations. It creates a governance challenge that neither content moderation (addressing individual items) nor algorithmic transparency (explaining individual recommendations) can adequately address.

The curated and adversarial categories are analytically separable by intent but structurally coupled in operation. Adversarial construction does not bypass curated systems; it parasitises them. An operator engineers content to satisfy the engagement signals that curated systems reward, so that the platform's commercial optimisation completes the adversarial design without any deliberate collaboration (cf. Yeung, 2018, on design-based ordering). In Table 1, these appear as distinct rows because their governance tools differ – transparency and user control for curated construction, prohibition and detection for adversarial – but a complete governance response must address their coupling, since tools aimed at one without the other leave the compound pathway open.

3.3. Institutional frameworks

Public sector organisations increasingly deploy AI systems for decision support in healthcare, criminal justice, welfare administration and regulatory enforcement. These systems interact with and potentially reshape existing institutional perceptual frameworks in ways that carry significant governance implications.

When judges repeatedly use recidivism risk assessment algorithms, they may gradually internalise the frameworks these tools construct – perceiving certain demographic or behavioural factors as more predictive of recidivism than warranted by evidence. Medical diagnosis support systems trained primarily on data from one demographic group construct frameworks where symptoms in that group appear more salient and clinically meaningful, potentially reshaping how physicians perceive patients from underrepresented populations. Welfare eligibility algorithms construct frameworks determining what factors appear relevant to need assessment, potentially displacing the nuanced professional judgment that experienced caseworkers bring to complex human situations.

The governance challenge extends beyond individual decision accuracy: AI systems can fundamentally alter the institutional frameworks through which public servants perceive their work, constituting what Yeung (2018) classifies as pre-emptive algorithmic regulation operating within organisational structures. Over time, the organic professional judgment built through years of training and practice may be displaced by algorithmically constructed frameworks that embed particular assumptions about what patterns matter

and what factors predict outcomes. When the displacement occurs gradually and unconsciously – when professionals do not recognise that their perceptual frameworks have shifted – the result is institutional framework capture: the organisation's perceptual infrastructure has been reshaped by the AI system without deliberate institutional choice.

3.4. Hybrid and contested frameworks

Real-world perceptual frameworks rarely fit neatly into single categories. A citizen's framework for perceiving public affairs combines organic elements developed through education and lived experience, curated elements constructed through algorithmic personalisation, institutional elements shaped by professional training and organisational norms, and technological elements determined by interface design and medium affordances. The interaction among these elements creates hybrid frameworks with complex governance characteristics.

Democratic deliberation characteristically occurs across contested hybrid frameworks, where different citizen groups may perceive fundamentally different information from shared data. When a municipal government proposes a zoning change, residents whose frameworks have been shaped by different algorithmic environments may perceive the proposal through incompatible lenses – one group perceiving community development opportunity, another perceiving displacement threat, another perceiving corruption. These are not merely different values brought to shared facts; they may be different factual perceptions arising from divergent frameworks.

Contested frameworks present acute governance challenges because the mechanisms designed for managing value disagreement – compromise, negotiation, deliberative forums, democratic voting – presuppose that participants perceive a shared factual landscape, even if they evaluate it differently. When the frameworks are so divergent that participants perceive different factual landscapes, these mechanisms lose their epistemic grounding. Governance in conditions of contested frameworks requires first addressing the perceptual divergence before conventional deliberative mechanisms can function effectively.

4. Regulatory gap analysis: The EU AI Act and perceptual infrastructure

The EU AI Act represents the most comprehensive regulatory framework for AI governance, establishing risk-based classification from unacceptable through high-risk to limited and minimal risk (AI Act, Recital 26–27). However, the Act's treatment of perceptual infrastructure reveals significant governance gaps that warrant systematic analysis.

4.1. Classification limitations

The Act classifies AI systems used for recommendation and algorithmic curation as limited-risk or minimal-risk applications, subject primarily to transparency obligations rather than the more rigorous requirements applied to high-risk systems in domains such as healthcare, criminal justice and employment. The classification reflects a focus on the risk posed by individual system interactions rather than the cumulative, long-duration effects of algorithmic framework construction.

Yadav (2025) observes that the temporal dynamics of algorithmic influence may extend beyond the Act's assessment frameworks – a concern, which the analysis substantiates. The construction of perceptual frameworks operates through years of micro-exposures that individually appear innocuous but cumulatively reshape the infrastructure through which citizens perceive democratic governance. The Act's static risk classification, which assesses systems at specific points in time, does not capture the dynamic, cumulative process. A recommendation system that scores as low-risk in any given interaction may, over years of sustained operation, fundamentally reshape the perceptual infrastructure of millions of citizens. The risk is not located in any individual algorithmic decision but in the aggregate pattern of decisions sustained over time – a temporal dimension that the Act's classification framework is not designed to capture.

4.2. Transparency deficits

The Act mandates transparency for general purpose AI models and imposes disclosure obligations on providers and deployers. However, transparency about perceptual framework construction requires fundamentally different mechanisms than transparency about individual outputs. Citizens need not merely to understand why a particular recommendation was made but to perceive how their informational environment has been systematically shaped over time – the aggregate patterns, source distributions, topic emphases, and narrative framings that constitute their perceptual infrastructure.

The Act's transparency provisions, oriented toward individual decisions and system-level documentation, do not require the longitudinal, aggregate disclosure. A platform could provide full transparency about why each individual piece of content was recommended – satisfying the Act's requirements – while the cumulative pattern of those recommendations constructs a perceptual framework that is never disclosed to the user and remains invisible to regulators.

4.3. The content-infrastructure gap in the EU regulatory architecture

The AI Act's provisions, complemented by the DSA's algorithmic transparency and risk assessment requirements for VLOPs, address content moderation and individual recommendation decisions. The DSA requires platforms to assess and mitigate systemic risks,

including risks to civic discourse and electoral processes. However, neither instrument addresses the construction of perceptual infrastructure as a distinct governance domain.

A platform could comply fully with both the AI Act and the DSA – moderating prohibited content, providing transparency about individual algorithmic decisions, conducting required systemic risk assessments – while systematically constructing divergent perceptual frameworks that undermine the epistemic foundations of democratic deliberation. The systemic risk assessments required by the DSA focus on identifiable harms (disinformation, electoral manipulation, negative effects on minors) rather than on the gradual construction of perceptual frameworks that may not produce identifiable harms at any single point in time but cumulatively reshape the democratic public's capacity for shared perception.

The gap reflects a broader conceptual limitation in contemporary AI governance: the predominant focus on outputs (what systems recommend, decide, or generate) rather than on the upstream construction of the perceptual frameworks that determine how outputs are received and processed. The EU AI Act's prohibition on AI systems that deploy subliminal techniques or exploit vulnerabilities to materially distort behaviour provides a starting point, but these provisions target discrete manipulative acts rather than the gradual, cumulative construction of perceptual frameworks through sustained algorithmic curation.

4.4. Omission or implementation deficit?

A natural objection is that perceptual-infrastructure harms might already be reachable under existing provisions, making the problem one of *implementation* rather than *omission*. The distinction matters for governance: an implementation deficit calls for enforcement resources and guidance, whereas an omission calls for new regulatory categories. The analysis here indicates the gap is *primarily* one of omission, with a *secondary* implementation dimension.

The case for omission rests on the object of regulation. The EU AI Act's risk classification (Article 6 and Annex III) attaches to *systems* and *use cases* assessed at defined points, and the DSA's systemic-risk regime (Articles 34–35) attaches to *identifiable* systemic risks – disinformation, electoral integrity, effects on minors, fundamental rights. Neither instrument contains a provision whose object is the cumulative, longitudinal construction of perceptual frameworks. The prohibition on subliminal and manipulative techniques in the AI Act reaches *discrete* manipulative acts that *materially distort behaviour;* it does not reach the gradual, individually innocuous micro-exposures whose harm is constituted only in aggregate over time (Yadav, 2025). The harm identified here is thus excluded not because regulators have failed to apply an existing rule but because no provision takes the relevant object – aggregate framework construction – as its subject. That is omission.

The secondary implementation dimension is real but narrower. Where the DSA's systemic-risk assessments *could,* in principle, capture framework effects – civic discourse risk under Article 34 is drafted broadly enough to admit a generous reading – the limiting factor is that the assessments target point-in-time identifiable harms and lack the

longitudinal, aggregate methodology that framework construction requires. Here, the deficit is partly implementational (the methodology does not yet exist) and partly definitional (the regulation does not require it). The governance implication is correspondingly two-part: extend the regulatory object to recognise perceptual infrastructure as a distinct domain (addressing the omission), and develop the longitudinal audit methodologies that would let existing systemic-risk mandates reach framework effects (addressing the deficit). Both are developed in Sections 5–6.

5. An accountability framework for perceptual infrastructure

Governing the algorithmic construction of perceptual frameworks requires mechanisms distinct from – though complementary to – existing content moderation and algorithmic transparency approaches.

5.1. Aggregate transparency

Traditional algorithmic transparency focuses on explaining individual decisions – why a particular recommendation was made or a specific piece of content surfaced. Governing perceptual infrastructure requires *aggregate transparency:* disclosure mechanisms showing how algorithmic systems construct perceptual frameworks over time.

Platforms should be required to provide users with longitudinal dashboards showing the distribution of sources they have been exposed to, topic emphasis patterns in their information environment, diversity metrics indicating framework homogeneity or heterogeneity, and historical changes in their informational landscape. These dashboards should enable citizens to perceive the aggregate patterns that constitute their perceptual frameworks – making visible what is by design invisible in current platform architectures. Platforms should also publicly disclose high-level principles governing framework construction: what metrics drive construction decisions, how personalisation algorithms balance individual preferences against exposure to diverse perspectives, and what values are embedded in construction processes.

Critically, transparency must be *meaningful* – accessible and actionable for citizens without technical expertise. It requires coupling disclosure with effective visualisation tools, plain language explanations, comparative views showing how individual frameworks differ from population-level patterns, and mechanisms allowing users to modify how their frameworks are being constructed. Transparency that is technically accurate but practically inaccessible fails the democratic purpose it is intended to serve.

5.2. Perceptual sovereignty as a citizen right

This paper proposes *perceptual sovereignty* – the principle that citizens have legitimate interests in controlling how their own perceptual frameworks develop. It does not mean

isolation from all external influence – human beings inevitably develop perceptual frameworks through social interaction, education and institutional participation – but establishes a presumption against covert, systematic, algorithmically driven manipulation of frameworks optimised for commercial rather than democratic ends.

Perceptual sovereignty requires several concrete mechanisms: opt-out options allowing citizens to receive chronological content without algorithmic curation; meaningful control over the principles governing framework construction, including source diversity requirements and the ability to set parameters for algorithmic exposure; periodic "reset" options preventing extreme framework homogenisation; access to multiple curated environments constructed according to different principles, enabling citizens to compare how different frameworks shape their perception; and enhanced protections for information environments directly relevant to democratic participation, including electoral periods and policy deliberation. Special protections are warranted for vulnerable populations, including minors and those with cognitive impairments, whose capacity to recognise and resist framework manipulation may be limited.

5.3. Proportional responsibility

Accountability for the algorithmic construction of perceptual frameworks should be proportionally distributed among stakeholders based on their capacity to shape outcomes and the magnitude of their influence.

Platform responsibilities: Platforms constructing frameworks at scale bear primary responsibility for aggregate transparency, meaningful user control, prevention of extreme framework homogenisation, protection against adversarial framework construction, regular third-party auditing of framework construction effects, and investment in research on the democratic implications of their systems. Platforms should also be required to implement automated "framework circuit breakers" – mechanisms that detect when a user's framework has become highly homogeneous and trigger deliberate exposure to diverse perspectives.

Public institutional responsibilities: Government organisations deploying AI decision support bear responsibility for understanding and monitoring how systems reshape institutional frameworks, maintaining organic professional judgment alongside AI-constructed frameworks through training and institutional design, ensuring that AI-constructed institutional frameworks align with democratic values and legal requirements, and conducting framework impact assessments before deploying new AI systems.

Regulatory responsibilities: Governments and regulatory bodies should establish aggregate transparency requirements extending beyond individual decisions to longitudinal framework effects, mandate meaningful user control mechanisms, prohibit covert adversarial framework construction, commission independent auditing of framework construction at scale, develop standards for framework construction in domains directly relevant to democratic governance, and invest in research on democratic implications. The EU AI Office provides an institutional model that should be extended to encompass perceptual infrastructure as a distinct governance domain.

User responsibilities: User responsibility should be understood as limited by the frameworks citizens inhabit. Holding individuals fully responsible for perceptions shaped by years of algorithmic curation is both normatively inappropriate and empirically unsupported. Users bear responsibility for engaging with transparency mechanisms where available and developing awareness of their own perceptual frameworks, but the responsibility is secondary to the obligations of those who construct the frameworks at scale.

5.4. Meta-perceptual literacy

Effective governance of perceptual infrastructure requires citizens capable of recognising that they perceive through frameworks susceptible to algorithmic construction. Traditional digital literacy emphasises evaluating individual sources and claims – a necessary but insufficient skill set. What may be termed *meta-perceptual literacy* extends it to evaluating the frameworks through which one perceives information: recognising how sustained algorithmic exposure patterns shape perception over time, identifying how systems construct frameworks through content selection, sequencing and presentation, deliberately adopting different frameworks to understand how other citizens perceive shared data, and recognising markers of adversarial framework manipulation.

The EU AI Act's literacy provisions, mandating Member States to promote AI literacy, provide a regulatory foundation for such initiatives. However, technical AI literacy – understanding how algorithms function at a mechanical level – is insufficient without the meta-perceptual dimension: understanding how sustained algorithmic exposure shapes one's own perceptual capabilities and recognising the gap between perceiving one's framework as natural or self-evident and recognising it as partially constructed. Educational curricula should include practical exercises in framework shifting – deliberately encountering information through alternative frameworks to develop awareness of one's own perceptual assumptions and the contingency of one's habitual perception patterns.

6. Implications for public administration

6.1. Public sector decision-making and institutional frameworks

As public organisations deploy AI decision support systems across healthcare, welfare, criminal justice and regulatory enforcement, they must attend to how these systems reshape institutional perceptual frameworks. *Framework impact assessments* are recommended – systematic evaluations conducted prior to AI deployment examining what frameworks the system is likely to construct among organisational users, whether these frameworks align with institutional values and democratic commitments, what training is needed for staff to maintain critical distance from AI-constructed

frameworks, and how the interaction between AI-constructed and organic professional frameworks will be managed over time.

Hybrid intelligence approaches – preserving space for professional judgment informed by organic institutional frameworks alongside AI-constructed frameworks – can prevent the wholesale displacement of professional perceptual competence. Organisations should maintain practices that exercise and develop organic professional judgment even as they integrate AI decision support, ensuring that algorithmic frameworks complement rather than replace human expertise. It requires deliberate institutional design: regular case reviews conducted without AI support, professional development activities that strengthen organic perceptual capacities, and rotation between AI-assisted and non-assisted work environments.

Concretely, public organisations should 1. require a *framework impact assessment* – modelled on data-protection impact assessments – as a precondition for procuring any AI decision-support system; 2. mandate periodic *AI-free case reviews* in which professionals exercise unassisted judgment on a sampled caseload, generating a baseline against which to detect framework drift; and 3. assign explicit organisational ownership of framework monitoring, rather than leaving it diffused across IT and compliance functions.

6.2. Regulatory capacity and institutional design

Governing the algorithmic construction of perceptual frameworks demands regulatory capacities that few public administrations currently possess. Regulators need analytical tools capable of auditing framework construction effects at scale, multidisciplinary expertise bridging computational systems, cognitive science and democratic theory, institutional structures enabling continuous monitoring rather than one-time compliance assessment, and the ability to assess cumulative effects of algorithmic exposure over time rather than evaluating individual system interactions in isolation.

The EU AI Office and the broader governance architecture established by the AI Act provide institutional models. However, governing perceptual infrastructure requires extending these models beyond their current mandates, developing new audit methodologies appropriate to long-duration framework construction rather than point-in-time system evaluation, and building regulatory capacity that can keep pace with rapidly evolving AI capabilities. National regulatory authorities will need to develop competencies that currently exist in fragmented form across data protection agencies, media regulators and competition authorities – suggesting the need for institutional coordination mechanisms or dedicated regulatory bodies with cross-cutting mandates.

6.3. Democratic resilience and perceptual common ground

Public administration scholarship has increasingly recognised resilience as a governance imperative – the capacity of public systems to absorb shocks, adapt to disruptions, and maintain core functions during periods of stress (Alibašić, 2024). Algorithmic reshaping

of perceptual infrastructure represents a novel and underexamined threat to democratic resilience.

Democratic systems depend upon citizens sharing sufficient perceptual common ground to engage in collective sense-making during periods of stress. When algorithmic systems fragment the common ground – constructing divergent frameworks through which different population segments perceive fundamentally different information from shared evidence – democratic systems lose adaptive capacity precisely when they most need it. A public health emergency requires citizens who can perceive shared evidence about disease transmission and treatment efficacy; an economic transition requires citizens who can recognise common challenges and evaluate proposed responses; a security threat requires citizens who can identify shared risks and support collective action. When perceptual infrastructure is fragmented, each of these collective capacities is compromised, and the democratic system's ability to mount coherent responses to crises deteriorates.

It has particular relevance for governance in contexts of democratic fragility, including the Central and Eastern European democracies that have experienced both rapid digitalisation and democratic stress. In these settings, the algorithmic construction of divergent perceptual frameworks intersects with existing social cleavages, institutional fragilities, historical mistrust of public institutions, and foreign influence operations. Societies where democratic institutions are still consolidating and where historical experience has generated justified scepticism toward state authority are particularly vulnerable to framework construction that systematically amplifies institutional distrust. The compound effect – algorithmic framework construction operating on populations with pre-existing institutional scepticism, in media environments with limited independent journalism capacity, and under conditions of active foreign influence operations – creates threats to democratic resilience that neither content moderation nor conventional media regulation can adequately address.

7. Limitations and future directions

Several limitations warrant acknowledgment. First, the framework developed here is primarily theoretical; empirical validation through longitudinal studies tracking framework development, experimental research on framework manipulation resistance, and field studies of governance interventions remains essential. The measurement of perceptual frameworks presents particular methodological challenges: most existing research relies on behavioural proxies (what content users engage with) or self-reports (what users say they perceive), neither of which directly captures the perceptual frameworks themselves. Developing more direct measurement methods constitutes a research priority.

Second, most existing research on algorithmic framework construction focuses on Western, English-speaking democracies with robust independent media ecosystems. How these dynamics operate across diverse political systems, languages, media environments and cultural contexts – including the Central and Eastern European contexts – requires

substantially more investigation. Governance mechanisms effective in one political and media environment may require significant adaptation for others.

Third, AI technologies continue evolving rapidly, with large language models, virtual and augmented reality, and multimodal AI systems raising new questions about framework construction that current analysis cannot fully anticipate. The emergence of AI systems capable of generating personalised audio, video and interactive content creates framework construction capabilities qualitatively different from text-based recommendation systems. Governance frameworks must remain adaptable as these capabilities emerge.

Future research should prioritise the development of measurement methods for perceptual frameworks and their algorithmic construction; cross-jurisdictional comparative analysis of regulatory approaches to algorithmic influence on democratic perception; investigation of how framework construction interacts with existing democratic institutions across different governance traditions; empirical assessment of intervention effectiveness including aggregate transparency mechanisms, perceptual sovereignty tools, framework circuit breakers and meta-perceptual literacy programs; and examination of how compound threats – algorithmic framework construction intersecting with foreign influence operations and democratic fragility – manifest in specific national contexts.

8. Conclusion

The argument moved from a conceptual account of information as perception given a framework (Section 2), through a taxonomy distinguishing curated, adversarial, institutional and hybrid construction (Section 3), to a provision-level demonstration that the EU architecture omits perceptual infrastructure as a regulatory object (Section 4); the accountability framework (Section 5) and administrative implications (Section 6) follow from that identified gap rather than standing as independent proposals.

The paper has argued that the algorithmic construction of perceptual frameworks constitutes a fundamental and inadequately governed challenge for democratic governance – a challenge that defines the post-factual governance environment in which public administration now operates. AI systems exercise power not merely through content selection but by reshaping the perceptual infrastructure through which citizens encounter public affairs – a mechanism operating below the threshold of conventional content moderation and existing regulatory instruments, including the EU AI Act and the Digital Services Act.

The theoretical framework grounded in the principle that information emerges through the interaction of perception and interpretive frameworks reveals why conventional responses – content moderation, individual algorithmic transparency, media literacy – are necessary but insufficient. These interventions address symptoms while leaving the underlying infrastructure unexamined. When the perceptual infrastructure of democracy is fragmented, fact-checking individual claims cannot restore shared perception, and transparency about individual recommendations cannot make visible the aggregate patterns that constitute citizens' perceptual frameworks.

The accountability framework proposed here – incorporating aggregate transparency, perceptual sovereignty as a citizen right, proportional responsibility and meta-perceptual literacy – provides a foundation for governing the challenge. However, implementing these principles requires extending existing regulatory instruments to recognise perceptual infrastructure as a distinct governance domain warranting dedicated institutional attention, regulatory capacity and enforcement mechanisms.

The stakes are fundamental. When algorithmic systems construct the frameworks through which citizens perceive democratic governance, they exercise a form of power that threatens the epistemic foundations of collective self-governance. Public administration scholarship must engage the challenge directly – developing both the theoretical understanding and practical governance mechanisms needed to preserve democratic capacity in the age of algorithmically mediated post-factual reality.

References

Alibašić, H. (2019). *Brexit* by Toby Haynes: Moral and Ethical Irrelevance and Naïve Expectations. *Public Integrity, 21*(5), 537–542. Online: https://doi.org/10.1080/10999922.2019.1626796

Alibašić, H. (2024). *The Post-Factual Polity. Ethical, Governance, Administrative, and Policy Crises in the Disinformation Era.* Information Age Publishing. Online: https://doi.org/10.1108/979-8-88730-661-2

Alibašić, H. & Rose, J. (2019). Fake News in Context: Truth and Untruths. *Public Integrity, 21*(5), 463–468. Online: https://doi.org/10.1080/10999922.2019.1622359

Bakshy, E., Messing, S. & Adamic, L. A. (2015). Exposure to Ideologically Diverse News and Opinion on Facebook. *Science, 348*(6239), 1130–1132. Online: https://doi.org/10.1126/science.aaa1160

Bateson, G. (1972). *Steps to an Ecology of Mind. Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology*. University of Chicago Press. Online: https://doi.org/10.7208/chicago/9780226924601.001.0001

Diakopoulos, N. (2015). Algorithmic Accountability: Journalistic Investigation of Computational Power Structures. *Digital Journalism, 3*(3), 398–415. Online: https://doi.org/10.1080/21670811.2014.976411

Dryzek, J. S. (2000). *Deliberative Democracy and Beyond. Liberals, Critics, Contestations.* Oxford University Press. Online: https://doi.org/10.1093/019925043x.001.0001

Dubois, E. & Blank, G. (2018). The Echo Chamber Is Overstated: The Moderating Effect of Political Interest and Diverse Media. *Information, Communication & Society, 21*(5), 729–745. Online: https://doi.org/10.1080/1369118X.2018.1428656

European Parliament & Council of the European Union (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union,* L, 2024/1689, 12.7.2024. Online: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

European Parliament & Council of the European Union (2022). Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a single market for digital services and amending Directive 2000/31/EC (Digital Services Act). *Official Journal of the European Union,* L 277, 1–102. Online: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2065

Floridi, L. (2011). *The Philosophy of Information.* Oxford University Press. Online: https://doi.org/10.1093/acprof:oso/9780199232383.001.0001

Gillespie, T. (2018). *Custodians of the Internet. Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. Yale University Press. Online: https://doi.org/10.12987/9780300235029

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., González-Bailón, S., Kennedy, E., Kim, Y. M., Lazer, D., Moehler, D., Nyhan, B., Rivera, C. V., Settle, J., Thomas, D. R., Thorson, E. A., Wang, X., Wojcieszak, M., Tucker, J. A. & Van Bavel, J. J. (2023). How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign? *Science, 381*(6656), 398–404. Online: https://doi.org/10.1126/science.abp9364

Habermas, J. (1996). *Between Facts and Norms. Contributions to a Discourse Theory of Law and Democracy* (W. Rehg, Trans.). MIT Press. Online: https://doi.org/10.7551/mitpress/1564.001.0001

Hosseinmardi, H., Ghasemian, A., Clauset, A., Mobius, M., Rothschild, D. M. & Watts, D. J. (2021). Examining the Consumption of Radical Content on YouTube. *Proceedings of the National Academy of Sciences, 118*(32). Online: https://doi.org/10.1073/pnas.2101967118

Landemore, H. (2013). *Democratic Reason. Politics, Collective Intelligence, and the Rule of the Many.* Princeton University Press. Online: https://doi.org/10.23943/princeton/9780691155654.001.0001

Levy, R. (2021). Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *American Economic Review, 111*(3), 831–870. Online: https://doi.org/10.1257/aer.20191777

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal, 27*(3), 379–423. Online: https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Statista (2024). Number of Daily Active Facebook Users Worldwide [Data set]. Online: https://www.statista.com/statistics/346167/facebook-global-dau/

Yadav, S. (2025). *The European Union A.I. Act 2024: Understanding the Context and Exploring its Future* (ORF Occasional Paper No. 486). Observer Research Foundation.

Yeung, K. (2018). Algorithmic Regulation: A Critical Interrogation. *Regulation & Governance, 12*(4), 505–523. Online: https://doi.org/10.1111/rego.12158

Zuboff, S. (2019). *The Age of Surveillance Capitalism. The Fight for a Human Future at the New Frontier of Power.* PublicAffairs.

Contents