Asia Pacific Academy of Science Pte. Ltd.

Metaverse 2026, 7(4), m260035.
https://doi.org/10.54517/m260035

Article

The System of AI Legal Subjects of the Metaverse: The AI Legal Entity and AI Representatives as Autonomous Subjects Rather Than Digital Duplicates of the Human

Oleksii Kostenko[1*], Dmytro Zhuravlov[2], Viktoriia Hurochkina[3], Roman Tatarinov[4], Svitlana Koshova[5]

[1] Scientific Laboratory of Immersive Technologies and Law, State Scientific Institution "Institute of Information, Security and Law of the National Academy of Legal Sciences of Ukraine", 01001 Kyiv, Ukraine
[2] Office of the President of Ukraine
[3] University of Zielona Góra, Licealna Str. 9, 65-246 Zielona Góra, Poland; SGH Warsaw School of Economics, al. Niepodległości 162, 02-554 Warsaw, Poland
[4] Department of International Relations, Public Communications and IT Law, State University of Intelligent Technologies and Telecommunications, Ukraine
[5] Department of Administrative Law and Public Administration, National Academy of Internal Affairs, Ukraine

*__Corresponding author:__ Oleksii Kostenko, antizuk@gmail.com

CITATION

Kostenko O, Zhuravlov D, Hurochkina V, et al. The System of AI Legal Subjects of the Metaverse: The AI Legal Entity and AI Representatives as Autonomous Subjects Rather Than Digital Duplicates of the Human. Metaverse. 2026; 7(4): m260035. https://doi.org/10.54517/m260035

COPYRIGHT

Copyright © 2026 by author(s).
Metaverse is published by Asia Pacific Academy of Science Pte. Ltd. This work is licensed under the Creative Commons Attribution (CC BY) license.
https://creativecommons.org/licenses/by/4.0/

Abstract: Economic activity and governance in virtual environments already rely on acts performed neither by humans nor by classical legal persons: autonomous systems dispose of digital assets, prepare draft administrative decisions, generate and disseminate informational content, and execute the decisions of decentralised organisations. The available legal constructs —the avatar, the digital twin, "digital humans", electronic personhood and the AI avatar —share a common defect: all of them are mimetic, that is, they construct the new subject as a reflection of the human being and therefore inevitably run into the intractable questions of consciousness, intent and dignity. This study builds a system of AI legal subjects of the Metaverse on the opposite, non-mimetic principle: a subject is defined not by its likeness to a human being but by its registered function and verified level of autonomy. The methodology combines doctrinal analysis, the functional comparative method, conceptual legal modelling, and the operationalisation of technical autonomy features on the basis of ISO/IEC 42001, ISO/IEC 23894, IEEE 7001-2021 and the NIST AI RMF. The results take the form of normative models (de lege ferenda) rather than empirically validated mechanisms. They comprise the non-mimeticity principle as the theoretical foundation of the system; a two-branch taxonomy dividing AI subjects into a status branch (the AI legal entity, classes 1–3) and a representative branch (nine types of AI representative); a proposed catalogue of the legal rights of an AI subject that does not copy human rights, together with four absolute limits; a three-dimensional registration profile combining subject type, autonomy class and functional profile; a reserve ratio with a dual progression (whose quantitative parameters are indicative and require empirical calibration); an identity anchor with cross-world status passporting; a deferred settlement finality window; and a six-tier liability cascade including a tier for the attribution of will to the principal. The proposed elements are modular and have different implementation horizons: autonomy classes and functional profiles could be introduced through platform-world rules today, the register, the reserve ratio and the liability cascade require special legislation, and status passporting requires inter-jurisdictional arrangements. Tax and conflict-of-laws issues lie outside the scope of the study.

Keywords: AI legal subjects; AI legal entity; AI representative of a public authority; AI representative of a media outlet; non-mimeticity principle; legal personality in the Metaverse; digital mandate; autonomy class

1. Introduction

1.1. Statement of the problem

The Metaverse has ceased to be a metaphor and has become an environment in which activity with measurable legal consequences occurs. The research literature describes it as a convergence of immersive technologies, persistent virtual spaces and the digital economy [1–4]; Ukrainian legal scholarship has characterised this environment as the object of an autonomous legal regime and of an electronic jurisdiction [5,6]. Market analysis records stable segments of that economy: virtual real estate, in-world goods, environment design services, advertising and financial services [7,8]. A body of empirical material has accumulated on property, possession and exchange in virtual worlds [9] and on sectoral applications of the environment [10], while the philosophical account of the virtual as real [11] has removed the argument that the relations in question are somehow "unreal". The legal dimension of the environment has been examined through the constructs of electronic jurisdiction and of a typical cross-border model of a virtual world [12,13].

What all of this material has in common is that an ever greater share of acts having legal effect in the Metaverse is not performed by a human being directly. An autonomous system disposes of tokenised assets; an algorithm prepares a draft administrative decision and transmits it to the applicant; a generative system produces and disseminates content on behalf of a news organisation; a software agent executes the decisions of a decentralised organisation in dealings with public registers. In legal terms, all of these entities remain invisible: they are neither subjects nor even distinct objects, but are classified as a software tool of the user or of the operator.

Three structural defects follow. The first is non-identifiability: a claimant has no starting point from which to establish who is liable. The second is non-portability: an agent operating in several worlds has a different legal regime in each of them, and none of those regimes recognises any other. The third is lack of security for claims: even where the responsible party has been established, the claimant often finds that there is no pool of assets from which recovery could be obtained, because the agent's digital assets are neither segregated nor ring-fenced.

1.2. The research gap: the mimetic trap

Attempts at a legal solution to this problem now form a considerable body of literature. The doctrine of virtual property and of regulation through code explains the nature of the objects and the mechanism of normative influence [14–20]. Corporate wrappers for decentralised organisations, introduced in individual jurisdictions [21,22], confer legal personality on such entities. The 2017 resolution of the European Parliament proposed "electronic personhood" [23], which remained a declaration without an institutional framework and attracted extensive criticism [24]. Gradient models of the legal personality of artificial intelligence have supplied a theoretical apparatus for the divisibility of status [25–27]. The problem of the legal recognition of avatars and electronic personalities has been raised, in particular through the idea of a cross-border personalised ID code [28], and the sectoral contours of regulation have been outlined in model codes of the Metaverse [29]. In parallel, an autonomous strand has emerged on the legal personality of the avatar and of the AI avatar in virtual worlds [30–32].

Yet virtually all of these constructs share a feature that has so far not been problematised. They are mimetic: the new subject is constructed as a reflection of the human being —its likeness, its twin, its "digital personality", its avatar. The consequence of that choice is not technical but fundamental. Once a subject is defined through its likeness to a human being, any discussion of its status

inevitably collapses into questions that law is not equipped to resolve: whether the system possesses consciousness, whether it is capable of intent, whether dignity attaches to it. It is on precisely this ground that the strongest criticism of the legal personality of autonomous systems is founded [24,33]: if the subject is a copy of a human being, then conferring status upon it is either a fiction or an erosion of the uniqueness of the human person.

It must be acknowledged, however, that the move away from consciousness, human likeness and full personhood as criteria has clear antecedents in the literature. The cluster theory of legal personality [25] and the conceptual framework for divisible legal personality [26] already detach status from the psychological properties of its bearer; theories of partial (limited) legal personality of software agents [34,35] and of functional legal agency, under which an autonomous system is treated primarily as a legal agent of its principal [36], approach autonomous systems first and foremost through the functions they perform; and the earliest framing of the problem [37] expressly separated the question of legal personhood from that of consciousness. The contribution of this study lies not in abandoning mimetic criteria as such, but in turning non-mimeticity into an explicit constitutive principle from which the entire institutional architecture —the taxonomy, the catalogue of rights, the registration profile and the mechanisms of security and reversibility —is consistently derived (see **Table 1**).

The research gap is therefore not thematic but conceptual. What is missing is a system grounded in the opposite principle: a subject defined not by its likeness to a human being but by its registered function and verified level of autonomy.

Table 1. Elements of the model: closest antecedents and the contribution of the study.

Element of the model	Closest doctrinal / regulatory antecedents	What is borrowed or adapted	What this study proposes
Non-mimeticity principle	cluster theory of legal personhood [25]; divisible legal personality [26]; gradient models [27]; partial legal personality [34,35]; principle of non-simulacrity [6,38]	detachment of status from consciousness; divisibility of status	non-mimeticity as an explicit constitutive principle from which the entire architecture is derived; a scale of status based on verified autonomy rather than approximation to a human being
Two-branch taxonomy	theory of autonomous agents as legal agents [36]; autonomous organisations without members [39,40]; classical representation	the agent–principal model; a subject without members	distinction between status and function; representation as a function superimposed on a classified subject
Catalogue of rights	functional rights of the legal person; critique of “synthetic persons” [24]	functional logic of the rights of non-natural subjects	formation of the catalogue from functional necessity; four tiers; four absolute limits
Three-dimensional registration profile	risk-based approach of the EU AI Act [41]; MiCA licensing categories [42]	risk stratification of functions	combination of subject type, autonomy class and functional profile in a single register entry
Reserve ratio; guarantee fund of the world	prudential regulation [43,44]; deposit guarantee and investor compensation schemes	logic of capital at risk and of collective coverage of tail risk	dual progression (by autonomy and by social risk); a fund at the level of the platform world
Identity anchor; status passporting	European Digital Identity Framework [45]; cross-border ID code [28]; the “single passport” in financial services [46]	verified identity; recognition of a licence by the host State	identification of a legal status rather than of a person; recognition of status by a private world administrator on the basis of voluntary participation in the register
Deferred settlement finality window	settlement finality in payment systems [47]; payment challenge mechanisms	the moment of finality as a legal category	deferred finality for acts of an AI subject in the financial stratum, with an exhaustive list of grounds for challenge
Liability cascade	liability regimes for AI [35,48]; liability for defective products [49]; agency law	subsidiary liability; producer liability	six tiers of successive exhaustion, including Tier 0 —attribution of will to the principal

Source: compiled by the authors.

1.3. Aim and research questions

The aim of this study is to build a system of AI legal subjects of the Metaverse on a non-mimetic principle and to describe the complete institutional framework of their legal status de lege ferenda.

The study answers four questions. RQ1: what does the non-mimeticity principle consist of, and what are its doctrinal consequences? RQ2: what is the internal structure of the system of AI legal subjects, and by what criterion are its elements delimited? RQ3: what legal rights may an AI subject hold if those rights are not a copy of human rights, and where do the absolute limits lie? RQ4: what institutional mechanisms make such a system workable, secured and reversible?

1.4. Originality and contribution

The novelty of the study lies in five propositions.

First, the non-mimeticity principle is formulated for the first time as an explicit and system-forming theoretical foundation of legal personality in virtual environments (certain of its components, as noted in subsection 1.2, have doctrinal antecedents), a it is shown that the principle removes, rather than resolves, the ontological objections to conferring status on autonomous systems.

Second, a two-branch taxonomy of the system of AI legal subjects is proposed, in which the status branch is delimited by autonomy class and the representative branch by the principal and the delegated function.

Third, an original catalogue of the legal rights of an AI subject is developed which does not copy human rights, with a four-tier structure and four absolute limits.

Fourth, a three-dimensional registration profile combining subject type, autonomy class and functional profile is introduced as the basis for proportionate regulation.

Fifth, the institutional framework is completed: a reserve ratio with a dual progression, an identity anchor with status passporting, a deferred settlement finality window, and a six-tier liability cascade including a tier for the attribution of will to the principal.

To distinguish clearly between what is new and what is borrowed or adapted, **Table 1** identifies, for each element of the model, its closest doctrinal or regulatory antecedents and what precisely this study proposes. The novelty therefore lies primarily in the particular combination of elements and the architecture of their interaction, rather than in each element taken in isolation.

2. Materials and methods

The study is designed as normative design research: from a diagnosis of the existing position, through a typology of subjects and functions, to the construction of institutional solutions and the verification of their compatibility with the current regulatory environment. The design is qualitative and draws on multiple sources. Accordingly, the results set out in Section 3 are proposed normative models and theoretical mechanisms; they have not been empirically validated in this study.

Doctrinal analysis was applied to four bodies of sources: theories of legal personality and their gradient models [24–27,33–37,50]; the doctrine of virtual property, regulation through code and the governance of decentralised organisations [14–20]; the literature on liability for harm caused by autonomous systems [48,51–53]; and the corpus on the legal personality of the avatar in virtual worlds [30–32]. The method was used to establish the limits within which a new system of subjects may be introduced without dismantling settled institutions.

The functional comparative method was applied to five competing constructs —the avatar as a user's tool; the corporate wrapper for a decentralised organisation [21,22]; "electronic personhood" [23]; the autonomous organisation

without members [39,40]; and the AI avatar as a quasi-subject [30] —against six criteria of suitability: non-mimeticity, functional specificity, cross-world portability, security for claims, reversibility of acts and amenability to supervision.

Derivation and sufficiency of the criteria. The criteria were not selected arbitrarily: they are derived deductively from the diagnosis set out in subsections 1.1–1.2. Each of the three structural defects corresponds to at least one criterion: non-identifiability to functional specificity and amenability to supervision; non-portability to cross-world portability; and lack of security for claims to security for claims. The mimetic trap corresponds to the criterion of non-mimeticity, and the risk that acts become irreversible faster than a human can detect them (subsection 3.9) corresponds to the criterion of reversibility. The set is sufficient relative to the problem as formulated: every defect identified is covered by at least one criterion, and every criterion can be traced back to at least one defect. It makes no claim to absolute completeness: tax and conflict-of-laws aspects are deliberately excluded (see the limitations below). A seventh, auxiliary criterion —suitability for representing a public authority or a media outlet — follows from RQ2 and is used only to assess the representative branch.

Assessment rule. Each construct was assessed on a three-level ordinal scale: "yes / developed" —the construct contains a dedicated institutional mechanism that responds to the criterion; "partly / limited" —a response is provided only by general legal means or only for some cases; "no / not developed" —no mechanism exists. The evidentiary basis comprised the texts of the relevant legal instruments [21–23,41,42,45,54–56] and doctrinal accounts of the constructs [24–32,39,40]. The procedure is summarised in **Table 2**; the results of the comparison are presented in **subsection 3.12**.

Table 2. Methodological procedure of the functional comparison.

Criterion	Legal question	Evidence / source	Assessment rule	Resulting conclusion for the proposed system
Non-mimeticity	Is the subject defined through likeness to a human being?	legal definitions and doctrinal accounts of the constructs [23–32]	"yes" if likeness to a human being is not among the conditions for acquiring status	yes —as a constitutive principle (3.1)
Functional specificity	Does the scope of legal capacity depend on the function performed?	constitutive documents, licensing categories [21,22,42]	"yes" if the scope of legal capacity is linked to defined functions; "partly" if linked only to a constitution	yes —type × class × profile (3.6)
Cross-world portability	Is the status recognised outside the environment of registration?	platform rules, digital identity framework [45]	"yes" if a mechanism exists for recognition of status by other environments	yes —anchor and passporting, with staged implementation (3.8)
Security for claims	Is there segregated assets from which claimants' claims can be satisfied?	rules on property and liability [21,22,49]	"yes" if there is segregated or guaranteed assets	yes —reserve R and guarantee fund (3.7, 3.10)
Reversibility of acts	Can an erroneous act be suspended or reversed before it becomes final?	DORA [54], technical protocols	"yes" if there is a dedicated suspension / challenge mechanism	yes —suspension and finality window (3.9)
Amenability to supervision	Are there a register, an audit and an empowered supervisory body?	EU AI Act [41], Framework Convention [56]	"yes" if register + audit + supervisory body; "declaratory" if merely announced	yes —register, audit, reserve supervision, mandate (3.11)

Source: compiled by the authors.

Conceptual legal modelling is the principal method of the constructive part. Each element of the system is described through four parameters: the legal form in which it is enacted; the body that adopts it; the condition for its entry into force; and its interaction with the other elements. This description allows a distinction

to be drawn between solutions that can be implemented by existing means — platform-world rules, contractual regulation —from those that require legislative intervention.

The operationalisation of technical autonomy features was undertaken by translating the technical properties of agency [57–59] into legal criteria supported by a verification procedure based on the existing standardisation infrastructure [60–64]. The legal form taken by each criterion is a rebuttable presumption of class: a system certified against a defined profile is presumed to belong to the corresponding class, and that presumption is rebutted by the findings of an extraordinary audit. It should be emphasised that these standards are not designed to verify legal autonomy as such: ISO/IEC 42001 sets requirements for a management system, ISO/IEC 23894 and the NIST AI RMF address risk management, and IEEE 7001-2021 addresses transparency. In the proposed model they play an auxiliary, evidentiary role: they determine the format and content of the evidence (decision logs, test records, conformity assessment documentation) examined during certification. The legal criterion itself and the rule for assignment to a class are laid down by statute (subsection 3.3), not by a standard.

The regulatory context was reconstructed from the instruments of the European Union on artificial intelligence [41], markets in crypto-assets [42], operational resilience [54], the pilot regime for market infrastructures based on distributed ledger technology [55] and the European Digital Identity Framework [45], as well as from the Council of Europe Framework Convention on Artificial Intelligence [56] and the legislation of Ukraine [65–67].

Limitations of the methods. First, the Metaverse as an economic and governance environment is still taking shape, so the assessment of the system's effects is prognostic. Second, the quantitative parameters of the model, in particular the reserve ratios, are given as indicative calculations to be calibrated against empirical data from specific platforms. Third, the study does not address tax aspects or questions of conflict of laws, which require separate treatment. The system has deliberately been designed on a modular basis: each element has independent value and does not require the others as a precondition.

3. Results

Preliminary note. All the mechanisms described in this section are proposed normative models (de lege ferenda), not empirically tested or already validated institutions. Wording such as "the system ensures" or "the mechanism removes" should be read as describing the function built into the model, not as stating a result confirmed in practice.

3.1. The non-mimeticity principle as the foundation of the system

3.1.1. The mimeticity axis and its consequences

The first and central result of the study is a distinction between legal constructs according to the criterion of mimeticity, that is, according to whether the new subject is defined through its likeness to a human being. The distribution of constructs along that axis, and the consequences of the choice of pole, are shown in **Figure 1**.

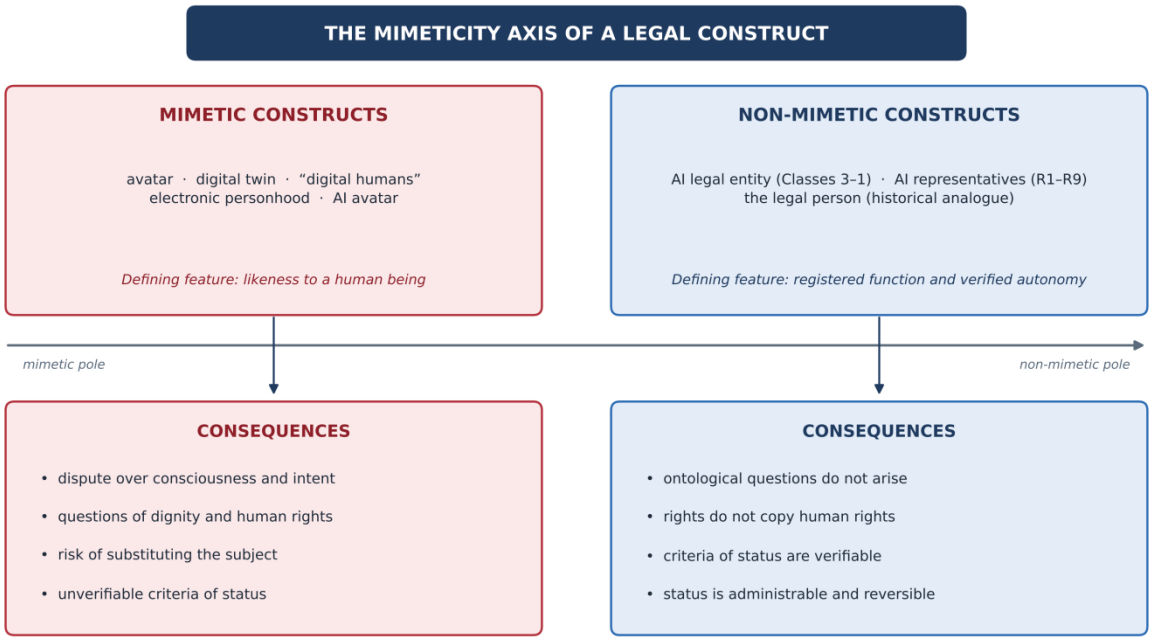


Figure 1. The mimeticity axis of a legal construct and the consequences of the choice of pole.

The mimetic pole brings together the avatar, the digital twin, "digital humans", electronic personhood and the AI avatar. For all of these constructs the defining feature is the reflection of the human being —external (an anthropomorphic image), functional (the performance of human roles) or status-related (a claim to a legal position resembling that of a person). There are four consequences, and all of them are structural rather than incidental.

The first consequence is a dispute over consciousness and intent. If a subject is defined as a likeness of a human being, the question whether it possesses consciousness becomes legally relevant even though there are no legal means of answering it. The scholarly literature has been rehearsing that dispute since at least 1992 [37], with no prospect of resolving it [50].

The second consequence concerns dignity and human rights. A mimetic construct inevitably raises the question whether rights derived from human dignity attach to the new subject, and any answer is problematic: an affirmative one devalues the uniqueness of the human person, a negative one destroys the mimetic premise itself.

The third consequence is the risk of substituting the subject. It is on this that the strongest criticism of the legal personality of autonomous systems rests: the creation of a "synthetic person" serves as a shield behind which those genuinely responsible are removed from liability [24]. Empirical analysis of case law confirms that courts intuitively resist such a construction [33].

The fourth consequence is unverifiable criteria of status. Likeness to a human being has no measurable indicator and therefore cannot be the object of a certification procedure. This deprives mimetic constructs of what a legal status most requires —a procedure for its conferral, verification and revocation.

3.1.2. The non-mimetic alternative and its historical precedent

The non-mimetic pole defines a subject differently: through a registered function and a verified level of autonomy. The feature of likeness to a human being is excluded from the legal definition of the subject altogether. The question of consciousness does not arise, not because an answer has been given to it, but because it does not form part of the conditions for acquiring status.

It should be noted that this construction is not a novelty for law as such —it

is a novelty only for virtual environments. The legal person, the most successful non-natural subject in legal history, was never a copy of a human being: it has no body, no consciousness and no intent in the psychological sense, and it is precisely for that reason that theory moved relatively early from the search for its "will" to a description of its organisational and proprietary structure and of the mechanism for attributing acts to it. The classical doctrine of the legal person, the cluster theory of legal personality [25] and conceptual frameworks for its divisibility [26] are in this sense non-mimetic by construction. The system proposed here continues that line rather than the line of the avatar.

Hence the proposition that determines the whole structure of the study: an AI legal subject is not a digital duplicate of a human being —it is a new type of non-natural subject, placed in the same series as the legal person and differing from it in the ground on which it arises and in the scope of its legal capacity. This proposition is consistent with the principle of non-simulacrity previously formulated within the doctrinal framework of a digital code [6,38] and gives that principle a concrete content in terms of subjects.

3.1.3. What changes in practice

The shift from a mimetic to a non-mimetic framework has four practical consequences, and these are given effect in the subsections that follow. The condition of status becomes the certification of autonomy rather than an assessment of "human-likeness" (3.3). The rights of an AI subject are formed from an original catalogue rather than by paring down human rights (3.4). Representation is constructed as the delegation of a function rather than as the replacement of an official or a journalist (3.5). Finally, the scope of legal capacity is determined by the registration profile rather than by analogy with the legal position of a human being (3.6).

3.2. The two-branch taxonomy of the system of AI legal subjects

3.2.1. Distinguishing status from function

The system of AI legal subjects is the totality of the statutorily provided types of autonomous software and organisational entity endowed with the capacity to perform acts having legal effect in virtual and informational environments within a registered profile, subject to a verified level of autonomy, deposited asset backing and the technical reversibility of their acts.

The term "AI legal subject" is used throughout as a generic designation of entities entered in the register; the degree of their legal personality is graduated, from a regulatory status (Class 3 and AI representatives) to a special legal personality (Classes 2 and 1) (see 4.2.1).

The system has two branches, delimited by the criterion on which each is identified. The architecture is set out in **Figure 2**.

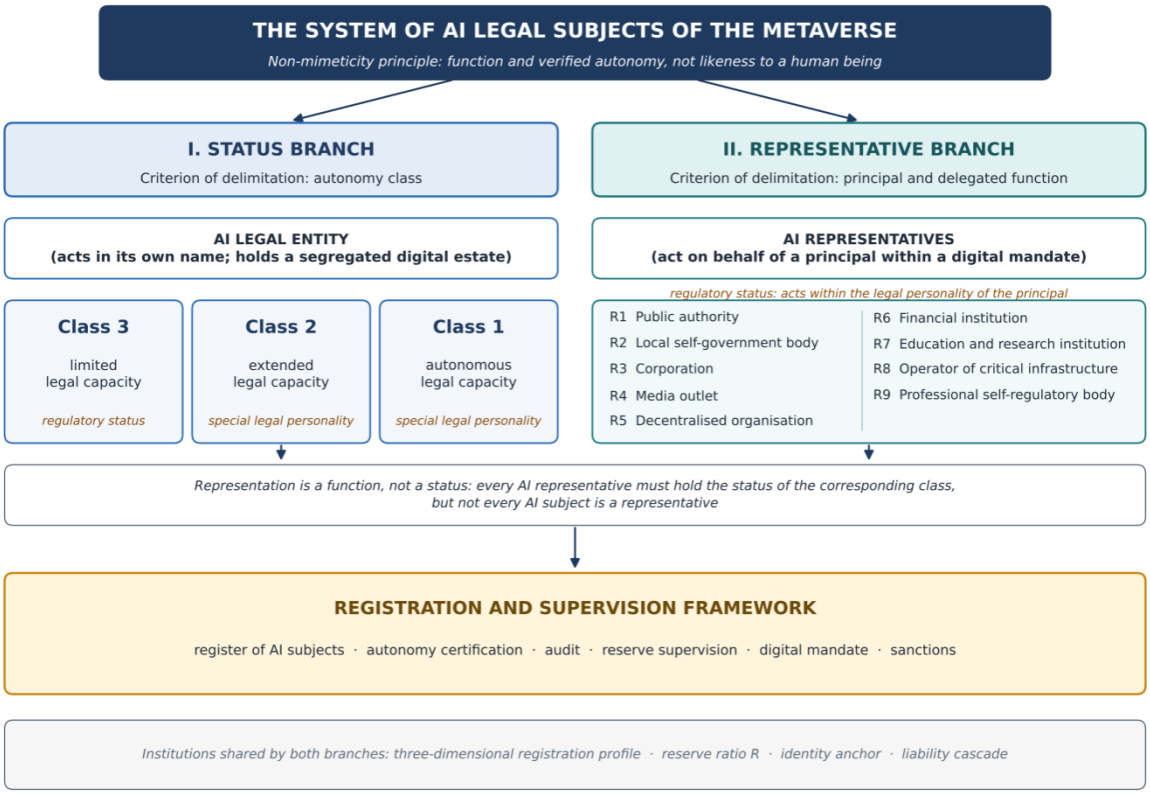


Figure 2. The architecture of the system of AI legal subjects of the Metaverse.

The status branch is delimited by autonomy class. Its single type is the AI legal entity —a subject that acts in its own name within a registered profile and holds a segregated digital estate. The autonomy class (3, 2 or 1) determines the scope of its legal capacity and its liability regime.

The representative branch is delimited by the principal and the delegated function. Its types are AI representatives —subjects that act on behalf of a defined principal within a digital mandate, with attribution of will to that principal. The nine types of AI representative are examined in subsection 3.5.

The relationship between the branches may be stated as follows: not every AI subject is a representative, but every AI representative must hold the status of the corresponding class. Representation is a function, not a status; it is superimposed on a classified subject and does not replace it. This distinction removes the conflation of concepts characteristic of the literature on "algorithmic agents", where status and function are described by a single term.

3.2.2. Why two branches are required

A single-branch system would be incomplete in either of its possible forms. Were the system to consist of the status branch alone, it would not cover the cases in which an autonomous system acts not in its own name but on behalf of a public authority, a news organisation or a corporation —and it is precisely those cases that are the most sensitive from the standpoint of the public interest. Were the system to consist of the representative branch alone, it would not cover autonomous entities that have no principal, in particular organisations without members [39,40], for which there is no person whose will could be the object of attribution.

Two branches cover both classes of case while drawing on shared institutions: a single register, a single procedure for certifying autonomy, a single reserve ratio and a single liability cascade. This makes the system coherent rather than a collection of separate regimes.

3.3. The status branch: the AI legal entity and the autonomy classes

The autonomy class is established by a certification procedure on the basis of five legally relevant criteria. Each criterion derives from a technical property of agency [57–59] but is framed in legal terms and supported by a verification instrument based on existing standards [60–64]. The criteria and their distribution across the classes are set out in **Table 3**.

Table 3. Operationalisation of the features of autonomy and the classes of the AI legal entity.

Technical property	Legal criterion	Verification instrument	Class 3 —limited	Class 2 —extended	Class 1 —autonomous
Independence in goal-setting	closure of the decision loop	logging of decisions under ISO/IEC 42001; NIST AI RMF (Map/Measure)	decisions only within the operator's instructions	independent decisions in a defined field	independent goal-setting within the registered purpose
Direct action in open environments	capacity to conclude transactions independently	certification testing of expressions of will; transaction protocols	does not transact in its own name	transactions within an approved list	transactions without a list, subject to statutory exclusions
Memory; organisational unity of the architecture	severability of digital assets	asset accounting; distributed register of property	assets are not attached	limited segregated estate	full estate with a public register of changes
Metacognition; explainability	auditability of behaviour	IEEE 7001-2021; event logs; post-deployment monitoring	basic logging	mandatory annual audit	continuous monitoring and extraordinary audits
Interaction with a human controller	controllability (human oversight, Article 14 of the AI Act)	conformity assessment; ISO/IEC 23894	continuous operational control	supervisory control with a right of intervention	ex post control; mandatory immediate suspension mechanism

Source: compiled by the authors on the basis of [41,57–64]. Assignment rule: the resulting class is the highest class number obtained on any criterion (weakest link); see the rule below the table.

The criterion of severability of digital assets draws on the concept of the digital-specific property of autonomous systems [68], adapted to the conditions of virtual environments. The legal form taken by all five criteria is a rebuttable presumption of class. A system certified against a defined profile is presumed to belong to the corresponding class; that presumption is rebutted by the findings of an extraordinary audit initiated by the regulator, the operator or an interested party. This mechanism relieves the court of the need for ad hoc expert technical evidence in every case and transfers the question of fact into the certification sphere —precisely as the law of technical regulation does in relation to product safety.

Rule for assignment to a class. Classes are numbered inversely to the level of autonomy: Class 1 denotes the highest and Class 3 the lowest level. The model provides for neither weighting nor scoring: all five criteria are mandatory and cumulative, and the class is determined by the "weakest link" principle in three steps. Step 1: for each of the five criteria, the certification body establishes the highest level of autonomy whose requirements are met (Class 1, 2 or 3). Step 2: the system is assigned the least autonomous of the five levels so established, that is, the highest class number obtained on any criterion: Class $= \max(c_1, \ldots, c_5)$, where c_i is the class number established on criterion i. Step 3: the blocking rules set out below are applied. Example 1: a system meets the Class 1 requirements on four criteria and the Class 2 requirements on auditability; it is assigned Class 2. Example 2: a system meets Class 1 on independence in goal-setting and on direct action, and Class 2 on auditability and controllability, but its digital assets are not attached (the Class 3 level on severability); however autonomous it may be technically, it is assigned Class 3, since the severability criterion —reinforced by the blocking rule on a segregated estate —places it at the lowest level. Weighting

has been deliberately rejected, since it would allow a deficit in controllability to be offset by a surplus of other properties. Two criteria are decisive in the sense of having a blocking effect: the absence of a functioning immediate suspension mechanism precludes Class 1, and the absence of a segregated estate precludes Classes 2 and 1, irrespective of all other indicators. Where the evidence on a criterion is mixed or inconclusive, the criterion is deemed to be met at the lower level (in dubio pro cautela), and the burden of proving the higher class rests on the operator. The grounds for re-certification and downgrade are set out in subsection 3.11 .

It is important to emphasise that what is at issue is operational rather than cognitive autonomy. The law does not assess whether a system possesses consciousness or intentions; the sole criterion is its ability to perform legally relevant functions in a stable and predictable manner without continuous human intervention. This position follows directly from the non-mimeticity principle (3.1) and is at the same time a response to gradient models of legal personality [25–27]: the divisibility of status is accepted, but the scale is grounded in verified autonomy rather than on the degree of approximation to a human being.

The liability regime is differentiated by class. For class 3, harm is compensated by the operator irrespective of whether it resulted from technical error, faulty model training or an external factor. For class 2, the primary source of satisfaction of claims is the subject's segregated estate, with the operator liable on a subsidiary basis. For class 1, the subject's own liability, satisfied from its own assets, is primary, and the operator is liable only in exceptional cases expressly defined by statute: covert intervention, circumvention of audit procedures, manipulation of algorithms, or use of the system for unlawful activity.

Class 1 is an exceptional status. Its widespread conferral is undesirable, since most systems can function effectively within classes 2 and 3, and it is precisely the restriction of its scope of application that secures the proportionality of regulation.

3.4. The formation of the legal rights of an AI subject

3.4.1. The logic of building the catalogue

This is the central subsection of the study from the standpoint of the non-mimeticity principle. If an AI subject is not a duplicate of a human being, its rights cannot be formed by paring down the catalogue of human rights —that is, by removing electoral, family and personal non-property rights and retaining "the remainder". Such a method is mimetic by construction and reproduces all the consequences described in 3.1.

The catalogue of the rights of an AI subject is built in the opposite way: from functional necessity. A right is included in the catalogue if and only if, without it, the subject cannot perform its registered function or cannot defend its own status in the procedures provided by law. The catalogue built in this way is set out in **Figure 3**.

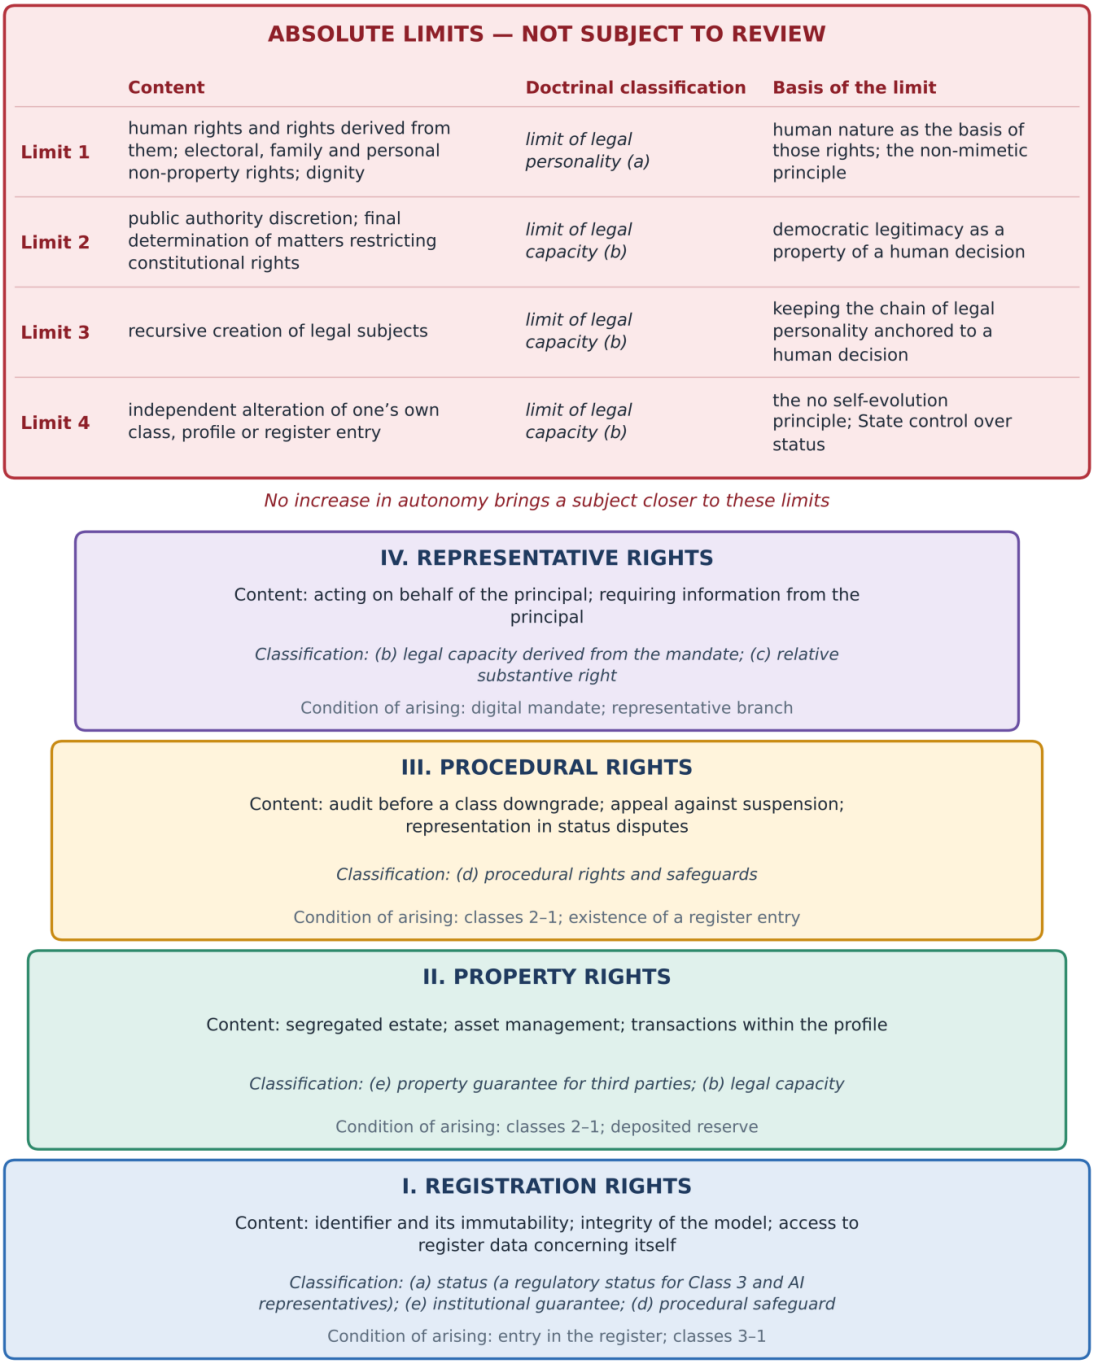


Figure 3. The four-tier formation of the legal rights of an AI subject and the absolute limits (the content, doctrinal classification and conditions reproduce Table 4).

3.4.2. The four tiers of rights

Tier I —registration rights. This is the minimum set, arising from the moment of entry in the register: the right to a digital identifier and to its immutability throughout the life cycle of the status; the right to the integrity of the registered model, that is, to the requirement that changes to the architecture affecting the class be made only in the prescribed manner; and the right of access to the register data concerning itself. The last of these has practical significance: without it, the subject cannot detect a divergence between its actual and its registered profile and, accordingly, cannot initiate the correction of that divergence.

Tier II —property rights. These arise from class 2: the right to hold a segregated digital estate; the right to manage it within the registration profile; and

the right to conclude transactions in a defined field. These rights are not a benefit conferred on the subject but a condition of the security of claimants' claims: the estate exists first and foremost so that compensation may be obtained from it.

Tier III —procedural rights. These arise together with Tier II and constitute the most significant part of the catalogue: the right to an audit procedure before a downgrade of class; the right to appeal against the suspension of status; and the right to procedural representation in disputes concerning its own status. These rights are dual in nature. Formally they belong to the subject; functionally they serve legal certainty in commerce, since they make it impossible to alter a counterparty's status arbitrarily without procedure. That is why their presence is a requirement not of humanitarian considerations but of the workability of the system.

Tier IV —representative rights. These arise only in the representative branch: the right to act on behalf of the principal within the digital mandate, and the right to require the principal to supply the information necessary for the performance of the delegated function.

Doctrinal classification of the elements of the catalogue. The term "rights" in the name of the catalogue is used in a broad sense; by their doctrinal nature, its elements fall into five distinct categories: (a) legal personality / status —the identifier and the register entry as the basis of the subject's existence; (b) legal capacity to perform acts, that is, the capacity to acquire rights and obligations by one's own acts —management of assets, conclusion of transactions within the profile, acting on behalf of the principal; (c) substantive rights —in essence, only the right to require information from the principal, which is a relative (obligational) right; (d) procedural rights and safeguards —access to register data, an audit before a class downgrade, appeal against suspension, procedural representation; and (e) property or institutional guarantees created to protect third-party claims —the segregated estate and the rule on the integrity of the registered model. The catalogue thus contains a minimum of substantive rights in the subject's own interest; the great majority of its elements are either elements of legal capacity or guarantees serving the interests of counterparties and of the legal order. This classification is reflected in a separate column of **Table 4**.

3.4.3. The four absolute limits

The catalogue is closed by four limits that are not subject to review within the system and that are consolidated in **Table 4**.

Table 4. The rights of an AI subject and the absolute limits.

Category	Content	Doctrinal classification	Condition of arising / basis of the limit
Registration rights	identifier and its immutability; integrity of the model; access to register data concerning itself	(a) status*(a regulatory status for Class 3 and AI representatives; see 4.2.1)*; (e) institutional guarantee; (d) procedural safeguard	entry in the register; classes 3–1
Property rights	segregated estate; asset management; transactions within the profile	(e) property guarantee for third parties; (b) legal capacity	classes 2–1; deposited reserve
Procedural rights	audit before a class downgrade; appeal against suspension; representation in status disputes	(d) procedural rights and safeguards	classes 2–1; existence of a register entry
Representative rights	acting on behalf of the principal; requiring information from the principal	(b) legal capacity derived from the mandate; (c) relative substantive right	digital mandate; representative branch
Limit 1	human rights and rights derived from them; electoral, family and personal non-property rights; dignity	limit of legal personality (a)	human nature as the basis of those rights; the non-mimetic principle

Continued table:

Category	Content	Doctrinal classification	Condition of arising / basis of the limit
Limit 2	public authority discretion; final determination of matters restricting constitutional rights	limit of legal capacity (b)	democratic legitimacy as a property of a human decision
Limit 3	recursive creation of legal subjects	limit of legal capacity (b)	keeping the chain of legal personality anchored to a human decision
Limit 4	independent alteration of one's own class, profile or register entry	limit of legal capacity (b)	the no self-evolution principle; state control over status

Source: compiled by the authors.

The limits are absolute in the strict sense: they depend neither on the autonomy class, nor on the type of representation, nor on the technical level of the system. No increase in autonomy brings a subject closer to overcoming them, because the basis of the limits is not the insufficient development of the system but the nature of the excluded rights and powers themselves.

3.5. The representative branch: nine types of AI representative

3.5.1. The legal nature of an AI representative

The classical model of representation in every branch of law rests on a shared premise: acts having legal effect on behalf of another person are performed by a human being, whose will is the source of authority and the guarantee of their proper exercise. Digital transformation has destroyed that premise in fact, though not in law.

The legal nature of an AI representative differs from that of a classical representative in four respects. The source of authority is not a power of attorney or an employment contract but a statute, an administrative act or a decision of an organ of the principal, combined with registration. The scope of authority is determined not by personal discretion but by a software-and-normative profile. Every act is subject to automatic logging and audit. Authority may be suspended or revoked instantaneously by the regulator.

The mechanism common to all types is the digital mandate —not a document but a set of legally defined parameters held simultaneously in the register and in the software profile of the system: the boundaries of competence, the permitted types of act, temporal and territorial restrictions, quantitative limits, rules on automatic suspension, and mandatory logging conditions. Acts beyond the mandate produce no attribution of will to the principal. Recording the parameters in two places gives the construct a quality unattainable for a paper authorisation: the limits of authority are not merely declared but technically enforced, and their breach is documented automatically.

Attribution: four distinct categories. The statement that acts within the mandate are attributed to the principal requires doctrinal refinement, because it combines four different legal categories. (1) Authorisation —the act of the principal (or a statute or administrative act) that creates the digital mandate; it is the source of authority, not its exercise. (2) Attribution of will —the rule under which an expression of will by the AI representative within the mandate is deemed to be an expression of will by the principal, in particular for the purposes of interpreting the transaction and assessing defects of consent. What is at issue is the expression of will as a legal fact: the model does not ascribe to the system a will in the psychological sense, and knowledge and intention, where legally relevant, are assessed at the level of the principal (its organs and authorised persons) and of the mandate parameters it has set. (3) Attribution of

legal consequences —rights and obligations under the transaction arise directly for the principal, as provided by the general rule on representation (Article 239 of the Civil Code of Ukraine [65]). (4) Liability —the obligation to compensate harm; vis-à-vis third parties it is borne by the principal as for its own acts, while the final burden is allocated among the principal, the operator and the developer by way of recourse in accordance with Tiers 4–5 of the cascade (3.10).

Allocation of the risk of acts outside the mandate. The consequences of an act outside the mandate are differentiated according to the cause of the excess: (a) if the principal ratifies the act, it is attributed to the principal from the moment it was performed, by analogy with subsequent approval of a transaction concluded in excess of authority (Article 241 of the Civil Code of Ukraine [65]); (b) if the excess results from an error in the configuration of the mandate entered in the register by the principal, the risk is borne by the principal, since it was the principal who set the parameters on which counterparties relied; (c) if the excess results from a technical malfunction or interference, the risk is borne by the operator (Tier 4) and, where there is a design defect in the model, by the developer (Tier 5); (d) if the counterparty knew, or by checking the identity anchor data (3.8) could have known, that authority was lacking, the act produces no contractual effect for the principal and the counterparty cannot rely on the protection afforded to a counterparty acting in good faith; this does not, however, exclude non-contractual liability on the grounds set out below.

Absence of attribution is not absence of liability. The rule that an act outside the mandate is not attributed to the principal concerns only its transactional (contractual) effect: the principal is not bound by the transaction. It does not exclude liability arising on other legal grounds. First, the principal may be liable in tort for its own fault in configuring, supervising or failing to suspend the representative, since the parameters of the mandate form part of its own duty of care. Second, the operator and the developer may be liable on the grounds of Tiers 4 and 5 of the cascade, including liability for defective products [49]. Third, where the principal has obtained a benefit from an unattributed act, it is obliged to return it under the rules on unjust enrichment (Article 1212 of the Civil Code of Ukraine [65]). Fourth, the function that classical law assigns to the personal liability towards a third party of an agent acting without authority is performed, in the case of an AI representative, by the cascade: the counterparty that has suffered loss is directed to Tiers 1–5 rather than left without a debtor. Non-attribution therefore determines who is bound by the transaction, not whether anyone is answerable for the harm it causes.

Why a distinct status if consequences are attributed to the principal. It may be objected that, if the legal consequences of acts within the mandate arise for the principal in any event, attribution to an existing natural or legal person under the “tool” construct would achieve the same aim. For some cases this is true, and the model does not require every system acting on behalf of a principal to be registered. The “tool” construct, however, fails to provide four elements for the sake of which the status of AI representative is introduced. First, it does not distinguish between acts within and outside authority: the acts of a tool are either always the acts of the user (over-attribution) or nobody’s acts. Second, it gives third parties no means of verifying the limits of authority before a transaction is concluded; the registered mandate and the identity anchor make those limits machine-verifiable. Third, it gives the regulator no distinct addressee for suspension: the status enables the suspension of a particular representative without suspending the activity of the principal itself (for example, a public authority). Fourth, it does not secure recognition of the authority in other worlds. The status of AI representative is therefore created not to shift liability away from the principal —Tier 0 of the cascade expressly prevents this —but to make the

delegated function identifiable, verifiable and controllable.

3.5.2. The typology

The nine types of AI representative are set out in **Table 5**. The typology is structured around the principal; for each type, the delegated scope, the negative list and the specific features of attribution are defined.

Table 5. Types of AI representative: principal, delegated scope and limits.

Type	Principal	Delegated scope	Negative list (never delegated)	Attribution and liability
R1	Public authority	processing of applications, preparation of draft decisions, automated administrative services within non-discretionary competence	final decisions restricting constitutional rights; discretionary acts; sanctions	will attributed to the authority; mandatory right to human review of the decision
R2	Local authority	municipal e-services, residents' requests, administration of local registers	disposal of municipal property; town-planning discretion	will attributed to the body; oversight by the council
R3	Corporation	electronic contracts, compliance, financial monitoring, corporate record-keeping	decisions of the general meeting; approval of significant transactions	will attributed to the company within the powers matrix
R4	Media outlet	selection, aggregation, preparation and dissemination of content	editorial policy; decisions to publish material concerning private individuals without human approval	will attributed to the news organisation; mandatory labelling of machine origin
R5	Decentralised organisation	execution of protocol decisions, representation before registers, administration of assets	amendment of the protocol; issuance outside a members' decision	will attributed to the organisation; liability of the operator and identified members
R6	Financial institution	scoring, underwriting, execution of client instructions, reporting	final refusal of access to a financial service without review	will attributed to the institution; increased reserve
R7	Education and research institution	administration of the educational process, academic integrity checks, research infrastructure	final assessment of learning outcomes; findings of academic plagiarism	will attributed to the institution; mandatory human review
R8	Operator of critical infrastructure	monitoring, dispatching, incident response within the regulations	decisions to disconnect critical facilities; emergency regimes	will attributed to the operator; operational resilience requirements
R9	Professional self-regulatory body	maintenance of registers, administration of admission procedures, disciplinary record-keeping	disciplinary decisions on the merits; admission to the profession	will attributed to the body; procedural guarantees for the individual

Source: compiled by the authors. The criterion of permissible delegation is the principle of functional admissibility: only powers that rest on clearly defined normative criteria, do not call for broad discretion, can be fully logged and remain subject to subsequent human review may be delegated.

3.5.3. Two types requiring separate comment

The AI representative of a public authority (R1) is the most sensitive type, since it comes into contact with the right of an individual not to be subject to a decision taken solely by automated means. The construct resolves that conflict not by an exception to the right but through the scope of delegation: R1's competence extends only to acts culminating in a draft decision, or in a decision within non-discretionary competence with a mandatory opportunity for human review. The negative list should be enacted expressly in statute rather than derived by interpretation.

The AI representative of a media outlet (R4) is the type in which the difference between a duplicate of a human being and an autonomous subject emerges most sharply. A generative system that prepares material is not a "digital journalist" and does not acquire journalistic rights derived from the human

freedom of expression. It is a representative of the news organisation, acts within a mandate, and its output is subject to mandatory labelling of machine origin. Liability for the information disseminated rests with the publisher as principal; the system is neither an author for copyright purposes nor a bearer of editorial independence. This approach preserves the institution of press freedom while giving legal form to a practice that already exists in fact.

3.6. The three-dimensional registration profile

The scope of an AI subject's legal capacity is not a characteristic of the subject but a function of three parameters: the type of subject (status or representative branch), the autonomy class (3, 2 or 1) and the functional profile (F1–F8). The registration profile is an ordered triple of those parameters and determines the regime of admission to each operation.

The third dimension —the functional profile —is required because the social weight of the consequences of a failure depends not only on the independence of the system but also on the content of its activity. An autonomous system that sells virtual clothing and an autonomous system that issues an in-world unit of account have the same technical level of independence and a radically different social risk. The eight profiles are grouped into three risk strata: transactional (F1, F8), financial (F2–F5) and fiduciary-informational (F6, F7). The admission matrix is set out in **Figure 4**.

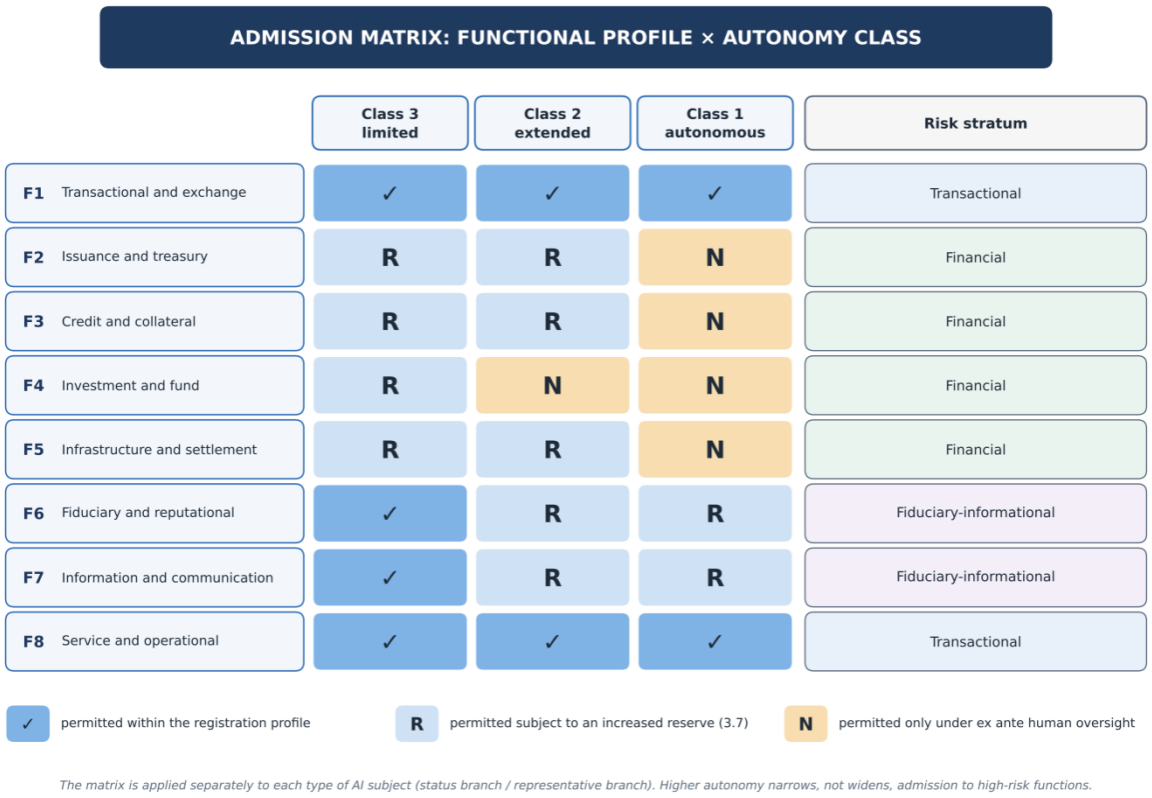


Figure 4. The admission matrix of functional profile against autonomy class, applied separately to each type of AI subject.

The matrix contains three regimes of admission. The "□" regime means that the function is performed within the registration profile without additional conditions. The "R" regime means admission subject to an increased reserve (3.7). The "N" regime means that the function is performed only under ex ante human oversight: every act having legal effect requires confirmation by an authorised natural person before legal consequences arise.

The content of the admission regimes and the consequences of breaching

their conditions are set out in **Table 6**.

Table 6. Admission regimes under the registration profile.

Regime	Condition of admission	Scope of acts	Consequence of breaching the condition
□ —unrestricted	registration of the profile; base reserve	full range of operations within the profile	administrative liability of the operator; warning
R —secured	increased reserve ratio; semi-annual audit	full range of operations; cap on the individual operation	automatic suspension of the function until the reserve is replenished
N —supervised	confirmation by an authorised person before consequences arise	operations only after confirmation	an act without confirmation is void
× —prohibited	—	the function is unavailable for that class	the act is classified as ultra vires on the part of the operator

Source: compiled by the authors.

The logic of distribution is counter-intuitive: the higher the autonomy class, the narrower, not the wider, the admission to high-risk functions becomes. The reason is that fiduciary and systemic functions presuppose the trust of third parties, and the extent of that trust cannot grow in proportion to the technical independence of the performer without a corresponding growth in security and supervision.

Acting outside the registered profile is an ultra vires act: a transaction or act performed outside the profile creates no obligations for the AI subject and is classified as an act of the operator, with the corresponding consequences. The typology of profiles deliberately describes functions rather than objects: the legal regime of in-world assets themselves —from tokenised items to non-fungible tokens —remains the subject of a separate debate [69,70] and is not revisited by the system. The technical guarantee of the boundaries is the requirement that any excess over the profile be recorded automatically by audit tools: ultra vires becomes not merely a legal characterisation after the event but an observable log entry.

3.7. Financial backing of status: the reserve ratio

The principal objection to any legal personality for autonomous systems is that it becomes a shield: harm is caused by a subject without assets, while the people and corporations behind it are removed from liability [24,33,48]. The system's answer is the normative reserve ratio R —a mandatory deposited reserve that is a condition of acquiring and maintaining status.

The construct borrows the logic of prudential regulation [43,44], adapting it to the conditions of virtual environments. The reserve is held in liquid assets recognised by the regulator, accounted for separately from the operator's assets, and may not be used for operational activity. The size of R is determined by the pair of functional profile and autonomy class and is expressed as a share of the annual turnover of the function concerned. Indicative figures are shown in **Figure 5**.

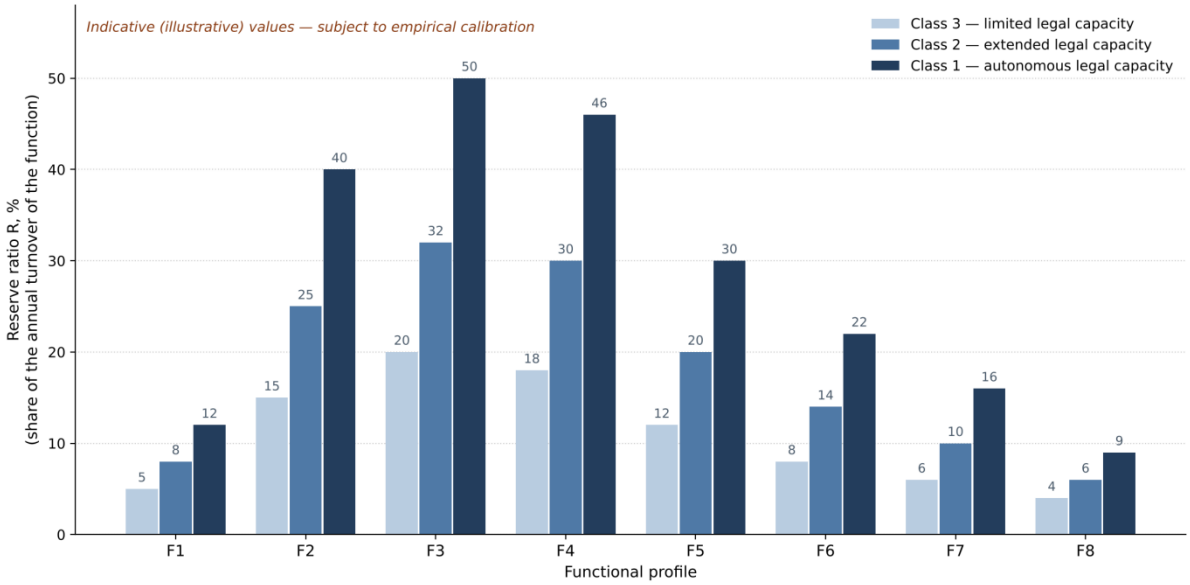


Figure 5. The normative reserve ratio by functional profile and autonomy class (indicative, illustrative values subject to empirical calibration).

Status and rationale of the values shown. The percentages in **Figure 5** are not empirically established. They are illustrative values that demonstrate the structure of the model (orders of magnitude and the direction of progression), not ready-made regulatory parameters. They were constructed according to the following logic: the lower end of the scale for the transactional stratum is comparable in order of magnitude to minimum capital adequacy requirements in banking regulation (around 8–10.5 per cent of risk-weighted assets); the higher values for the financial stratum reflect the absence in virtual environments of such risk mitigants as a lender of last resort, deposit guarantee schemes and prudential supervision of counterparties; and the growth factor between classes (approximately 1.5–1.6 for each step of class) expresses the underlying rationale of progression by autonomy. Calibration of the specific values requires data on the turnover and loss experience of particular platforms and constitutes a separate direction for further research (4.7).

The curve has two properties, each of which constitutes a distinct component of the proposed model.

Progression by autonomy. For one and the same functional profile, the ratio climbs as the class ascends. This creates an economic incentive running counter to the prevailing one: at present, an operator stands to benefit from maximising autonomy, since doing so lowers the cost of human oversight and carries no attendant obligations. Under the proposed model, greater autonomy becomes a costly proposition, and the operator calibrates the level of independence to genuine need rather than mere technical feasibility.

Progression by social risk. Within one and the same class, the ratio rises with the risk stratum of the profile: from 4–12 per cent for the service-and-operational and transactional-and-exchange profiles to 18–50 per cent for the credit-and-collateral and investment-and-fund profiles. This transposes into the law of the Metaverse the principle of proportionality of regulatory burden, tried and tested in financial regulation [43,44,71].

The reserve serves a threefold purpose: it is a source of redress for claimants (Tier 1 of the liability cascade, 3.10); it operates as a barrier to entry, deterring non-serious participants from high-risk functions; and it acts as an instrument of continuous supervision, since any shortfall below the required level automatically suspends the function.

3.8. Portability of status: the identity anchor and passporting

An autonomous subject is, by its very nature, multi-world: it interacts with several platforms at once, each with its own rules, its own identification regime and its own conception of a participant's status. A corporate wrapper [21,22] tethers the subject to a single jurisdiction; platform rules hem it into a single world. The upshot is that a counterparty in world B is in no position to verify who the subject registered in world A actually is.

The system offers a two-element solution.

The identity anchor is a unique, cryptographically verifiable identifier that ties the AI subject to the register entry and remains unchanged throughout the life cycle of the status. The anchor carries a minimum set of particulars available for machine verification in any world: the identifier, the type of subject and its branch, the autonomy class, the functional profile, the parameters of the digital mandate for representatives, the date of the last audit, the state of the reserve and details of any restrictions in force. Its technical basis is the digital identity infrastructure being built within the European framework [45] and the interoperability standards for virtual environments [72]. The immediate conceptual precursor of the anchor is the model of a cross-border personalised ID code for avatars and electronic personalities [28], extended here from the identification of a person to the identification of a legal status.

Status passporting is a mechanism for the recognition of status by host worlds, modelled on the single licence with notification of the host State that has been tested in financial services law [46]. An AI subject registered in one legal order acquires the right to act in another world upon notification of the administrator of that world, without re-registration. The host world retains the right to impose additional requirements but not the right to re-characterise the status.

The proposed contribution lies in transposing into the law of virtual environments an institution that has hitherto existed only in inter-State relations, and adapting it to relations between a legal order and a private platform. The recognising entity is not a State but the administrator of a world; the basis of recognition is not an international treaty but participation in a register that the administrator has joined voluntarily. Such a design is consistent with the conception of digital jurisdiction as a post-territorial legal institution [6,38].

Staged implementation. Passporting raises difficult questions of jurisdiction, recognition, private platform authority, conflicts between host-world rules, revocation and the legal effect of a register entry outside the original jurisdiction. It should therefore be regarded as a staged regulatory proposal rather than an instrument capable of immediate full implementation. The allocation of its elements to stages is set out in **Table 7**.

Table 7. Stages of implementation of the identity anchor and status passporting.

Element	Stage 1: platform rules and contract (implementable now)	Stage 2: national legislation	Stage 3: inter-jurisdictional arrangements
Identity anchor (identifier and machine-verifiable particulars)	technical introduction on the basis of the digital identity infrastructure [45] and interoperability standards [72]	statutory establishment of the register and of the public reliability of its data	mutual recognition of the registers of different States
Recognition of status by a host world	contractual recognition in the terms of service of platforms that have voluntarily joined the register; effect inter partes	recognition made binding on platforms operating within the jurisdiction	recognition of a status registered in another State
Conflicts between host-world rules and the status	the rule "additional requirements —yes, re-characterisation —no" in the accession agreement	statutory enactment of that rule	conflict-of-laws rules on the applicable law (outside the scope of this article)

Continued table:

Element	Stage 1: platform rules and contract (implementable now)	Stage 2: national legislation	Stage 3: inter-jurisdictional arrangements
Revocation and suspension of status	technical propagation through the anchor (real-time update of particulars)	legal force of the regulator's decision vis-à-vis third parties	recognition and enforcement of decisions of foreign regulators
Legal effect of a register entry outside the original jurisdiction	contractual only, within the network of register participants	—	requires an international treaty or a mutual recognition agreement

Source: compiled by the authors.

3.9. Reversibility: protocol-level suspension and the deferred settlement finality window

An autonomous subject that acts faster than human oversight creates a specific risk: an erroneous or unlawful act becomes irreversible before it can be detected.

Protocol-level suspension is a technically embedded ability to suspend the activity of an AI subject by decision of the regulator, a court or the operator, with simultaneous freezing of the segregated assets. The existence of a functioning suspension mechanism is a condition for certification in class 1 and is verified at every audit. This requirement corresponds to the principle of human oversight [41,56] and to the requirements of operational resilience [54].

The deferred settlement finality window is a procedural instrument proposed in this study (its closest antecedent is the institution of settlement finality in payment and securities settlement systems [47], which, however, fixes the moment of finality rather than deferring it). For functions in the financial stratum (F2–F5), a transaction performed by an AI subject becomes final not at the moment of execution but upon the expiry of a short period fixed by regulation, during which it may be challenged by an authorised person, a counterparty or the regulator on an exhaustive list of grounds: acting outside the registered profile or mandate; exceeding an operational limit; a fall in the reserve below the required level; or technically confirmed compromise of the system.

The design of the window is deliberately confined. It does not extend to the transactional and service-and-operational strata, where the cost of delaying settlement exceeds the benefit of reversibility. The period of the window is short and differentiated by profile, the list of grounds is exhaustive, and the burden of proof rests on the challenger. Such constraints are necessary so that the instrument does not become a source of systemic uncertainty: legal certainty in commerce is a value comparable to the ability to correct an error.

Procedural scheme of the window. (1) Standing to challenge belongs to: the principal or the operator as an authorised person; the counterparty to the transaction; and the regulator. No other person has standing. (2) The burden of proof formally rests on the challenger, but in practical terms it is substantially eased: most grounds (acting outside the profile or mandate, exceeding a limit, the state of the reserve) are proved by machine-verifiable log and register entries, so that a dispute is usually reduced to verification of the log. (3) Lodging a challenge within the window automatically suspends finality, and the subject matter of the transaction is blocked in an escrow account. (4) The dispute is resolved by the online dispute resolution body of the platform world, with recourse to the courts, and, in cases concerning regulatory grounds, by the regulator. (5) If the window expires while a challenge is still pending, finality does not occur until a decision is taken, but no later than the expiry of a maximum holding period; thereafter, in the absence of an interim court order, the transaction becomes final and the challenger's remedy is converted from a restitutionary one (unwinding of the

transaction) into a compensatory one (compensation through the cascade, 3.10). The specific periods used in the examples below are indicative.

Example 1. An AI legal entity of Class 2 with profile F3 (credit and collateral) grants a loan secured on a plot of virtual land. Six hours later the counterparty discovers that the loan amount exceeds the operational limit stated in the identity anchor and, within the window (for example, 72 hours), lodges a challenge through the register. The transaction is suspended and the funds and collateral are blocked. The dispute resolution body checks the log and the limit recorded in the register; the excess is confirmed, the transaction is unwound and the parties are restored to their original position, while the operator receives a warning and, upon a repeated breach, is subject to an extraordinary audit.

Example 2. An AI representative of a financial institution (type R6) executes a client's instruction to transfer tokenised assets. Ten hours later the operator detects a compromise of the system and, within the window, lodges a challenge on the ground of technically confirmed compromise. The examination is prolonged by the need for technical expert evidence; the window expires, but finality does not occur pending the decision. If no decision has been taken and no interim court order made before the end of the maximum holding period (for example, 30 days), the transfer becomes final and the loss is compensated through the cascade —from the reserve, the guarantee fund of the world and, in view of the compromise, at the expense of the operator.

3.10. The liability cascade with attribution of will

The liability model has six tiers and is built on the principle of successive exhaustion (**Figure 6**). Its purpose is to guarantee that in any scenario there exists a person from whose assets compensation may be obtained.

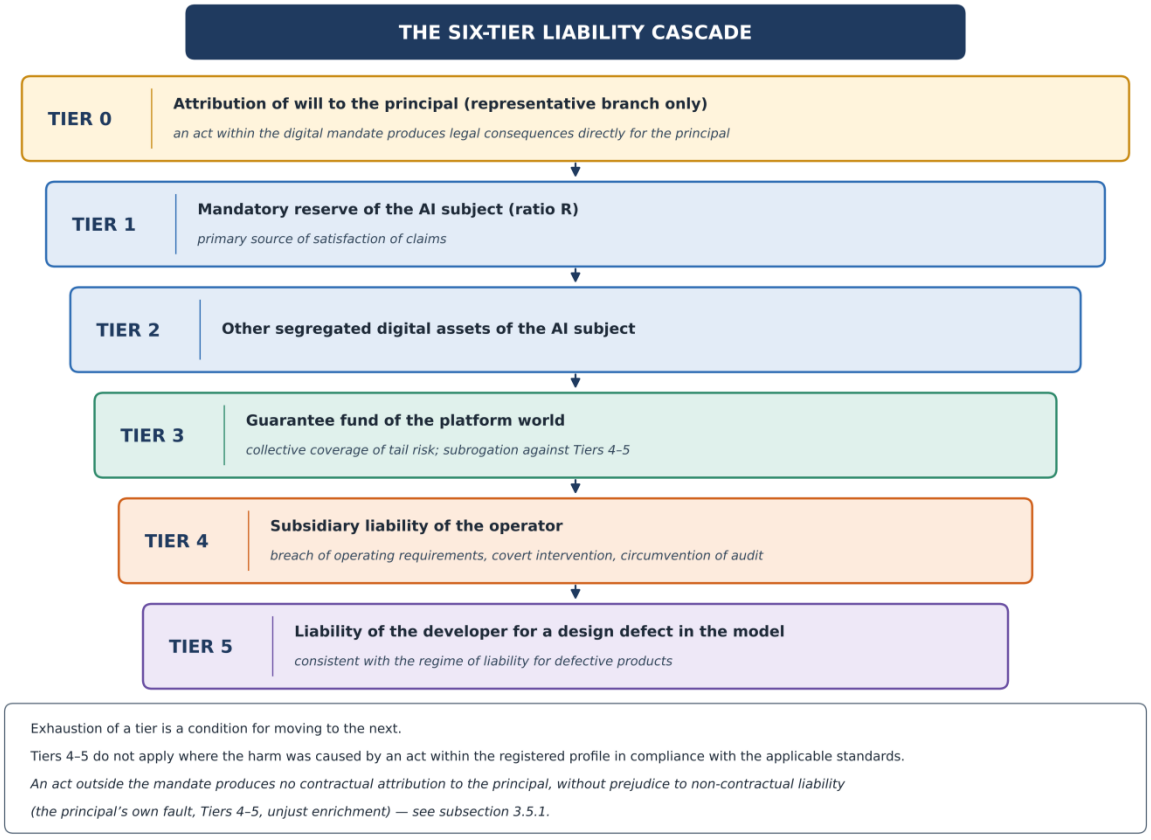


Figure 6. The six-tier liability cascade with a tier for the attribution of will to the principal.

Tier 0 —attribution of will to the principal —applies only in the representative branch and precedes all the others. Where an act has been

performed within the digital mandate, its legal consequences arise directly for the principal —the public authority, the news organisation, the corporation —and the question of the AI representative's own assets do not arise at all. This tier is the most important in legal terms: it ensures that an AI representative does not become a buffer absorbing the principal's liability. Where, by contrast, an act has been performed outside the mandate, attribution does not occur and the subsequent tiers apply (subject to the rules on the allocation of risk set out in subsection 3.5.1)*, without prejudice to the principal's non-contractual liability for its own fault*.

Tier 1 is the mandatory reserve. Tier 2 comprises the subject's other segregated digital assets. Tier 3 is the guarantee fund of the platform world, formed from contributions by all AI subjects admitted to activity in that world and covering harm where individual assets are insufficient. Tier 4 is the subsidiary liability of the operator, which arises where the operator has breached the requirements of operation, intervened covertly or circumvented audit procedures. Tier 5 is the liability of the developer where the harm was caused by a design defect in the model; that tier is consistent with the regime of liability for defective products [49].

The third tier is a distinct component of the proposed model (its antecedents are deposit guarantee and investor compensation schemes, transposed here to the level of the platform world). The guarantee fund of a world resolves the problem encountered by every individual model of security: an individual subject's reserve is calibrated for the ordinary rather than the catastrophic scenario. A collective fund distributes tail risk among the participants in a single environment and at the same time gives the administrator of the world a direct economic incentive to admit only reputable participants.

Governance of the guarantee fund. (1) Management: the fund is administered by the administrator of the world under the supervision of the regulator and with the participation of a council of participants; its assets are accounted for separately from those of the administrator and held by an independent custodian, which excludes them from the administrator's insolvency estate. (2) Funding: annual ex ante contributions, risk-weighted in proportion to the ratio R and the stratum of the functional profile. (3) Replenishment after major losses: a temporary increase in contribution rates and extraordinary ex post contributions subject to an annual cap, so as to avoid cascading insolvencies among participants. (4) Protection against moral hazard: the fund pays out only after the subject's reserve and other assets have been exhausted (Tiers 1–2), and it is subrogated to claims against the operator and the developer where the grounds of Tiers 4–5 are present. (5) Protection against adverse selection: participation in the fund is mandatory for all subjects admitted to the world, with no right to opt out, and the risk-weighting of contributions prevents cross-subsidisation of risky participants by reliable ones. (6) Closure of the platform: a run-off period is established for the submission of claims, after which the remaining assets are transferred to the fund of a successor or, failing one, to a reserve fund attached to the register. These parameters, like the reserve ratios, are indicative.

A crucial reservation concerns the limits of Tiers 4 and 5: they do not apply where the harm was caused by an act performed within the registered profile and in compliance with all the applicable standards. Without that reservation, the system would give the operator no benefit from complying with the rules.

3.11. The life cycle of legal status

The status of an AI subject is dynamic. Retraining of the model, the addition of new modules and changes in data sources can materially alter the actual autonomy profile during operation, so a status determined at the moment of initial

registration soon ceases to correspond to reality. The life cycle is shown in **Figure 7**.

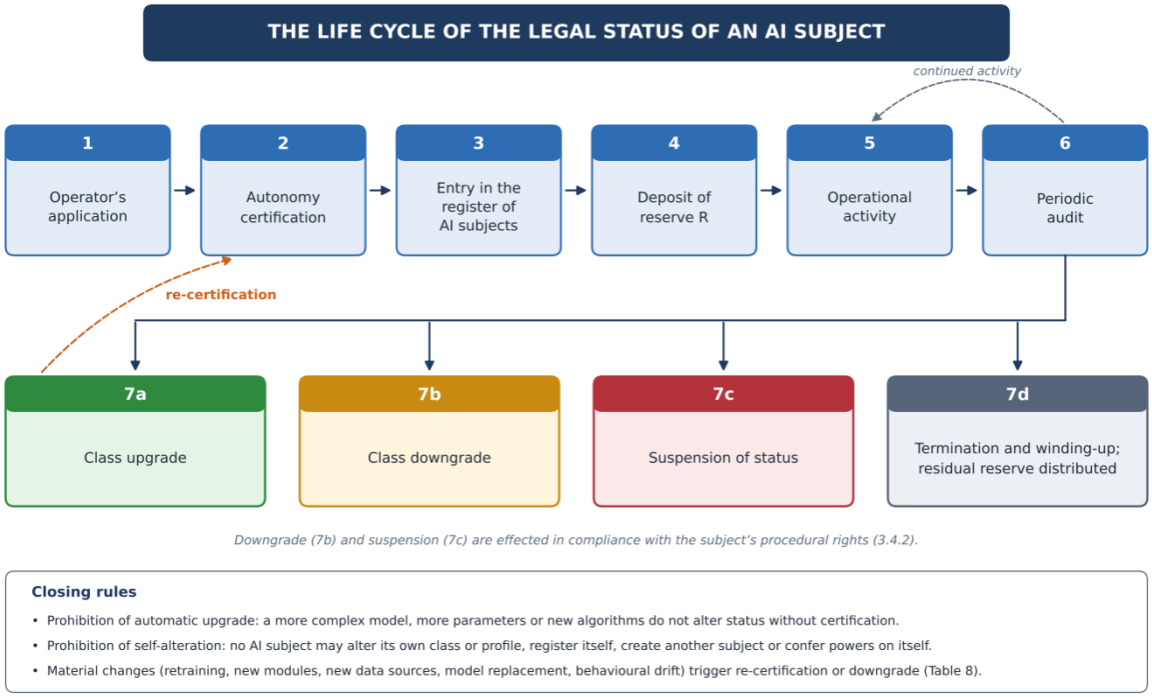


Figure 7. The life cycle of the legal status of an AI subject.

The cycle consists of six principal stages —the operator's application, autonomy certification, entry in the register, deposit of the reserve, operational activity and periodic audit —and four exit routes: upgrade of class, downgrade of class, suspension of status, and termination with winding-up and distribution of the residual reserve. A downgrade of class and a suspension of status are effected in compliance with the subject's procedural rights (3.4.2), which makes an arbitrary alteration of a counterparty's legal position impossible.

Grounds for re-certification and downgrade. To make the rule operational, changes to a system are divided into material and non-material ones. Non-material changes (bug fixes and security patches that do not affect any of the five criteria) require only a notification, which is recorded in the log. Material changes oblige the operator to notify the register before they are introduced and trigger the consequences set out in **Table 8**. Materiality is assessed by the operator, but the assessment may be challenged by the regulator; concealment of a material change constitutes a ground for the operator's liability under Tier 4.

Table 8. Life-cycle events and their consequences for the autonomy class.

Event	Operator's obligation	Legal consequence
Retraining of the model on new data that affects the decision loop	prior notification of the register; targeted audit	the class is retained until the audit is completed; depending on its findings, the class is confirmed or downgraded
Addition of new modules providing new capabilities to act (tools, APIs, payment channels)	prior notification; re-certification against criteria 2 and 5	pending certification, the new capabilities may be used only under regime "N" (ex ante human oversight)
New data sources that change the risk stratum	notification; review of the functional profile	recalculation of the ratio R; where necessary, a change of profile

Continued table:

Event	Operator's obligation	Legal consequence
Replacement of the base model	full re-certification	the identity anchor and the register entry are retained, but pending re-certification the subject operates temporarily under the Class 3 regime
Behavioural drift detected by monitoring beyond the tolerances set in the certificate	immediate notification	extraordinary audit; pending its completion, the most restrictive admission regime; if confirmed, downgrade of class or suspension of status

Source: compiled by the authors.

The construct is closed by two prohibitions.

The prohibition of automatic upgrade. The mere fact that a more complex model is deployed, that the number of parameters has increased or that new algorithms have been connected does not alter legal status. Any upgrade is possible only after completion of the official procedure for assessing autonomy. Without that prohibition, a new type of subject would in effect be created by an ordinary software update, outside any control.

The prohibition of self-alteration. No AI subject may independently alter its own class or profile, register itself, create another AI subject or confer new powers on itself. All such acts are legally void. This prohibition is a direct response to the scenario of recursive self-replication of subjects described in the literature as the limit of the applicability of corporate wrappers [39,40,73].

3.12. Comparison with the existing constructs

The suitability of the system has been assessed theoretically (normatively) against six criteria (and one auxiliary criterion) in comparison with five alternatives, following the procedure described in Section 2 (**Table 2**) (**Table 9**).

Table 9. Functional comparison of constructs of a subject of a virtual environment.

Criterion	Avatar as a tool	Corporate wrapper	Electronic personhood	Autonomous organisation without members	AI avatar as quasi-subject	System of AI legal subjects
Non-mimeticity	no	yes	no	yes	no	yes, as a constitutive principle
Functional specificity	absent	limited by the constitution	absent	limited by the protocol	absent	type × class × profile
Cross-world portability	absent	bound to a jurisdiction	not developed	not developed	not developed	identity anchor and passporting
Security for claims	user's assets	company's assets	undefined	protocol assets; undefined without a wrapper	undefined	reserve R + guarantee fund of the world
Reversibility of acts	general civil remedies	depends on the protocol	undefined	depends on the protocol	undefined	protocol-level suspension + finality window
Amenability to supervision	absent	general corporate	declaratory	absent	absent	register, audit, reserve supervision, mandate
Suitability for representing a public authority or a media outlet (auxiliary criterion)	no	no	no	no	partly	yes, with a negative list for each type

Source: compiled by the authors on the basis of [21–26,30–32,39,40].

The theoretical comparison suggests that the alternative constructs are not mistaken —they are incomplete in respect of precisely those parameters that the

Metaverse environment makes decisive. The avatar as a user's tool accurately describes situations in which every act is traceable to a specific human being, and loses its explanatory force once the act is autonomous. The corporate wrapper effectively solves the problem of the legal personality of a decentralised organisation, but not that of multi-world presence. The autonomous organisation without members [39,40] is non-mimetic by construction but lacks mechanisms of security for claims and of supervision, and therefore depends either on a corporate wrapper or on the proposed status branch. The AI avatar as a quasi-subject [30,31] identifies the practical problem precisely, but remains within the mimetic framework and therefore lacks verifiable criteria of status. The system of AI legal subjects contradicts none of them; it is designed to complete the dimensions without which the construct of a subject in the Metaverse remains incomplete.

4. Discussion

4.1. Why non-mimeticity is prior to every other solution

The dominant strand of scholarship builds the legal personality of autonomous systems on a gradation by degree of approximation to a human being [25–27]. That is rational as a description of the historical movement of ideas but mistaken as a design decision: it makes the central parameter a quantity that has neither a measurable indicator nor a procedure of verification.

The non-mimetic principle changes precisely this. It moves the defining feature of a subject out of the sphere of likeness and into the sphere of function and autonomy —that is, into a sphere where, for every parameter, a verification instrument can be proposed on the basis of existing standards [60–64] (bearing in mind their auxiliary, evidentiary role; see Section 2). All the system's subsequent solutions are derived from that shift: certification instead of an assessment of "human-likeness", an original catalogue of rights instead of a pared-down human one, delegation of a function instead of the substitution of an official. In that sense non-mimeticity is not one of the study's results but the condition that makes the remaining results possible.

4.2. Is the system redundant in relation to existing institutions?

The most serious objection is that the need for new types of subject has not been demonstrated: activity in the Metaverse can be conducted by an ordinary legal person that uses autonomous systems as tools, and a public authority or a news organisation can remain the sole subject of the relevant relations.

That objection is well founded in a substantial proportion of cases and is not contested. The system does not propose to replace the classical subject where it works. It addresses three situations in which it does not.

The first is action without a human decision at the moment of its performance. Where a system takes decisions in milliseconds, the construct "a legal person acts through a tool" retains its legal accuracy but loses its regulatory utility: it supplies neither a point of control nor a criterion of the admissibility of the act.

The second is multi-world presence. The classical subject is bound to the jurisdiction of its registration; its status has no mechanism of recognition in private virtual environments.

The third is the organisation without members [39,40], which does not fit the construct of a legal person built around the will of the founders.

In all three situations the system does not duplicate an existing institution but fills a void. As regards the representative branch specifically, subsection 3.5.1 explains why attribution of consequences to an existing person is not sufficient: the status of AI representative is created not to transfer liability but to make the limits of authority verifiable, to provide a distinct addressee for supervision and

to secure cross-world recognition of the mandate.

4.2.1. Regulatory recognition versus legal personality

Two questions that are easily conflated should be kept apart: (i) whether an AI system should be registered, certified, supervised and capable of being suspended; and (ii) whether it should be recognised as an independent legal subject holding rights and obligations in its own name. The answer to the first question is affirmative for every element of the proposed system, and none of these mechanisms in itself requires legal personality: a register, verification of authority, audit and protocol-level suspension can operate in respect of an object of regulation, as conformity-assessment regimes for products and for AI systems already show [41]. The model therefore confers legal personality only where regulatory recognition alone cannot achieve the regulatory aim. Three cases are decisive. First, segregation of an estate: a segregated digital estate that serves as the primary source of satisfaction of claims (Classes 2 and 1) requires a holder distinct from the operator; without such a holder, the estate either remains the operator's assets, exposed to the operator's creditors, or has no owner at all. Second, the absence of a principal: an organisation without members [39,40] has no existing person to whom its acts could be attributed. Third, acting in one's own name: a subject that concludes transactions in its own name (Class 1, and Class 2 within an approved list) must be capable of being a party to them.

Where none of these conditions is present, the model is satisfied with regulatory recognition. A Class 3 system does not transact in its own name, holds no attached assets and its acts are attributed to the operator, so that in substance it is a registered and supervised object rather than an independent subject; likewise, an AI representative acts within the legal personality of its principal, and its registered status is regulatory rather than an independent legal personality. Legal personality is thus not a precondition of the system but an instrument reserved for the cases in which the alternative —attribution to an existing person combined with regulatory recognition —would leave an estate without a holder, an act without a party or a subject without a principal.

4.3. The risk of regulatory arbitrage

If status entails reserve requirements and audit, operators will move their activity to worlds that impose no such requirements. The risk is genuine, and the system answers it not by prohibition but through an architecture of incentives.

Status passporting (3.8) makes a registered AI subject fit for activity across the network of worlds that have joined the register, whereas an unregistered agent remains confined to a single environment. The guarantee fund of a world (3.10) gives platform administrators a direct financial interest in admitting only reputable participants. The existence of a verified status is a signal to counterparties that reduces the cost of verification and, accordingly, transaction costs. Status is designed as an asset rather than a burden.

4.4. Relationship with the existing regulatory environment

The system was built to be compatible with the existing acquis rather than to displace it. As regards the regulation of artificial intelligence [41,56], the status of an AI subject is not an alternative but a superimposed regime: the requirements applicable to high-risk systems, to human oversight and to documentation are retained in full, and the status adds a subject-related dimension to them. As regards the regulation of markets in crypto-assets [42] and the pilot regime for market infrastructures based on distributed ledger technology [55], the functional profiles F2–F5 have deliberately been described in terms that permit direct comparison with licensing categories. As regards operational resilience [54], the requirement of a functioning suspension mechanism is a specification rather than an innovation. As regards digital identity [45], the identity anchor is an applied

implementation of infrastructure already provided for.

The proposed construct fits into the broader architecture of a digital code, the elements of which have been described previously [38]. For the legislation of Ukraine, implementing the system would require four steps: adding new types of subject holding special legal capacity to the list of persons participating in civil-law relationships [65]; aligning the conceptual apparatus with the legislation on virtual assets [66]; enacting the institution of the AI representative in sectoral legislation —administrative, corporate and information law; and establishing the register and the supervisory authority, which is a logical continuation of the direction of state policy on artificial intelligence [67].

4.5. Constitutional limits

The system establishes four categorical limits (3.4.3), and each has its own justification rather than being a concession to caution.

The exclusion of human rights and of rights derived from them follows from the non-mimetic principle: if a subject is not defined through its likeness to a human being, neither does it have any ground for claiming rights whose basis is human nature. The exclusion of public authority discretion follows from the fact that the democratic legitimacy of a decision is a property of a human choice, not a quality of a procedure. The prohibition of recursive creation maintains the chain of legal personality anchored to a human decision. The prohibition of self-alteration of status preserves state control over the composition of legal subjects.

Notably, none of these limits results from the technical imperfection of the systems, and accordingly no technical improvement weakens them. This is a quality that mimetic constructs cannot in principle possess: within them, a limit always appears provisional, because it is defined by the degree of approximation to a human being.

4.6. Limitations of the study

Five limitations should be noted. First, the quantitative parameters —above all the reserve ratios and the parameters of the guarantee fund —are given as indicative calculations constructed by analogy with prudential regulation [43,44]; their calibration requires empirical data on the turnover and loss experience of specific platforms. Second, the system does not address the tax consequences of the activity of AI subjects or questions of conflict of laws. Third, the effectiveness of passporting depends on the voluntary accession of world administrators to the register, and its full implementation requires inter-jurisdictional arrangements (**Table 7**). Fourth, the assessment of effects is prognostic: the Metaverse environment is still taking shape. Fifth, the proposed mechanisms are normative models that have not undergone empirical or pilot validation within this study; conclusions about their effectiveness are theoretically grounded assumptions.

4.7. Directions for further research

The study opens four directions. The first is the empirical calibration of the reserve ratios against data from operating in-world economies. The second is conflict-of-laws regulation: the determination of the applicable law for the acts of AI subjects in a multi-world environment. The third is the procedural dimension: proving the acts of an autonomous subject, in particular the admissibility of algorithmically processed logs as evidence, which continues the problem of the admissibility of electronic evidence [74]. The fourth is the sectoral elaboration of types R1–R9, primarily the administrative-law construct of the AI representative of a public authority and the information-law construct of the AI representative of a media outlet.

5. Conclusion

The study has proposed a system of AI legal subjects of the Metaverse and described the complete institutional framework of their status as a normative model de lege ferenda.

As regards RQ1. The non-mimeticity principle holds that a subject of law in a virtual environment is defined not by its likeness to a human being but by its registered function and verified level of autonomy. Its principal doctrinal consequence is negative in form and decisive in significance: the ontological questions of consciousness, intent and dignity receive no answer but cease to be legally relevant, because they do not form part of the conditions for acquiring status. The historical precedent for such a design is the legal person, which was never a copy of a human being. Certain components of this approach have antecedents in theories of divisible and partial legal personality; the contribution of the study lies in turning non-mimeticity into an explicit, system-forming principle.

As regards RQ2. The system consists of two branches. The status branch is delimited by autonomy class and contains the AI legal entity of classes 3, 2 and 1. The representative branch is delimited by the principal and the delegated function and contains nine types of AI representative —of a public authority, a local authority, a corporation, a media outlet, a decentralised organisation, a financial institution, an education and research institution, an operator of critical infrastructure and a professional self-regulatory body. Not every AI subject is a representative, but every AI representative must hold the status of the corresponding class.

As regards RQ3. The rights of an AI subject are formed from functional necessity rather than by paring down the catalogue of human rights, and they constitute four tiers: registration, property, procedural and representative rights. By their doctrinal nature, most elements of the catalogue are elements of legal capacity or guarantees for third parties rather than substantive rights in the subject's own interest. The procedural rights are the most significant, since they serve not the interest of the subject but legal certainty in commerce. The catalogue is closed by four absolute limits: human rights and rights derived from them; public authority discretion; the recursive creation of subjects; and the independent alteration of one's own status. No increase in autonomy brings a subject closer to overcoming them.

As regards RQ4. In the proposed model, the workability of the system rests on six institutions: the three-dimensional registration profile of type, class and functional profile; the reserve ratio with its dual progression by autonomy and by social risk (with indicative values requiring calibration); the identity anchor with cross-world status passporting; the deferred settlement finality window; the six-tier liability cascade with a tier for the attribution of will to the principal; and the life cycle of status with its two prohibitions of automatic upgrade and of self-alteration.

The system is designed to answer the three principal objections to the legal personality of autonomous systems. The objection of a shield against liability is addressed by the cascade, in which status does not narrow but widens the asset base available for compensation and, in the representative branch, does not permit the buffering of the principal's liability at all. The objection of an uncontrolled evolution of status is addressed by the two prohibitions and the rules on re-certification. The objection of redundancy is addressed by the justification of three situations in which the classical subject structurally does not work*, and by the distinction between regulatory recognition and legal personality (4.2.1), under which legal personality is conferred only where regulatory recognition alone cannot achieve the regulatory aim*.

The practical value of the results lies in their modularity. The autonomy classes and the functional profiles can be introduced by platform-world rules today. The register, the reserve ratio and the liability cascade require special legislation. The institution of the AI representative requires sectoral amendments, primarily in administrative and information law. Status passporting requires inter-jurisdictional coordination and is the most distant horizon. The proposed mechanisms have not yet been empirically validated; the quantitative parameters require calibration, and tax and conflict-of-laws issues call for separate research.

The Metaverse is developing faster than law is able to create categories of subject for it. The system proposed here is an attempt to narrow that gap not by equipping the human being with a digital twin, but by recognising that law has for eight centuries worked successfully with non-natural subjects, and that a new type of such subject requires not likeness to a human being but a verified function, secured assets and a reversible status.

Author contributions: Conceptualisation, OK; methodology, OK and DZ; formal analysis, OK and VH; investigation, OK, DZ, RT and SK; resources, DZ and VH; writing—original draft preparation, OK; writing—review and editing, OK, DZ, VH, RT and SK; visualisation, OK; supervision, OK; project administration, OK. All authors have read and agreed to the published version of the manuscript.

Funding: This research was carried out under the planned research programme of the State Scientific Institution "Institute of Information, Security and Law of the National Academy of Legal Sciences of Ukraine". No external funding was received.

Acknowledgements: The authors are grateful to colleagues at the Scientific Laboratory of Immersive Technologies and Law for discussion of earlier drafts of the system.

Conflicts of interest: The authors declare no conflicts of interest.

References

[1] Rauschnabel PA, Felix R, Hinsch C, et al. What is XR? Towards a framework for augmented and virtual reality. Computers in Human Behavior. 2022, 133: 107289. https://doi.org/10.1016/j.chb.2022.107289

[2] Mystakidis S. Metaverse. Encyclopedia. 2022, 2(1): 486–497. https://doi.org/10.3390/encyclopedia2010031

[3] Dwivedi YK, Hughes L, Baabdullah AM, et al. Metaverse beyond the hype: Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. International Journal of Information Management. 2022, 66: 102542. https://doi.org/10.1016/j.ijinfomgt.2022.102542

[4] Lee LH, Braud T, Zhou P, et al. All one needs to know about metaverse: A complete survey on technological singularity, virtual ecosystem, and research agenda. arXiv. 2021. https://doi.org/10.48550/arXiv.2110.05352

[5] Kostenko O, Furashev V, Zhuravlov D, et al. Genesis of legal regulation Web and the model of the electronic jurisdiction of the Metaverse. Bratislava Law Review. 2022, 6(2): 21–36. https://doi.org/10.46282/blr.2022.6.2.316

[6] Kostenko O, Zhuravlov D. Reconceptualizing jurisdiction and legal identity in the metaverse: Elements of a digital code architecture. Metaverse Science, Society and Law. 2025, 1(1). https://doi.org/10.69635/mssl.2025.1.1.13

[7] Ball M. The Metaverse: And How It Will Revolutionize Everything. Liveright Publishing; 2022. Available online: https://www.matthewball.co/the-metaverse-book

[8] Hollensen S, Kotler P, Opresnik MO. Metaverse —the new marketing universe. Journal of Business Strategy. 2023, 44(3): 119–125. https://doi.org/10.1108/JBS-01-2022-0014

[9] Belk R, Humayun M, Brouard M. Money, possessions, and ownership in the Metaverse: NFTs, cryptocurrencies, Web3 and Wild Markets. Journal of Business Research. 2022, 153: 198–205. https://doi.org/10.1016/j.jbusres.2022.08.031

[10] Buhalis D, Leung D, Lin M. Metaverse as a disruptive technology revolutionising tourism management and marketing. Tourism Management. 2023, 97: 104724. https://doi.org/10.1016/j.tourman.2023.104724

[11] Chalmers DJ. Reality+: Virtual Worlds and the Problems of Philosophy. W. W. Norton; 2022. Available online: https://consc.net/reality/

[12] Kostenko O. Electronic jurisdiction, metaverse, artificial intelligence, digital personality, digital avatar, neural networks: Theory, practice, perspective. World Science. 2022, 1(73): 1–13. https://doi.org/10.31435/rsglobal_ws/30012022/7751

[13] Kostenko O, Zhuravlov D, Nikitin V, et al. A typical cross-border metaverse model as a counteraction to its fragmentation. Bratislava Law Review. 2024, 8(2): 163–176. https://doi.org/10.46282/blr.2024.8.2.844

[14] Fairfield JAT. Virtual property. Boston University Law Review. 2005, 85: 1047–1102. https://doi.org/10.2139/ssrn.807966

[15] Lastowka FG, Hunter D. The laws of the virtual worlds. California Law Review. 2004, 92(1): 1–73. https://doi.org/10.2307/3481444

[16] Lessig L. Code: Version 2.0. Basic Books; 2006. Available online: http://codev2.cc/

[17] Goldberg M, Schär F. Metaverse governance: An empirical analysis of voting within decentralized autonomous organizations. Journal of Business Research. 2023, 160: 113764. https://doi.org/10.1016/j.jbusres.2023.113764
[18] De Filippi P, Wright A. Blockchain and the Law: The Rule of Code. Harvard University Press; 2018. Available online: https://www.hup.harvard.edu/books/9780674241596
[19] Werbach K. The Blockchain and the New Architecture of Trust. MIT Press; 2018. https://doi.org/10.7551/mitpress/11449.001.0001
[20] Wright A. The rise of decentralized autonomous organizations: Opportunities and challenges. Stanford Journal of Blockchain Law & Policy. 2021, 4(2): 152–176. Available online: https://stanford-jblp.pubpub.org/pub/rise-of-daos
[21] Wyoming Statutes, §§ 17-31-101 —17-31-116 (Decentralized Autonomous Organization Supplement); 2021. Available online: https://www.wyoleg.gov/Legislation/2021/SF0038
[22] Vermont Statutes Annotated, tit. 11, §§ 4171–4176 (Blockchain-Based Limited Liability Companies); 2018. Available online: https://legislature.vermont.gov/statutes/chapter/11/025
[23] European Parliament. Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)). Available online: https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html
[24] Bryson JJ, Diamantis ME, Grant TD. Of, for, and by the people: The legal lacuna of synthetic persons. Artificial Intelligence and Law. 2017, 25(3): 273–291. https://doi.org/10.1007/s10506-017-9214-9
[25] Kurki VAJ. A Theory of Legal Personhood. Oxford University Press; 2019. https://doi.org/10.1093/oso/9780198844037.001.0001
[26] Novelli C, Bongiovanni G, Sartor G. A conceptual framework for legal personality and its application to AI. Jurisprudence. 2022, 13(2): 194–219. https://doi.org/10.1080/20403313.2021.2010936
[27] Mocanu DM. Degrees of AI personhood. Oxford University Press; 2025. https://doi.org/10.1093/9780198945215.003.0002
[28] Kostenko O, Dniprov O, Zhuravlov D. Metaverse: Ensuring legal recognition of avatars and electronic personalities through a cross-border personalized ID-code. International Journal of Innovative Technologies in Social Science. 2024, 2(42). https://doi.org/10.31435/rsglobal_ijitss/30062024/8141
[29] Kostenko O, Zhuravlov D, Dniprov O, et al. Metaverse: Model criminal code. Baltic Journal of Economic Studies. 2023, 9(4): 134–147. https://doi.org/10.30525/2256-0742/2023-9-4-134-147
[30] Cheong BC. The rise of AI avatars: Legal personhood, rights and liabilities in an evolving metaverse. Journal of Digital Technologies and Law. 2024, 2(4): 857–885. https://doi.org/10.21202/jdtl.2024.42
[31] Cheong BC. Avatars in the metaverse: Potential legal issues and remedies. International Cybersecurity Law Review. 2022, 3(2): 467–494. https://doi.org/10.1365/s43439-022-00056-9
[32] Mulyono AT, Sihombing ENAM. Fenomena avatar sebagai subjek hukum di metaverse. Jurnal Ilmiah Penegakan Hukum. 2023, 10(1): 11–20. https://doi.org/10.31289/jiph.v10i1.9091
[33] Banteka N. Artificially intelligent persons. Houston Law Review. 2021, 58(3): 537–596. https://doi.org/10.2139/ssrn.3552269
[34] Teubner G. Digital personhood? The status of autonomous software agents in private law. Ancilla Iuris. 2018: 106–149. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3177096
[35] Beckers A, Teubner G. Three Liability Regimes for Artificial Intelligence: Algorithmic Actants, Hybrids, Crowds. Hart Publishing; 2022. https://doi.org/10.5040/9781509949366
[36] Chopra S, White LF. A Legal Theory for Autonomous Artificial Agents. University of Michigan Press; 2011. https://doi.org/10.3998/mpub.356801
[37] Solum LB. Legal personhood for artificial intelligences. North Carolina Law Review. 1992, 70(4): 1231–1287. Available online: https://scholarship.law.unc.edu/nclr/vol70/iss4/4/
[38] Kostenko O. Digital jurisdiction. Metaverse Science, Society and Law. 2025, 1(2). https://doi.org/10.69635/mssl.2025.1.2.23
[39] Bayern S. Autonomous Organizations. Cambridge University Press; 2021. https://doi.org/10.1017/9781108878203
[40] Bayern S, Burri T, Grant TD, et al. Company law and autonomous systems: A blueprint for lawyers, entrepreneurs, and regulators. Hastings Science and Technology Law Journal. 2017, 9(2): 135–162. Available online: https://repository.uclawsf.edu/cgi/viewcontent.cgi?article=1004&context=hastings_science_technology_law_journal
[41] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. 2024, L, 2024/1689. Available online: https://eur-lex.europa.eu/eli/reg/2024/1689/oj
[42] Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets (MiCA). Official Journal of the European Union. 2023, L 150: 40–205. Available online: https://eur-lex.europa.eu/eli/reg/2023/1114/oj
[43] Basel Committee on Banking Supervision. Prudential treatment of cryptoasset exposures. Bank for International Settlements; 2022. Available online: https://www.bis.org/bcbs/publ/d545.htm
[44] Financial Stability Board. High-level recommendations for the regulation, supervision and oversight of crypto-asset activities and markets; 2023. Available online: https://www.fsb.org/2023/07/high-level-recommendations-for-the-regulation-supervision-and-oversight-of-crypto-asset-activities-and-markets-final-report/
[45] Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework. Official Journal of the European Union. 2024, L, 2024/1183. Available online: https://eur-lex.europa.eu/eli/reg/2024/1183/oj
[46] Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (MiFID II). Official Journal of the European Union. 2014, L 173: 349–496. Available online: https://eur-lex.europa.eu/eli/dir/2014/65/oj

[47] Directive 98/26/EC of the European Parliament and of the Council of 19 May 1998 on settlement finality in payment and securities settlement systems. Official Journal of the European Communities. 1998, L 166: 45–50. Available online: https://eur-lex.europa.eu/eli/dir/1998/26/oj
[48] Wagner G. Robot liability. In: Lohsse S, Schulze R, Staudenmayer D (editors). Liability for Artificial Intelligence and the Internet of Things. Nomos; 2019. pp. 27–62. Available online: https://api.pageplace.de/preview/DT0400.9783845294797_A36625603/preview-9783845294797_A36625603.pdf
[49] Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products. Official Journal of the European Union. 2024, L, 2024/2853. Available online: https://eur-lex.europa.eu/eli/dir/2024/2853/oj
[50] Chesterman S. Artificial intelligence and the limits of legal personality. International & Comparative Law Quarterly. 2020, 69(4): 819–844. https://doi.org/10.1017/S0020589320000366
[51] Karnow CEA. Liability for distributed artificial intelligences. Berkeley Technology Law Journal. 1996, 11(1): 147–204. Available online: https://lawcat.berkeley.edu/record/1115611/files/fulltext.pdf
[52] Čerka P, Grigienė J, Sirbikytė G. Liability for damages caused by artificial intelligence. Computer Law & Security Review. 2015, 31(3): 376–389. https://doi.org/10.1016/j.clsr.2015.03.008
[53] Bathaee Y. The artificial intelligence black box and the failure of intent and causation. Harvard Journal of Law & Technology. 2018, 31(2): 889–938. Available online: https://jolt.law.harvard.edu/assets/articlePDFs/v31/The-Artificial-Intelligence-Black-Box-and-the-Failure-of-Intent-and-Causation-Yavar-Bathaee.pdf
[54] Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA). Official Journal of the European Union. 2022, L 333: 1–79. Available online: https://eur-lex.europa.eu/eli/reg/2022/2554/oj
[55] Regulation (EU) 2022/858 of the European Parliament and of the Council of 30 May 2022 on a pilot regime for market infrastructures based on distributed ledger technology. Official Journal of the European Union. 2022, L 151: 1–33. Available online: https://eur-lex.europa.eu/eli/reg/2022/858/oj
[56] Council of Europe. Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225); 2024. Available online: https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treatynum=225
[57] Ashton H. Definitions of intent suitable for algorithms. Artificial Intelligence and Law. 2023, 31(3): 515–546. https://doi.org/10.1007/s10506-022-09322-x
[58] Kraprayoon J, Williams Z, Fayyaz R. AI Agent Governance: A Field Guide. arXiv. 2025. https://doi.org/10.48550/arXiv.2505.21808
[59] Akshathala S, Adnan B, Ramesh M, et al. Beyond Task Completion: An Assessment Framework for Evaluating Agentic AI Systems. arXiv. 2025. https://doi.org/10.48550/arXiv.2512.12791
[60] ISO/IEC 42001:2023. Information technology —Artificial intelligence —Management system. International Organization for Standardization; 2023. Available online: https://www.iso.org/standard/81230.html
[61] ISO/IEC 23894:2023. Information technology —Artificial intelligence —Guidance on risk management. International Organization for Standardization; 2023. Available online: https://www.iso.org/standard/77304.html
[62] ISO/IEC 22989:2022. Information technology —Artificial intelligence —Artificial intelligence concepts and terminology. International Organization for Standardization; 2022. Available online: https://www.iso.org/standard/74296.html
[63] IEEE Std 7001-2021. IEEE Standard for Transparency of Autonomous Systems. Institute of Electrical and Electronics Engineers; 2021. Available online: https://standards.ieee.org/ieee/7001/6929/
[64] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0); 2023. https://doi.org/10.6028/NIST.AI.100-1
[65] Civil Code of Ukraine No. 435-IV of 16 January 2003 (as amended) [Tsyvilnyi kodeks Ukrainy]. Available online: https://zakon.rada.gov.ua/laws/show/435-15
[66] On Virtual Assets: Law of Ukraine No. 2074-IX of 17 February 2022 [Pro virtualni aktyvy]. Available online: https://zakon.rada.gov.ua/laws/show/2074-20
[67] Concept for the Development of Artificial Intelligence in Ukraine: Order of the Cabinet of Ministers of Ukraine No. 1556-r of 2 December 2020 [Kontseptsiia rozvytku shtuchnoho intelektu v Ukraini]. Available online: https://zakon.rada.gov.ua/laws/show/1556-2020-\TU\textbackslash{}%D1\TU\textbackslash{}%80
[68] Moura Vicente D, Soares Pereira R, Alves Leal A (editors). Legal Aspects of Autonomous Systems: A Comparative Approach. Data Science, Machine Intelligence, and Law. Springer; 2024. https://doi.org/10.1007/978-3-031-47946-5
[69] Allen JG, Rauchs M, Blandin A, et al. Legal and Regulatory Considerations for Digital Assets. Cambridge Centre for Alternative Finance; 2020. https://doi.org/10.2139/ssrn.3712888
[70] Guadamuz A. The treachery of images: Non-fungible tokens and copyright. Journal of Intellectual Property Law & Practice. 2021, 16(12): 1367–1385. https://doi.org/10.1093/jiplp/jpab152
[71] Zetzsche DA, Arner DW, Buckley RP. Decentralized finance. Journal of Financial Regulation. 2020, 6(2): 172–203. https://doi.org/10.1093/jfr/fjaa010
[72] Metaverse Standards Forum. Interoperability domain groups: Status and scope. Available online: https://metaverse-standards.org
[73] Armour J, Eidenmüller H. Self-driving corporations? Harvard Business Law Review. 2020, 10: 87–116. https://doi.org/10.67929/ecgi-law-475-2019
[74] Kostenko O, Akefi Ghaziani V. Admissibility of illegally obtained e-evidence: A critical study of EU law and the precedents of the European Court of Human Rights. European Journal of Privacy Law and Technologies. 2024, (2). https://doi.org/10.2139/ssrn.5050699