EUROPEAN CONFERENCE

Conference Proceedings

XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

August 24–26, 2026
Prague, Czech Republic

PRIORITY AREAS OF SCIENTIFIC RESEARCH IN THE MODERN WORLD

Abstracts of XXXIV International Scientific and Practical Conference

Prague, Czech Republic
(August 24–26, 2026)

SOVEREIGN AI AND EPISTEMIC RESILIENCE: GOVERNING NATIONAL ALGORITHMIC INFRASTRUCTURES

Prokopenko Serhii Olexandrovych
Doctor of Philosophy, Senior Lecturer
Simon Kuznets Kharkiv National Economic University, Department of Marketing,

The accelerating diffusion of artificial intelligence systems into public administration, media ecosystems, security governance, and economic coordination has fundamentally reconfigured the epistemic foundations upon which contemporary societies build shared understanding. As algorithmic systems increasingly mediate access to information and actively structure the interpretive frames through which citizens, officials, and institutions make sense of the world, they simultaneously shape what is perceived as credible knowledge, legitimate expertise, and plausible political and economic futures. This paper problematizes the concept of "sovereign AI" through the analytical lens of epistemic resilience, arguing that the strategic pursuit of algorithmic autonomy, while often framed as a defensive response to geopolitical competition and technological dependence, carries within it the latent risk of institutionalizing national biases and producing new forms of epistemic closure.

Sovereign AI is typically presented in policy discourse as a straightforward strategic good: nationally controlled large language models, domestically owned compute infrastructure, and state-backed algorithmic ecosystems are positioned as necessary correctives to asymmetric technological dependence on a small number of foreign providers. The European debate on digital sovereignty, the ongoing technological rivalry between the United States and China, and the proliferation of national AI strategies among middle and smaller powers all reflect this trajectory [6],[9].

Yet this framing obscures a deeper paradox. Systems designed to reduce external dependence may simultaneously reproduce and amplify domestic epistemic biases, since machine learning models are trained on culturally situated corpora shaped by linguistic dominance, historical narrative conventions, and prevailing political discourse. The optimization functions and reinforcement signals used to align these systems encode particular patterns of relevance and credibility that risk becoming institutionalized as national frames of interpretation, effectively naturalizing a specific epistemic order under the guise of technical neutrality [1,3,4,5]

As AI becomes embedded in decision-support architectures across public governance, defense planning, media moderation, and economic policy, such embedded biases risk shaping the very boundaries of national epistemic reality.

The central analytical problem, therefore, is not merely misinformation understood as discrete false claims circulating within an otherwise neutral informational environment, but rather the structural conditioning of what counts as information in the first place. This reframing shifts the governance question away from narrow content

moderation toward the design of the underlying cognitive infrastructure itself, and it implies that oversight regimes calibrated solely to detect and remove false statements will systematically miss the subtler, and arguably more consequential, effects of algorithmic systems on the boundaries of the sayable and the thinkable within a given national context.

Drawing on structuration theory and sociotechnical systems scholarship [10,11] this paper conceptualizes AI systems as recursive infrastructures that both draw upon and reproduce the epistemic conditions of their training environments. Algorithmic outputs recursively reinforce dominant narratives through feedback loops, platform integration, and institutional uptake, such that early biases are not merely preserved but actively amplified across successive iterations of model training and deployment. From an institutional perspective, AI systems function as epistemic institutions in their own right: they codify norms of relevance, coherence, and credibility, and their outputs measurably influence sensemaking processes at both the organizational and societal levels, often without the interpretive labor being visible to those who rely on them.

The concept of digital sovereignty conventionally emphasizes autonomy in technological capacity, data governance, and infrastructure control. However, this paper argues that algorithmic nationalism emerges as a distinct and under-theorized risk when national AI architectures privilege domestic narratives, selectively filter external information, or align their outputs with strategic state interests. This produces a dual and interlocking risk structure. On one hand, external vulnerability persists as reliance on foreign AI systems continues to expose societies to epistemic influence and data extraction by dominant technology providers. On the other hand, internal homogenization emerges as sovereign systems narrow epistemic diversity domestically and institutionalize culturally specific biases as though they were universal or neutral outputs. Sovereign AI, in other words, does not resolve the epistemic vulnerability of dependence; it relocates that vulnerability inward.

To address this dual risk, the paper advances epistemic resilience as a systemic capacity to detect, counterbalance, and pluralize algorithmic bias across national AI architectures. Drawing on resilience theory as developed in organizational and socio-technical systems research [12], resilience is understood here not as a static property but as the capacity to anticipate, adapt, and transform under disruption. Applied to AI governance, epistemic resilience denotes the systemic ability to detect bias, maintain a plurality of interpretive frames, and sustain critical reflexivity within algorithmic environments. Rather than pursuing an unattainable ideal of algorithmic neutrality, epistemic resilience emphasizes adaptive pluralism and the deliberate construction of institutional safeguards that can surface and correct systemic distortions over time.

Methodologically, the paper adopts a qualitative multi-case research design examining national AI governance frameworks across selected jurisdictions, including the European Union, the United States, China, and emerging sovereign AI initiatives in smaller states. Data sources include national AI strategy documents, regulatory proposals and legislative texts, technical documentation of publicly known sovereign AI initiatives, policy speeches and official discourse, and, where available, secondary expert interviews. The analysis combines discourse analysis to identify epistemic

framing and sovereignty narratives, comparative institutional analysis of governance architectures across jurisdictions, and conceptual modeling to develop the proposed framework of epistemic resilience in AI systems. The objective of this design is theory-building rather than empirical generalization, and the resulting framework is intended to be applicable across diverse national contexts rather than tied to the particularities of any single case.

The discussion advances three governance implications that follow directly from the epistemic resilience framework. First, plural model architecture should be actively encouraged, favoring multi-model ecosystems over monolithic national models that risk becoming single points of epistemic failure. Second, cross-border epistemic interoperability should be maintained through structured exposure to external knowledge systems, ensuring that sovereign infrastructures do not evolve into closed epistemic loops insulated from correction by outside perspectives. Third, institutionalized reflexivity should be established through independent audit bodies and transparency mechanisms explicitly tasked with monitoring epistemic bias, rather than relegating this function to voluntary industry self-assessment.

If such mechanisms are absent, sovereign AI risks functioning as a soft totalizing system, one that operates not through overt coercion but through the subtle conditioning of cognitive horizons over time. This is arguably a more durable and less visible form of epistemic capture than earlier, more overt forms of state information control, precisely because it operates through infrastructure that is widely perceived as a neutral technical utility rather than a site of contestation over meaning and legitimacy.

The paper concludes that the central policy question is not whether sovereign AI should exist, given that the geopolitical and economic pressures driving its adoption are unlikely to abate, but rather how sovereign AI can be embedded within resilient epistemic ecosystems that actively preserve plurality, reflexivity, and democratic cognition.

This reframes AI governance debates, which have historically centered on technological competitiveness and regulatory compliance, as fundamentally epistemic resilience challenges requiring sustained institutional attention. Future research should extend this framework through empirical case studies of specific national AI architectures, particularly in smaller states whose sovereign AI initiatives remain comparatively under-examined in the existing literature, and through longitudinal analysis of how epistemic biases embedded in early model generations propagate or attenuate across subsequent iterations of national algorithmic infrastructure.

From a practical governance perspective, several concrete instruments could operationalize the epistemic resilience framework proposed here. National regulators could require sovereign AI developers to publish periodic epistemic diversity assessments alongside existing safety and bias audits, disclosing the range of interpretive frames represented in the training corpus and the degree to which model outputs converge on singular narrative positions across politically or culturally contested topics. International cooperation frameworks could establish shared benchmarks for cross-border epistemic interoperability, analogous to existing interoperability standards in data protection and cybersecurity, allowing independent

researchers to evaluate whether sovereign systems maintain meaningful openness to external knowledge production rather than operating as closed epistemic loops. Academic and civil society institutions could be granted structured access to sovereign AI systems for the specific purpose of independent epistemic auditing, mirroring the role that independent media and academic institutions have historically played in checking the epistemic authority of state institutions more broadly. None of these instruments alone would fully resolve the tension between sovereignty and pluralism identified in this paper, but taken together, they would begin to institutionalize the reflexive correction mechanisms that epistemic resilience requires, shifting sovereign AI governance from a narrow focus on technical safety toward a broader concern with the health of national and transnational epistemic ecosystems as such.

Ultimately, the stakes of this governance challenge extend well beyond the technology sector narrowly conceived. As algorithmic systems become further embedded in the everyday infrastructure of public reasoning, from search and recommendation systems to administrative decision support and educational tools, the cumulative effect of unaddressed epistemic bias in sovereign AI architectures compounds over time in ways that are difficult to reverse once entrenched. Treating epistemic resilience as a first-order governance objective, rather than a secondary consideration subordinate to strategic autonomy or economic competitiveness, offers the most promising path toward sovereign AI systems that strengthen rather than erode the pluralistic foundations of democratic epistemic life.

The theoretical foundations of this argument draw on five interconnected streams of scholarship, each of which addresses a component of the problem without, on its own, integrating sovereign AI and epistemic resilience into a unified governance challenge. The first stream, AI bias and fairness research, demonstrates how training data, model architecture, and deployment context jointly produce measurable demographic and cultural bias in algorithmic outputs. This literature has been instrumental in establishing that bias is not an incidental technical defect to be patched after the fact, but a structural effect of how socially embedded models are constructed, trained, and validated. The second stream, digital sovereignty and the geopolitics of technology, examines strategic autonomy in AI development and control, situating national AI policy within broader debates over weaponized interdependence and the uneven distribution of infrastructural power in the global technology stack. The third stream, institutional theory and structuration, explains how technological systems recursively shape and are shaped by social structures, offering the conceptual vocabulary needed to understand AI not as an exogenous tool imposed upon institutions but as a co-constitutive element of institutional life itself [10,11]. The fourth stream, resilience and systems theory, conceptualizes adaptive capacity under uncertainty and systemic shock, providing the analytical bridge that allows this paper to move the discussion from a static diagnosis of "AI bias" toward a dynamic framework of "epistemic resilience" capable of guiding ongoing governance practice [12]. The fifth stream, epistemic governance and disinformation studies, highlights how cognitive infrastructures shape public reasoning and democratic legitimacy, connecting the technical question of algorithmic design to normative questions about

epistemic autonomy and human rights in an information environment increasingly mediated by machine intermediaries [13].

Taken together, these five streams suggest that the governance of sovereign AI cannot be reduced to a purely technical exercise in bias mitigation, nor to a purely geopolitical exercise in strategic autonomy. Instead, it requires an integrated framework that treats algorithmic infrastructure as a constitutive element of national epistemic order, subject to the same requirements of pluralism, accountability, and reflexive correction that democratic societies have historically demanded of other epistemic institutions, including the press, the academy, and independent judiciaries. The comparative case material examined in this study illustrates how different jurisdictions have approached this challenge in markedly different ways.

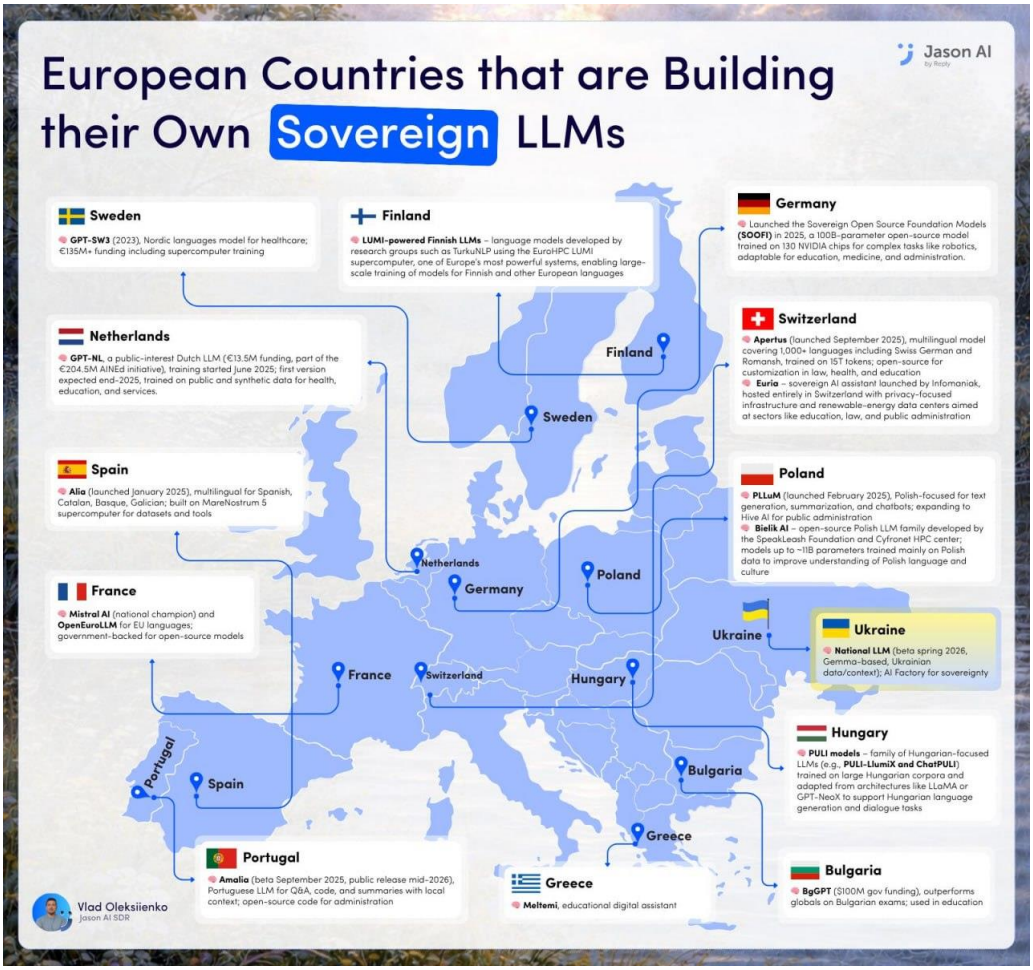

Picture source: Oleksiienko V. While everyone debates AI ethics, Europe invested millions in sovereign LLMs. LinkedIn, 2026

The European Union has pursued a predominantly regulatory strategy, embedding risk-based obligations for high-impact AI systems within a broader framework of fundamental rights protection, though the epistemic pluralism implications of this approach remain underdeveloped relative to its data protection and safety provisions. The United States has favored a more market-driven approach, relying on competitive pressure among a plurality of large private developers to generate a degree of de facto epistemic diversity, albeit without any explicit institutional mandate for pluralism as

such. China's national AI development strategy illustrates the opposite pole of the spectrum, in which state-directed alignment between algorithmic outputs and official narratives is treated as an explicit policy objective rather than a risk to be mitigated, offering the clearest illustration of the internal homogenization risk this paper identifies.

Emerging sovereign AI initiatives among smaller and middle powers, meanwhile, face a distinctive constraint: limited compute and data resources often force reliance on foundation models developed elsewhere, even as national policy discourse frames these initiatives in the language of digital sovereignty, producing a hybrid condition of partial autonomy that has received comparatively little scholarly attention to date.

These comparative observations reinforce the paper's central claim that sovereign AI should not be evaluated solely in terms of strategic autonomy or economic competitiveness, since doing so risks overlooking the ways in which algorithmic infrastructure reshapes collective cognition regardless of the underlying strategic rationale for its adoption. A resilience-oriented governance agenda would treat the three implications discussed above, plural model architecture, cross-border epistemic interoperability, and institutionalized reflexivity, not as optional best practices but as necessary structural safeguards against the soft totalizing tendency inherent in any sufficiently centralized epistemic infrastructure, however benign its stated policy intentions.

The paper acknowledges several limitations that suggest productive directions for future work. As a theory-building contribution based on qualitative multi-case analysis of publicly available policy documents, the framework proposed here has not yet been tested against systematic empirical measurement of epistemic diversity within deployed sovereign AI systems, a methodological challenge that will require the development of new indicators capable of capturing pluralism at the level of interpretive framing rather than merely at the level of surface-level content diversity.

Further, the comparative case selection, while illustrative of a meaningful range of governance postures, cannot capture the full diversity of national approaches emerging across the Global South, where sovereign AI initiatives are often shaped simultaneously by resource constraints, geopolitical alignment choices, and distinct cultural conceptions of epistemic authority that merit dedicated future investigation.

References

1. FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. Pages 610 - 623.https://doi.org/10.1145/3442188.3445922
2. Zammit, F. (2021). Race After Technology Abolitionist Tools for the New Jim Code. Share Magazine - Issue 16.
3. Benjamin, R. (2019). Race after technology: Abolitionist tools for the new Jim code. Polity Press.
4. Noble, S. U. (2018). Algorithms of oppression: How search engines reinforce racism. New York University Press.

5. Floridi, L. (2019). Establishing the rules for building trustworthy AI. Nature Machine Intelligence, 1(6), 1–2. https://doi.org/10.1038/s42256-019-0055-y

6. Bradford, Anu, The Brussels Effect: How the European Union Rules the World (New York, 2020; online edn, Oxford Academic, 19 Dec. 2019), https://doi.org/10.1093/oso/9780190088583.001.0001, accessed 16 Aug. 2026.

7. Gatrell, C., Muzio, D., Post, C. and Wickert, C. (2024), Here, There and Everywhere: On the Responsible Use of Artificial Intelligence (AI) in Management Research and the Peer-Review Process. J. Manage. Stud., 61: 739-751. https://doi.org/10.1111/joms.13045

8. Anthropic. Claude (Sonnet 5 version) [Large language model], 2026. https://claude.ai

9. Farrell, H., & Newman, A. L. (2019). Weaponized Interdependence: How Global Economic Networks Shape State Coercion. International Security, 44(1), 42–79. https://www.jstor.org/stable/26777882

10. Author, A. A. (2026). Structuration theory: The constitution of society – by Anthony Giddens (1984). In R. Spiller, C. Rudeloff, & T. Döbler (Eds.), Key works: Theories in communication studies. Springer Fachmedien Wiesbaden. https://doi.org/10.1007/978-3-658-50422-9_23

11. Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. Organization Science, 3(3), 398–427. https://doi.org/10.1287/orsc.3.3.398

12. Walker, B., C. S. Holling, S. R. Carpenter, and A. Kinzig. 2004. Resilience, adaptability and transformability in social–ecological systems. Ecology and Society 9(2): 5. [online] URL: http://www.ecologyandsociety.org/vol9/iss2/art5/

13. Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.

Scientific publications

MATERIALS

The XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

Prague, Czech Republic
(August 24–26, 2026)