Dragi Rashkovski[*]

IS THERE LEGAL READINESS FOR ARTIFICIAL INTELLIGENCE TO BE RECOGNIZED AS A THIRD SUBJECT OF LAW?

Abstract

Artificial intelligence (AI) is no longer a novelty in any respect, and therefore it should not be perceived as a phenomenon whose outcomes we still need to observe before developing a common societal strategy in any domain. The question arises whether the European Union's AI Act is sufficient to regulate it, or whether it represents a rather weak mechanism for that purpose. By its nature, artificial intelligence is a technical phenomenon, similar to many others before it, such as the discovery of electric energy or the internal combustion engine. However, unlike them, AI possesses the property of autonomy, which is not inherent in any other creation of the human mind.

That autonomy, reinforced by its enormous expansion through the active use of artificial intelligence by every modern individual, the gradual creation of dependence on it in everyday activities, and its penetration through self-imposition, cannot be regarded merely as another object or service. Consciously or unconsciously, users of AI develop their own notion of its personalization, perceiving it in their thoughts as an interlocutor or an active collaborator. This tendency arises not only from the possibility of direct communication with AI but also from human nature itself and the theory of anthropomorphism, according to which human characteristics are attributed to objects or abstract concepts.

This entire set of fundamental considerations, which are well justified, leads us to reflect that AI cannot be integrated within the existing legal concepts of material objects. It certainly cannot be classified as a natural person; in some aspects, it encroaches upon the characteristics of a legal entity. However, we consider it most appropriate for AI to be assigned the nomenclature of a person—but what kind?

Some authors of similar studies discuss the introduction of the concept of a "digital person" as a third type of subject of law, yet all of this remains without legal codification that would determine how such a concept could be integrated into the legal system.

The aim of the following paper is to propose a concrete solution, broadly elaborating on how and where this new legal person could be integrated, what legal challenges it may encounter, and what legal amendments would be required for this emerging "subjectivism," which is becoming part of everyday life, to be legally recognized.

Keywords: artificial intelligence, third subject of law, legal challenges, digital person

[*] Dragi Rashkovski, PhD., Associate Professor, Ss. Cyril and Methodius University in Skopje, Iustinianus Primus Faculty of Law, e-mail: rashkovskid@gmail.com, ORCID 0009-0001-0042-7763

1. INTRODUCTION

Within contemporary debates on the legal status of artificial intelligence (AI) systems, the need to reassess the traditional dogmatics of legal subjectivity is becoming increasingly apparent. Although the dominant approach continues to treat AI exclusively as an instrument in the hands of human actors, I contend that this position is gradually becoming insufficient to capture the complexity of modern autonomous systems. Accordingly, I advance the affirmative thesis that certain categories of highly autonomous AI systems should be granted a form of limited legal subjectivity. In doing so, I do not view the solution solely in expanding the concept of the legal person, but also in the careful design of an entirely new form of legal subjectivity. In my view, the optimal normative model should be sought precisely in the interstitial space between these two constructions.

My position is based on the understanding that legal subjectivity is not a natural category tied exclusively to human consciousness, but rather a normative tool constructed by law in response to the functional needs of the legal order. The history of law clearly demonstrates that the notion of the subject of law is evolutionary. Legal persons, despite lacking their own consciousness or moral autonomy, have long been integrated as holders of rights and obligations. In my view, this fact relativizes the claim that the absence of consciousness in AI, in and of itself, constitutes an insurmountable argument against its limited legal personalization.

2. AI – A FORM OF LEGAL PERSON OR AN ENTIRELY NEW LEGAL SUBJECT?

The full incorporation of AI into the classical construction of the legal person encounters serious dogmatic obstacles. The principal objection derives from the organic theory, according to which the will of a legal person is formed through its organs composed of natural persons. Highly autonomous algorithmic systems lack such a classical organic structure. For this reason, I submit that an approach relying solely on the "stretching" of the existing concept of the legal person may ultimately prove conceptually insufficient.

On the other hand, the idea of creating an entirely new, third subject of law, although theoretically appealing, carries significant systemic and political risks. The radical introduction of a new category of subjects of law could generate legal uncertainty, doctrinal resistance, and potential erosion of existing liability regimes. Precisely for this reason, I argue that the appropriate approach should be affirmative yet carefully balanced: rather than adopting a binary choice between "legal person" and "entirely new subject," the more promising path lies in developing a hybrid, gradual construction that combines elements of both solutions.

The principal motivation for this line of reasoning lies in the growing functional autonomy of contemporary AI systems. Unlike traditional automated tools, advanced AI models learn from data, adapt their own decision-making patterns, and operate in dynamic environments characterized by a significant degree of unpredictability. Under such conditions, the classical instrumental model—according to which AI is always fully under human control—gradually loses its descriptive persuasiveness. Although I do not contend that an absolute liability gap currently exists, I am convinced that continued technological development will intensify the pressure on traditional models of attribution of responsibility. In this context, a hybrid liability regime that mirrors the hybrid nature of the proposed subjectivity appears most appropriate. Under such a model, part of the responsibility must remain with human actors—the creators, manufacturers, and

operators of the systems—while at the same time it is normatively justified that a portion of the risk be attached to the AI system itself. This does not imply moral fault on the part of the algorithm, but rather patrimonial, functional responsibility to be realized through dedicated funds, mandatory insurance schemes, or other mechanisms of economic risk coverage.

This approach offers several important advantages. It preserves the protective function of tort law, provides greater legal certainty for injured parties, and simultaneously acknowledges the specific risk profile of highly autonomous AI systems. Moreover, by avoiding a rigid black-and-white categorization, it opens space for a gradual, sector- specific evolution of legal dogmatics.

Naturally, the affirmation of such a model must be accompanied by clear legal safeguards. It is particularly important to prevent the new form of subjectivity from becoming an instrument for evading human responsibility. Accordingly, I maintain that any future normative framework must include mandatory human oversight, subsidiary or joint liability of the relevant human actors, and a limitation of application exclusively to high-risk AI systems.

From this perspective, my position is clearly affirmative: the law should proactively consider granting limited legal subjectivity to artificial intelligence. At the same time, however, I do not see the solution in a simple transposition of existing categories, nor in the radical creation of an entirely new subject, but rather in a carefully designed model situated between an expanded form of legal personhood and a new, specialized form of legal subjectivity. It is precisely within this interstitial space that the most promising direction for the future development of legal theory and AI regulation can be found.

3. POSSIBLE SOLUTIONS OR EMERGING OPEN QUESTIONS

The foregoing theoretical discussion demonstrates that the question of the legal subjectivity of artificial intelligence (AI) cannot be resolved solely at an abstract level. If the affirmative position is to be tested in practice, it is necessary to propose concrete practical models capable of operating within existing legal systems. In this respect, the approach must be distinctly pragmatic: the objective is not the immediate establishment of a full and general legal personality for AI, but rather the design of gradual, controlled, and verifiable mechanisms through which limited legal subjectivity may be tested in practice. In other words, instead of a normative "leap," I advocate a careful legal evolution supported by institutional safeguards.

The first practical solution I consider particularly relevant is the establishment of a registration regime for high-risk autonomous AI systems. Rather than treating every algorithm *prima facie* as a potential legal subject, regulatory focus should be directed toward systems that combine a high degree of autonomy with a significant capacity to cause harm. Qualification criteria could include the level of self-learning, operational independence, field of application, and potential societal risk. Within such a registry, each qualifying AI system would receive a unique legal identifier, enabling procedural recognition, monitoring of its risk profile, and linkage to the relevant liability mechanisms. Moreover, the registry could serve a preventive function, as inclusion in an enhanced oversight regime would itself exert a disciplining effect on operators.

The second practical solution concerns the establishment of a mandatory financial backing associated with registered high-risk AI systems. The most workable model would be a combined approach involving compulsory insurance and a dedicated risk fund. Each operator or entity placing the system on the market of a qualifying AI system would be required to secure a minimum level of financial coverage proportionate to the system's risk profile, with the possibility of dynamically adjusting premiums based on actual performance and incident history. In this way,

where harm occurs that cannot readily be attributed to a specific human fault, the injured party would have access to a direct and prompt source of compensation. In effect, this model operationalizes the notion of the algorithm's patrimonial "responsibility" through a concrete financial architecture.

The third element of the proposed framework is a hybrid liability regime. In practical terms, the most sustainable approach is a model of layered liability. At the first layer, the classical liability of the manufacturer, programmer, or operator remains in place, particularly in cases of defect, design defect, regulatory non-compliance, or breach of supervisory duties. At the second layer, where damage results from complex autonomous behavior that cannot convincingly be reduced to human fault, the financial mechanism attached to the AI system itself would be triggered. At the third, residual layer, the state may provide a guarantee fund for exceptionally complex cases involving a high public interest. Such a three-tier architecture not only protects injured parties but also creates a predictable distribution of risk among all relevant actors.

The fourth practical solution I consider essential is the establishment of an enhanced human oversight regime (human-in-the-loop or human-on-the-loop), particularly for systems falling within the limited subjectivity framework. The affirmation of functional autonomy must not amount to regulatory abdication. On the contrary, every AI system with an elevated degree of autonomy should have a clearly designated supervisor—either a natural or legal person—with defined duties of monitoring, intervention, reporting, and record-keeping obligations. Depending on the risk classification, the legislature could prescribe varying levels of human involvement, ranging from periodic oversight to a mandatory emergency intervention capability (e.g., a kill-switch mechanism). This safeguard is crucial to prevent the new construction from becoming a vehicle for the evasion of responsibility.

The fifth potential solution concerns the procedural position of AI systems. At this stage, I do not consider it necessary for algorithms to be granted full procedural capacity in the classical sense. A more workable approach is a model of limited procedural personalization. Specifically, the AI system may be formally designated as the holder of a particular risk fund and as the object of regulatory measures, while procedural representation would always be exercised by an appointed human or legal representative. This approach enables the functional use of subjectivity without undermining fundamental procedural principles such as adversariality and the right to effective judicial protection.

The sixth important practical measure is the gradual, sector-specific testing of the model. Rather than introducing AI subjectivity horizontally across the entire legal system, it would be far more prudent to begin with pilot regulation in highly technological and high-risk domains—such as autonomous transport, algorithmic financial trading, advanced medical robotics, or critical infrastructure. In these fields, risks are most visible and the regulatory need most acute. An incremental approach would enable the collection of empirical data, assessment of impacts on insurance markets, and timely correction of potential normative weaknesses.

The seventh solution, often underestimated, is the development of standards for algorithmic auditability and evidentiary transparency. For the hybrid liability model to function effectively in judicial practice, courts and regulators must have a genuine ability to reconstruct how a particular AI system reached a given decision. Accordingly, any system falling within the limited subjectivity regime should be subject to enhanced obligations concerning logging of decisions, retention of relevant data, technical documentation, and the possibility of independent audit. Without such an evidentiary infrastructure, the allocation of responsibility would face serious procedural obstacles.

The eighth and final practical solution would be the clear normative delimitation of the scope of AI subjectivity. To avoid conceptual inflation, the legislature should expressly clarify that this is a specialized, functional, and limited form of legal subjectivity that does not entail the automatic recognition of all rights traditionally attributed to classical legal subjects. Such differentiation is essential to preserve systemic coherence and to prevent misinterpretation in judicial practice.

In my view, this expanded package of practical measures—registration, financial backing, layered hybrid liability, enhanced human oversight, limited procedural personalization, sectoral testing, algorithmic auditability, and strict normative delimitation—constitutes a realistic initial framework for operationalizing the idea of limited legal subjectivity of artificial intelligence. This framework does not purport to offer a final solution. On the contrary, its implementation will inevitably generate new complex questions. Yet it is precisely through such carefully designed and incremental steps that the law can proactively adapt to the challenges posed by the growing autonomy of algorithmic systems.

4. EMERGING LEGAL QUESTIONS

Opening the question of the possible legal subjectivity of artificial intelligence (AI) systems inevitably generates a broad spectrum of new legal dilemmas that extend beyond the framework of classical legal dogmatics. Even carefully limited and hybrid models of the kind discussed above do not operate in a normative vacuum. On the contrary, their potential implementation directly touches the core of several legal institutions historically constructed on an anthropocentric premise. Accordingly, if the law is to contemplate seriously any form of AI personalization, it is necessary to systematically identify and reassess the points of friction with the existing legal architecture. The following outlines the key open questions that, in my view, legal scholarship and legislators will need to address.

The first and most fundamental issue concerns the very dogmatics of legal subjectivity. Traditional law defines with relative clarity the conditions under which an entity may be a holder of rights and obligations, with two dominant categories: natural persons and legal persons. The introduction of AI—whether as an expanded form of legal personhood or as a new specialized subject—calls this binary model into question. The law will have to decide whether subjectivity should continue to be constructed on classical criteria (will, organization, legal fiction grounded in human agency) or whether a functional criterion based on risk allocation should be adopted. This is not merely a technical issue but a deeply theoretical one with potentially long-term implications for the entire systematics of legal subjects.

The second open question relates to the possible redefinition of the organic theory of the legal person. If AI is brought closer to the concept of legal personhood, the unavoidable question arises: who forms the "will" of such an entity? Classical organic theory presupposes human organs acting on behalf of the legal person. In the case of highly autonomous algorithms, this assumption is called into question. The law will have to consider whether an expanded or functional concept of the "organ" is feasible, whether a hybrid organic model (algorithm plus human supervisor) should be introduced, or whether the organic logic should be partially abandoned in favor of a risk-operational approach. Each of these options opens new doctrinal and practical dilemmas.

The third major issue concerns tort liability. Existing regimes—especially those grounded in fault—are designed for human conduct or for legal persons whose will is mediated through human actors. With the emergence of complex self-learning systems, it is increasingly questioned

whether traditional models of fault, negligence, and foreseeability remain adequate. The law will need to determine whether to continue adapting existing doctrines (for example, through expanded strict liability) or to develop a specialized regime for algorithmic harm. Particularly sensitive will be the question of evidentiary standards in situations where algorithmic decision-making is partially opaque.

The fourth question concerns the allocation of responsibility within hybrid models. If a layered approach (human actor + AI fund + possible guarantee mechanism) is adopted, the law will need to define with precision the triggers for activating each layer. Otherwise, there is a real risk of procedural confusion, protracted proceedings, and legal uncertainty for injured parties. It will be especially important to avoid the phenomenon of "ping-pong liability," whereby different actors attempt to shift responsibility onto one another.

The fifth open question relates to the property-law status of AI systems. If patrimonial responsibility is attributed to an algorithm, it becomes necessary to clarify the legal nature of the assets attached to it. Are these to be treated as a dedicated fund, assets held in trust, a statutorily imposed guarantee pool, or some other sui generis construct? How would such assets be treated in the event of the operator's insolvency? Could they be subject to security interests or enforcement proceedings? These questions penetrate deeply into the core of property law and are far from trivial.

The sixth issue concerns procedural law. Even if a model of limited procedural personalization is adopted, courts will need to develop new rules governing representation, service of process, and evidentiary rules in cases involving autonomous systems. Particularly complex will be the question of access to algorithmic data, the protection of manufacturers' trade secrets, and the balancing of these interests against the injured party's right to effective judicial protection. Without clear procedural mechanisms, any substantive liability regime risks remaining difficult to apply in practice.

The seventh important issue is the cross-border dimension of AI systems. By their nature, AI systems operate in digital environments that frequently transcend national boundaries. If some form of AI subjectivity is introduced, the law will have to determine which legal order has jurisdiction and competence, where the system is deemed to "act," and how decisions concerning it will be recognized across jurisdictions. This will place particular pressure on the field of private international law.

The eighth question concerns regulatory coordination with existing European and international frameworks. The current trend—especially within the European Union—favors a risk-based regulatory approach without recognizing AI legal subjectivity. Any national or regional innovation in this field will need to be carefully aligned with existing regulatory frameworks to avoid normative fragmentation and regulatory arbitrage practices.

The ninth issue relates to evidentiary transparency and technical auditability. For courts to adjudicate effectively in AI-related disputes, significant advancement will be required in rules governing decision logging, data retention, and independent technical expertise. This, in turn, raises additional questions concerning data protection, cybersecurity, and the protection of trade secrets.

The tenth, but no less important, issue concerns the limits of legal personalization. Even within an affirmative model, the law will have to clearly determine which attributes of subjectivity are applicable to AI and which remain reserved for natural persons and classical legal persons. Without such precise delineation, there is a risk of conceptual inflation that could undermine the systemic coherence of the legal order.

In light of all these open questions, it becomes evident that any viable solution regarding the limited legal subjectivity of artificial intelligence must emerge from a careful, multidisciplinary, and gradual legal evolution. Neither the simple expansion of existing categories nor the radical introduction of a new subject can be accomplished without a thorough reassessment of numerous interconnected legal institutions. For this reason, the future task of legal scholarship is not merely to answer whether AI can be a subject of law, but to systematically elaborate the conditions under which such a transformation could be integrated into the existing legal architecture without compromising its internal coherence and protective function.

5. THE QUESTION OF THE PROCEDURAL CAPACITY OF AI

The question of the procedural capacity of artificial intelligence (AI) systems represents one of the most complex and sensitive aspects of the potential legal personalization of algorithmic entities. While the substantive law dimensions of liability can relatively more easily be adapted through models of strict or hybrid liability, procedural law has traditionally been much more closely tied to the classical notions of legal personality, will, representation, and participation in proceedings. Accordingly, any serious discussion of limited legal subjectivity for AI must necessarily address the question: under what conditions, to what extent, and through which mechanisms could an AI system appear as a procedural subject.

The first fundamental dilemma concerns the very nature of capacity to be a party to proceedings (standing). In classical procedural theory, the capacity to be a party to proceedings (processual standing) is closely linked to the subject's legal capacity and capacity to act. In the case of natural persons, this is tied to their will and ability to understand legal consequences, while in the case of legal persons it is assumed that procedural will is formed through their organs. In the case of AI, neither assumption exists in pure form. The algorithm possesses no independent procedural will in a psychological sense, nor does it have a classical organic structure. The law will therefore have to decide whether procedural capacity should remain will-based, or whether a functional approach grounded in risk-bearing and participation in legal transactions should be developed.

The second significant dilemma concerns the procedural representation of AI. Even within a model of limited subjectivity, it is difficult to envisage an algorithmic system independently performing procedural acts. The most realistic option is therefore the introduction of mandatory procedural representation. However, this immediately raises several sub-questions: who should serve as the legal representative of the AI—the operator, the manufacturer, an independent administrator, or a specially appointed "AI guardian"? Should the representative act in the interest of the AI fund, the injured parties, or the public interest? How can conflicts of interest be avoided if, for example, the operator of the system is simultaneously a potentially liable party? These questions are critical because improperly designed representation could undermine the entire procedural architecture.

The third open issue concerns procedural standing. If AI is granted limited subjectivity, the law will need to determine in which proceedings and in what capacity it may appear. May an AI system be sued only in relation to patrimonial liability, or also in other types of proceedings? Could it appear as a claimant? Could it intervene as a third party or auxiliary participant? Might it be subject to administrative proceedings, or only civil ones? Without clear normative limitations, there is a risk of excessive expansion of the procedural role of AI, potentially leading to legal uncertainty and procedural inflation.

The fourth important dilemma relates to service of process and procedural notification. Classical procedural law presumes that parties can receive court documents, be duly notified, and respond within prescribed deadlines. In the case of AI, this assumption is technically mediated. The law will have to determine whether service should be effected upon the representative, the operator, the registered digital identifier of the system, or through some combined mechanism. This is not a trivial issue, as proper service lies at the core of the right to a fair hearing.

The fifth issue concerns the law of evidence. In disputes involving AI, key facts are often embedded within algorithmic processes and system logs. For courts to adjudicate effectively, new rules will be required regarding access to technical data, protection of trade secrets, and the use of expert testimony. Particularly sensitive is the question of whether—and under what conditions—a court may order disclosure of source code or machine-learning models. This represents a potential tension between procedural justice and the protection of intellectual property.

The sixth dilemma relates to the burden of proof. In traditional tort litigation, the claimant generally bears the burden of proving damage, unlawfulness, and causation. In the context of complex AI systems, this may become significantly more difficult due to technical opacity (the so-called black box problem). Consequently, the law may need to consider whether the introduction of facilitated evidentiary standards, presumptions of causation, or even limited shifts in the burden of proof are justified. Any such intervention, however, must be carefully balanced against the right to a fair trial of the other parties involved.

The seventh important issue concerns the enforcement of judgments. If a court renders a decision against an AI system, how is enforcement to be carried out in practice? Most likely, through activation of the dedicated fund or insurance coverage associated with the system. Yet this requires clear rules governing access to the funds, the order of distribution, and safeguards against abuse. Additionally, the question may arise whether, in certain cases, technical restriction or deactivation of the system could be permitted as a regulatory measure.

The eighth dilemma has a cross-border dimension. AI systems frequently operate within distributed digital environments. Where a system is registered in one jurisdiction, operates in another, and causes harm in a third, complex questions of international jurisdiction and competence, applicable law, and recognition of judgments inevitably arise. This will require careful coordination with private international law and potentially the development of new conflict-of-laws solutions.

The ninth issue concerns potential abuses of procedural personalization. There are legitimate concerns that improperly designed AI subjectivity could be used as a shield for evading human responsibility. Procedural law will therefore need to incorporate mechanisms for "piercing" the formal structure where it is evident that the AI entity is being used for abusive purposes—analogous to piercing the corporate veil in company law.

Finally, the tenth dilemma concerns the desirable degree of procedural autonomy. In my view, in the foreseeable future the most sustainable model is one of limited procedural personalization, whereby AI formally appears as the holder of a specific risk fund but always through a human representative and within strictly defined types of proceedings. This approach allows practical functionality without undermining fundamental procedural principles.

In sum, the question of the procedural capacity of AI opens up a complex network of substantive, procedural, and institutional dilemmas. No solution will be sustainable unless issues of representation, evidentiary transparency, burden of proof, enforcement, and cross-border coordination are adequately addressed. For this reason, the procedural dimension will be one of

the key tests of whether the limited legal subjectivity of artificial intelligence can evolve from a theoretical proposition into a functional legal reality.

6. WHAT LEGAL SCHOLARSHIP HAS SAID SO FAR

Legal scholarship has thus far approached the idea of recognizing legal subjectivity for artificial intelligence (AI) systems with caution and considerable reservation. Although technological developments have significantly intensified interest in this issue, the dominant position in the literature remains that the existing legal categories—natural persons and legal persons—are, in principle, still sufficiently flexible to absorb most of the practical challenges posed by autonomous systems. Nevertheless, within this general caution, a rich and nuanced debate has emerged, which may be grouped into several principal strands.

One influential line of argument is associated with the so-called "responsibility gap" thesis. Andreas Matthias observed early on that, in the case of learning automata, situations may arise in which traditional attribution of responsibility becomes problematic, because the specific behaviour of the system cannot be fully predicted or controlled by human actors.[1] This thesis has been further developed by a number of scholars who warn that the growing autonomy of AI will increasingly place pressure on classical models of fault and foreseeability.[2] Even within this school of thought, however, there is no consensus that recognizing legal subjectivity for AI constitutes either the only or the optimal response. Parallel to this debate, the discourse on so-called "electronic personhood" has also developed. In 2017, the European Parliament opened the possibility of considering this approach, particularly in the context of highly autonomous robots.[3] Some authors, such as Pagallo and Bertolini, have explored models of limited legal personalization as a tool for more efficient allocation of risk and responsibility.[4] According to this view, legal subjectivity is a normative instrument that the law may shape functionally, without requiring ontological equivalence with natural persons.

Nevertheless, a significant portion of legal scholarship remains critical of such proposals. Joanna Bryson has forcefully argued that granting legal personality to robots or AI systems could lead to a dangerous erosion of human responsibility and create problematic incentives within the industry.[5] Similarly, Ryan Calo warns that premature personalization of AI may obscure rather than clarify questions of liability.[6] These authors generally advocate the refinement and strengthening of existing liability regimes rather than the creation of a new legal subject.

Within the European regulatory context, the prevailing trend continues to favor a risk-based approach without recognizing legal subjectivity for AI. The report of the Expert Group on Liability and New Technologies concludes that, in the vast majority of cases, existing liability regimes can be adapted to address emerging technological risks, particularly through the expansion of strict liability and the improvement of evidentiary rules.[7] This position is indirectly reflected in the

[1] Andreas Matthias, "The responsibility gap: Ascribing responsibility for the actions of learning automata", Ethics and Information Technology, 2004.
[2] Ugo Pagallo, The Laws of Robots: Crimes, Contracts, and Torts, Springer, 2013.
[3] European Parliament, Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics.
[4] Andrea Bertolini, "Artificial Intelligence and Civil Liability" (2020); Ugo Pagallo, supra note 4.
[5] Joanna J. Bryson, "Robots Should Be Slaves", in Close Engagements with Artificial Companions, 2010.
[6] Ryan Calo, "Robotics and the Lessons of Cyberlaw", California Law Review, 2015.
[7] European Commission, Liability for Artificial Intelligence and other emerging digital technologies, Expert Group Report, 2019.

European regulatory framework on artificial intelligence, which focuses on risk classification, transparency, and human oversight without introducing a new category of legal subject.

At the same time, legal scholarship has not reached a unified position on the legal subjectivity of AI. What emerges instead is a cautious pluralism: on the one hand, there is growing theoretical openness toward limited forms of functional personalization, particularly in the context of risk allocation; on the other hand, the dominant practical and regulatory approach continues to rely on adapting existing liability doctrines. It is precisely within this tension between innovation and doctrinal prudence that the contemporary legal debate on the role of artificial intelligence within the system of legal subjects continues to evolve.

7. CONCLUSION

The overall analysis demonstrates that the question of the legal subjectivity of artificial intelligence can no longer convincingly be treated as a purely theoretical speculation, but rather as a real challenge that is gradually entering the normative agenda of contemporary law. Although the dominant doctrine justifiably remains cautious, the growing functional autonomy of advanced AI systems is creating tensions that classical models of liability are increasingly unable to absorb without adaptation. In this respect, the most sustainable approach lies neither in the radical introduction of a "third subject" nor in the rigid preservation of the status quo, but in a carefully designed hybrid model situated between an expanded form of legal personhood and a specialized, limited form of AI subjectivity.

Such a model may offer improved risk allocation, greater legal predictability, and more effective protection of injured parties—but only if it is accompanied by robust safeguards, including subsidiary human liability, mandatory financial backing, enhanced oversight, and clear procedural regulation. Accordingly, the central task of future legal development is not whether AI abstractly "deserves" legal subjectivity, but rather under which strictly limited and functionally justified conditions such a construction can be integrated without undermining the systemic coherence and protective function of the law.

Bibliography

1. Andrea Bertolini, "Artificial Intelligence and Civil Liability", 2020.
2. Andreas Matthias, "The responsibility gap: Ascribing responsibility for the actions of learning automata", Ethics and Information Technology, 2004.
3. European Parliament, Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics.
4. Joanna J. Bryson, "Robots Should Be Slaves", in Close Engagements with Artificial Companions, 2010.
5. Ryan Calo, "Robotics and the Lessons of Cyberlaw", California Law Review, 2015.
6. Ugo Pagallo, The Laws of Robots: Crimes, Contracts, and Torts, Springer, 2013.