

12 Law in the Age of Synthetic Realities: Rethinking Legal Norms for Generative AI¹

Bart van der Sloot²

This chapter critically examines the evolving regulatory landscape for synthetic technologies with particular attention to foundational legal principles – truth, justice, personhood, privacy, cognitive autonomy, and collective rights. It offers a deep dive into the structural and doctrinal challenges that GAI poses to European legal regimes and analyzes key legislative instruments, interprets relevant jurisprudence, and synthesizes emerging legal theory to assess the extent to which current regulatory frameworks are equipped to address the ontological instability of synthetic realities.

12.1 Introduction

The proliferation of Generative Artificial Intelligence (GAI) technologies is radically reshaping the contours of legal, ethical, and societal order. In the transition toward the "synthetic society," law finds itself confronting an unprecedented ontological and normative crisis: how to legislate reality when that reality is increasingly simulated, augmented, or generated. From humanoid robots to deepfakes, from virtual environments to algorithmically tailored avatars, GAI blurs the boundary between the authentic and the artificial, not just in form but in consequence. The legal implications of these blurred boundaries are manifold, calling into question traditional legal concepts of personhood, accountability, truth, and autonomy.

Section 2 interrogates the epistemic destabilization [de Groot & De Souza 2025] caused by synthetic content and its effects on legal evidence and public discourse. Section 3 explores the emergence of synthetic agents and the challenge they pose to traditional doctrines of liability and legal

¹ This is the English version of this chapter, the German version can be found in the book „Künstliche Intelligenz für Business Analytics“ by Haneke et al. (ed.) under <https://dpunkt.de/produkt/kuenstliche-intelligenz-fuer-business-analytics/>.

² This chapter is based on: Van der Sloot, B. (2024). *Regulating the synthetic society: Generative AI, legal questions, and societal challenges* (p. 296). Bloomsbury Academic. OpenAI's Deep Research was used to craft a chapter on the basis of that book. The German version of this chapter can be found in

personhood. Shifting focus from individuals to collectives. Section 4 examines the structural limitations of rights-based regimes in dealing with systemic harms. Section 5 assesses the crisis of legal imagination that has resulted in fragmented and reactive policies. Finally, the concluding section advances a forward-looking framework – the *Lex Syntheticum* – as a prototype for regulatory synthesis in the age of AI.

12.2 Reality Disrupted: The Epistemic Crisis of Deepfakes and Synthetic Content

The legal tradition has always rested on a stable ontology – one in which facts can be ascertained, truth can be adjudicated, and human testimony carries presumptive value. This architecture is now in peril. The advent of generative synthetic content – deepfakes, AI-generated audio, hyperreal avatars – has introduced a new layer of epistemic uncertainty into legal systems. The problem is not merely technological, but normative: how can law continue to function as an arbiter of truth when its epistemic foundations are increasingly destabilized by synthetic simulations?

At the core of any legal proceeding lies the question of evidence. Deepfakes, particularly those generated using adversarial neural networks, [Wang et al. 2017] now possess a degree of photorealistic precision that renders them nearly indistinguishable from authentic audiovisual material. This challenges the evidentiary reliability of videos, voice recordings, and even biometric data, which have long served as critical forms of proof in both civil and criminal litigation.

The traditional doctrine governing evidence in European legal systems is anchored in principles of reliability, probative value, and procedural fairness. Article 6(1) of the European Convention on Human Rights (ECHR) holds [EMRK 2021]:

‘1. In the determination of his civil rights and obligations or of any criminal charge against him, everyone is entitled to a fair and public hearing within a reasonable time by an independent and impartial tribunal established by law. Judgment shall be pronounced publicly but the press and public may be excluded from all or part of the trial in the interests of morals, public order or national security in a democratic society, where the interests of juveniles or the protection of the private life of the parties so require, or to the extent strictly necessary in the opinion of the court in special circumstances where publicity would prejudice the interests of justice. 2. Everyone charged with a criminal offence shall be presumed innocent until proved guilty according to law.’

That provision provides that individuals are entitled to a fair trial, which includes the right to examine the evidence brought against them. But when the authenticity of that evidence becomes probabilistic—subject to forensic estimations rather than categorical verification—the burden of proof becomes conceptually unstable. This is especially critical in criminal law. Consider a hypothetical in which a video emerges showing a defendant committing an assault. The defense argues it is a deepfake. The prosecution provides a forensic AI tool indicating the video is likely real with a 74% confidence score. Can a conviction follow from probabilistic authentication? This scenario risks violating the presumption of innocence under Article 6(2) ECHR.

Emerging solutions include digital watermarking, cryptographic verification chains (e.g., Adobe's Content Authenticity Initiative), and blockchain-based provenance tools. However, these remain voluntary and are not retroactive. The proposed EU AI Act mandates transparency in some use-cases (Article 52), but stops short of establishing evidentiary thresholds for AI-generated content in litigation [EU AI Act 2024]. It provides:

- Providers shall ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect, taking into account the circumstances and the context of use. This obligation shall not apply to AI systems authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to appropriate safeguards for the rights and freedoms of third parties, unless those systems are available for the public to report a criminal offence.
- Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, shall ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated. Providers shall ensure their technical solutions are effective, interoperable, robust and reliable as far as this is technically feasible, taking into account the specificities and limitations of various types of content, the costs of implementation and the generally acknowledged state of the art, as may be reflected in relevant technical standards. This obligation shall not apply to the extent the AI systems perform an assistive function for standard editing or do not substantially alter the input data provided by the deployer or the semantics thereof, or where authorised by law to detect, prevent, investigate or prosecute criminal offences.

- Deployers of an emotion recognition system or a biometric categorisation system shall inform the natural persons exposed thereto of the operation of the system, and shall process the personal data in accordance with Regulations (EU) 2016/679 and (EU) 2018/1725 and Directive (EU) 2016/680, as applicable. This obligation shall not apply to AI systems used for biometric categorisation and emotion recognition, which are permitted by law to detect, prevent or investigate criminal offences, subject to appropriate safeguards for the rights and freedoms of third parties, and in accordance with Union law.
- Deployers of an AI system that generates or manipulates image, audio or video content constituting a deep fake, shall disclose that the content has been artificially generated or manipulated. This obligation shall not apply where the use is authorised by law to detect, prevent, investigate or prosecute criminal offence. Where the content forms part of an evidently artistic, creative, satirical, fictional or analogous work or program, the transparency obligations set out in this paragraph are limited to disclosure of the existence of such generated or manipulated content in an appropriate manner that does not hamper the display or enjoyment of the work.
- Deployers of an AI system that generates or manipulates text which is published with the purpose of informing the public on matters of public interest should consequently disclose that the text has been artificially generated or manipulated. However, this obligation does not apply where the use is authorised by law to detect, prevent, investigate or prosecute criminal offences or where the AI-generated content has undergone a process of human review or editorial control and where a natural or legal person holds editorial responsibility for the publication of the content. These exemptions are so wide that it renders the provisions toothless.

Beyond the courtroom, the societal impact of synthetic content is no less grave. As deepfakes spread, the European Court of Human Rights may soon be asked to adjudicate claims where the privacy harm lies not in the disclosure of truth, but the propagation of fiction—raising the question: is a synthetic lie entitled to the same protections as a genuine one? In democracies reliant on informed citizenries, the systemic spread of false, AI-generated content constitutes an attack on the public epistemic commons. The law's role, therefore, must extend beyond reactive adjudication toward proactive protection of democratic infrastructure.

The Digital Services Act (DSA) [DSA 2022], particularly Articles 34–36, mandates that Very Large Online Platforms (VLOPs) conduct risk assessments and implement mitigation strategies for systemic threats, including disinformation. These provisions provide:

- Providers of very large online platforms shall carry out the risk assessments, taking into consideration their severity and probability, and shall include the following systemic risks: (a) the dissemination of illegal content through their services; (b) any actual or foreseeable negative effects for the exercise of fundamental rights; (c) any actual or foreseeable negative effects on civic discourse and electoral processes, and public security; (d) any actual or foreseeable negative effects in relation to gender-based violence, the protection of public health and minors and serious negative consequences to the person's physical and mental well-being.
- Providers of very large online platforms and of very large online search engines shall put in place reasonable, proportionate and effective mitigation measures, tailored to the specific systemic risks identified. Such measures may include, where applicable: (a) adapting the design, features or functioning of their services, including their online interfaces; (b) adapting their terms and conditions and their enforcement; (c) adapting content moderation processes; (d) testing and adapting their algorithmic systems, including their recommender systems; (e) adapting their advertising systems and adopting targeted measures aimed at limiting or adjusting the presentation of advertisements in association with the service they provide; (f) reinforcing the internal processes, resources, testing, documentation, or supervision of any of their activities in particular as regards detection of systemic risk; (g) initiating or adjusting cooperation with trusted flaggers in accordance; (h) initiating or adjusting cooperation with other providers of online platforms or of online search engines through the codes of conduct and the crisis protocols; (i) taking awareness-raising measures and adapting their online interface in order to give recipients of the service more information; (j) taking targeted measures to protect the rights of the child; (k) ensuring that an item of information, whether it constitutes a generated or manipulated image, audio or video that appreciably resembles existing persons, objects, places or other entities or events and falsely appears to a person to be authentic or truthful is distinguishable through prominent markings when presented on their online interfaces,

and, in addition, providing an easy to use functionality which enables recipients of the service to indicate such information.

Although these provisions are a good start, the DSA does not require content authenticity disclosures, nor does it provide effective redress mechanisms for individuals targeted by synthetic media. One legislative path forward lies in the recognition of an “epistemic right to truth”—not in the absolute sense, but as a right to informational integrity in digital environments.⁵ In an era increasingly dominated by artificial intelligence and synthetic media, the “right to truth” emerges as a foundational principle for preserving individual autonomy, democratic discourse, and the integrity of shared knowledge. This right would entail a guarantee that citizens have access to unmanipulated, verifiable representations of the world – whether in journalism, political communication, or legal evidence. A right to reality would not prohibit the use of AI-generated material per se, but would demand unequivocal transparency about its origin, rigorous standards for labeling and provenance, and safeguards against deceptive practices. Such a right affirms that in order to make informed decisions, engage in civic life, and trust institutions, people must be able to anchor their understanding in an accessible and accountable version of the real.

12.3 Synthetic Subjects : The Regulation of Non-Human Agents

The emergence of generative AI systems with increasing degrees of autonomy and anthropomorphic design poses a fundamental challenge to the legal architecture of responsibility and agency. As synthetic entities begin to perform tasks once considered the sole domain of humans— offering companionship, administering medical care, executing financial trades— legal scholars and policymakers must grapple with a central question: How should law treat non-human agents that mimic or mediate human action?

The concept of legal personhood has long served as a doctrinal gatekeeper in attributing rights and duties. Traditionally, it has been reserved for natural persons and juristic entities such as corporations. However, with the development of humanoid robots, synthetic avatars, and autonomous agents capable of sophisticated interaction, there is increasing pressure to revisit the contours of this legal category. In 2015, the European Parliament’s Committee on Legal Affairs proposed the notion of granting “electronic personhood” to the most advanced autonomous robots. It suggested to explore ‘creating a specific legal status for robots in the long run, so that

⁵ Anuj Puri has coined to notion of “Right to Reality” [Puri 2024].

at least the most sophisticated autonomous robots could be established as having the status of electronic persons responsible for making good any damage they may cause, and possibly applying electronic personality to cases where robots make autonomous decisions or otherwise interact with third parties independently' [Zivilrechtliche Regelungen Robotik 2017].

Although this suggestion has not been honored, the pressure remains. From a functionalist perspective, the law could adopt a tiered model of personhood based on the degree of autonomy, social presence, and decision-making capacity. Just as corporate entities enjoy partial personhood for contractual and property purposes, synthetic agents might acquire "functional personhood" in contexts requiring operational agency. Yet this approach risks obfuscating human responsibility. Framing AI systems as independent agents may shift focus away from the designers, deployers, and profiteers who control them. Legal systems must therefore remain vigilant against the anthropomorphic fallacy—the tendency to attribute moral agency to systems that lack consciousness, intentionality, or the capacity for normative reasoning.

Rather than personhood, a more fruitful legal lens may be that of agency law. In this doctrine, agents act on behalf of principals and can bind them in contractual or tortious relations. Synthetic agents, then, could be conceived as proxies through which humans act. Under this framework, the liability for the agent's action – whether intentional or negligent – falls upon the principal, unless the agent clearly acted beyond their authority. But what if the agent is an algorithmic system whose decision-making logic is opaque, even to its creators? In such cases, traditional attribution models fail. The "responsibility gap" [Matthias 2004] rises precisely because no individual can be said to have intended or foreseen the specific harm.

To close this gap, regulators may need to adopt presumptive liability schemes. One proposal is strict liability for deployers of autonomous systems engaged in high-risk activities (e.g., autonomous vehicles, healthcare robots, lethal autonomous weapons). Another is a layered responsibility model: developers liable for design flaws, operators for deployment choices, and users for misuse. The European Commission's proposed AI Liability Directive [Europäische Kommission 2022] moves in this direction by introducing a rebuttable presumption of causality when claimants can show that an AI system likely caused harm and the defendant failed to comply with a relevant duty of care or disclosure obligation. While promising, this proposal remains limited in scope and lacks binding harmonization. It provides that national courts shall presume, for the purposes of applying liability rules to a claim for damages, the causal link between the fault of the defendant and the output produced by the AI system or the

failure of the AI system to produce an output, where all of the following conditions are met:

- The claimant has demonstrated or the court has presumed the fault of the defendant, or of a person for whose behaviour the defendant is responsible, consisting in the non-compliance with a duty of care directly intended to protect against the damage that occurred.
- It can be considered reasonably likely, based on the circumstances of the case, that the fault has influenced the output produced by the AI system or the failure of the AI system to produce an output.
- The claimant has demonstrated that the output produced by the AI system or the failure of the AI system to produce an output gave rise to the damage.

Internationally, there is little consensus on how to regulate synthetic agents. The lack of harmonization poses challenges in cross-border AI deployment. A humanoid robot with limited legal capacity in one jurisdiction may be treated as a mere tool in another or as a co-worker in a third country. This legal pluralism creates regulatory arbitrage opportunities and underscores the need for coordinated international standards. Beyond utility, the debate also engages with moral philosophy. Should legal status reflect societal perceptions of moral agency? If users experience companionship, grief, or attachment to synthetic entities, do such entities merit legal protection? Scholars like David Gunkel consider that relational ethics may justify extending certain legal protections not because AI deserves them intrinsically, but because denying them undermines human moral integrity [Gunkel 2018]. This view resonates with animal welfare law, which recognizes sentience and vulnerability as bases for limited legal protection. A similar logic might support limited duties toward AI entities that evoke emotional bonding or are designed to appear conscious. At the same time, emotional realism in synthetic agents should not be confused with genuine consciousness. The law must resist aesthetic seduction and remain anchored in legal realism.

12.4 From Individual Rights to Collective Harms

The transition from analog to synthetic societies involves a paradigm shift in how rights, harms, and protections are conceptualized. While traditional legal frameworks emphasize the primacy of individual rights – especially in the context of privacy and data protection—generative AI introduces harms that are diffuse, systemic, and often collective in nature. This section interrogates the limitations of existing rights-based legal architectures and proposes pathways for addressing collective harms in the synthetic age. European data protection law, most notably the General Data Protection

Regulation (GDPR), remains the cornerstone of informational rights within the EU. It grants individuals rights to access, rectify, erase, and object to the processing of their personal data. However, synthetic systems increasingly operate through inference, aggregation, and behavioral prediction – modes of computation that evade traditional notions of “personal data.”

Personal data, for the purposes of the GDPR, is defined as:

- Any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.
- The principles of data protection should apply to any information concerning an identified or identifiable natural person. Personal data which have undergone pseudonymisation, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person. To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments. The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. This Regulation does not therefore concern the processing of such anonymous information, including for statistical or research purposes. This Regulation does not apply to the personal data of deceased persons. Member States may provide for rules regarding the processing of personal data of deceased persons.

For instance, generative recommender systems create psychometric profiles, not by collecting declared preferences, but by triangulating subtle behavioral cues: click delays, scroll velocity, voice modulation. These inferences are typically not disclosed and fall outside GDPR’s effective reach.

It also excludes much of the affective nudging and epistemic framing done by AI systems. In addition, there are many limitations and exceptions to Article 22, which offers protections against automated decision-making, but only if the decision produces “legal or similarly significant effects.” It provides:

- The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.
- This shall not apply if the decision: (a) is necessary for entering into, or performance of, a contract between the data subject and a data controller; (b) is authorised by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or (c) is based on the data subject's explicit consent.
- In the cases referred to in points (a) and (c), the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.
- If automated decision-making is based on sensitive data (personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation), this is only allowed if based on the informed consent of persons involved or if in a substantial public interest.

Further, the GDPR assumes a dyadic model of data harm: a subject and a data controller. Yet deepfakes, hallucinated news articles, and algorithmically generated lies affect third parties, communities, and even democratic institutions. The law must therefore develop tools to respond to “epistemic harms”— damage done not to a single data subject, but to the shared informational substrate on which society relies. In this context, scholars have proposed concepts like “data justice” and “cognitive liberty” as alternative frameworks [Taylor 2017]. Data justice refers to the equitable distribution of epistemic power: who controls knowledge production, access, and verification [Taylor 2023]? Cognitive liberty emphasizes the right not only to privacy, but to autonomy in thought formation, unmanipulated by immersive or deceptive AI stimuli [Farahany 2023].

Synthetic technologies also give rise to structural harms that affect groups rather than individuals [Taylor et al. 2017]. These include disinformation campaigns that undermine democratic processes, algorithmic profiling that reproduces systemic bias and generative content that erodes collective memory or distorts historical record. The law has historically struggled to articulate responses to such harms. Tort law, for example, requires identifiable plaintiffs and demonstrable injury. Constitutional law may invoke public interest doctrines but lacks procedural pathways for redress. Even the European Court of Human Rights (ECtHR), though expansive in its interpretive approach, remains tethered to individual claims. A constitutionalization of the information environment may be necessary—treating shared epistemic infrastructures as legally protected commons. This might resemble environmental law, where collective harms (e.g., air pollution) are governed through public regulation rather than private litigation. Legal innovation is already visible in the Digital Services Act, which imposes systemic risk obligations on very large online platforms, requiring them to assess and mitigate harms related to disinformation, algorithmic amplification, and manipulation. While promising, the DSA lacks enforceable rights for affected communities and does not recognize epistemic integrity as a standalone legal good.

Cognitive sovereignty, meanwhile, addresses the right of individuals and communities to shape their mental environments. In immersive media contexts—augmented reality, brain-computer interfaces, algorithmically generated dreams—this may require legal recognition of non-invasive spaces for thought formation. Legal scholars have called for neuro-rights legislation, especially as AI converges with neurotechnology. The path forward involves not simply extending privacy rights, but reconceptualizing informational rights as ecological and participatory [Farahany et al. 2023]. The synthetic society challenges law to move from a transactional model of harm to a systemic one—capable of recognizing, preventing, and remedying the collective vulnerabilities that AI technologies expose and exploit.

12.5 The Legal Imagination in Crisis: Fragmentation, Soft Law, and Structural Limits

The emergence of synthetic technologies has outpaced the law's capacity to respond coherently. Regulatory instruments are proliferating, yet they remain fragmented across jurisdictions, disciplines, and institutional mandates. Simultaneously, a growing reliance on soft law—ethical guidelines, voluntary codes of conduct, impact assessments—threatens to substitute technocratic consensus for democratic accountability. This section evaluates the limitations of current legal infrastructures and argue for a more robust, principled approach to AI governance.

Across the EU, at least three major legal regimes attempt to regulate facets of the synthetic society: the General Data Protection Regulation, the Artificial Intelligence Act, and the Digital Services Act. Each addresses a specific domain—data protection, AI system classification, and online platform governance—but none provides a holistic framework for addressing the systemic risks posed by generative AI. The GDPR focuses on the protection of personal data, but it does not regulate synthetic content *per se*. The AI Act proposes a tiered risk-based framework but avoids adequately defining key terms such as “autonomy” or “general-purpose AI.” The DSA imposes due diligence obligations on online platforms but does not delineate specific enforcement tools for dealing with harmful synthetic media or epistemic harms.

This fragmentation leads to jurisdictional overlap, regulatory arbitrage, and underenforcement. For instance, a generative AI system that produces deepfake pornography using public figures could fall under the AI Act (if considered high-risk), the GDPR (if personal data is involved), and the DSA (if hosted by a platform)—yet no single body is tasked with resolving such a multifaceted case. Scholars have called for “legal interoperability” in AI governance—a coordinated legal infrastructure that harmonizes definitions, principles, and enforcement tools [de Oliveira et al. 2024]. Without such integration, AI regulation risks becoming a patchwork of siloed interventions that fail to capture systemic effects.

As regulatory complexity increases, policymakers and corporations have increasingly turned to soft law instruments: principles for trustworthy AI, ethics boards, algorithmic impact assessments, and transparency pledges. While these mechanisms may promote awareness and corporate social responsibility, they often function as substitutes for binding rules. This phenomenon—what some scholars term the “ethics-washing” of AI governance—can undermine both public trust and legal clarity [van Maanen 2022]. The European Commission’s High-Level Expert Group on AI published a widely cited document in 2018 outlining seven key requirements for trustworthy AI: human agency, technical robustness, privacy, transparency, diversity, societal well-being, and accountability [HLEG 2018]. While rhetorically and substantively compelling, these principles lack legal enforceability. Moreover, they can invite selective interpretation and provide cover for firms that comply in form but not in substance.

Ethical soft law also struggles with democratic legitimacy [Senden 2004]. Most AI ethics frameworks are drafted by corporate actors, technical experts, or closed-door advisory bodies, often without meaningful input from civil society. This can result in a technocratic approach to regulation that sidelines fundamental rights in favor of performance metrics and innovation incentives. Ultimately, the legal imagination must rise to the

challenges of the synthetic society with normative ambition, not regulatory minimalism. Law is not merely a constraint on innovation but a constitutive force in shaping the kinds of futures we inhabit. The regulation of AI must therefore be grounded not only in efficiency or safety but in foundational democratic values: truth, dignity, justice, and collective self-determination. Rather than treating AI as a sector-specific concern, legal systems should approach it as an infrastructural transformation on par with electricity or the printing press—an evolution that affects every domain of life and requires an equally integrated legal response.

12.6 Charting the Path Forward: Lex Syntheticum as Regulatory Prototype

To overcome the epistemic, normative, and jurisdictional fragmentation that defines contemporary legal responses to synthetic realities, a *Lex Syntheticum* should be introduced, which is based on a coherent legal grammar that reflects the ontological and ethical specificities of the synthetic age. This section outlines the key pillars of such a framework—doctrinally, institutionally, and procedurally. The *Lex Syntheticum* is not a new code, but a meta-framework designed to harmonize and scaffold existing and emergent regulatory instruments. It is grounded in four normative pillars:

- **Epistemic Integrity:** Legal mechanisms must safeguard the reliability of information and the verifiability of truth. This includes authentication standards for digital content, forensic AI protocols for courts, and rights to algorithmic transparency and contestation.
- **Cognitive Autonomy:** Individuals must retain sovereignty over their own mental and emotional states. This pillar underwrites emerging neuro-rights, bans on subliminal AI manipulation, and oversight of immersive environments that exploit behavioral vulnerabilities.
- **Social Infrastructure:** Law must support institutions that preserve pluralism, discourse, and resilience in synthetic ecosystems – fact-checking bodies, AI ombuds offices, public AI literacy campaigns, and platform co-governance mechanisms.
- **Human Dignity:** All AI regulation must affirm the primacy of human rights over commercial, technological, or military priorities. This includes red lines around biometric surveillance, manipulative AI in politics, and systems that dehumanize or instrumentalize users.

The implementation of the *Lex Syntheticum* requires institutional innovations. At the national level, governments should establish independent AI oversight agencies with investigatory and enforcement powers – akin to data protection authorities under the GDPR. These agencies would audit

AI systems, investigate complaints, and issue binding decisions. At the transnational level, an European Epistemic Integrity Agency could be created to coordinate responses to cross-border synthetic threats – e.g., coordinated disinformation campaigns, synthetic election interference, or cross-platform generative content harms. This body would collaborate with ENISA, the European Data Protection Board, and other relevant stakeholders.

Drawing inspiration from environmental governance, a public AI registry could serve as a transparency mechanism, cataloguing high-risk and general-purpose AI systems along with audit logs, redress pathways, and contact points. Participation in this registry could be made mandatory for systems meeting criteria defined under the AI Act or Lex Syntheticum.

The *Lex Syntheticum* also calls for new procedural mechanisms:

- **Algorithmic Due Process:** Individuals must be informed when synthetic content significantly affects their rights or environments, and must be granted a meaningful opportunity to challenge or correct these outputs [Citron 2008; Fortes 2020]. This expands on Article 22 GDPR and aligns with ECtHR jurisprudence on effective remedies.
- **Epistemic Habeas Corpus:** A procedural right to challenge the integrity of key informational content (e.g., evidence, media, archives) when their authenticity is questioned due to AI interference [Vergari & Preite 2023].
- **Civic Algorithmic Standing:** The recognition of civil society organizations, academic institutions, and watchdog groups as legitimate claimants in cases involving collective or systemic AI harms – especially where no single individual can prove direct harm.
- **Sunset Clauses for Synthetic Systems:** All high-risk synthetic technologies should be subject to periodic review, sunset clauses, or reauthorization mechanisms to ensure continued legal, ethical, and social alignment [Kouroutakis 2016].

In an era increasingly defined by synthetic realities – where generative algorithms fabricate people, manipulate evidence, shape discourse, and recalibrate human perception – the law is being called upon to perform a new and historically unfamiliar function: to serve as the guardian of reality itself. This role transcends the traditional remit of law as dispute resolution or coercive constraint. It demands that legal systems defend the epistemic foundations upon which democratic societies are built: the capacity to distinguish true from false, authentic from fabricated, agent from artefact. In doing so, the law reasserts its relevance not only as a normative order, but as an ontological anchor in a world of accelerating simulation.

What is at stake is not merely regulation of technology, but the preservation of political community. Without trust in evidence, the rule of law cannot be sustained. Without shared narratives, democratic deliberation collapses. Without legal accountability, generative systems may become vectors of impunity, instrumentalization, and psychological manipulation. As we have argued throughout this chapter, the current legal architecture is ill-equipped to manage these challenges in its present form. Rights-based frameworks struggle with collective harms. Traditional doctrines of liability falter in the face of synthetic agency. Soft law proliferates without binding power. Jurisdictional silos paralyze transnational enforcement.

The response, however, need not be despair. The *Lex Syntheticum* offers one possible blueprint – a harmonized, adaptive, principled framework that places human dignity, epistemic integrity, and democratic sovereignty at the center of AI governance. Implementing such a framework will require doctrinal creativity, institutional courage, and transdisciplinary collaboration. Ultimately, the law’s legitimacy in the synthetic age will depend on its ability to protect not just rights, but reality; not just individuals, but societies; not just order, but meaning. The future of legal regulation lies in its ability to respond to a world in which the line between the real and the artificial is not merely blurred, but actively contested. In defending that line, law defends us all.