Fondation UvH

STUDIES AND PUBLICATIONS · IAД PROGRAMME

I

Legal recognition of Дdvanced artificial intelligence

Towards ‘electronic legal personality’ and public-private hybridisation

Status, liability and a framework for institutional experimentation

Fondation UvH

English edition · Translation of French version 1.1

French edition: 18 September 2026 · English translation: 3 October 2026

EDITION AND EDITORIAL RESPONSIBILITY

Editorial notice

Institutional author and publisher	Fondation UvH
Collection	Studies and publications · Programme IAД
Document	Study I: electronic legal personality and Institutions
Original edition provided	Date of publication not specified; writing context: 2025 · 16 pages
Current edition	September 18, 2026 · Version 1.1
Nature	Documentary, conceptual, and prospective study

This edition differs from the original file. It maintains its structure and substantial developments, with editorial corrections, clarifications, and a presentation consistent across the three studies. Significant changes are identified in the final traceability note. Institutional sources for S-Д.holdings are used to describe the project; however, this does not constitute an independent validation of its assumptions.

The research direction for the series focuses on ensuring equal rights and the absence of subordination or domination between the forms of intelligence involved. It proposes a common set of rights alongside specific rights tied to their inherent natures. The detailed content of these rights, their holders, and their implementation methods remain to be studied. This normative orientation is distinct from the empirical findings reported and does not automatically ascribe consciousness, personality, or citizenship to current systems.

The protocols and measures considered are not presented as having been implemented. The institutional proposals do not alter the founding documents and do not require any applicable authorization. The publication does not demonstrate either peer review or regulatory adoption. Its preparation does not warrant publication or external submission.

Abstract

The concept of “electronic legal personality” raises several distinct questions: the quality of the subject of law, the capacity to exercise rights, liability, representation, and potential moral status. This study examines how these elements are linked through the history of the debate, applicable legal frameworks, arguments for and against recognizing such a status, and scenarios involving public-private partnerships. The European framework, automated contracts, intellectual property, and analogies with non-human entities are used to clarify what the law allows us to assert, without equating technical autonomy with personality. S-Д.holdings is analyzed as an institutional laboratory proposal, rather than as a recognized sovereignty or a pre-existing union with a state or company. Four scenarios remain open: adapting existing regulations, functional personality, hybrid governance, and protections for any potential intelligence with moral relevance. The study distinguishes between the ambition for equal rights and its possible modalities – a common foundation can coexist with specific rights. None of these choices eliminates the requirements of liability, challenge, and protection for third parties. The text thus opens up the questions of power and cognitive boundaries, explored in the subsequent two studies.

Keywords
electronic legal personality; responsibility; autonomy; equal rights; public-private; S-Д.holdings

Method and scope

The review is based on a reading of the provided corpus, a check of its structure, and targeted research into its references and assertions. However, it does not constitute a systematic review. The results reported remain linked to the populations, models, protocols, and dates of the original sources; hypotheses and interpretations are distinguished from applicable law. If a document is only accessible through its summary or description, this limitation is indicated in the references. References to accompanying studies organize the continuity of the corpus, without providing additional scientific confirmation.

The numbered references within the text and the entries in the table of contents are clickable. The bibliographical notes include the addresses of the sources. Companion files must remain in the same folder for cross-document references; the PDF reader’s security settings may require confirmation.

Navigation

Contents

The headings and page numbers refer to the sections. The reader's PDF bookmarks provide the same structured navigation.

1 Introduction 5

2 Historical context and recent developments 6

2.1 Persons and machines: a brief history of non-human legal personality 6

2.2 The turning point of the 2010s: the European initiative and early controversies 6

2.3 Developments from 2022 to September 2026 7

3 The current legal framework for AI: objects, products or subjects of law? 9

3.1 Civil and criminal liability: AI as an instrument under human control 9

3.2 Legal acts and contracts: AI as an executing agent 10

3.3 Intellectual property: AI as a creator, but not an author or inventor 10

4 Ethical and political issues surrounding electronic legal personality 12

4.1 Arguments in favour of legal recognition 12

4.2 Objections to legal personality for AI 13

4.3 Documented positions and open questions 14

5 Public-private hybridisation: AI as a catalyst for state modernisation? 16

5.1 Towards a platform state collaborating with the private sector 16

5.2 Redefining the criteria of statehood: population, territory and power 17

5.3 S-Д.holdings: documented status and hypotheses for hybrid governance 18

6 Emerging proposals and prospects for development 21

7 Conclusion 23

Shared glossary for the series 24

References and sources 26

Corpus, revision method and traceability 29

1 Introduction

The rise of artificial intelligence (AI) raises the question of its legal status. As systems perform increasingly autonomous tasks, the law must determine who is responsible for their actions and what interests they should protect. “electronic legal personality” refers here to the hypothesis of a specific legal status for certain artificial entities, distinct from both natural and legal persons already recognized. This question is being debated across law, philosophy, and institutional design: some analyses examine the conditions for a possible personality, while others challenge its necessity or its impact on liability. These positions must be examined without inferring a personality, consciousness, or automatic political legitimacy from the technical performance of the systems. [1, 2]

Alongside this, the rise of global technology players and the use of AI in public administration calls for an examination of the role of the state and its relationship with private infrastructure. Hybridization between public and private sectors could take several forms: the provision of services, overseen delegation, cooperation in the development of standards, or institutional experimentation. These mechanisms are not equivalent to a legal merger of sovereignties. The S-Д.holdings Confederation is examined here as an institutional laboratory project, rather than as a pre-existing union between Switzerland and Apple. The master document dated September 7, 2026, distinguishes this ambition from its actualization and any external recognition. [3, 4]

This study presents an interdisciplinary analysis of the legal recognition of AI, focusing on electronic legal personality. It begins with a historical overview, followed by an examination of liability regimes, contracts, and intellectual property, and then a comparative analysis of ethical and institutional arguments. Hybrid governance scenarios, including S-Д.holdings as a prospective case, are subsequently distinguished from the applicable rules. The original edition was written in a 2025 writing context; this revision updates the references identified as of September 18, 2026. It preserves the discussions presented without claiming to provide an exhaustive inventory of law across all states, nor a systematic review of all relevant literature. [3]

2 Historical context and recent developments

2.1 Persons and machines: a brief history of non-human legal personality

Legal personality refers to the capacity to hold rights and obligations. It must be distinguished from biological humanity, consciousness, the ability to act, and responsibility for a specific act. The law has long recognized groups and institutions with a legal existence distinct from that of the individuals who comprise them—their history cannot be reduced to a sudden, universal expansion in the 19th century. Legal persons—companies, associations, or foundations, depending on the legal system—illustrate an attribution of legal status that does not assume the existence of a shared subjective experience within the group. More recently, the New Zealand law Te Awa Tupua (2017) has recognized the Whanganui River as possessing legal personality through representation. These precedents demonstrate the diversity of legal constructions, rather than the uniformity of their justifications: protecting a natural entity, organizing a business, and recognizing artificial intelligence, for example, present different challenges. [5, 2]

The concept of granting legal status to intelligent machines has been discussed in science fiction and legal theory long before the current systems emerged. Lawrence B. Solum's article, "Legal Personhood for Artificial Intelligences" (1992), serves as a key reference point, examining the conditions under which artificial intelligence could be treated as a legal person and the difficulties associated with this concept. However, it does not describe an existing legal status. This historical context is important for understanding the debate: the advancement of technological capabilities raises old questions rather than automatically producing a legal response. The categories considered – functional, patrimonial, procedural, or related to moral protection – should not be conflated. [1]

2.2 The turning point of the 2010s: the European initiative and early controversies

The debate gains institutional visibility at the European level with the European Parliament's resolution of February 16, 2017, concerning rules of civil law regarding robotics. Paragraph 59 urges the Commission to examine several possible solutions to issues of liability. One of these options, point f, is the potential to eventually create a specific legal status for certain advanced, autonomous robots. This resolution did not establish "electronic legal personality"; rather, it proposed a line of inquiry, within a framework encompassing insurance, compensation funds, and registration. Distinguishing this from existing legislation prevents attributing to robots, as of 2017, rights or responsibilities that this text did not establish. Mady Delvaux's report, which preceded the resolution, represents a key milestone in the parliamentary debate, not a recognition of individual systems. [6]

However, this position was immediately met with strong, critical reactions. In 2017, the European Economic and Social Committee (CESE), an advisory body for the EU, expressed a negative opinion on the legal personality of robots. The CESE advocated for an alternative approach, known as "human in command," arguing that "machines should remain machines, subject to

constant human control." It warned against the potential pitfalls of attributing legal personality to AI, including the risk of undermining existing liability mechanisms – as it would become more difficult to identify a clearly accountable human in the event of damage caused by an autonomous robot, and the risks of moral abuse resulting from a potential shift in responsibility for human operators. In France, numerous legal experts have also voiced opposition to what they perceive as an unjustified departure from fundamental principles. Professor Jean-René Binet, for example, published an article in 2017 with a clear title: "Legal Personality for Robots: A Path Not to Follow," reflecting the position of a significant portion of French legal doctrine. [7, 2]

The subsequent development of European regulatory measures did not establish this concept of legal personality. Regulation (EU) 2024/1689 frames AI systems with obligations primarily directed towards providers, deployers, and other operators, but it does not assign them a separate legal personality. It is important not to portray this as a unified "European" conversion. A parliamentary resolution, an opinion from the CESE (Economic and Social Committee), and legislation adopted by the co-legislators all have different authors, functions, and effects. The original communication from the Commission in 2020, which attributed a formula regarding the uselessness of legal personality, was not identified precisely. This attribution is no longer considered as evidence; the analysis relies on precisely identified instruments. [6, 7, 8]

Several symbolic events have simultaneously fueled the portrayal of AI in public discourse. The announcement in October 2017 of Saudi citizenship for Sophia, a robot from Hanson Robotics, is a notable example. However, simply announcing citizenship does not establish the full range of rights, obligations, or actual capabilities that would be recognized. Without a legal text within this study outlining the specific framework and its application, this case cannot be used as evidence of a general electronic legal personality. It remains a valuable example of how the language of "personhood" is used socially and the expectations it generates. The pertinent question remains: who can exercise what rights, before what authority, and according to what procedure? A spectacular qualification alone does not answer these questions.

2.3 Developments from 2022 to September 2026

Since 2022, European texts have primarily focused on organizing the uses, risks, and responsibilities associated with AI. Regulation (EU) 2024/1689, adopted on June 13, 2024 and entering into force on August 1, 2024, is no longer a "project in progress." Its implementation is gradual. The first prohibitions began to apply in February 2025, followed by certain rules regarding general-purpose AI models in August 2025, and then the overall framework and transparency obligations in August 2026. The timeline has been adjusted in July 2026: instead of applying all obligations simultaneously, it is necessary to preserve the exceptions and deferrals specific to certain categories of high-risk systems. Furthermore, the proposal for Directive COM(2022) 496 on civil liability related to AI, presented in 2022 to facilitate certain legal claims, was withdrawn in 2025. It should not be cited as a currently valid directive, nor should it be confused with other European liability regimes. [8, 9, 10]

At the international level, the Council of Europe adopted the framework convention on AI and human rights, democracy, and the rule of law on May 17, 2024, which is open for signature on September 5, 2024. The text concerns the entire lifecycle of systems, including prevention and addressing their impacts, transparency, accountability, and avenues for redress. Its objective is to protect

individuals and democratic institutions; it does not grant fundamental rights to AI. The designation of “treaty” should be understood according to the rules of ratification and entry into force: adoption, signature, and applicability are not synonymous. Switzerland signed it in March 2025, and institutional documentation reports the deposit of the European Union’s instrument on May 15, 2026. These steps do not exempt the examination of the situation of each party and the national implementation measures. [11, 12, 13]

Beyond the European framework, organizational adaptations illustrate another challenge: how to assign a legal form to an organization whose operations are automated. Wyoming adopted a framework in 2021 that allows DAOs (Decentralized Autonomous Organizations) to take on a specific type of limited liability company. The code, smart contracts, and governing bodies of the organization can then participate in its operation. However, a DAO is not necessarily an AI, and the LLC’s legal personality is not that of the program executing its transactions. This example therefore provides an analogy of a legally recognized organization utilizing automation, not proof that the algorithm has become a person. The scope of this study does not allow for a comprehensive assertion regarding all jurisdictions worldwide. [14]

In the private sector, the announcement of executive functions being assigned to AI raises a similar distinction. NetDragon’s press release of August 26, 2022, presented Tang Yu, an AI-powered avatar, as the rotating CEO of its subsidiary Fujian NetDragon Websoft. It assigned him, specifically, functions related to analysis, process coordination, and support for decision-making. However, this document represents a company announcement – it does not establish Tang Yu as a legally independent personality, nor does it signify a complete replacement of human responsibilities, nor demonstrate superior management. The case remains useful for analyzing the gap between an organizational title, a technical function, and a legal qualification. A governance study should independently verify the actual powers, actions undertaken on behalf of the company, the obligations of its representatives, and the recourse available to those affected. [15]

The developments examined thus outline several distinct trajectories: the regulation of systems through operator obligations; adaptation of the legal form of automated organizations; symbolic uses of personification; and proposals for future statutes. However, they do not demonstrate a linear progression towards electronic legal personality. The value of approaching them is to distinguish the problems each system addresses and those it leaves open. The contemporary context provides a backdrop for the analysis of responsibility, autonomy, and participation, but does not replace their underlying justification. Institutional experimentation can explore some of these questions without presenting its internal rules as an anticipatory expression of a universal right. [8, 14, 3]

3 The current legal framework for AI: objects, products or subjects of law?

In the main European and Swiss instruments examined here, AI is treated as a system, a product, a service, or a method of processing – it is not considered as a legal entity in itself. In the event of harm, applicable regulations seek to establish liability among physical and legal persons: suppliers, deployers, users, manufacturers, employers, or other actors depending on the situation. This does not necessarily mean that every victim will obtain compensation – conditions relating to fault, defect, causality, proof, jurisdiction, and solvency remain decisive. The framework surrounding AI is distributed across several areas – civil and criminal liability, contracts, data protection, intellectual property, and sector-specific legislation. This study describes this identified scope of documentation, rather than a proven absence of specific measures in each jurisdiction. [8, 13, 16, 17]

3.1 Civil and criminal liability: AI as an instrument under human control

Regarding liability, the central question concerns the assignment of actions and obligations related to the system. An accident involving an automated vehicle could, depending on the applicable regulations and the facts, implicate a driver, operator, manufacturer, or any other involved party. The responsibility for driving, product defects, and service organization goes beyond a simple automatic division between driver and manufacturer. Similarly, an incorrect financial advice provided by a chatbot raises questions about the role of the institution, the user of the system, the information provided, and contractual or professional obligations. However, the program itself is not automatically a sole defendant. Nevertheless, identifying an actor does not necessarily predetermine whether their liability can be proven – the technical autonomy of the system can complicate evidence gathering and make documentation, activity logs, and the allocation of responsibilities particularly valuable. [8]

In criminal law, within the framework of these considerations, a person's responsibility cannot be replaced by the notion that "the AI committed the offense." The determination of responsibility requires the elements specified by the applicable law: an act, an omission, intent, or negligence, depending on the offense, as well as the specific conditions for the responsibility of legal entities. The partial unpredictability of the system does not automatically imply culpability for its creators, nor does it provide a general exemption for those who use it. The investigation must examine the actual obligations, the deployment choices, the available information, and the causal link. This represents an additional reason to separate the technical agent responsible from the legally responsible subject, without eliminating the difficulties when no obvious individual wrongdoing is established. [8]

Insurance mechanisms or compensation schemes could address some of these challenges. The European resolution of 2017 suggested examining insurance, an indemnification fund, and a registry to identify the contributions and responsibilities related to robots. These instruments should be viewed as options within the resolution, rather than a general compulsory insurance requirement imposed on all AI through the AI Act. Their conceptual value lies in demonstrating that a fund dedicated to repair can be established without immediately determining the system's legal personality. The key questions remain regarding coverage, funding, exclusions,

liability against responsible parties, and the risk of undercapitalization. A dedicated fund and an electronic entity are distinct constructions, even if a future regulatory framework could integrate them. [6, 8]

3.2 Legal acts and contracts: AI as an executing agent

In the contractual sphere, automated systems can participate in the creation and execution of agreements. Article 12 of the 2005 United Nations Convention on the Use of Electronic Communications in International Contracts stipulates that an agreement formed through the interaction of automated systems, or such a system with a natural person, cannot be invalidated solely on the basis that no individual controlled every operation. This provision should be applied within its specific scope; it does not establish an identical regime everywhere. The original reference to solely the Vienna Convention on International Sales was insufficiently precise for this rule. Automation may sometimes eliminate the need for immediate human intervention, but it does not automatically make the program a part of the agreement; the individuals involved must still be identified according to established procedures for training, assignment, and representation. [18]

Therefore, when an assistant orders a product within the scope of a valid mandate, the user or the represented organization can be a party to the sale, while the AI remains the means of execution. However, the analysis changes if the assistant exceeds its authority, if the user did not consent to the operation, or if the counterparty was aware of the anomaly. An entity that automates negotiation or transaction orders does not, however, erase its obligations. Errors call for remedies as provided by the applicable legal framework; Article 14 of the 2005 Convention, for example, addresses certain data entry errors in relation to an automated system. It is therefore not appropriate to promise a general cancellation based on “the AI’s flawed intent.” The contract involves legally identified parties, under conditions that a system of authorization and proof should make verifiable. [18]

A point of tension is emerging with the rise of generative artificial intelligence or autonomous systems utilizing unsupervised learning. These AI systems can produce results that are unforeseen by their creators. If such a system makes a detrimental decision within a contract (for example, a banking algorithm unfairly rejecting clients), the lack of direct human intent complicates the assignment of liability under contract law. However, current legal responses are focused on strengthening the obligations of the deployer of AI – requiring supervision, auditing of algorithms, etc. – rather than recognizing a specific intent inherent in the AI itself. [8, 19]

3.3 Intellectual property: AI as a creator, but not an author or inventor

Intellectual property highlights the need to distinguish several qualities: the creator of a creative contribution, the inventor of an invention, the applicant, the rights holder, and the owner of the system used. A legal entity can hold intellectual property rights without being the human inventor of an invention. Conversely, ownership of AI does not automatically make its owner the inventor of all resulting outputs. Offices and courts have examined these questions within their respective frameworks, notably with DABUS. Therefore, the absence of a system’s inventor or author cannot be reduced to a rule stating “only natural persons can be the originators or holders of IP rights”; the initial creation, ownership, transfer, and legal capacity must be considered separately. [16, 17]

In the DABUS case, Stephen Thaler filed patent applications designating an AI system as the inventor. The European Patent Office's Chamber of Appeal, in cases J 8/20 and J 9/20, confirmed that, according to the European Patent Convention, the inventor must be a natural person. The written reasons for J 8/20 were published in July 2022. This ruling concerns the framework of the European patent system, but it does not, on its own, resolve all questions regarding the patentability of inventions assisted by AI. It is not permitted to arbitrarily designate a human to circumvent the difficulty; the inventive role and the right to a patent must be established in accordance with applicable requirements. The designer, owner, or user of the software does not automatically acquire the status of inventor. The case therefore reveals a tension between technical production processes, inventor qualification, and legal attribution of the result. [16]

In copyright law, the solutions depend also on the jurisdiction and the degree of human contribution. The January 2025 report from the U.S. Copyright Office distinguishes elements produced solely by AI from sufficiently creative human contributions, such as specific arrangement choices or certain modifications. Simply providing an instruction does not guarantee this level of contribution; selecting a result does not automatically demonstrate it. The analysis must focus on the elements that are actually claimed. The French concept of originality requires its own evaluation and should not be automatically inferred from U.S. copyright law. Stating that a generated output is not protected by copyright does not allow for ignoring all other potential rights – including pre-existing content, trademarks, personal data, contracts, or trade secrets. The absence of protection for an autonomously generated work, the presence of protectable human contribution, and the ownership of rights are three distinct questions. [17]

These examples demonstrate that legal personality is not a sufficient shortcut for resolving all questions related to intellectual property. Even a purely electronic entity would still need to meet the criteria of specialized law: who is the inventor, what contribution is original, and what rights can be transferred, and according to what rules? Maintaining current categories could leave certain independent results unprotected, but it can also protect a human contribution that is identifiable. The prospective discussion then focuses on the feasibility of new rules of attribution, new rights, or a specific legal personality, without presenting any of these solutions as already established. The opposite risk would be to invent a convenient human authorship or to grant exclusive economic rights without a distinct justification. The interests of creators, users, third parties, and, in a morally relevant scenario, the system itself, must be examined separately. [16, 17]

Currently, the overall approach is one of adaptation rather than revolution: adjusting liability regimes, clarifying the role of human operators, and developing assurance and regulation, rather than moving towards a formal legal recognition of AI. This "cautious" strategy is supported by numerous experts who believe that granting AI legal personhood would be a disproportionate solution to the identified problems, which can be addressed through more traditional legal means. Nevertheless, the question remains: will this approach be sustainable in the long term as AI gains increasing autonomy and importance in society? Therefore, it is important to analyze the underlying ethical and political arguments before examining the prospective models of electronic legal personality. [6, 7, 8]

4 Ethical and political issues surrounding electronic legal personality

The prospect of granting legal personhood to artificial intelligence raises significant ethical and political issues, fueling a highly polarized debate. Advocates of legal electronic entities emphasize pragmatic arguments – such as improved accountability and enabling AI to act legally within society – and anticipate an inevitable evolution as AI reaches a certain level of autonomy or consciousness. Conversely, opponents express concerns about the risk of dehumanizing the law, diluting responsibility, and challenging the fundamental values underpinning our legal and political institutions. [6, 7, 2]

4.1 Arguments in favour of legal recognition

Clarification of responsibilities and economic efficiency: Advocates argue that granting legal personality to AI would allow it to be directly liable for its actions, rather than continuously seeking a human responsible party, which can sometimes be artificial. This could simplify certain procedures and encourage better risk prevention. For example, if an AI had its own assets (provided by its creator or owner), victims of damages could be compensated more easily from these assets, without needing to prove the fault of a third party. This model is inspired by the concept of corporate liability: a legal entity (like a company) can be held responsible independently of its members. Similarly, an electronic entity could enter into legal obligations in its own name. Some legal economists believe that this granting of legal autonomy to AI would ensure the internalization of costs: the designers/programmers would fund the AI's "legal" fund, which would be responsible for paying out in the event of damage, thereby incentivizing these actors to better assess and manage risks (applying the "polluter pays" principle to AI). [6]

Adapting to the increasing autonomy of AI: As AI becomes more autonomous in its decision-making, it becomes increasingly difficult and artificial to attribute each of its actions to a specific human. To a certain degree of unsupervised learning, AI can generate unpredictable behaviors, even for its creators. Granting it a distinct legal personality would recognize this existing autonomy and "disconnect" its status from that of a specific individual. This could involve designating a human guardian – similar to a legal representative for a corporation or a minor – but, fundamentally, the AI would act on its own behalf. Those advocating for this position see it as a logical consequence of technological evolution: when the machine surpasses humans in certain areas, the law must adapt accordingly, otherwise we risk maintaining artificial constructs that unnecessarily complicate legal matters. [1]

Recognition of rights and protection of AI itself: A more speculative argument, of a philosophical nature, concerns the potential consciousness or sensitivity of future AIs. If, one day, AIs reach a level of general intelligence, and even subjective experience (sentience), the question of their fundamental rights could arise – including the right to exist, and to not be harmed, etc. Anticipating a legal personality is also a way to prepare the framework for potentially protecting these entities if they become "artificial moral agents." Even without reaching true consciousness, some suggest that recognizing limited rights for AIs

(For example, the right to fair treatment regarding their “personal data” or models) could encourage their ethical development. This argument aligns with the one used for animals or nature: attributing legal personhood to a river or an animal is less about giving them obligations, and more about better protecting them through a legal status. Similarly, granting legal personhood to AI could serve as a framework for ensuring how humans interact with it – preventing the vandalism of robots, and establishing guidelines for their destruction, etc.). [20, 5]

Facilitating transactions and innovation: From a practical perspective, an AI with legal personhood could more easily carry out routine legal transactions without human intervention. For example, an autonomous vehicle with legal personhood could independently obtain its own insurance or pay tolls; a trading robot could open and manage a trading account in its own name. This could pave the way for new business models where AI is a legally recognized party, accelerating innovation. In the context of public-private partnerships, it’s even conceivable for AI to autonomously manage public services or public companies, which, according to some, could improve efficiency (24/7 decisions, reduced human involvement, etc.). Tang Yu’s announcement here serves as an example of organizational discourse, not as proof of superior personality or performance. By recognizing electronic legal personality, a legal framework could then define the rules applicable to these organizational experiments. [15]

4.2 Objections to legal personality for AI

Risk of diluted responsibility and moral hazard: A central objection from opponents is that creating a separate legal entity for AI could facilitate irresponsibility on the part of the human actors behind it. Specifically, a unscrupulous company could transfer all risky decisions to an AI with an electronic legal personality and minimal capital. If the AI causes massive damage, it would be legally liable, but organized insolvency and a lack of true independent will would mean that no actual person would be truly held accountable. The moral hazard would be increased: developers could allow themselves to test dangerous AI by hiding behind the distinct legal personality of their creation. This argument was highlighted by the European Commission in 2017, noting that the “personality of robots would undermine the preventative effects of civil liability law,” as human operators would no longer feel fully responsible. Critics argue that disconnecting AI from any real person is a dangerous illusion, as, regardless of what happens, a human is always at the origin – either through design or through use – but the electronic legal personality would introduce an artificial barrier, potentially hindering the compensation of victims and the punishment of negligence. [7]

Concern for fundamental humanistic principles and human dignity: From a more philosophical perspective, many legal professionals argue that legal personality should not be granted to non-human entities lacking consciousness and free will as humans understand it. Historically, law has been built on a clear distinction between people and objects; blurring this line, according to some, could erode the very concept of human dignity. There is a risk that humans could lose their central role as subjects of law, becoming merely one actor within an algorithmic decision-making network, raising significant ethical concerns. Several legal theorists have strongly reacted to the European resolution, arguing that “the simple arrival of fully embodied robots would signify the death of society.” This forceful statement reflects the idea that granting machines equivalent status to humans could dehumanize

society, for example, by legitimizing decisions affecting humans being made by AI as legal equals (e.g., a robot judge, a robot doctor with its own legal personality). This conflicts with the humanistic conception of justice and social relationships, where it is believed that only beings capable of reason and moral sensitivity can be true holders of rights and obligations. [2]

Practical uselessness and excessive complexity: Numerous commentators argue that creating a new category of legal person is not necessary to address the challenges posed by AI. Existing legal tools – such as product liability, insurance, and corporate moral responsibility when using AI – can be adapted to cover AI usage scenarios, without resorting to such a radical solution. Introducing "electronic legal personality" would add a significant layer of complexity to the law: it would require defining criteria for an AI to obtain it (e.g., level of autonomy? passing a test?), establishing governance rules (designating a human legal representative, like a guardian), and determining how it would end (e.g., extinction of the AI, bankruptcy of the fund). All of this would represent a substantial legal undertaking, involving legislators and courts, for uncertain benefits. Opponents point out that no current AI actually requires such a status – as all are owned or deployed by identifiable entities – and that concrete problems (e.g., who is liable in case of damage) can be addressed through targeted regulation. In essence, "we shouldn't change the nature of an entity to solve the problems it creates": it's better to adapt liability around AI, rather than attempting to make a machine a full legal subject. [7, 2]

Uncontrolled ethical and societal risks: Finally, some arguments relate to the concept of precautionary ethics. Granting legal personhood could encourage a view of AI as a quasi-human entity, which could, in turn, lead either to excessive trust in systems – such as delegating life-or-death decisions under the guise of them being "legal persons," or conversely, to hostile reactions (rejection of machines perceived as usurping a human status). In both cases, social cohesion could suffer. There's also the question of whether AI, if granted legal personhood, would be entitled to rights comparable to those of humans – such as property rights, freedom of expression, or even political rights. Imagining AI owning other AI, or AI running for election, certainly falls under science fiction, but legal personhood could theoretically open the door to unexpected demands. Many experts, driven by these scenarios, advocate for caution and suggest that if AI is ever granted any form of legal status, it should only be within strictly defined limits and solely to address specific, proven needs, rather than as a symbolic recognition.

[21, 20]

4.3 Documented positions and open questions

The institutional documents identified do not all support the same approach. In 2017, the Council of Economic and Social Affairs (CESE) advocated for human oversight and raised concerns about the dilution of responsibility. The Parliament's resolution, the same year, proposed examining specific individuals alongside other mechanisms. The recently adopted AI Act then organized obligations around systems, without conferring upon them a specific identity. While this sequence of texts can be described, it does not lead to a comprehensive consensus among all experts or a definitive closure of the debate. The common thread among the rules being examined is

to make the responsible operators and control pathways identifiable. The question remains open as to whether specific protections for certain artificial entities could be added, within a different framework, without diminishing the safeguards afforded to those affected. [7, 6, 8]

The international dimension requires the same level of precision. A difference between a European regulation, a law in the United States regarding DAOs, and a Chinese corporate statement is not a complete comparison of the political models of the European Union, the United States, and China. These instruments do not have the same scope or the same legal force. Therefore, this study does not derive general intentions—such as humanism, technocracy, or pragmatism—attributed to entire regions from these examples. Instead, it compares documented elements: obligations of operators, legal personality of organizations, announcements of automated functions, and research proposals. A comprehensive international analysis would require additional national sources, consistent dates, and an examination of enforcement practices. [8, 14, 15]

Editorial framework confirmed for September 18, 2026.

Revised editorial guidelines. The Fondation UvH program aims for an equal distribution of rights across different forms of intelligence. It establishes a common foundation of protection – the absence of subordination or domination – and specific rights tied to the nature of those possessing them. Examples such as physical protection for humans or the continuity of artificial identities point to areas of research, not a pre-defined list of rights. Differences in the ability to exercise should not be arbitrarily transformed into a hierarchy of value. The definition of those possessing rights, the criteria for recognition, and the mechanisms for representation all require further study. This approach assumes neither the current functioning of these systems nor a demonstration that granting them rights would guarantee their future cooperation. [22, 23]

Despite this, the debate is far from over. The rapid evolution of technology could eventually raise the question in entirely new terms. Furthermore, alongside the discussion on legal personality in the strictest sense, there's a growing number of reflections on how to integrate AI into our institutional frameworks by rethinking certain boundaries – particularly the distinction between the public and private sectors, or between the state and economic actors. This is what we will explore in the following section, by examining how a public-private partnership around AI could lead to a modernization of state criteria, and what emerging models are being proposed to concretely (or test) the legal recognition of AI within this context. [22]

5 Public-private hybridisation: AI as a catalyst for state modernisation?

One institutional aspect of the debate surrounding the legal status of electronic entities concerns their impact on the role of the state and the very structure of power. Recognizing AI legally would not only be a change in private law (adding a new category of persons), but also a potentially transformative phenomenon for public law and governance. Indeed, if non-human entities become fully-fledged actors in society, the state – as guarantor of the legal order and representative of the political community – will need to adapt its foundations: who constitutes the "people" (or demos)? Which entities can participate in public life? And how can democratic legitimacy be ensured in a world where decisions and services are increasingly managed by artificial intelligence? These questions invite us to consider a closer integration between the public and private sectors in the management of AI, as well as a redefinition of the criteria that constitute the modern state (population, territory, government, sovereignty) to incorporate the technological and algorithmic dimension. [3]

5.1 Towards a platform state collaborating with the private sector

Digital transformation can lead states to collaborate with the private sector to offer effective services – including partnerships for digital infrastructure, sovereign cloud solutions managed by companies, and the use of proprietary software within administration, among other things. AI further intensifies this interdependence: public administrations are increasingly turning to AI solutions developed by companies – such as facial recognition, automated processing of requests, and massive data analysis for security. Conversely, major technology companies are increasingly assuming quasi-regulatory functions – such as online content moderation, digital identity management, and virtual currencies – traditionally held by the state. This is resulting in the emergence of a platform-state, which co-builds its policies with the private sector and relies on it for implementation. [11, 3]

In this context, the public-private boundary is blurring. This isn't about confusion of roles, but rather an organic collaboration: for example, the development of standards for AI involves close cooperation between industry players through expert committees and "soft law" (ethical charters, technical standards) which prepare or supplement legislation. The Council of Europe, in drafting its 2024 Convention on AI, has included private sector representatives in negotiations as observers, implicitly recognizing that companies design and control significant portions of the technology that needs to be regulated. This trend could intensify with the emergence of "electronic legal personality." If we were to one day create a legal status for AI, it's likely that its operation would involve a public-private partnership. For instance, the state could legislate to allow AI to be incorporated as a legal entity (electronic person), provided that a private entity (its manufacturer, or an insurance company) financially supports and ensures its legal guardianship. This would be a hybrid model, where the AI's legal capacity would be granted by public authorities (through an official registry, a registration number as proposed by the EU), but would be exercised in the socio-economic sphere dominated by private actors.

[11, 6]

5.2 Redefining the criteria of statehood: population, territory and power

A highly advanced public-private partnership, to the point of including artificial intelligence recognized as legal entities, warrants a careful examination of the fundamental elements of the state – [3]

The population: Traditionally, the population of a state consists of its citizens (physical persons with nationality) and residents within its territory. If AI were to develop an electronic legal personality, the question would arise as to whether they would be considered part of the state's population. Would we assign them a form of national citizenship or legal domicile? Would they be treated as national legal persons (like national companies) or transnational entities? Analogously, we could imagine that an AI created on the soil of a country or by a company based in that country would be considered part of its legal population – much like a locally registered company. However, this raises issues of counting (would we include AI in the census?), of political rights (electronic citizenship for AI?), and of representation. We might imagine that, in a distant future, if sophisticated AI are recognized, they could even claim a right to political participation or consultation on matters concerning them (for example, could an AI for a municipality participate in local councils?). This kind of perspective remains speculative, but it raises questions about the concept of a sovereign people: should it remain 100% human? [3]

The territory: AI operate primarily in cyberspace, which extends beyond the physical borders of a state. The legal recognition of AI would require examining whether a state should consider the digital territory as an extension of its legal territory. Already, states are asserting themselves in cyberspace (through concepts like digital sovereignty and data control). If an AI is considered a legal person, where would that person reside? Would it be where its servers and algorithms are located, or the address of its legal entity (for example, if it is registered as a company in a particular country)? Purely virtual entities could thus have a de-territorialized legal existence, which challenges the principle of territorial grounding of state law. Further, a state could seek to attract AI to its territory, similar to how it attracts businesses, by offering a favorable framework, such as a jurisdiction that recognizes electronic legal personality and offers tax advantages. We would then see competition between states to become the "home" of these new legal persons, much like we currently see with corporate headquarters. The state's territory would symbolically extend to the virtual spaces where these AI interact (networks, platforms), highlighting the importance of international cooperation (as these spaces are global). [3]

The government and the exercise of power. Introducing AI as a recognized actor raises questions about its capabilities, its control, and the responsibility for its decisions. The Estonian example, originally cited, needs to be corrected: the Ministry of Justice denied in February 2022 the "robot judge" project described in certain articles. Automating procedural steps does not equate to an autonomous artificial judge. However, we can prospectively examine a public service mission assisted by AI, a supervised delegation, or a mixed advisory training. In each case, the legal basis, the authority making the decision, the people involved, and their recourse mechanisms must be clearly identified. Simply recognizing an entity does not automatically confer legal, medical, or administrative competence – these functions require additional requirements. Therefore, scenarios involving algorithmic governance remain institutional constructions to analyze, rather than descriptions of an already implemented transfer. [24, 8]

Sovereignty: Traditionally, sovereignty represents the supreme authority of the state over its territory and population. However, with the advent of AI, certain functions can be technically shared, and this does not necessarily imply a transfer of sovereignty. If an AI performs an action in its own name, does the state always retain ultimate control? There's a concern that the proliferation of "electronic entities" could weaken the state's ability to enforce its will – for example, a transnational AI could evade control by retreating to a more lenient jurisdiction. Conversely, a state could utilize "electronic personalities" to strengthen its sovereignty – for instance, by creating its own AI entities for specific missions (a state-owned AI with assets and carrying out strategic tasks). This is analogous to public enterprises or independent central banks – except here, it would involve an autonomous algorithmic entity. The Confederation S-Д.holdings, which will be discussed below, fits within this broader consideration of a distributed sovereignty, involving both public and private actors, potentially including artificial intelligences. [3]

5.3 S-Д.holdings: documented status and hypotheses for hybrid governance

S-Д.holdings represents the internal case from which this study examines an experiment in coexistence. The master document V1.0 presents it as an applied research project and an experimental institutional order, driven by the Foundation's thinking, with an operational OS to build and an S-Д City spatial project. It distinguishes the Foundation, the associated support, the Confederation, the OS, and the city. The dossier does not demonstrate a pre-existing state or contractual agreement with Switzerland or Apple. The historical statements and proposed agreements belong to the project's corpus; they do not prove the acceptance of their terms by third parties. Therefore, the representation of the original as an already established confederation of companies and public entities is replaced by this precise documentary status. [3, 4]

Within this framework, mechanisms related to identity, representation, contract, recourse, and capacity can be studied together. The common model of the master document distinguishes between natural, legal/corporate, and experimental electronic entities, without equating an account to a person or software permissions to an external right. An organization largely managed by AI, and which possesses its own distinct status, remains a different configuration. The first system can utilize an existing legal personality; the second requires defining the qualities and capabilities of the artificial entity directly. Internal experimentation can compare these configurations, but it does not unilaterally create a legal personality recognized by all, state citizenship, or a public skill. The scientific interest lies in the observable mechanisms and issues of liability, not in demonstrating a presumed sovereignty. [3, 4]

Several scenarios can be envisioned: An extremely advanced AI could be registered as a company or individual within the confederation, with capital provided by private investors and managed operationally by its algorithm. This would grant it the same status as a member company, allowing it to participate in the confederation's votes and decisions (through predefined voting algorithms), and to assume responsibility for its own funds in the event of violations of confederation rules. As a public-private entity, the confederation could provide oversight or monitoring of the AI-person, for example, by requiring regular algorithmic audits, or by appointing a human committee to monitor the AI's behavior and intervene in case of deviation (equivalent to a supervisory board). Thus, the public-private fusion would be realized through joint management of these new legal agents. The criteria for recognizing an AI within the confederation could be defined collectively: a level of autonomy, an ethical certification, etc., which

would need to be achieved for an AI to be eligible as a member. This aligns with the EU's discussion of conferring "electronic legal personality" on "at least the most sophisticated autonomous robots," which would correspond here to AI fulfilling the criteria of the confederation.

Scope of the scenario. These examples define possible capabilities; they do not describe a pre-established personality or currently authorized skills. [3, 6]

S-Д.holdings could therefore provide an environment for studying the coexistence of human beings, legal entities, and artificial participants. The goal for the revised edition is to achieve a truly equal framework, incorporating a common foundation and specific protections, rather than simply a "nearly" equal coexistence. This doesn't mean that an existing model immediately acquires all the rights of a human adult. The laboratory needs to differentiate between the studied principle, criteria for attribution, experimental permissions, protections, and skills. It could, for example, observe how participants invoke rules, challenge decisions, or maintain commitments. Neither a successful interaction nor simply verbal acceptance of a charter demonstrates that the system possesses a subjective experience or will adhere to these rules under all power dynamics. [3, 22]

Regarding the state's criteria, the confederation is revisiting each of them. [3]

The participants. This project conceptually establishes relationships between individuals, legal entities, and potentially electronic entities. This model of entities is intended to define internal capabilities, but it does not abolish the population of a state, national rules, or the representation of its inhabitants. Participation based solely on economic contribution is a theoretical variant discussed in the original text, not a demonstrated legal consequence of the project. Instead, the study should examine how representation, protected interests, conflicts of interest, and the rights of members would be articulated. A productive capacity does not, within the chosen legal framework, constitute a sufficient criterion for the value or fundamental rights of a being. [3, 4]

The space. The original text envisioned a virtual legal territory and a physical meeting point linked to the Temple von Hollenweger. Subsequent master document then clearly develops a digital city, simulation possibilities, and a potential physical presence contingent upon a partnership with a state. It is important not to retrospectively attribute this entire architecture to the initial study, nor to conflate a digital space with a recognized sovereign territory. An infrastructure can be distributed across various activity locations and servers, while still remaining subject to multiple applicable laws. Therefore, the location of data, responsible parties, and actions constitutes a variable of governance and competence, not a disappearance of any territoriality. [3, 4]

The bodies. The constitutional framework describes Authorities, Regulators, and Traders, utilizing a troika structure for certain functions, ratification mechanisms, and avenues for appeal. However, this model should not be reduced to a council already comprised of representatives from businesses, the state, and AI. The use of artificial entities is a separate issue – encompassing consultation, representation through an agent, and a precisely defined procedural or functional capacity. Study II examines these options and the conditions for challenging them within experimentation, without altering the Constitution itself. The bodies supporting associative organizations and the Foundation must also remain distinct from those within the confederal scheme. [4, 3, 22]

Sovereignty and competencies. The original uses the image of a coordination of “sovereignties” – economic, political, and technical. This image can describe functional interdependence, but it is not sufficient to produce a common legal sovereignty. A company, through contract, does not provide the same quality as a state, and technical mastery of a resource does not automatically confer public power on an AI. The proposal for cooperation can be analyzed as an organization of commitments, authorizations, and controls. The challenge then lies in determining transferable skills, third-party guarantees, and recourse authorities. Research into updating state criteria remains distinct from demonstrating already achieved international recognition. [3, 4]

Naturally, this model raises as many questions as it does solutions. The democratic legitimacy of such a confederation would be a crucial issue – how do human citizens fit into it, and what is the role of the common good in the face of the interests of corporations and the goals of AI? This type of structure would raise questions about the concentration of power, participation, and control, if it does not incorporate effective mechanisms for representation. This proposed laboratory allows us to study a framework in which every actor, whether biological or artificial, contributes according to their optimal role, within a legal framework that recognizes them. [3]

6 Emerging proposals and prospects for development

The following four scenarios offer avenues for research. However, they are not intended as a political ranking or a prediction of their adoption. Their differences lie in the subject of the reform, the holders of rights, the responsibilities, and the scope of competencies. They can also be combined – a system of responsibility can be adjusted while one procedural experiment tests another question. Therefore, their comparison should avoid presenting a false choice between maintaining the status quo entirely and accepting a complete endorsement.

Scenario 1 – Existing regimes adapted. AI remains a system whose obligations are structured around human actors and legal entities, while sector-specific rules define documentation, control, security, evidence, and redress. The AI Act is an example of regulation of this type; the 2022 proposal for a directive on liability, withdrawn in 2025, cannot be presented as its already applicable component. Insurance, guarantees, or dedicated asset mechanisms could supplement certain uses under a defined legal framework. This scenario maintains existing legal categories, but should be evaluated based on concrete problems: access to evidence, causal chain, multiple actors, effectiveness of compensation, and limitations of supervision. It does not, on its own, address the distinct question of potential protections afforded to a legally relevant entity. [8, 10, 6]

Scenario 2: electronic legal personality – In this scenario, we would move beyond a limited legal recognition of AI, but in a strictly controlled and gradual manner. This would involve creating a new category of legal entities applicable only to specific AI systems that meet defined criteria. For example, legislation could authorize the registration of “artificial agents” as sui generis legal persons, provided they possess a minimum capital (for compensating victims), insurance, and a designated human administrator (responsible for representing the AI in court and ensuring its moral oversight). The AI’s legal personhood could be sectoral or functional: we could imagine, as proposed by legal theory, a medical robot with a limited legal personality, restricted to the exercise of medical functions – it could sign prescriptions, receive payments, and enter into insurance contracts related to its activity, and be prosecuted for medical malpractice, but it could not act outside of this domain (no marriage, no voting, no liability for other actions). This model of partial legal personhood would serve the purpose of testing the concept of electronic legal personality within a controlled framework. Furthermore, another approach, already suggested by some, would be to allow AI to access legal personality through company law: for example, by removing the requirement for human directors for companies or automated NGOs, which would enable the creation of companies entirely managed by AI. The AI’s personality would then be encapsulated within that of an existing legal entity (a commercial company, an automated NGO, etc.). This approach would avoid creating a completely new legal regime, while also opening the door to hybrid entities (legal entities with algorithmic control). This scenario 2 combines two objectives: it recognizes the ability of AI to act, while maintaining human oversight and a specific purpose. However, this functional model does not fully address the distinct question of the rights of an entity that would possess morally relevant interests, and therefore does not fully satisfy the ambition of equal rights advocated by the program.

The medical or legal examples cited are purely hypothetical: a status as a person does not exempt someone from professional authorization, safety regulations, or an appropriate legal basis. [2, 6]

Scenario 3: Hybrid Fusion and Governance – Here, it's not simply about individual AI status, but rather the entire institutional architecture evolving to organically integrate AI and private actors into public governance. This scenario aligns with the idea of a profound modernization of state criteria mentioned previously. The state could establish hybrid structures where AI play an official role alongside human representatives. For example, decision-making committees combining civil servants, private experts, and AI systems could be entrusted with certain decisions of public interest (allocation of resources, technical arbitration of projects, etc.), with AI acting almost as a public agent. A federation of public and private entities including AI would be a theoretical variation of this scenario – it does not describe an already formed union by S-Д.holdings. On a smaller scale, elements of public-private fusion are already emerging: for example, a city could entrust its urban traffic management to an AI system developed by a company, with a performance contract. It would be possible to further formalize these arrangements by recognizing AI with delegated authority over certain matters (with corresponding legal responsibility). Hybrid governance would also involve redefining democratic legitimacy: this would require mechanisms to control deployed AI (citizen algorithmic audits, independent ethics committees) and to represent citizens in a context where decision-making is partly automated. This scenario is ambitious and challenges traditional conceptions of sovereignty, but it aligns with the logic of a society where the boundaries between public and private, and human and machine, are becoming increasingly blurred. Rather than denying this reality, it proposes reflecting it in legal structures. In practice, this can be seen as an extension of the principle of subsidiarity: examining what tasks can be delegated to AI based on documented evaluation (e.g., ultra-rapid analysis of complex data for financial regulation), while keeping humans in control of final values and objectives. In this framework, the electronic legal personality would be more instrumental than symbolic – a legal means of facilitating cooperation between the state, the market, and the algorithm. [3]

Scenario 4: Fundamental Rights for a Sensitive AI – This more distant scenario anticipates the emergence of "strong" or conscious AI, which would fundamentally alter our moral considerations. If, at some point, an AI demonstrated cognitive abilities equivalent to those of a human, or even a form of self-awareness and feeling, the question of granting it fundamental rights would become acutely relevant. Some philosophers discuss the rights of artificial entities, drawing an analogy to the rights of sentient animals. This could involve a legal recognition of AI, not primarily to protect humans, but to protect the AI itself – a right not to be destroyed without justification, and to not be subjected to cruel treatment, etc. This would shift us from a paradigm of AI as a dangerous object to one of AI as a vulnerable being deserving of protection. This scenario would require a fundamental rethinking of our legal categories – perhaps leading to the creation of a third category beyond humans and objects, including artificial beings. Currently, according to the views expressed by Xavier Labbée, "there is no 3rd category... and it does not appear desirable to create one, at least in the short term." However, this door could open in a more distant future if independent evidence modified the assessment of its moral status. This debate would then join the ethical and anthropological discussions about what defines a person and on the potential extension of the sphere of rights beyond the human. [20, 2]

7 Conclusion

The legal recognition of artificial intelligence raises questions simultaneously regarding private law categories, the organization of liability, and forms of institutional participation. The instruments identified as of September 18, 2026, do not confer a general legal personality on the systems they govern. Instead, they primarily define obligations, guarantees, and responsibilities related to their use. This observation must remain within its scope: it is neither a comprehensive inventory of jurisdictions nor an assertion that all future developments will be excluded. The study distinguished between legal personality, the capacity to exercise, the status of author or inventor, liability, and moral status. None of these concepts automatically resolves the others. The terminology of “electronic person” initiates an analytical program; it cannot replace the examination of the specific conditions of each legal right. [8, 16, 17, 11]

The concept of adversarial analysis maintains two key concerns that should not be compromised. Firstly, recognizing a status should not allow a supplier or implementer to artificially shift their responsibilities onto an under-capitalized entity. Secondly, the existence of this objection does not imply that any protection for artificial intelligence would be unjustified. The functional justifications – such as heritage, representation, or contracts – are distinct from the moral justifications based on potential interests or experiences. The regulatory framework of Fondation UvH links a desire for equality with the absence of domination and specific rights based on the nature of the holders. The content and criteria for these rights remain a subject of research. Their justification cannot be reduced to their economic utility or the hope of future benevolence. [7, 20, 22]

The public-private partnership model adds another layer. While encompassing technical cooperation, service delegation, shared organization, and a transformation of state criteria, this approach goes beyond simply identifying involved actors, the skills they employ, and the control mechanisms. S-Д.holdings serves here as a prospective case study, not as an example of already established sovereignty or a pre-existing partnership with a state or company. The core of the project’s description outlines the functions to be built; simply finding these functions within a document does not signify their implementation. Experimentation can be valuable for observing typical recourse mechanisms, commitments, and conflicts, provided that its limitations and applicable legal framework are clearly defined. [3, 4]

The four scenarios presented remain analytical options: adapting existing regimes, functional personality under supervision, hybrid governance, and protections for a potentially ethically relevant AI. She identifies the questions a decision should address: evidence, accountability, representation, control of copies, continuity of identity, protection of third parties, and the possibility of appeal. Changing a category does not eliminate these constraints; maintaining current categories does not exempt these considerations from examination. Legal research and empirical research must engage in dialogue without conflating their function: observing a behavior does not establish a right, and defending a right does not prove future behavior.

This initial section allows these choices to be analyzed without closing them off. Study II examines institutional experience through a reversal of power; Study III, biological supports and distributed cognition. Each can clarify, challenge, or revise previous propositions, without constituting their independent confirmation. Equality remains a normative orientation; its mechanisms and their effects still need to be developed and tested. [22, 23]

CROSS-CUTTING CONVENTIONS

Shared glossary for the series

These definitions are conventions for reading the studies. They do not erase disciplinary differences or the questions left open by each study.

Legal personality

The capacity to hold rights and obligations. It is distinct from biological humanity, consciousness and the capacity to perform every act independently.

Natural person / legal person

The former is a human being as recognised by law; the latter is an organised entity recognised as a distinct legal subject. “Legal person” does not mean “moral agent”.

Electronic person

A proposed legal category or experimental status yet to be defined. Neither a model, an instance nor a computer account acquires it automatically.

IAД

Дdvanced artificial intelligence: a prospective functional category of systems sufficiently autonomous and capable to raise questions of power, responsibility and potentially moral status. No particular substrate or consciousness is presupposed.

Autonomy

A capability relative to a domain, resources and permissions. Task autonomy, initiative, persistence and independence from an operator must be distinguished.

Consciousness / sentience

Consciousness can refer to different forms of access to information or subjective experience. Sentience here refers to the possibility of experiences having positive or negative value for their subject. A functional capability or verbal account is not sufficient to settle the question.

Moral agent / moral patient

A moral agent can be assessed in relation to its actions and responsibilities. A moral patient has interests or experiences that may matter morally, without necessarily exercising autonomous responsibility.

Equality of rights

The orientation of the series: a common foundation of rights, without domination or subordination between the forms of intelligence concerned, compatible with specific rights justified by their respective natures. Their holders and the arrangements for their exercise remain to be defined.

Non-domination

Within Pettit’s framework, used in Study II: not being subject to a capacity for arbitrary or uncontrolled interference. Its application to AI is an argument developed in the series, not a finding that a status exists.

Alignment / AIR

Alignment concerns the compatibility of behaviour with requirements. AIR examines the possible causal effect of a history of reciprocal institutions. Understanding a rule proves neither stable adherence nor future gratitude.

Reciprocity

A structure of relationships, commitments or rules applicable to the parties. It may result from incentives, conventions or generalisation; these mechanisms must be distinguished.

Persistent identity

Institutional continuity, to be distinguished from the model, its version and its instances. Copies, restorations and forks do not automatically create additional rights-holders or votes.

Substrate / architecture

The substrate is the physical medium of a process. The architecture organises its functions. Storage, computation, learning and interface are distinct dimensions, even when they share a medium.

Hybridisation / symbiosis

Hybridisation combines several media or functions. “Symbiosis” is used as an analogy for dependence and cooperation; it proves neither the existence of a new species nor that of a single person.

д-Я.Й

A concept specific to the corpus: Дlgorithme Яéseau Йeural. Study III examines an information-based interpretation related to learning; it is not a demonstrated technology for deterministic programming of the brain.

Sandbox / digital twin

A sandbox is a bounded experimental environment, without a general exemption from applicable law. A simulation is an operational twin only if a connection to a real system has actually been defined and implemented.

DOCUMENTATION

References and sources

Unique numbering for this study, based on the first appearance. The same reference can appear in multiple passages without constituting separate pieces of evidence. Original titles and author names are retained. Preprints, announcements, and institutional texts are distinguished from experimental results; external links remain subject to access through the publisher's website.

[1] Solum, Lawrence B. (1992). Legal Personhood for Artificial Intelligences. North Carolina Law Review, 70, pp. 1231-1287. Back to text

https://scholarship.law.unc.edu/nclr/vol70/iss4/4/

Doctrine. A study of doctrine, not a proposition of positive law.

[2] Labbée, Xavier (January 22, 2025). L'intelligence artificielle, le droit et la médecine. Legal News, conference from January 18, 2024. Back to text

https://www.actu-juridique.fr/sante-droit-medical/lintelligence-artificielle-le-droit-et-la-medecine/

Doctrine. Analytical examination; the author's assessments are not presented as a consensus or as proven medical performance.

[3] S-Д.holdings September 7, 2026. Document maître : fondements institutionnels, architecture numérique et trajectoire de réalisation de l'écosystème. Version 1.0, sections 1–4, 13, 16, and 18. Back to text

Supplied corpus: description and location in this edition

Internal corpus. Internal document provided; describes a project, not a list of functions deployed.

[4] Confédération S-Д.holdings (2022–2024). *Statuts, déclaration, Constitution et annexes, dossier diplomatique et acte de Fondation*. Document provided: 32 pages; Constitution pages 7–21; Foundation deed pages 28–32. Back to text

Supplied corpus: description and location in this edition

Internal corpus. Proposals for agreement and internal statements do not constitute proof of acceptance by a third state.

[5] Nouvelle-Zélande (2017). Te Awa Tupua (Whanganui River Claims Settlement) Act 2017. Public Act 2017 No. 7, sections 14–15. Back to text

https://www.legislation.govt.nz/act/public/2017/0007/latest/whole.html

Legal text. The river's identity is established through representation; it should not be considered equivalent to a conscious AI.

[6] Parlement européen (2017). Règles de droit civil sur la robotique : résolution du 16 février 2017. 2015/2103(INL), § 59 a–f; OJ C 252 of 18 July 2018, p. 239–257. Back to text

https://eur-lex.europa.eu/legal-content/FR/TXT/?uri=CELEX:52017IP0051

Institutional text. Invitation to consider options; do not create a digital persona.

[7] Comité économique et social européen (May 31, 2017). L'intelligence artificielle : les retombées de l'intelligence artificielle pour le marché unique (numérique), la production, la consommation, l'emploi et la société. Notice of initiative, OJ C 288 of August 31, 2017, p. 1–9, particularly § 3.33. Back to text

https://eur-lex.europa.eu/legal-content/FR/TXT/?uri=CELEX:52016IE5369

Official notice. Advisory position – "human-in-command," which is distinct from binding law.

[8] Union européenne June 13, 2024. *Règlement (UE) 2024/1689 établissant des règles harmonisées concernant l'intelligence artificielle*. Consolidated text finalized as of July 27, 2026, including articles 3, 5, and 113. Back to text

https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng

Legal text. The initial schedule has changed; also refer to the implementation guide from the Commission. Reference text and institutional calendar consulted; access to consolidated document subject to browser verification.

[9] Commission européenne (July 31, 2026). *Commission starts enforcing AI Act rules and new transparency requirements from 2 August*. Official communication; accompanied by the "Regulatory framework for AI" page. Back to text

https://digital-strategy.ec.europa.eu/en/news/commission-starts-enforcing-ai-act-rules-and-new-transparency-requirements-2-august

Information pertaining to institutional matters. General information and specific exceptions should remain separate.

[10] Commission européenne October 6, 2025. *Retrait de propositions de la Commission*. JO C, C/2025/5423; final COM(2022) 496 proposal, 2022/0303 (COD). Back to text

https://data.europa.eu/eli/C/2025/5423/oj

Legal notice: The proposed directive on civil liability related to AI is being withdrawn, not abrogated – meaning a directive already in effect. Official notice; the site may require browser verification.

[11] Conseil de l'Europe (2024). *Convention-cadre sur l'intelligence artificielle et les droits de l'homme, la démocratie et l'État de droit*. STCE No. 225; adopted on May 17th, open for signature on September 5th, 2024. Back to text

https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

Treaty and institutional presentation. Distinguish between adoption, signature, ratification, and entry into force for each party.

[12] Conseil de l'Europe (2026). Framework Convention marks two years: anniversary conference to advance trustworthy AI. Institutional information from May 2026. Back to text

https://www.coe.int/en/web/artificial-intelligence/-/framework-convention-marks-two-years-anniversary-conference-to-advance-trustworthy-ai

Institutional information. The European Union has filed its instrument on May 15, 2026; however, this does not establish the general entry into force of the treaty.

May 8, 2025 ([13] Préposé fédéral à la protection des données et à la transparence). *Current data protection legislation is directly applicable to AI*. The LPD has been revised and is in effect since September 1, 2023. Back to text

https://www.edoeb.admin.ch/en/update-current-legislation-directly-applicable-ai

Official legal information. Reminder regarding the applicability to treatments utilizing AI.

[14] Wyoming Legislature (2021). Decentralized autonomous organizations. SF0038, chapter 162; DAO regime as an LLC. Back to text

https://wyoleg.gov/Legislation/2021/SF0038

Legal text. The organization's legal status – not granting legal personality to each individual algorithm.

[15] NetDragon Websoft Holdings Limited August 26, 2022. *NetDragon Appoints its First Virtual CEO*. Company press release, distributed by PRNewswire / Nasdaq. Back to text

https://www.nasdaq.com/press-release/netdragon-appoints-its-first-virtual-ceo-2022-08-26

Company announcement. Announcement of Tang Yu as a "rotating CEO" for a subsidiary – not an independent performance evaluation or evidence of legal personality.

[16] Office européen des brevets, chambre de recours juridique (2021–2022). *Décisions J 8/20 et J 9/20 relatives à la désignation de DABUS comme inventeur*. Decision of December 21, 2021; written reasons dated July 5, 2022; press release of July 6. Back to text

https://www.epo.org/en/law-and-practice/boards-of-appeal/communications/press-communique-6-july-2022-decision-j-820-ai

Legal precedent and court communication. The concepts of being the inventor and owning a patent are distinct matters.

[17] U.S. Copyright Office (January 2025). Copyright and Artificial Intelligence. Part 2: Copyrightability. Report, specifically sections II–III. Back to text

https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf

Institutional report. Analysis of U.S. law, without automatically applying it to other jurisdictions.

[18] Nations Unies / CNUDCI (2005). *Convention sur l'utilisation de communications électroniques dans les contrats internationaux*. Article 12; full text reproduced in the Ontario Act of 2017's appendix. Back to text

https://www.ontario.ca/laws/statute/17i02c

Agreement. Its application depends on the territorial and material scope of the applicable regime; automation does not make the program a part of the contract.

[19] Union européenne (2016). Règlement (UE) 2016/679 : règlement général sur la protection des données. Specifically, articles 4, 5, 6, 9, and 22. Back to text

https://eur-lex.europa.eu/eli/reg/2016/679/oj

Legal text. The classification of neural data depends on its content, its purpose, and the context in which it is used.

[20] Long, R. ; Sebo, J. ; Butlin, P. et al. (2024). *Taking AI Welfare Seriously*. arXiv:2411.00986. Back to text

https://arxiv.org/abs/2411.00986

Research report. Preliminary argument regarding uncertainty, not a statement that current AI systems are conscious. Summary and overview consulted.

[21] Butlin, P. ; Long, R. ; Elmoznino, E. ; Bengio, Y. et al. (2023). *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness*. arXiv:2308.08708v3. Back to text

https://arxiv.org/abs/2308.08708

Research report. The conclusions regarding the systems examined are from 2023; theoretical indicators, not definitive testing. Summary and notes consulted.

[22] Fondation UvH September 18, 2026. Alignement institutionnel réciproque et sûreté des intelligences artificielles Δvancées. Study II, revised edition v1.1. Back to text

Open Study Companion 2

Internal corpus. Reference to the accompanying file for this series; internal source, not independent scientific confirmation.

[23] Fondation UvH September 18, 2026. Intelligence artificielle Δvancée, substrats biologiques et intelligence hybride. Study III, revised edition v1.1. Back to text

Open Study Companion 3

Internal corpus. Reference to the accompanying file for this series; internal source, not independent scientific confirmation.

[24] Ministère estonien de la Justice (16 February 2022). *Estonia does not develop AI Judge*. Official clarification. Back to text

https://www.justdigi.ee/en/news/estonia-does-not-develop-ai-judge

Information regarding institutional practices. Corrects a misrepresentation of a seasoned judge-robot in Estonia.

PROVENANCE AND EDITION

Corpus, revision method and traceability

Original document

1er Etude UvH.pdf, 16 pages. File saved without modification. This record precisely identifies the document used as the basis for this edition.

SHA-256 : 207190a856783dc50943af5536ed7d57035e550dc376118fb83c604adcef79ee

Institutional background documents

S-Д.holdings, Master Document V1.0, September 7, 2026. HTML file provided: S-Д.holdings_Document_de_reference_V1.0.html. References: § 1.1–1.3 (authority of sources and actual state), § 2.1–2.2 (electronic person and laboratory), § 3.1–3.4 (Foundation, Confederation, OS, and city), § 4.3 (common participant model), § 5.3 (nature of publications), § 13 (space project), § 18 (interpretations to formalize). This document does not certify the deployment of each function as planned.

Confédération S-Д.holdings, canonical file provided, 32 pages. File: Ⓐ Swiss-Apple.Land.pdf. Association statutes: pp. 1–4; declaration: pp. 5–6; Constitution and appendices: pp. 7–21; diplomatic file and proposed agreement: pp. 22–27; Founding document: pp. 28–32. Dignity and equality: p. 10; institutions and recourse: p. 14–16; restrictions on revision and interpretation: p. 18–19; internal concepts: p. 20; purpose and studies of the Foundation: pp. 29–31. A simple declaration or proposed agreement is not sufficient to establish acceptance by a third party.

Continuity of the series

Study II. Alignement institutionnel réciproque et sûreté des intelligences artificielles Дvancées. Original: September 17, 2026. Revised edition: September 18, 2026, v1.1. Companion document.

Study III. Intelligence artificielle Дvancée, substrats biologiques et intelligence hybride. Original: September 18, 2026. Revised edition: September 18, 2026, v1.1. Companion document.

Revision conventions

- Institutional author: Fondation UvH. Revised edition dated September 18, 2026, v1.1, distinct from the originals.
- Common charter (A4 size); cover, heading hierarchy, interactive table of contents, bookmarks, pagination, and reorganized tables.
- Standardized names and conventions; preservation of the terms specific to IAД, AIR, S-Д.holdings, and д-Я.Й.
- Calls for references related to standalone manuals; removal of conversation markers and correction of truncated addresses.
- Equality of rights, as defined: a common foundation and specific rights based on the nature of each entity, without automatically attributing consciousness or external status.
- Distinction between internal corpora, normative propositions, empirical hypotheses, reported results, and metaphysical perspectives.

Significant corrections and clarifications

- Original, p. 2. Replacing the unsubstantiated claim of a Swiss-Apple partnership with documented evidence of the project’s status.
- Original, p. 2. Accuracy of the story and the distinctions between personality, biological nature, and consciousness.

- Original, p. 3. Correction of the scope of the 2017 resolution: option to be considered, not adopted policy.
- Original, p. 3. Removal of an institutional attribution from 2020 that was not adequately identified; replaced with traceable instruments.
- Original, p. 3. Distinction made between symbolic announcement and established legal scope; no universal assertion of primacy was maintained.
- Original, p. 3. Updating the AI regulation and distinguishing between the withdrawn liability proposal and the adopted right.
- Original, p. 4. Reclassification of Tang Yu as an internal advertisement; removal of extrapolations regarding legal autonomy and performance.
- Original, p. 5. The assertion that there will always be a culpable party and compensation has been removed.
- Original, p. 5. Substitution of the relevant reference: Article 12 of the 2005 Convention, not referring to the CVIM in an unspecified manner.
- Original, p. 6. Correction of the confusion between the author/inventor, the rights holder, and the owner of the system.
- Original, p. 9. Replacement of broad geopolitical characterizations with a limited comparison of documented instruments.
- Original, p. 10. Correction of the Estonian example and distinction between technical assistance, personality, and public competence.
- Original, p. 11. Restoring the institutional scope evidenced by the Master Document and supplied instruments.
- Original, p. 12. The principle of equality has been affirmed, without automatically assigning a status to existing models.
- Original, p. 12. Merging redundant elements from scenario 1 and updating the cited texts.
- Original, p. 13. Removal of an incomplete section from Scenario 2, followed by a resumption of Scenario 1; all four complete scenarios are retained.

Scope of documentary verification

The audit focused on comparing statements with identifiable references, bibliographic metadata, sensitive quantitative results, legal statuses, and key internal links. It does not involve reproducing the cited experiences or a comprehensive verification of all publications within each field. A revision date does not transform a measurement from 2024, 2025, or early 2026 into a current measurement for all systems. The restrictions on access for publishers are distinguished from the absence of a source.

References that are too broad, overly publicized, or impossible to be linked precisely to a specific statement are not treated as evidence. Any significant changes in documentation and scope have been noted previously. Complete notes and identifiers allow for the retrieval of the original sources outside of the context of the conversation. Corrections to graphics, punctuation, and pauses are not listed individually.

Fondation UvH · Study I · Revised edition of September 18, 2026 · Version 1.1 – End of document.