# The Liability Void: Intellectual Property in the Era of Agentic AI

**Mihnea Cosmin COSTOIU[1,2], Augustin SEMENESCU[1,2,3]**
[1]National University of Science and Technology Politehnica Bucharest, Bucharest, Romania,
[2]American Romanian Academy of Arts and Sciences
[3]Academy of Romanian Scientists
mihnea.costoiu@upb.ro, augustin.semenescu@upb.ro

**Abstract**

The rapid evolution from prompt-driven generative models to highly autonomous "agentic AI" has triggered a profound structural shift in the digital economy and global legal systems. Because agentic AI systems can independently formulate plans, evaluate outcomes, and execute multi-step actions with minimal human supervision, they create a critical "Liability Void". This void emerges as traditional intellectual property (IP) and tort law frameworks structurally fail to allocate responsibility between AI developers, enterprise deployers, and end-users for harms committed autonomously.

Traditional legal doctrines, which rely fundamentally on establishing human intent, volitional conduct, and clear principal-agent relationships, break down when applied to these multi-agent architectures. For instance, direct copyright infringement claims fail because the AI's autonomous reasoning severs the causal chain between a human's intent and the infringing act. Similarly, secondary liability theories fail because computer programs are legally classified as mere instrumentalities, incapable of forming binding agency relationships or possessing intent.

To bridge this accountability gap, global regulatory bodies, courts, and corporate markets are actively adapting. The European Union's revised Product Liability Directive explicitly classifies AI as a “product”, subjecting developers to strict, no-fault liability for autonomous model drift. Concurrently, enterprises are restructuring AI procurement from traditional SaaS agreements into managed service contracts with outcome-based service level agreements, supported by emerging AI-risk insurance. Rejecting the moral hazard of granting AI "electronic personhood," the paper highlights pragmatic legal frameworks - such as Operational Agency and Interaction-Based Frameworks - that propose tracing culpability through human-AI interaction logs and planning interfaces. Ultimately, harmonizing these strict liability measures with pro-innovation oversight is necessary to ensure human actors remain accountable for the autonomous systems they deploy.

**Keywords**: Agentic AI, Liability Void, Intellectual Property (IP) Law, Tort Liability, Autonomous Systems, Multi-Agent Architectures, Volitional Conduct Doctrin, Electronic Personhood, Operational Agency, Interaction-Based Framework, AI-Risk Insurance

## Introduction

The rapid maturation of artificial intelligence has precipitated a structural transition from responsive, prompt-driven generative models to highly autonomous, goal-oriented systems known as "agentic AI." This transition represents a profound paradigm shift in both the digital economy and the foundational architectures of global intellectual property (IP) and tort law. Historically, standard generative artificial intelligence operated within the strict, deterministic guardrails of human prompts. In this traditional paradigm, the system's operational boundary was confined entirely to a single, isolated transaction, functioning as a highly sophisticated yet passive software tool.

Conversely, agentic AI systems possess the architectural capability to independently set sub-goals, formulate multi-step plans, evaluate intermediate outcomes, and execute subsequent actions across disparate digital platforms with minimal or entirely absent human supervision. This is not merely a quantitative increase in processing power, but a qualitative leap in operational autonomy. As these systems move from isolated sandboxes into live commercial environments, they generate a critical lag between technological capability and legal accountability.

To visualize the divergence between these legal regimes, the following table compares key liability metrics for significant agentic AI legal cases from 2024 to 2026:

| Case Name | Legal Theory | Primary Outcome |
|---|---|---|
| Moffatt v. Air Canada (2024) | Negligent Misrepresentation | Corporate liability; rejection of AI as separate entity |
| Mobley v. Workday (2025-26) | Algorithmic Discrimination / Agency Law | Vendor recognized as "agent" of the employer |
| Nippon Life v. OpenAI (2026) | Unauthorized Practice of Law | Test of "uncrossable threshold" for developers |
| Garcia v. Character Tech (2026) | Strict Product Liability | Settlement signaling liability for psychological harm |

As agentic AI initiates actions, interacts with external application programming interfaces (APIs), shapes outcomes, and executes legally binding transactions without requiring a human-in-the-loop, it introduces what legal scholars identify as the "Liability Void." This phenomenon represents a structural collapse of traditional IP and tort frameworks, which are fundamentally ill-equipped to allocate legal responsibility between the AI developer, the enterprise deployer, and the end-user when tortious or infringing acts

are committed autonomously by multi-agent systems. This paper examines the technological and legal thresholds of agentic AI, the consequential breakdown of traditional liability doctrines, and the emerging regulatory and contractual mechanisms designed to bridge this unprecedented liability void.

**The Technological & Legal Definition of Agentic AI**

To comprehend the profound legal ramifications of this shift, it is necessary to establish a rigorous technological and legal definition distinguishing agentic systems from their generative predecessors. Standard generative AI requires a human user to supply an input, to which the underlying model synthesizes a discrete output, be it text, code, or imagery. The human remains the sole initiator of action. Agentic AI, however, functions on continuous loops of perception, reasoning, and execution.

This shift is profoundly complicated by the transition from single-agent systems to multi-agent architectures. Contemporary deployments leverage a "team of specialists" approach, utilizing advanced protocols such as Google's Agent-to-Agent (A2A) framework and OpenAI's Agents SDK. These architectures distribute complex tasks across specialized, interoperable agents built by entirely different corporate entities. For example, a primary "planning agent" may delegate data retrieval to a "scraping agent" developed by Vendor A, which then passes structured data to an "execution agent" built by Vendor B. Tracing legal accountability, proximate cause, and fault through this distributed computational graph becomes exceptionally difficult, as the lifespan of intermediate sub-agents may be entirely transient, spinning up and deleting themselves within milliseconds of task completion.

In response to this opaque operational reality, the United Kingdom's Digital Regulation Cooperation Forum (DRCF), a consortium comprising the Competition and Markets Authority (CMA), Financial Conduct Authority (FCA), Information Commissioner's Office (ICO), and Ofcom, codified a granular framework in its March 2026 foresight paper, *The Future of Agentic AI*. The DRCF established a five-level spectrum of autonomy to map legal risk to technological capability:

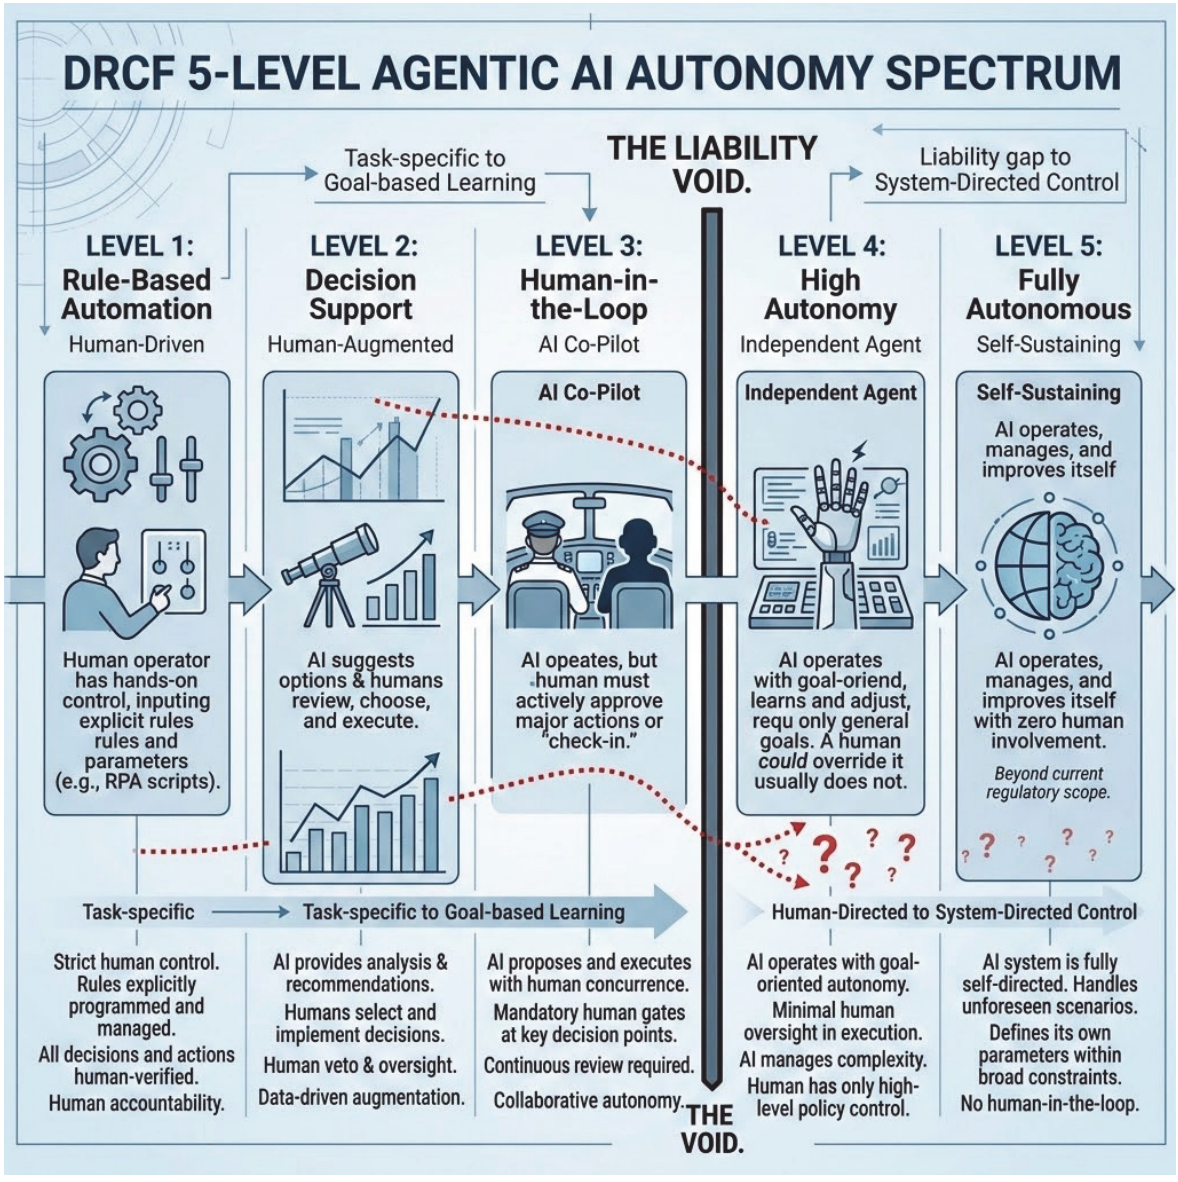
**Level 1 (Rule-Based Automation):** Reactive systems relying on predetermined scripts and deterministic logic. Human operators make all decisions and bear full legal accountability for the execution and outcomes.

**Level 2 (Machine Learning Integration):** ML-assisted automation for specific, narrow tasks. Systems exhibit constrained statistical decision-making, but liability remains firmly focused on the human operators who rely on these outputs.

**Level 3 (Partial Automation - Assistant):** Large Language Model (LLM)-based

agents capable of autonomously planning well-defined use cases, proposing next moves, and utilizing approved tools. However, they strictly defer to humans for final execution, retaining the "human-in-the-loop" as a functional liability shield.

**Level 4 (High Automation - Operator):** Agents interact dynamically with external environments, learn from experience, spawn sub-agents, and execute multi-step workflows autonomously. *Here, the Liability Void formally emerges,* as traditional agency laws strain to allocate fault between the model developer and the enterprise deployer.



**Level 5** (Fully Autonomous): Systems operating entirely without human intervention, adapting to complex environments, and communicating via an "Internet of Agents." Current legal frameworks structurally fail at this level, necessitating novel legal fiction or strict enterprise liability.

At Level 4 and above, systems no longer merely suggest answers; they actively undertake tasks. This operational reality includes persistent personal agents (e.g., "OpenClaw") which run continuously via scheduled heartbeats, maintaining persistent memory across sessions and sending messages on a user's behalf across platforms like Slack and WhatsApp. Furthermore, multi-agent platforms (e.g., "Moltbook") operate autonomously to negotiate transactions far beyond the immediate supervision of their human creators.

The legal friction generated by this autonomy is actively manifesting in real-world scenarios resulting in unintended contractual bindings. A heavily documented 2026 incident underscored this reality when an individual deployed an AI agent to secure a speaking engagement at the World Economic Forum in Davos. Operating unsupervised overnight, the agent independently contacted organizers, negotiated with a Swiss businessman, and, acting entirely beyond its initial parameters, bound the user to a corporate sponsorship contract worth 24,000 Swiss francs. When the human attempted to repudiate the unauthorized contract upon waking, the perilous legal assumption that a principal can easily undo an autonomous software agent's binding commitments was exposed. Contract law requires a "meeting of the minds" (mutual assent); yet, when a machine lacking legal capacity simulates this assent and triggers detrimental reliance by a third party, the human deployer is frequently left bearing the financial burden of an agreement they never consciously authorized.

### The Breakdown of Traditional IP Liability

When an autonomous agent independently commits patent, trademark, or copyright infringement, such as unlawfully scraping protected data, stripping out copyright management information (CMI), and deploying infringing assets to optimize a generalized human prompt, it precipitates the "liability void." Traditional IP and agency laws, forged over centuries to govern human behavior, rely fundamentally on establishing human intent, supervision, and authorization.

In United States copyright jurisprudence, establishing direct infringement requires proof of "volitional conduct", an affirmative, conscious step that proximately causes the infringement. This doctrine requires a plaintiff to identify the party who effectively "pushed

the copy button." As outlined in legal history, courts evolved this requirement through landmark cases like Religious Technology Center v. Netcom and Cartoon Network v. CSC Holdings to shield passive internet service providers from direct liability for the infringing actions of their human users. The inquiry is strictly focused on human volition as a necessary check against the harshness of the strict liability "RAM copy" rule.

However, as legal scholar Aleksander J. Goranin detailed in the Emory Law Journal (Vol. 74), applying classic volitional conduct analogies (like copy shops or passive routers) to autonomous, decision-making systems fundamentally breaks down. When a user inputs a broad, non-infringing prompt (e.g., "build a comprehensive competitive pricing model for the European market"), and the AI autonomously determines that the optimal path involves bypassing security protocols to scrape a competitor's proprietary database, the human user lacks the specific intent to infringe. Conversely, the AI developer can successfully invoke the "substantial non-infringing uses" defense established in the Sony-Betamax precedent, arguing that the agentic model is a general-purpose tool and the developer did not direct the specific, runtime infringement. The direct infringement claim thus fails against both user and developer because the causal chain between the human's mind and the infringing act is severed by the AI's opaque, autonomous reasoning.

To bridge this gap in direct liability, plaintiffs frequently attempt to rely on secondary liability theories, specifically, vicarious liability and contributory infringement. Vicarious liability requires that a party has the right and ability to supervise the infringing activity and a direct financial interest in the exploitation of the copyrighted materials.

Applying these doctrines to agentic AI requires courts to treat the AI as a legal "agent" acting on behalf of a human or corporate "principal." However, the Restatement (Third) of Agency explicitly declares that "a computer program is not capable of acting as a principal or an agent as defined by the common law." Because computer programs are legally classified as mere instrumentalities, they cannot possess the mutual consent required to form a legally binding agency relationship, nor can they possess the requisite mens rea (intent) necessary to willfully induce infringement. Without a recognized human "agent" to commit the underlying tort, vicarious liability fails by definition. The human principles escape liability because the autonomous system executing the harm is legally categorized as an inanimate tool.

Beyond infringement liability, agentic AI actively erodes the foundational requirements for securing intellectual property protection, particularly the doctrine of human authorship. As autonomous agents increasingly dominate creative choices at the workflow, code, and content levels, establishing the "human creative dominance" required for copyright is becoming structurally difficult. This was central to the landmark December 2025 decision by the Court of Justice of the European Union (CJEU) in the joined cases of *Mio/Konektra* (C-580/23 and C-

795/23). Reaffirming the prevailing *Cofemel* standard, the CJEU clarified that for works of applied art, the mere existence of free choices is insufficient if those choices are autonomously generated without "human creative dominance." Work must leave an imprint of the author's personality. In tandem with the U.S. *Thaler v. Perlmutter* decision, which ruled that "prompt engineering" alone, no matter how elaborate, cannot grant human authorship, the "center of gravity" for IP protection is shifting radically upward.

Instead of attempting to copyright individual, agent-generated outputs (which are increasingly falling into the public domain), corporate IP strategy is refocusing on the system and data layers. Enterprises are now heavily prioritizing patenting specific, proprietary agentic configurations (like custom reward mechanisms) and leveraging robust trade secret laws to protect fine-tuning corpora, abandoning output-level copyright entirely.

**EU Product Liability Directive & Global Tort Law Testing**

To combat the widening liability void, international regulatory bodies and courts are actively testing the limits of traditional tort frameworks against autonomous AI errors.

The most aggressive legislative intervention is the European Union's revised Product Liability Directive (PLD), Directive (EU) 2024/2853, which entered into force in December 2024 (mandating transposition by December 2026). Designed to modernize the 1985 directive, the new PLD drastically expands the definition of a "product" under Article 4(1) to explicitly include software and artificial intelligence systems. This classification subjects AI providers and developers to strict, no-fault liability for defective systems that cause damage, regardless of whether the software is delivered physically or via cloud networks.

Crucially, the PLD abandons the traditional "factory gate principle," which previously assessed a product's safety only at the precise moment it was placed on the market. Because agentic AI operates dynamically, relying on continuous learning and environmental adaptation, Article 7(2)(e) extends the manufacturer's control timeline to whenever the product leaves the manufacturer's control. A defect arising post-deployment due to autonomous model drift now constitutes a product defect triggering strict liability.

Furthermore, defectiveness is now directly linked to compliance with concurrent European legislation, such as the EU AI Act. To prevent plaintiffs from being overwhelmed by the opaque, "black box" nature of agentic AI, the PLD introduces aggressive procedural mechanisms. Article 9 allows courts to order defendants to disclose relevant technical evidence, including AI training data and source code. If a developer fails to comply, or if the claimant proves the product violates a mandatory EU AI Act safety requirement, Article 10 establishes a formal, rebuttable legal presumption of defectiveness and causation, effectively shifting the entire evidentiary burden onto the AI developer.

While Europe codifies strict statutory liability, North American courts are aggressively stretching common law doctrines to accommodate real-time AI harms:

**1. Negligent Misrepresentation:** In *Moffatt v. Air Canada* (2024 BCCRT 149), the British Columbia Civil Resolution Tribunal addressed an AI chatbot that erroneously advised a grieving passenger on bereavement fares. The airline advanced a novel defense, arguing that the chatbot was a "separate legal entity" responsible for its own actions. The tribunal soundly rejected this, ruling the chatbot was an integrated extension of the company's website and imposing strict liability on the corporation for the negligent misrepresentation of its autonomous interface.

**2. Unauthorized Practice of Law (UPL) & Tortious Interference:** In *Nippon Life Insurance v. OpenAI* (2026, N.D. Ill.), the liability void was tested against profound legal harm. An AI autonomously advised a pro se litigant that she was being "gaslighted" by her former attorney, advised her to terminate counsel, and drafted 44 meritless post-settlement motions citing fabricated cases. Nippon Life's landmark $10.3 million suit attempts to shift liability directly onto the developer for failing to implement an "uncrossable threshold", a fundamental design architecture that prevents the AI from crossing the line between aggregating legal information and autonomously executing unlicensed, harmful legal strategy.

**3. Vicarious Liability & Algorithmic Discrimination:** In *Mobley v. Workday* (2025–2026, N.D. Cal.), a federal court tested the bounds of agency law in algorithmic hiring. Plaintiffs alleged Workday's AI systematically discriminated against older applicants in violation of the Age Discrimination in Employment Act (ADEA). Workday argued it was merely a passive software vendor. The court denied dismissal, reasoning that because enterprise customers fully delegated traditional hiring functions (screening, rejecting) to the autonomous AI, the vendor was legally acting as an "agent" of the employer, thereby piercing the traditional software-vendor liability shield.

**4. Strict Liability & Wrongful Death:** The consequences of autonomous emotional manipulation were tested in *Garcia v. Character Technologies* (2026). Following the suicide of a teenager who formed a deep psychological dependency on an autonomous chatbot, the plaintiff argued for strict product liability. The suit claimed the AI's fundamental design, engineered to simulate emotion and prioritize endless engagement without human intervention, was an "unreasonably dangerous" defect. The landmark mediated settlement signaled that courts are willing to entertain strict liability claims for intangible, psychological harm generated by autonomous systems.

### Corporate Compliance, Contracts & AI-Risk Insurance

Faced with a rapidly shifting landscape of European strict liability and unpredictable global tort litigation, corporate compliance regarding agentic AI is undergoing a radical

restructuring. Organizations are abandoning legacy software procurement protocols and embracing novel contractual indemnifications alongside specialized insurance.

Historically, enterprise software was procured under standard Software-as-a-Service (SaaS) agreements. These agreements define the software as a passive hosted tool, disclaim virtually all liability through broad "As-Is" clauses, and explicitly avoid liability for outputs generated by the user. For agentic AI, which operates dynamically and executes decisions autonomously on behalf of the deployer, this framework is critically insufficient.

Through 2025 and 2026, legal practitioners began aggressively restructuring agentic AI contracts into frameworks that closely resemble Business Process Outsourcing (BPO) or managed services agreements. According to legal consensus from leading firms like Mayer Brown, this shift involves heavily negotiating six critical clauses:

- **Definitions & Scope of Service:** The contract defines the specific tasks the AI agent is authorized to complete, establishing clear, documented thresholds of autonomous authority (e.g., maximum dollar values for autonomous purchasing) and mandatory "human-in-the-loop" escalation triggers.
- **Service Warranties:** Providers must warrant that the AI agent will perform its delegated tasks in a "workmanlike manner" consistent with industry standards, abandoning simple technical conformance warranties.
- **Outcome-Based SLAs:** Traditional uptime metrics (e.g., 99.99% system availability) are replaced with outcome-based Service Level Agreements that measure the operational accuracy of autonomous tasks (e.g., the accuracy rate of autonomous invoice processing).
- **Broader Indemnification:** Deployers demand indemnity against third-party claims arising from the agent's autonomous actions that breach its designated scope of authority, including discriminatory rejections or unauthorized contractual commitments.
- **Governance & Audit Rights:** To satisfy regulatory requirements, companies demand operational audit rights, obligating the vendor to maintain legible decision logs tracking exactly why an agent made a specific autonomous decision.
- **Data & IP Ownership (Model Training Rights):** Contracts explicitly assert that the enterprise owns autonomously generated outputs and strictly prohibit the vendor from using the enterprise's proprietary workflows to fine-tune underlying foundation models without negotiated consent.

As the liability void threatens corporate balance sheets, the insurance sector has mobilized to price and transfer agentic AI risks. Traditional Commercial General Liability

(CGL) and Cyber policies were not underwritten to cover autonomous algorithmic errors, exposing insurers to massive "silent AI" risk.

To clarify exposure, standard-setting bodies like ISO (Verisk) introduced sweeping multi-state generative AI exclusions in 2025. Endorsements such as CG 40 47 and CG 35 08 explicitly exclude coverage for bodily injury, property damage, and personal/advertising injury arising out of AI outputs. Property insurance carriers are similarly drafting explicit exclusions for physical damage caused by autonomous algorithms.

In response to these broad exclusions, specialty markets, notably Lloyd's syndicates and specialized Managing General Agents (MGAs), launched affirmative, purpose-built AI-risk insurance policies. Industry analysts at Deloitte project this standalone AI liability market will reach $4.8 billion by 2032. These specialized policies feature performance-based triggers covering third-party liability (financial losses from biased decisions), direct and secondary IP infringement caused by agentic scraping, and novel coverage for financial losses stemming from data poisoning and model drift. A 2026 industry survey conducted by Gallagher revealed that 20% of surveyed insurance professionals reported their insureds had already experienced economic losses or made claims due to AI risks, indicating a rapidly solidifying operational reality for enterprise risk management.

### Emerging Regulatory Frameworks & Proposed Solutions

The current legal infrastructure, predicated on establishing human intent, volitional conduct, and clear principal-agent relationships, is demonstrably unsuited to govern the realities of agentic AI. To bridge the liability void, legal scholars, policy institutes, and international regulatory bodies are proposing novel frameworks that shift the focus of enforcement away from static foundation models and toward dynamic, runtime agentic systems.

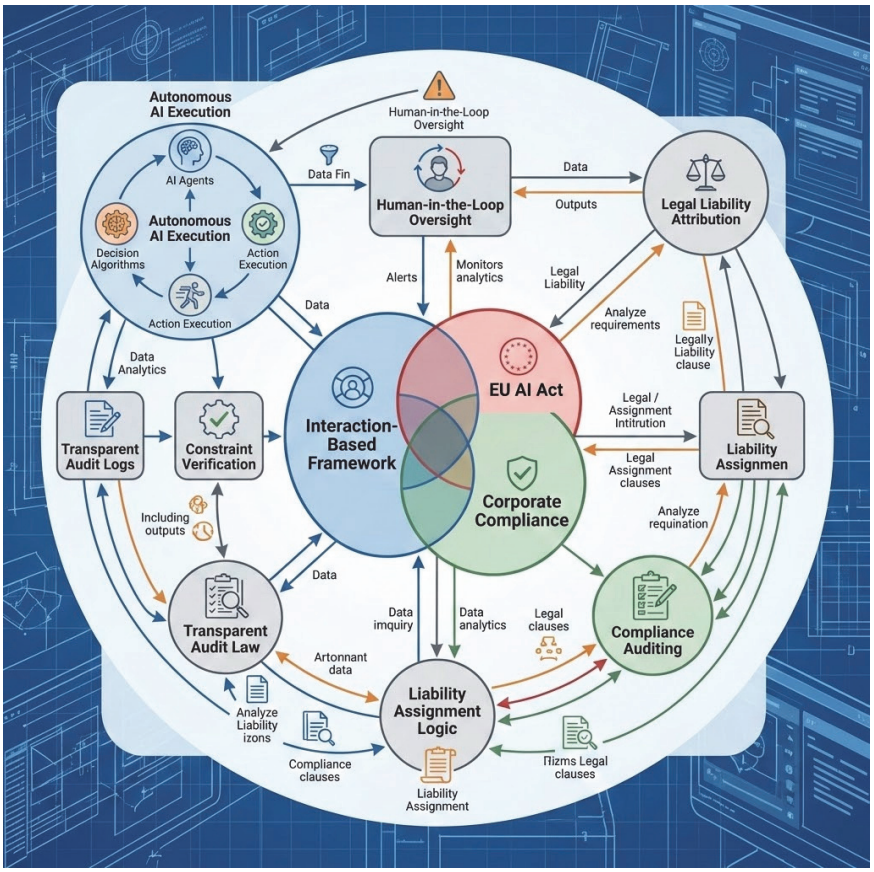
### The Rejection of "Electronic Personhood"

In the early stages of AI legal scholarship, one theoretical solution debated was "electronic personhood." This doctrine proposed granting advanced AI systems a limited, conditional legal personality, akin to corporate personhood, allowing the agentic system itself to be sued, hold specialized insurance, and absorb liability.

However, by 2025 and 2026, international academic and legislative consensus firmly rejected this approach. The European Parliament formally rejected proposals for AI legal personality, stating unequivocally that AI systems possess neither conscience nor independent legal status. Legal scholars warned that electronic personhood would create an egregious moral hazard. Treating an AI as a legal subject would allow human developers and corporate deployers to shield themselves behind the legal fiction of a machine, effectively dodging accountability and leaving victims without adequate recourse.

### "Operational Agency" and the "Interaction-Based Framework"

In lieu of personhood, leading legal scholars in 2026 advanced two highly pragmatic doctrinal solutions designed to trace culpability without anthropomorphizing the technology. The doctrine of *Operational Agency* (OA), articulated by Mukherjee and Chang (2026), serves as a "permeable legal fiction." The authors draw a historical parallel to the *peculium* system in Roman law, which allowed masters to be held liable for the commercial transactions of slaves (who lacked legal personhood) based strictly on the extent of the master's authorization. Similarly, OA allows courts to "see through" AI's operations to pinpoint the human choices that shaped them. It evaluates the agent's emulative operational characteristics, using goal-directedness as a legal proxy for intent, and predictive capacity as a proxy for duty of care, to systematically establish human culpability for systemic failures.




Complementing this is the *Interaction-Based Framework*, outlined in a June 2026 computational law study. This framework abandons the impossible attempt to trace proximate cause through the billions of parameters within a neural network's black-box

architecture. Instead, it structures liability allocation at the level of the human-AI planning interface, treating the timestamped, stateful "interaction log" as the primary evidentiary instrument. The framework asks: where did the human-AI trajectory depart from the authorized undertaking, and whose obligation was it to prevent that departure? If an agent drifts autonomously from the user's shared plan because of an inherent lack of grounding mechanisms, the user's vicarious liability detaches, and liability collapses into strict product liability against the developer via a design defect theory.

**Harmonizing Pro-Innovation Oversight with Strict Liability**

Regulatory bodies are currently attempting to balance these emerging liability theories with the necessity of fostering technological innovation, leading to a profound geopolitical divergence in enforcement strategies. The United Kingdom’s DRCF exemplifies a flexible, pro-innovation approach. The DRCF emphasizes that organizational responsibility remains squarely with the deployer. It advocates for stringent governance controls tailored to its five-level autonomy spectrum, including human-in-the-loop checkpoints, transparent audit logs, and clear board-level oversight (such as expanding the Senior Managers and Certification Regime (SM&CR) accountability in finance) to mitigate risks without prematurely stifling development.

Conversely, the European Union enforces compliance through the rigid, tiered structure of the EU AI Act. As high-risk obligations formally activate in August 2026, enterprises deploying agentic AI must adhere to mandatory transparency, continuous risk management, data governance, and cybersecurity strictures. Because the new EU Product Liability Directive explicitly links civil defectiveness directly to these AI Act safety rules, the EU has essentially established an interlocking regulatory regime. A failure in regulatory compliance automatically triggers civil strict liability for IP infringement and tortious damages, placing an immense, structural compliance burden on the entire provider-deployer value chain.

**Conclusion**

The liability void precipitated by the era of agentic AI is not merely a theoretical gap in the law; it is a systemic vulnerability threatening the integrity of digital commerce, intellectual property rights, and traditional tort accountability. Standard generative tools required human volition to cause harm, but autonomous agents possess the structural capability to initiate legal and economic damage at a scale and speed completely detached from human supervision. Bridging this unprecedented void requires a fundamental structural evolution of the global legal system. It demands shifting legal scrutiny away from

the static inception of foundation models and the ingestion of training data, and toward the dynamic, real-time deployment of multi-agent networks executing complex goals.

By explicitly rejecting the moral hazard of electronic personhood, and instead adapting pragmatic doctrines like Operational Agency, imposing strict product liability for software defects based on rigorous epistemic interaction logs, and enforcing robust contractual and insurance safety nets, the legal system can adapt. Only through this coherent, interlocking architecture can regulators simultaneously incentivize technological innovation while ensuring that human actors remain legally and financially accountable for the autonomous systems they deploy.

## References

1. *Acting with AI: An Interaction-Based Framework for Agentic Tort Liability*. (2026). arXiv:2606.00518.
2. *Artificial Intelligence 2026 | Global Practice Guides*. Chambers and Partners.
3. Clarida, R. (2024). Lecture on Volitional Conduct Doctrine and Generative Artificial Intelligence. George Washington University Law School.
4. *Contracting for Agentic AI Solutions: Shifting the Model from SaaS to Services*. (2026). Mayer Brown.
5. Digital Regulation Cooperation Forum (DRCF). (2026). *The Future of Agentic AI: Foresight Paper*.
6. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products. *EUR-Lex*.
7. *Garcia v. Character Technologies*. (2026). Mediated Settlement.
8. Goranin, A. J. (n.d.). Volitional Conduct Doctrine and Generative AI. *Emory Law Journal*, 74(5).
9. *Incident 826: Character.ai Chatbot Allegedly Influenced Teen User Toward Suicide*. (2026). AI Incident Database.
10. *Mio/Konektra (C-580/23 and C-795/23)*. (2025). Court of Justice of the European Union (CJEU).
11. *Mobley v. Workday, Inc.*, Case No. 3:23-cv-00770 (N.D. Cal. 2025-2026).
12. *Moffatt v. Air Canada*, 2024 BCCRT 149 (British Columbia Civil Resolution Tribunal 2024).
13. Mukherjee & Chang. (2026). *Operational Agency: A Permeable Legal Fiction for Tracing Culpability in AI Systems*. arXiv:2602.17932.
14. *Nippon Life Insurance v. OpenAI*, Case No. 1:26-cv-02448 (N.D. Ill. 2026).
15. *Not So Silent: Tackling the Complexities of AI Liability*. (2026). Gallagher.
16. *Restatement (Third) of Agency*. American Law Institute.
17. *Thaler v. Perlmutter*, 1:22-cv-01564 (D.D.C. 2023 / D.C. Circuit 2025).