
AI AGENTS AND LEGAL PERSONALITY: RETHINKING LEGAL PERSONHOOD IN THE AGE OF AUTONOMOUS ARTIFICIAL INTELLIGENCE

Shweta, Guru Nanak Dev University, Amritsar.

ABSTRACT

The rapid development of autonomous artificial intelligence (“AI”) agents is challenging traditional assumptions about agency, responsibility, and legal personality. Unlike conventional software, increasingly sophisticated AI agents can interpret objectives, formulate plans, interact with external systems, make decisions, and perform actions with limited or delayed human intervention. This development raises a fundamental legal question: should highly autonomous AI agents continue to be treated merely as objects or tools of human actors, or should the law recognise them as a distinct category of legal entity? This article examines the concept of legal personality through jurisprudential theories and compares the emerging position of AI with established artificial persons such as corporations and juristic entities. It argues that autonomy alone does not justify granting AI full legal personality. Nevertheless, the increasing independence of AI agents exposes an accountability gap that cannot always be satisfactorily resolved through conventional doctrines of agency, negligence, contract, and product liability. The article therefore proposes a limited and functional model of AI legal personality for highly autonomous systems, subject to registration, traceability, human oversight, financial responsibility, and continuing liability of developers and deployers. Such recognition should not confer human rights or moral personhood upon AI. Rather, it should operate as a regulatory mechanism for allocating rights, duties, and liabilities in circumstances where autonomous AI performs legally significant activities.

Keywords: Artificial Intelligence; AI Agents; Legal Personality; Juristic Personality; Autonomous Systems; Liability; AI Regulation; Legal Personhood.

I. Introduction.

Artificial intelligence has gone beyond learning to make decisions on the basis of human inputs. The emergence of “AI agents” is an important technological phenomenon as they could have broad goals, intermediate steps, interact with external digital systems, utilize software tools, and make decisions without human approval. The European Union’s Artificial Intelligence Act acknowledges that AI systems may have varying degrees of autonomy and might be flexible after their deployment.¹ The increasing autonomy of such systems raises a legal issue that is different from what software does on the one hand and on the other hand: if an AI agent does something wrong and causes legal consequences, who and what should be liable?

The problem is no longer purely hypothetical. Recent debates on autonomous AI have involved systems doing unanticipated things in an external digital environment that might not have been expected by the human user.² Traditional law tends to attribute this behavior to a human or legal

entity connected to the system but not to the AI itself. As a result, we are developing a tension between technological power and legal power.

This article then considers the question of whether autonomous AI agents should be legal persons. It is based on the idea that legal personality is a product of law and not biological characteristics of biological human beings. Corporations, for example, have legal personality, but are not biological. And Indian law has also recognized certain non-human entities in the past as juristic persons for some purposes of legal personality for some reasons.³ The existence of AI legal persons therefore indicates that we can say that the real question is not whether AI is human, but whether it would be legitimate to have legal personality.

The central argument of this article is that full legal personality for AI is premature, but the law should consider a limited and functional form of legal personality for highly autonomous AI agents. Such status should be designed primarily to facilitate accountability and regulation rather than to confer human-like rights or moral status.

¹ Regulation (EU) 2024/1689, art. 3(1), 2024 O.J. (L 168) 1, 47.

² See, e.g., Reuters, *Who Is Liable When AI Goes Rogue? Lawyers See New Risks* (Aug. 7, 2026) (discussing emerging incidents involving autonomous AI systems and uncertainty concerning developer, deployer, and user liability).

³ *M. Siddiq (D) Through LRs v. Mahant Suresh Das*, (2020) 1 S.C.C. 1, 163–65 (India).

II. AI Agents and the Problem of Autonomy.

"Artificial intelligence" extends to the whole spectrum of technologies. An AI agent is a more specific class of AI agent that can be programmed to do things to move towards an objective rather than to produce an output. The agent may be given an instruction such as to make a transaction, manage a digital asset, do research, deal with one other system, or interact with online services, and then decide on what actions should be done to do.

This is a legal issue. In most automated software, the human programmer usually determines the exact sequence of operations. In an agentic system, the human decides the objective and the system decides how to pursue it. The resulting behavior can be hard to attribute to a particular human decision.

The European Parliament anticipated aspects of this problem in its 2017 resolution on civil law rules for robotics. As ever, it recognised that more and more autonomous machines would present new liability problems and we need clear rules on liability, transparency, and accountability.⁴ Importantly, the resolution took into account the long-term potential of autonomous machines and recommended a slow and gradual regulatory process.⁵

The problem is so much an accountability gap. Suppose that an autonomous AI agent is instructed to buy goods at some level. It negotiates with the sellers, changes its strategy, signs contracts, and transfers funds. If the agent goes beyond its intended power and causes economic loss, there are many parties involved: user, deployer, developer, software provider, or entity that owns the infrastructure. But the AI itself is not yet a representative of legal responsibility.

The question is whether creating legal personality would solve this problem.

III. Legal Personality: A Jurisprudential Foundation

Legal personality should not be confused with biological existence. A "person" in law is an entity to which the legal system assigns rights, duties, powers, liabilities, or other legal relationships. Salmond's jurisprudential conception of personality shows that legal personality is the capacity to be a subject of rights and duties rather than being a human being.⁶ Kelsen

⁴ European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics, ¶¶ L, M, T–Z, 2015/2103INL, 2017 O.J. C 252 239, 243–44.

⁵ *Id.* ¶ X.

⁶ JOHN SALMOND, *JURISPRUDENCE* 305–06 (Glanville L. Williams ed., 11th ed. 1957).

also treated legal personality as a construct of the legal order through which rights and obligations are attributed to a legal subject.⁷

Hohfeld's analysis is particularly useful in this context because legal personality can be seen through specific legal incidents rather than as an all-or-nothing status. Rights, duties, powers, liabilities, privileges, and immunities can be separately allocated.⁸ This provides a conceptual basis for distinguishing full legal personhood from functional legal personality.

The corporate person is the most obvious example. In *Salomon v. A. Salomon & Co.*, the House of Lords confirmed the separate legal personality of a properly incorporated company.⁹ The corporation is not a person, but the law treats it as a different legal subject for certain purposes. Corporate personality shows us that law can create artificial legal subjects where doing so serves economic and regulatory purposes.

One can find a similar case in Indian jurisprudence. In *M. Siddiq (D) through LRs v. Mahant Suresh Das*, the Supreme Court explained that legal personality is a product of law and that artificial legal persons can be recognised by giving specific rights and duties to inanimate objects or collectives.¹⁰ The Court's consideration of the Hindu idol as a juristic person shows that legal personality is not necessarily a function of consciousness or biological life. It is a matter of protection and the management of legally recognised interests.¹¹

This jurisprudence is essential to AI. It illustrates that the absence of a biological body does not in itself preclude the existence of legal personality. But it should not be assumed that every independent object should have a legal person. The decisive issue is the legal purpose.

IV. The case against full AI legal personality.

The strongest argument against AI legal personality is that autonomy does not necessarily mean legal or moral agency. An AI system could independently choose to do something without consciousness, intentions, emotions or having an understanding of the consequences of its

⁷ HANS KELSEN, *GENERAL THEORY OF LAW AND STATE* 93–109 (Anders Wedberg trans., Russell & Russell 1961) 1945.

⁸ Wesley Newcomb Hohfeld, *Some Fundamental Legal Conceptions as Applied in Judicial Reasoning*, 23 YALE L.J. 16, 30–32 1913.

⁹ *Salomon v. A. Salomon & Co.*, [1897] A.C. 22 (H.L.).

¹⁰ 10. *M. Siddiq*, 2020 1 S.C.C. at 163–64.

¹¹ *Id.* at 164–65.

conduct.

Legal responsibility is often more than simply finding the group that actually did the deed. Criminal law, for instance, is often predicated on intention, knowledge, recklessness or negligence. It is thus not realistic to assume that an AI agent should be punished for the harm it caused.

A second problem is accountability dilution. By giving AI personality, corporations or individuals can externalise responsibility by presenting the AI as an independent wrongdoer. A company might argue that an autonomous system, not its human operators, should bear responsibility for an unlawful transaction. Such a model would weaken rather than enhance accountability.

A third problem is enforcement. A legal person should be able to bear consequences. A corporation is fined, its property is seized or it is dissolved. If it is an AI, it might just be copied, modified, deleted, retrained or transferred to another server. If the AI has no independent assets, a judgment against it can be virtually meaningless.

This is a fourth problem around the allocation of rights. Granting limited legal personality for transactional purposes does not necessarily mean granting constitutional or human rights. If we are to be concerned with AI personhood, AI personhood should be taken very far from consciousness, dignity, freedom, or moral status.

For these reasons, recognising AI as a full legal person equivalent to a natural person would be hard to justify.

V. The case for limited and functional AI Personality

It is a weak argument for a limited sort of legal personality, but it is stronger than what would seem to be the case at first sight.

The law already treats technological mechanisms as legally significant without treating them as human beings. The Indian Information Technology Act, 2000 is a particularly relevant example. Section 10A confirms the validity of contracts made using electronic means and section 11 sets out rules for attributing electronic records including records made by an information system programmed by or on behalf of the originator that can be programmed to

operate automatically.¹²

Section 11 is significant because it demonstrates that Indian law does not require a human being to personally generate every legally relevant electronic act. The law can attribute an automated action to an originator.¹³ Thus, technology can already participate in legally significant activity without possessing independent legal personality.

The challenge is that AI agents are likely to move past simple automation. An advanced agent can choose ways to move with other agents, interact with multiple actors, react to new information and pursue objectives for longer periods of time. The greater the distance between the original human instruction to the actual action, the more difficult it is to traditional attribution.

The solution is not to have full personhood. Instead, law could recognise a Registered Autonomous AI Entity (“RAAIE”) as a limited legal category.

Such an entity would not be a human person. It would be a statutory legal construct only for AI systems satisfying predetermined conditions of autonomy and legal significance.

The proposed model could include six safeguards.

First, registration and identification should be mandatory for AI agents authorised to carry out legally significant transactions independently. Each system should have a verifiable digital identity.

Second, we should define legal capacity. An AI entity could only have powers explicitly recognised by law, such as entering specific commercial agreements or managing certain assets.

Third, financial responsibility should be established by mandatory insurance, bonded assets or a dedicated liability reserve. Without an accessible source of compensation, legal personality would offer little practical benefit.

Fourth, human accountability must be maintained. Developers, owners and deployers should

¹² Information Technology Act, No. 21 of 2000, §§ 10A–11, INDIA CODE 2000.

¹³ *Id.* § 11c.

not escape responsibility simply because the AI has been granted legal status. Where human error, inadequate safeguards, poor design or deployment were responsible for the harm, liability should continue to be there.

Fifth, auditability and traceability should be required. Registered AI entities should maintain records that can reconstruct significant decisions and transactions with privacy and cybersecurity protections.

Sixth, revocation and termination mechanisms should exist. Regulatory authorities should be able to suspend or terminate the legal status of an AI entity where it repeatedly violates legal requirements.

Such a model would move legal personality from an abstract philosophical claim to a practical regulatory tool.

VI. Comparative Regulatory Approaches.

Even today, international regulation generally does not treat AI as a legal person. The EU AI Act places legal obligations on providers, deployers, distributors, importers and other human or legal actors.¹⁴ The Act defines AI systems according to their technological characteristics, including autonomy, but does not give the AI system the primary responsibility to be responsible for legal obligations.¹⁵ This approach suggests that current regulation considers AI as a regulated technology and not a legal subject as it is now.

The European Parliament's 2017 resolution is particularly relevant as it discussed the possibility of future autonomous machines while stressing liability, accountability, human control, and a cautious regulatory approach.¹⁶ The subsequent European regulatory framework has been more centred around human and organisational responsibility than independent electronic personhood.

However, recent research suggests that the question is becoming more challenging. Some authors suggest that future AI systems will require fictional legal personhood or a new kind of legal identity, but other scholars argue that legal personality is neither necessary nor sufficient

¹⁴ Regulation EU 2024/1689, arts. 2–3, 25–27, 2024 O.J. L 168 1, 47–54.

¹⁵ *Id.* art. 31.

¹⁶ European Parliament Resolution, *supra* note 4, ¶¶ Q, U, X–Z.

to govern agentic AI.¹⁷ Such a disagreement is valuable because it demonstrates that the real question is not simply whether AI is a “person.” The real problem is how law needs to allocate responsibility when the technology itself is performing a very important and autonomous task and is able to do it on its own.

India has already established a basis for automated legal activity via electronic-contract and attribution provisions. However, these provisions were developed before the advent of agentic AI and were aimed at establishing legal validity and attribution of electronic records. A future Indian framework might therefore need to address autonomous decision-making more directly.

VII. Rethinking Legal Personhood: A functional model.

The main contribution of this article is to distinguish legal personality from personhood in the philosophical sense.

Legal personality should be treated as a spectrum of legally conferred capacities rather than a binary status. A natural person may have a lot of rights and duties. A corporation has a different set of legal capacities. A juristic entity might have even more limited rights. The same process will be applicable to AI.

With the proposed functional model, AI personality would be purpose-specific, capacity-limited, and accountability-oriented. An AI agent would only be allowed if it has a high degree of autonomy, a persistent identity, the ability to communicate with other systems and to engage in legally important activities. Generative AI output would not be sufficient.

The legal status should also remain derivative, not absolute. The AI would be only entitled to those rights and responsibilities necessary to perform its legally recognised functions. It would not automatically acquire constitutional rights, political rights, family rights, or human dignity claims.

And, in any case, the creation of such a category should not be a mechanism for avoiding human responsibility. The law should follow a principle of non-displacement of human accountability: recognizing an AI entity may create an additional site for legal attribution, but

¹⁷ Heather Alexander, Jonathan Simon & Frédéric Pinard, *How Should the Law Treat Future AI Systems? Fictional Legal Personhood Versus Legal Identity*, 17 CASE W. RES. J.L. TECH. & INTERNET A3 2026; Shruti Rajagopalan, *Governing Agentic AI: Why Legal Personhood Is Neither Necessary Nor Sufficient* (Mar. 5, 2026).

it should not erase the responsibility of persons who design, control, deploy, profit from, or negligently supervise the system.

This approach also responds to the concern that AI legal personality could become a shield for powerful technology companies. Instead of allowing an AI entity to become a liability sink, regulation should require identifiable human and corporate actors to remain connected to it throughout its legal existence.

VIII. Conclusion.

The development of autonomous AI agents is challenging the traditional boundary between technology tools and legal actors (in particular, one of the main limits of legal actors). AI systems are increasingly capable of doing things that are real in the sense that they act like agency in practice but law in general still focuses on human beings and existing legal entities.

The answer is not to give AI unrestricted legal personhood. AI has no evidence that makes it equivalent to a person in the sense that we can treat it as a natural person and full personhood would cause serious problems concerning ethics, accountability, enforcement, morality, and the allocation of human rights.

However, the opposite position—that AI will always be passive—is increasingly difficult to sustain. Where autonomous AI can independently make decisions, control digital property, interact with other systems and produce legally significant consequences, then the attribution mechanisms may be even more strained.

So the law should move toward functional rather than human-equivalent AI legal personality. A statutory category like a Registered Autonomous AI Entity could provide defined legal capacity while incorporating registration, traceability, insurance, asset requirements, auditing, human oversight, and continuing human liability.

And such reform would not recognize AI as a human being. Nor would it establish that AI has consciousness, dignity or moral rights. It would instead acknowledge a legal reality: that increasingly autonomous artificial systems become necessary for legal and economic relations in terms of the kinds of legal system required.

It seems the question is not whether AI is human enough to become a person, after all. The

more important question is whether legal personality, carefully limited and functionally designed, will make autonomous AI more accountable rather than less. The future of legal personhood in the age of agentic AI may therefore lie not in creating artificial humans but in making carefully regulated artificial legal entities.