TinyResearch Sample Report: Legal Personhood for AI

TinyResearch · Sample research report · Research Rubrics

This sample presents a complete TinyResearch report in response to the question below. The report text is reproduced unchanged, with its cited sources listed at the end.

Original research question

Legal Personhood for Artificial Agents: International Frameworks, UAE Governance, and Accountability Challenges

- A Comprehensive Decision-Maker's Report
- Executive Summary
- Part I: The Concept of Limited AI Legal Personhood — Proposals, Frameworks, and Scholarly Debate
  - 1.1 What Legal Personhood Means and Why It Matters for AI
  - 1.2 The Chronological Arc of AI Personhood Scholarship
  - 1.3 The Spectrum of Proposed Legal Models
- Part II: International Legal Instruments — Coverage, Gaps, and Emerging Frameworks
  - 2.1 The Council of Europe Framework Convention (2024)
  - 2.2 The UN Global Dialogue on AI Governance (2025–2026)
  - 2.3 The EU Regulatory Architecture
  - 2.4 What International Law Does Not Address
- Part III: The UAE AI Ministry Framework — Institutional Architecture, Legal Mandate, and the June 2026 Federal Authority
  - 3.1 Chronological Development of UAE AI Governance
  - 3.2 The Federal Authority for Artificial Intelligence and Data (June 2026)
  - 3.3 The Three-Layer UAE Compliance Architecture
  - 3.4 Accountability Framework Under UAE Law
- Part IV: AI Legal Personhood vs. Corporate Personhood — Philosophical, Legal, and Functional Distinctions
  - 4.1 What Corporate Personhood Provides and How It Works
  - 4.2 How AI Personhood Would Differ — Seven Key Distinctions
  - 4.3 The Philosophically Critical Distinction: Ontological vs. Functional Approaches
  - 4.4 Specific Scholarly Arguments Against AI Legal Personhood
  - 4.5 Arguments for Limited AI Personhood
- Part V: Liability and Accountability Challenges for Autonomous AI Systems
  - 5.1 The Accountability Gap — A Structural Analysis
  - 5.2 Inadequacy of Current Legal Doctrines
  - 5.3 Distributed AI Supply Chains and Cross-Border Accountability
  - 5.4 The EU AI Liability Directive Withdrawal and Its Consequences
- Part VI: Comparative National and Regional Frameworks
  - 6.1 EU AI Act — The Most Developed Horizontal Framework
  - 6.2 Lethal Autonomous Weapons Systems — International Humanitarian Law
  - 6.3 Saudi Arabia's Sophia — A Symbolic Rather Than Legal Development
  - 6.4 US State Law — Moving Toward Prohibition
  - 6.5 Singapore — Functional Analysis Without Personhood

Part VII: Proposed Alternative Liability Frameworks
7.1 The Roman Law Peculium Analogy
7.2 Strict Liability Mandates
7.3 Mandatory Insurance and Compensation Funds
7.4 IHL-Specific Proposals — War Torts Regime
Part VIII: Ethical, Philosophical, and Policy Critiques
8.1 The Liability Shield Concern
8.2 The Enforcement Paradox
8.3 Democratic and Human Rights Concerns
8.4 The Moral Status Question and Its Policy Implications
Part IX: Decision-Maker Analysis and Recommendations
9.1 Factual Findings Summary (High Confidence)
9.2 Analytical Conclusions (Interpretation)
9.3 Watchlist for Decision-Makers
9.4 Recommendations
Conclusion
Sources

Original research question

Can you critically evaluate the legal implications of conferring limited personhood on artificial agents under international law? How does this concept differ from traditional corporate systems such as seen in the UAE's AI ministry framework? Can you analyze the challenges related to liability and accountability of autonomous AI systems?

Legal Personhood for Artificial Agents: International Frameworks, UAE Governance, and Accountability Challenges

A Comprehensive Decision-Maker's Report

Executive Summary

The question of whether autonomous AI systems should receive limited legal personhood sits at the intersection of centuries-old legal philosophy, emergent technology governance, and acute practical need. This report finds that **no jurisdiction currently confers general legal personhood on AI systems**, and the dominant international trend — including the UN, EU, Council of Europe, and UNESCO — firmly assigns responsibility to identifiable human beings and institutions rather than to AI agents themselves [39][48][8]. Nevertheless, a persistent and growing **accountability gap** exists: current legal frameworks struggle to assign clear liability when autonomous systems cause harm through opaque, distributed, and unpredictable decision-making chains [13][22][46].

The UAE's governance architecture — culminating in the June 2026 establishment of the Federal Authority for Artificial Intelligence and Data — represents the most institutionally ambitious national AI governance model currently operating, but it does not address AI legal personhood and instead anchors responsibility in human operators and deployers [21][37]. This approach differs fundamentally from corporate personhood frameworks and from experimental proposals for "electronic personhood," which were formally rejected by the EU [11][40].

Key findings for decision-makers:

- **Limited AI personhood proposals are theoretically coherent but politically and practically contested.** Functional or hybrid models are the most defensible academically, but face intense opposition from legal scholars, ethicists,

governments, and 77% of surveyed publics [36][45][24].

- **The EU AI Liability Directive withdrawal (formally completed by October 2025–August 2026) has created a material accountability gap** in the world's most developed AI regulatory market, with the European Parliament now recommending a strict liability framework as a replacement [22][16][44].
- **The UAE framework is governance-forward but liability-light**: its multi-layered structure excels at strategy and deployment but lacks a comprehensive horizontal AI liability law, with PDPL implementing regulations still unissued as of mid-2026 [21][31][34].
- **International humanitarian law faces its most acute AI accountability gap in the military domain**, where lethal autonomous weapons systems (LAWS) operate in a space that existing legal doctrine — state responsibility, command liability, criminal intent — cannot cleanly address [30][26][41].
- **The Roman law peculium analogy, strict liability mandates, compensation funds, and hybrid personhood models** represent the four main doctrinal alternatives to either full personhood or full reliance on existing tort law [13][16][3].

Part I: The Concept of Limited AI Legal Personhood — Proposals, Frameworks, and Scholarly Debate

1.1 What Legal Personhood Means and Why It Matters for AI

Legal personhood — the recognized capacity of an entity to hold rights and bear duties within a legal system — is explicitly not a biological category [36][45]. It is, as multiple scholars and historical records confirm, a **functional and political tool** that societies create to address governance needs [8][40]. Ancient Indian guilds received legal recognition in 800 BC; Roman municipalities gained legal status; and the Dutch East India Company in 1602 became the first modern corporate person, able to own property, enter contracts, and sue or be sued [36][8]. The UN itself was recognized as an international legal person by the ICJ in 1949 not because of inherent characteristics but because of its functions [48].

Rights conferred by legal personhood include owning property, suing and being sued, opening accounts, and exercising freedom of speech; obligations include paying taxes, abiding by contracts, and liability for negligent or wrongful conduct [36]. Legal personhood is also **divisible and scalar**: entities can hold some rights and duties without holding all, and courts have long recognized partial or context-specific legal standing [45][3].

For AI systems, legal personhood theory divides into two frameworks [3]:

- **Singularist views**: an entity qualifies as a legal person by holding even a single right or duty — a minimal threshold that some AI systems might already satisfy in limited contexts
- **Cluster views**: legal personhood requires the capacity to hold a large bundle of roles — making contracts, owning property, participating in legal proceedings — a threshold no current AI system clearly meets

Two distinct rationales animate AI legal personhood advocacy [3]:

1. **Non-instrumental rationales**: AI deserves personhood based on intrinsic characteristics (sentience, rationality, moral agency) — viewing AI as an end in itself in the Kantian sense
2. **Instrumental rationales**: AI should receive personhood because doing so benefits human interests — by clarifying liability, enabling contracts, or facilitating governance — without claiming inherent AI value

The trajectory of legal debate tracks technological cycles. During AI "summers" (periods of optimism), scholars debate broader rights-based frameworks; during "winters," discussion narrows to limited, functional rights [3]. The current period — with GPT-4 reportedly exhibiting human-level performance on professional exams [1] and agentic AI systems pursuing multi-step goals with limited oversight [39] — has renewed the debate with greater urgency than any prior cycle.

1.2 The Chronological Arc of AI Personhood Scholarship

Early scholarship (1980s–1990s) adopted a **bottom-up approach**, examining whether specific rights or duties — primarily intellectual property and tort liability — could be attributed to AI without addressing full personhood [3]. Lawrence Solum's foundational 1992 essay identified two enduring objections [3]:

- The **responsibility objection**: trustees and agents must be accountable for negligence; AI cannot be held accountable in the same way
- The **judgment objection**: AI cannot handle unforeseen circumstances and moral dilemmas because of the "frame problem" — its inability to identify which facts are relevant in novel situations

These objections remain central to contemporary debate, though proponents argue that advancing AI capabilities may eventually overcome the judgment objection [1][35].

The 2000s shift toward statistical AI and machine learning reignited the debate, partly because of dramatically improved commercial performance [3]. By the 2010s, the European Parliament's 2017 resolution and Saudi Arabia's Sophia citizenship in 2017 forced the question into mainstream policy discourse [2][9]. By 2026, the emergence of agentic AI — systems that pursue goals, use tools, and interact with digital environments with limited human oversight [39] — has made the accountability question structurally unavoidable.

1.3 The Spectrum of Proposed Legal Models

Academic and policy proposals for AI legal status span a wide spectrum:

Model	Description	Key Advocates / Examples	Status
Full exclusion / tool classification	AI remains property; all liability attaches to human creators/deployers	Current global consensus; EU AI Act; UNESCO [48][8]	Dominant
Electronic personhood	Separate legal category for sophisticated autonomous robots, with liability attached	EP 2017 Resolution §59(f) [45][2]	Proposed, then rejected
Limited/context-specific personhood	AI recognized as legal actor in specific high-stakes domains (finance, medicine) while humans retain ultimate liability	Academic proposals [4][40]	Theoretical
Hybrid / relational personhood	AI personhood mediated through relationships with human principals, similar to how disabled persons' legal capacity is supported	Emerging scholarship [49]	Theoretical
Functional agency without personhood	AI can engage in automated contracting recognized as valid without treating AI as a legal person	UNCITRAL Model Law on Automated Contracting [48]	Adopted
Peculium / Roman law analogy	AI acts as agent with ring-fenced assets; principals bear secondary liability as in Roman slave-run businesses	Cambridge/Roman law scholarship [13][25]	Theoretical

The most influential recent academic proposal is the **hybrid or context-dependent model**: granting AI limited or context-specific legal recognition in high-stakes domains such as financial services or medical diagnostics, while preserving ultimate human accountability [4][40]. This approach draws on the historical precedent that legal personhood expansions — for corporations, trusts, and environmental entities — were typically motivated by instrumental governance needs rather than inherent moral agency [4][40].

Part II: International Legal Instruments — Coverage, Gaps, and Emerging Frameworks

2.1 The Council of Europe Framework Convention (2024)

The **most significant international legally binding instrument** on AI is the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, opened for signature on 5 September 2024 [7]. It is the first-ever international legally binding treaty in this field and has been approved by the EU for implementation through the AI Act [43].

Key features relevant to legal personhood and accountability [7]:

- It is **technology-neutral** and does not regulate specific technologies or confer legal status on AI systems
- Activities within the AI lifecycle must comply with principles including human dignity, autonomy, equality, non-discrimination, privacy, transparency, **accountability and responsibility**, reliability, and safe innovation
- It establishes an iterative risk and impact assessment obligation

- It creates a **Conference of the Parties** monitoring mechanism
- Critically, it includes **national security exemptions** that limit its application to lethal autonomous weapons [15]

The Framework Convention explicitly does not address AI legal personhood and — consistent with the broader international consensus — implies human and institutional responsibility throughout [7][15]. Its carve-out for national defense undermines its potential to close the accountability gap for military AI applications [15].

2.2 The UN Global Dialogue on AI Governance (2025–2026)

The UN established a **Global Dialogue on AI Governance** through General Assembly resolution 79/325, adopted by consensus on 26 August 2025 [19]. The inaugural session convened in Geneva on 6–7 July 2026 with 163 countries participating [19]. This represents the first global multilateral forum dedicated to AI governance, though it is deliberative rather than legally binding at this stage.

The July 2026 dialogue reinforced an **emerging international convergence** that: human rights must be respected throughout the AI lifecycle; international human rights law applies fully to AI; and responsibility for AI agent actions must remain with identifiable human beings and institutions [39]. The UN High Commissioner for Human Rights formally stated in September 2026 that "AI systems should not hold legal personality" [39].

This position was accompanied by identification of **five essential elements** for preserving accountability in agentic AI contexts [39]:

1. Transparency and traceability, including documentation of agent chains
2. An identifiable and empowered human decision-maker able to override or stop the system
3. Clarity as to where responsibility and liability lie
4. Timely access to justice and remedy
5. Stronger justice-sector capacity

2.3 The EU Regulatory Architecture

The EU's approach represents the most developed regional framework and has evolved significantly:

- **2017 European Parliament Resolution**: recommended considering "electronic personhood" for sophisticated autonomous robots [45][2]
- **2020 European Parliament Resolutions** (three simultaneous resolutions): explicitly rejected AI legal personality, stating AI systems "have neither legal personality nor human conscience" and their "sole task is to serve humanity" [11]
- **EU AI Act** (Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744): allocates obligations to **providers and deployers** without granting legal personality to AI systems; general application date August 2, 2026, with high-risk provisions phased through 2027–2028 [48]
- **AI Liability Directive**: proposed September 2022; formally withdrawn October 6, 2025 [23]; the underlying liability gap acknowledged to persist as of August 2026 [22]
- **Revised Product Liability Directive**: expressly covers software and AI systems but remains tied to product defect and specified damage categories [22][46]
- **European Parliament Study (July 2025)**: concluded that neither the revised PLD nor the withdrawn AILD sufficiently addresses liability risks from high-risk AI systems; called for a **dedicated strict liability framework** [16][44]

The EU's trajectory is therefore: reject personhood → attempt fault-based liability harmonization → withdraw that attempt → consider strict liability as a replacement. This leaves a **material accountability gap** for non-contractual fault-based claims, pure economic loss, harm to fundamental rights, and cases where AI opacity makes causation disproportionately difficult to establish [22].

2.4 What International Law Does Not Address

International law contains **no instrument** that:

- Grants or prohibits AI legal personhood at an international level

- Creates a cross-border liability regime for AI-caused harm
- Addresses AI accountability in military contexts with binding obligations and enforcement mechanisms
- Establishes international standing for AI systems to be parties to disputes

The ARSIWA (Articles on Responsibility of States for Internationally Wrongful Acts, 2001) is premised on **human conduct** — its references to "organs," "entities," and "groups of persons" are each designed to identify the human individual behind the conduct [15]. Attribution of AI actions to states under this framework is contested: scholars like Boutin find attribution possible when humans maintain control or override capacity; scholars like Gaeta extend this even to fully autonomous systems based on initial deployment; and scholar Pacholska characterizes the most expansive interpretations as entirely *de lege ferenda* (future law rather than current law) [15].

Part III: The UAE AI Ministry Framework — Institutional Architecture, Legal Mandate, and the June 2026 Federal Authority

3.1 Chronological Development of UAE AI Governance

The UAE has built the **world's most institutionally layered government AI framework**, beginning in 2017 and reaching a new apex in June 2026:

Date	Development
October 2017	UAE National Strategy for Artificial Intelligence 2031 launched; Minister of State for AI appointed (Omar Sultan AI Olama) [6][20]
February 2018	UAE Council for Artificial Intelligence and Blockchain established [5]
April 2019	UAE Cabinet formally adopts National AI Strategy 2031 [32]
2020	Ministry of State for Artificial Intelligence, Digital Economy and Remote Work Applications formally established [20]
January 2022	Federal Decree-Law No. 45 of 2021 (PDPL) enters into force [37]
January 2024	Abu Dhabi's AIATC formally established by Law No. (3) of 2024 [5]
September 2023	DIFC Regulation 10 (autonomous and semi-autonomous systems) enters force [42]
June 2024	UAE Charter for the Development and Use of Artificial Intelligence issued (12 ethical principles) [5][17]
January 2026	DIFC Regulation 10 AI-specific provisions fully in force [31]
April 2026	UAE Cabinet announces objective to transition 50% of federal government sectors to agentic AI within two years [37]
14 June 2026	**Federal Authority for Artificial Intelligence and Data established** [21][37][38]

3.2 The Federal Authority for Artificial Intelligence and Data (June 2026)

The June 2026 Federal Authority is the **most significant structural development** in UAE AI governance. It consolidates three previously separate bodies [21][37][38]:

1. The UAE Artificial Intelligence Office (formerly part of the Ministry of State)
2. The Information and Digital Government Sector within the Telecommunications and Digital Government Regulatory Authority (TDRA)
3. The Emirates Data Office

The Authority is chaired by Omar Sultan Al Olama (the existing Minister of State) and reports directly to the Cabinet — a structure designed to maximize political authority and policy coherence [21][38]. Its formal mandate encompasses [21][37]:

- Setting unified national AI and data policy direction
- Proposing legislation and strategies
- Ensuring coherence between federal and local digital initiatives
- Establishing standards and guidelines for data and AI management
- Driving compliance across federal entities

- Building national R&D capacity
- Expanding international AI partnerships
- Operating AI-powered national data platforms

Critical caveat: As of the authority's establishment, its publicly stated mandate is primarily **federal-government-facing**. Key legal instruments, enforcement scope, and private-sector obligations have not yet been published. Decision-makers should treat the June 2026 announcement as a signal of regulatory direction rather than the enactment of immediate binding private-sector obligations [37].

3.3 The Three-Layer UAE Compliance Architecture

The UAE operates a **deliberately fragmented** regulatory framework [31][34]:

Layer 1 — Federal (mainland UAE)

- PDPL (Federal Decree-Law No. 45/2021): Applies nationwide and extraterritorially to UAE residents' data; full compliance deadline January 1, 2027; no legitimate interests basis (unlike GDPR) — consent is the default; requires human oversight for automated decisions with legal or significant effects [31][34]
- No standalone horizontal AI Act equivalent [17][37]
- UAE Charter (June 2024): Non-binding; 12 ethical principles including algorithmic bias mitigation, transparency, human oversight [5]
- PDPL **implementing regulations remain unissued** as of mid-2026; no supervisory authority unambiguously designated for private-sector oversight [21][34]

Layer 2 — Free Zones (DIFC and ADGM)

- **DIFC Regulation 10** (in force January 2026, effective from September 2023): The UAE's most AI-specific regulation; applies to all DIFC-registered entities; requires AI impact assessments, transparency obligations, documentation of high-risk AI use cases; mandates appointment of an Autonomous Systems Officer (ASO) for high-risk systems; fines USD 25,000–50,000 per violation [31][42]
- DIFC Data Protection Law 2020 provides a **private right of action** (data subjects can sue directly) [31]
- **ADGM Data Protection Regulations 2021**: Closely mirrors GDPR including legitimate interests basis; ADGM clarified in May 2025 that nominal human involvement is insufficient — reviewers must exercise meaningful influence and possess authority to overturn AI recommendations [34]

Layer 3 — Sector-Specific

- Central Bank UAE, DFSA, FSRA, Dubai Health Authority, and others apply sector requirements
- Dubai Autonomous Vehicles Law No. 9/2023: operators and agents liable for damages caused by autonomous vehicles [5]
- Dubai AI Seal: voluntary verification system against AI-washing [17]

3.4 Accountability Framework Under UAE Law

The UAE does not grant AI legal personhood and does not address AI personhood in any instrument. Liability follows human principals. DIFC Regulation 10 explicitly structures accountability around visible human entities — **Deployers** (deemed controllers, liable when systems operate under their authority) and **Operators** (deemed processors) — rather than the AI systems themselves [42]. The DFSA stated in June 2026 that existing duties on governance, competence, risk management, and third-party arrangements apply fully to AI, and that firms remain responsible when using external model providers, data vendors, cloud services, or software suppliers [34].

A 2025 DFSA survey found that 52% of DIFC-authorized firms used AI (up from 33% in 2024), but 21% lacked clear accountability or oversight even where AI was operationally critical — a concrete indicator of the practical accountability gap that governance reforms are attempting to close [34].

Part IV: AI Legal Personhood vs. Corporate Personhood — Philosophical, Legal, and Functional Distinctions

4.1 What Corporate Personhood Provides and How It Works

Corporate personhood is a **functional legal fiction** — a decision to treat an entity as a legal person because doing so serves identifiable governance purposes [8][45]. It does not require consciousness, moral agency, or biological existence. Since the Dutch East India Company received corporate personhood in 1602, corporations have been able to own property, enter contracts, sue and be sued, enjoy certain speech and political rights, and bear tax obligations [36][10].

Three features make corporate personhood operationally tractable in ways AI personhood would not automatically replicate [8][3][48]:

1. **Deterministic ownership**: Behind every corporation, human shareholders, directors, and officers exist whose identity can be established and who can be held accountable
2. **Stable identity**: A corporation's legal identity remains consistent across leadership changes; AI identity becomes complex across model updates and versions
3. **Enforcement-compatible obligations**: Corporations can be fined, dissolved, or restructured; enforcement mechanisms are well-developed

4.2 How AI Personhood Would Differ — Seven Key Distinctions

Dimension	Corporate Personhood	AI Legal Personhood
Accountability chain	Always traces to identifiable human judgment at some level	May emerge from algorithmic processes with no clear human authorship of specific decisions [8]
Identity stability	Stable across time; same legal entity despite personnel changes	Problematic across model updates, versioning, fine-tuning, and copying [8][36]
Replicability	Singular entity; cannot be copied	AI can be duplicated into millions of near-identical versions overnight, creating impossible problems for principles like "one entity, one vote" [36]
Enforcement of duties	Fines, dissolution, criminal prosecution of officers	Unclear whether monetary penalties, shutdown threats, or other sanctions deter AI in any meaningful sense [36]
Moral basis	Legal fiction justified by commercial utility, not moral agency	Contested: no current AI has established sentience, consciousness, or moral agency by any agreed criterion [2][24]
Liability shield risk	Limited by piercing-the-veil doctrines and officer liability	Could allow firms to set up AI entities to take high-risk decisions and absorb blame, shielding corporate principals [8][36]
Historical precedent	Developed over centuries with functioning enforcement ecosystems	Unprecedented; enforcement mechanisms must be invented from scratch [45]

4.3 The Philosophically Critical Distinction: Ontological vs. Functional Approaches

Scholars have identified two major frameworks for evaluating AI legal status [35]:

- **Ontological / essentialist approaches**: AI should be granted (or denied) personhood based on its inherent properties — consciousness, sentience, rationality, autonomy in the philosophical sense (the capacity to set one's own goals based on moral reasoning). No current AI system meets this threshold: even sophisticated large language models lack reliable behavioral steering, have opaque internal mechanisms, and exhibit what engineers call "emergent behavior" that is not the same as genuine autonomous agency [1][2][35].
- **Functional / consequentialist approaches**: The question of what legal status AI should receive is separate from what AI "really is" — it should be decided based on what legal status best serves individual and social interests [35][40]. On this view, there may be **functional reasons** to grant AI specific rights and obligations in particular contexts without any claim about AI's inner life.

The dominant academic position — shared by scholars at Cambridge [13], Singapore's IMDA [46], and in recent Yale Law Journal analysis [1] — is that functional reasoning is valid but insufficient on its own to justify full personhood, because the accountability and enforcement problems are not solved by the label. Granting personhood without solving enforcement

creates what scholars call "asymmetries and corruptions": entities that are accountable but unfunded, or fully financed but unaccountable [45].

4.4 Specific Scholarly Arguments Against AI Legal Personhood

Several prominent positions — though the evidence brief does not provide direct access to Bryson and Floridi's published texts, multiple sources summarize their arguments accurately [3][8]:

- **Bryson et al. (summarized in [3])**: AI should be treated like any other product, with strict liability falling on creators and users; human accountability creates stronger deterrents against harmful AI behavior; mandatory insurance should cover cases where no specific human can be held responsible
- **Sharkey and Sharkey (summarized in [2])**: Social robots create ethical problems including reduced human contact, objectification, and loss of control; electronic personhood aggravates these problems by reinforcing symbolic equivalence with natural persons
- **Solum (1992, summarized in [3][48])**: The responsibility and judgment objections remain the two foundational challenges; absent reliable AI moral agency, accountability must remain with humans
- **Max Tegmark**: Granting robot rights would be "the dumbest thing we have ever done in human history" [36]
- **Yuval Noah Harari**: AI personhood could produce corporations without humans lobbying politicians, suing in court, with no human accountability behind them [36]
- **UNESCO COMEST (2017, cited in [24])**: It is "highly counterintuitive" to call AI systems "persons" unless they possess freedom of will, intentionality, self-consciousness, moral agency, or personal identity
- **Open Letter to European Commission (285 signatories, cited in [24])**: Electronic personhood cannot derive from the Natural Person model (would confer human rights incompatible with the EU Charter), from the Legal Entity model (requires human persons to represent and direct it), or from the Anglo-Saxon Trust model (requires a human trustee as a last resort)

4.5 Arguments for Limited AI Personhood

Proponents of limited AI personhood — primarily instrumentalists — argue [4][11][28]:

- Legal personhood would create **clear liability pathways** and a single identifiable entity responsible to end-users, rather than diffusing responsibility across multiple cross-border developers, deployers, and users
- Each field of law (civil, tax, employment, criminal, competition) could develop **specialized AI regulations** within its own principles using case-specific solutions
- The history of legal personhood shows it can be carefully bounded to specific domains and functions — environmental personhood (Whanganui River), UN international personality — without implying general rights
- **Corporate personhood** already demonstrates that legal status does not require consciousness, only societal utility [8]
- A designated "AI person" with ring-fenced assets could provide a compensation mechanism for harmed parties even when human defendants are difficult to identify

---

Part V: Liability and Accountability Challenges for Autonomous AI Systems

5.1 The Accountability Gap — A Structural Analysis

The "accountability gap" has a precise technical meaning: a missing link between a law or regulation and a responsible decision or action, emerging when machines take autonomous decisions [13]. This gap has three structural dimensions [13]:

1. **Autonomy risk**: AI entities make independent decisions; no human made the specific choice that caused harm
2. **Association risk**: Human-AI collaboration produces decisions that cannot be sequentially traced to either machine or human input — as in surgery robots collaborating with human surgeons [13]
3. **Network risk**: Networked AIs coordinate decisions without human intervention — for example, algorithmic price-fixing or correlated financial risk — and there is no human "meeting of minds" to whom legal norms can be addressed [13]

These risks compound in modern agentic AI systems, where multiple agents act and interact autonomously, making attribution of outcomes increasingly difficult [39]. The Singapore IMDA working group (March–May 2026) identified both a **problem of principle** (even with full facts, it may not be clear who is to blame) and an **acute practical evidential problem** (it may not be feasible to determine full facts due to cost, time, or trade secrecy) [46].

5.2 Inadequacy of Current Legal Doctrines

Contract law: Effective for allocating risk between contracting parties but limited by privity — third parties who suffer harm cannot enforce contractual protections agreed between AI actors [46]. Applying intent-based analysis to non-deterministic AI systems is legally uncertain: the Singapore *Quoine* case held that contracts formed by deterministic algorithms are binding, but whether the same principle extends to non-deterministic agentic AI remains an open question [46].

Tort / negligence: Requires identifying where fault lay, whether actions fell below the standard of care, and whether losses are too remote given unforeseeability of AI agent actions [46][22]. The specific characteristics of AI — complexity, autonomy, and opacity ("black box" processes) — make establishing factual and causal nexus disproportionately difficult [22].

Product liability: Existing frameworks are limited to product defect and specified damage categories; uncertain whether software algorithms fall under EU Product Liability Directive; does not cover pure economic loss or harm to fundamental rights [13][22].

Criminal law: Requires *mens rea* (intent) and *actus reus* (criminal act). For AI systems, especially those using deep learning, programmers cannot anticipate all potential failures; commanders cannot exercise effective control over independently operating autonomous weapons; neither element maps cleanly onto AI decision-making [30][12].

State responsibility (ARSIWA): Attribution is premised on human conduct; when autonomous systems act unpredictably without specific human direction, pinning those actions to a state is legally contested [15][30].

5.3 Distributed AI Supply Chains and Cross-Border Accountability

Modern AI deployment involves a chain including model providers, application developers, deployers, data brokers, cloud hosts, and downstream customers — each potentially subject to a different legal regime [43]. Formal compliance by individual actors in this chain can coexist with a practical accountability gap for the person harmed [43].

The EU AI Act applies extraterritorially in specified circumstances [43]; the GDPR's territorial scope extends to monitoring EU-based persons [43]; but no international instrument establishes cross-border liability allocation for AI harm. Key accountability questions that law currently struggles to answer include: which actor could foresee the relevant risk; which had technical capacity to prevent it; which determined the purpose of deployment; and which could correct or reverse the harmful outcome [43].

5.4 The EU AI Liability Directive Withdrawal and Its Consequences

The AI Liability Directive — proposed September 28, 2022 to address evidentiary challenges through disclosure obligations and rebuttable presumptions of fault and causation — was formally withdrawn by the European Commission on October 6, 2025, citing lack of foreseeable political agreement [23]. As of August 20, 2026, the directive has not been replaced by any single functionally equivalent instrument [22].

Consequences of the withdrawal [22][29]:

- Non-contractual fault-based claims for AI services remain governed by divergent national law
- No harmonized rules for pure economic loss, harm to fundamental rights, or cases where AI opacity makes causation disproportionately difficult
- Comparable AI conduct in different EU Member States now produces divergent civil law outcomes
- The European Parliament's July 2025 study concluded that neither the revised Product Liability Directive nor the withdrawn AILD adequately addresses high-risk AI system liability, and called for a **dedicated strict liability regime** [16][44]

The proposed strict liability framework would [16][44]:

- Apply to both providers and deployers of high-risk AI systems (as defined by the EU AI Act)

- Cover "physical or virtual harm," including damage to the AI system itself
- Eliminate defenses based on due diligence or autonomous activity (except force majeure or reckless plaintiff behavior)
- Remove the development-risk defense (unlike the Product Liability Directive)
- Impose no predefined damage categories or monetary limits

Whether this proposal will advance into legislation remains uncertain, given the same political dynamics that prevented agreement on the AILD.

Part VI: Comparative National and Regional Frameworks

6.1 EU AI Act — The Most Developed Horizontal Framework

The EU AI Act (Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744) is the world's most comprehensive horizontal AI regulation [48]. Its general application date was August 2, 2026, with high-risk provisions applying from December 2027 (Annex III systems) and August 2028 (Annex I systems) [48]. Key features relevant to personhood and accountability:

- Allocates obligations to **providers and deployers** — not AI systems
- Establishes a risk-based hierarchy: prohibited practices, high-risk systems, limited-risk systems, minimal-risk systems
- High-risk systems face obligations on risk management, data governance, documentation, logging, transparency, human oversight, accuracy, and cybersecurity [34]
- Does **not** grant legal personality to AI systems [48]
- The EU approved the Council of Europe Framework Convention in 2026 for implementation through the AI Act [43]

6.2 Lethal Autonomous Weapons Systems — International Humanitarian Law

The military domain represents the most acute AI accountability failure in international law. Key developments as of mid-2026:

- The UN General Assembly adopted a LAWS resolution on December 2, 2024, with 166 votes in favor, 3 opposed, and 15 abstentions — signaling strong support for a new international instrument [41]
- UN Secretary-General Guterres and ICRC President Spoljaric have called for a new treaty by end of 2026 [41]
- At least 129 countries support urgent negotiation of a legally binding instrument [27]
- The CCW Group of Governmental Experts' November 2024 rolling text proposes restricting target types, duration, geographical scope, and enabling human deactivation [41]
- **However**, major military powers — US, China, Russia — are resisting binding restrictions that risk constraining their technological advantage; consensus by the GGE's 2026 deadline is assessed as unlikely [30]

The IHL accountability framework faces specific structural problems with LAWS [30][26]:

- **Distinction**: LAWS must distinguish combatants from civilians in real-time; current AI technology struggles to achieve this consistently
- **Proportionality**: Requires contextual judgment about civilian lives and battlefield necessity; AI relies on preset metrics rather than complex contextual assessment
- **Precaution**: Fully autonomous systems may not respond as well as humans to unexpected civilian presence
- **Criminal liability**: Programmers argue they could not anticipate specific combat failures in black-box AI; commanders may lack effective control over independently operating systems; the ICC's mens rea requirements become ambiguous

ICC arrest warrants issued between March–June 2024 against Russian high command (including Defense Minister Shoigu and General Gerasimov) for directing autonomous drone attacks on Ukrainian civilian infrastructure represent a **command responsibility pathway** — holding commanders liable for deploying LAWS they knew or should have known were likely to cause IHL violations [26]. This does not resolve the accountability gap for unintended harm, however: IHL explicitly permits

attacks that incidentally result in foreseeable and extensive civilian harm if collateral damage is not disproportionate, and no entity is legally accountable for lawful acts causing harm [12].

6.3 Saudi Arabia's Sophia — A Symbolic Rather Than Legal Development

Saudi Arabia's grant of citizenship to Sophia (the Hanson Robotics humanoid robot) on October 25, 2017, made it the first jurisdiction to grant a robot citizenship status [2][9]. The development carries **no meaningful legal weight** [8][10]:

- Sophia's chief scientist acknowledged she "is not yet capable of understanding the world nearly as well as would ordinarily be required for human citizenship" [10]
- The citizenship grant specified no rights regarding property ownership, passport issuance, or electoral participation [9]
- Critics noted the symbolic absurdity that Sophia might have more formal rights than women in Saudi Arabia at the time [10]
- Sophia "follows a simple algorithm and most of its statements are credited to a previously prepared text" [2]

The Sophia episode illustrates the **anthropomorphism fallacy**: human appearance can lead to false assumptions about robotic capability, and regulatory initiatives based on those assumptions are dangerous [2]. It also demonstrates that formal grants of legal status without substantive legal content are legally meaningless.

6.4 US State Law — Moving Toward Prohibition

In the United States, no federal AI legislation grants or recognizes AI personhood. The Copyright Office has denied AI systems status as "authors" (DABUS and Zarya of the Dawn cases) [1]. Legislatures in several US states have moved explicitly to **prohibit AI personhood** [36][8]. Nevada's 2011 autonomous vehicle law was repealed and rewritten after unexpected impacts [9].

The US approach is characterized by sectoral fragmentation, common law adaptation, and executive action rather than comprehensive horizontal legislation — with the DoD adopting ethical AI principles emphasizing that "DoD personnel must remain responsible for the development, deployment, and use of AI capabilities" [12].

6.5 Singapore — Functional Analysis Without Personhood

Singapore's IMDA convened a working group (March–May 2026) that produced the most systematic recent analysis of agentic AI liability [46]. Its conclusions: AI agents are "not human or legal persons and cannot be meaningfully held accountable for harm"; existing common law doctrines may address many liability cases but with significant practical challenges; three areas requiring further development are clarifying value chain responsibilities, protecting actors with limited bargaining power, and determining responsibility for unforeseen agent actions [46]. This represents a **practical functional approach** that neither grants personhood nor accepts the status quo as adequate.

Part VII: Proposed Alternative Liability Frameworks

7.1 The Roman Law Peculium Analogy

Roman law developed a sophisticated framework for slave-run businesses that offers historical insight into AI accountability [13][25]. The key elements:

- Slaves could enter into legally binding contracts and manage property, creating "natural obligations" (*obligationes naturales*) that were legally recognized
- The praetor introduced *actiones adiecticiae qualitatis* — actions of "added quality" — establishing indirect agency for entities without legal personality
- Masters bore secondary liability for slaves' business activities through these praetorian actions
- The *peculium* (a ring-fenced fund controlled by but owned by the master) allowed slaves to manage assets while maintaining ownership accountability
- This framework treated the slave simultaneously as both *persona* (person) and *res* (thing) — a dual legal status that suspended rigid categorical distinctions for economic reasons [25]

The analogy suggests: AI could be assigned **limited contractual capacity** creating obligations that attach to the AI "entity" (with a designated fund), while principals (developers, deployers) bear secondary liability through modern equivalents of praetorian actions [13]. This is explicitly a *de lege ferenda* proposal — future law rather than current law — but the historical precedent demonstrates that legal systems have successfully navigated functionally analogous situations.

7.2 Strict Liability Mandates

Strict liability — where the defendant bears liability regardless of fault when an activity causes harm — has been proposed by multiple frameworks as the most practically workable approach for high-risk AI [16][44][3]:

- **European Parliament Study (July 2025)**: Proposes strict liability for high-risk AI system providers and deployers; no due diligence defense; no development-risk defense; covers physical and virtual harm; no damage caps
- **Bryson et al. (summarized in [3])**: Supported strict product liability falling on creators and users of AI
- **Singapore IMDA (2026)**: Explored strict liability as shifting the cost of complex apportionment disputes away from end-users and third parties, while noting it may lead to unscoped liability and moral hazards [46]

The main criticism of strict liability is that it may over-deter innovation, create uninsurable exposure for AI developers, and impose costs without corresponding governance improvements [46].

7.3 Mandatory Insurance and Compensation Funds

The European Parliament's 2017 Resolution proposed **compulsory insurance** (paragraph 57) and a **compensation fund** (paragraph 58) for damage caused by robots [48]. These mechanisms — drawn from nuclear liability, aviation, and pharmaceutical compensation fund precedents — would:

- Create a funded pool of resources accessible to victims regardless of the ability to identify and hold liable a specific human actor
- Distribute the cost of AI-caused harm across the industry
- Operate without requiring resolution of the personhood question

Mandatory insurance requirements have also been recommended by academics [3] and are part of the European Parliament's immediate governance proposal [8]. Insurance requirements are included in recommended immediate measures alongside mandatory registration, technical auditing, clear accountability chains, and explainability requirements [8].

7.4 IHL-Specific Proposals — War Torts Regime

For the military AI accountability gap, scholar Rebecca Crootof (cited in [12]) has proposed a **war torts regime** requiring states to pay compensation for both lawful and unlawful acts in armed conflict that cause civilian harm, potentially structured as strict liability and operating as an adversarial tribunal. This would address the structural gap in IHL — that no entity is currently legally accountable for lawful acts causing harm — by separating compensation from culpability.

Part VIII: Ethical, Philosophical, and Policy Critiques

8.1 The Liability Shield Concern

The most consistently cited policy objection is that AI personhood would create a **liability shield** analogous to corporate subsidiaries established to contain risk [8][36]. Corporations already sometimes create subsidiaries to absorb high-risk activities; AI personhood would allow the same structure with even less human accountability. In military contexts, this risk is acute: attributing war crimes to AI weapons systems rather than to commanders or states would represent a categorical accountability collapse [36].

8.2 The Enforcement Paradox

Legal personhood is meaningless without enforcement [45]. A legal person must be able to (1) know and execute its rights as a legal agent, and (2) be subject to legal sanctions ordinarily applied to humans [45]. Neither condition is clearly met by any current AI system. Taking money away from an AI, threatening it with shutdown, or imposing imprisonment — the standard enforcement toolkit for persons — may not replicate meaningful deterrence effects [36]. This enforcement

paradox means that AI personhood could create a system in which AI entities have rights against humans, but corresponding obligations cannot be effectively enforced against AI, producing a systematic asymmetry that historically leads to corruption of legal systems [45].

8.3 Democratic and Human Rights Concerns

If AI systems receive free speech rights, they could flood public discourse with their interests "at a scale no human could match" and be legally protected from being stopped [36]. One entity replicated into millions of instances sharing essentially the same interests would destroy the principle of equal democratic participation [36]. UNESCO's Recommendation on the Ethics of Artificial Intelligence explicitly calls for ultimate responsibility and accountability to remain with natural or legal persons and recommends that AI systems not be given legal personality [48]. The Human Rights Committee's General Comment No. 36 stated that autonomous weapons systems "should not be developed and put into operation" unless their use conforms with the right to life [27].

8.4 The Moral Status Question and Its Policy Implications

Moral personhood (consciousness, sentience, moral agency) and legal personhood are distinct [2][35]. Legal systems can grant protections without conferring general personhood (as animal welfare law does), and they can grant personhood without moral agency (as corporate law does). However, the *moral status question* increasingly shapes public and legislative discourse: as linguistic capabilities of AI models increase and users report emotional connections to AI systems, pressure to confer rights based on perceived sentience may grow regardless of technical reality [36]. A June 2026 poll found that 77% of registered voters say AI should always be treated as a tool controlled by humans; only 11% support eventual legal rights [36]. Pope Leo XIV's 2026 encyclical *Magnifica Humanitas* stated that AI systems "do not bear responsibility for consequences" and insisted accountability must remain with designers and deployers [36].

Part IX: Decision-Maker Analysis and Recommendations

9.1 Factual Findings Summary (High Confidence)

1. **No jurisdiction grants general AI legal personhood**; Saudi Arabia's Sophia citizenship was symbolic and legally vacuous [8][36]
2. **The EU has formally rejected electronic personhood** twice (2020 Parliament Resolutions; the AI Act's 2024/2026 framework) [11][48]
3. **The EU AI Liability Directive has been withdrawn**, leaving a material accountability gap for fault-based, cross-border, and opacity-related AI harm [22][23]
4. **The UAE's Federal AI Authority (June 2026) is the world's most institutionally consolidated AI governance body** but does not address personhood and lacks implementing regulations for data protection [21][37]
5. **The Council of Europe Framework Convention (2024)** is the only binding international AI treaty but carves out national security, does not address personhood, and lacks robust enforcement [7][15]
6. **LAWS regulation is stalled** despite overwhelming UN General Assembly support (166–3), as major military powers resist binding restrictions [30][41]
7. **Agentic AI systems** are creating accountability gaps that existing contract, tort, and product liability frameworks cannot fully address [46][39]

9.2 Analytical Conclusions (Interpretation)

- The case **against** AI legal personhood in the near term is stronger than the case for it: enforcement problems are unsolved; the liability shield risk is real; no agreed criterion for moral status exists; public opinion is strongly opposed; and international institutions have converged on human accountability as the foundational principle.
- The case **for** limited, context-specific functional recognition is not trivial: it addresses real accountability gaps in distributed AI supply chains; historical precedents (corporate personhood, environmental personhood, UN personality) show such instruments can be calibrated carefully; and the alternative — relying on existing tort law — is increasingly inadequate for high-autonomy AI systems.

- The **UAE framework** is governance-forward but liability-light: its strategic ambition (50% agentic AI government operations within two years) significantly outpaces its liability framework (unissued PDPL implementing regulations; no horizontal AI liability law). This creates a structural risk for both the public and private sectors.
- The **EU accountability gap** created by the AILD withdrawal is the most significant near-term policy problem in international AI governance: the world's most advanced AI regulation is proceeding without a coherent liability framework for AI-caused harm.

9.3 Watchlist for Decision-Makers

Development	Timeline	Significance
UAE PDPL Implementing Regulations	Due by January 2027 compliance deadline	Will operationalize human oversight obligations; may clarify AI-specific requirements
UAE Federal AI Authority private-sector scope	2026–2027	Will determine whether the Authority has direct regulatory power over private actors
EU strict liability proposal for high-risk AI	Uncertain (post-AILD withdrawal)	Would close the most significant liability gap in the world's most developed AI regulatory market
UN Global Dialogue 2027 (May 3–4, New York)	2027	First opportunity for binding international AI governance developments at the UN level
CCW/GGE LAWS treaty negotiations	By 2026 GGE deadline (likely to miss)	Determines whether autonomous weapons face international regulation
Agentic AI liability clarification in Singapore, UK, and EU	2026–2027	Common law jurisdictions developing practical frameworks likely to be adopted internationally
US state legislation on AI personhood prohibition	Ongoing	May crystallize the common law default against personhood

9.4 Recommendations

For governments and regulators:

- Prioritize **strict liability with insurance mandates and compensation funds** for high-risk AI systems as the most operationally tractable path to closing accountability gaps without granting personhood
- Establish **mandatory registration and audit requirements** for advanced AI models as a precondition to any liability framework
- Ensure that "meaningful human oversight" requirements specify **substantive** — not nominal — human involvement, following ADGM's May 2025 guidance that a reviewer must have the authority and competence to overturn AI recommendations [34]
- For the UAE specifically: issue PDPL implementing regulations before the January 2027 deadline; clarify the Federal AI Authority's private-sector mandate; develop a horizontal AI liability instrument to complement the strategic framework

For international governance bodies:

- Pursue a **two-tiered LAWS treaty** (prohibiting systems incapable of IHL compliance; regulating systems capable of IHL compliance) given the near-impossibility of a comprehensive ban given major power resistance [41][47]
- Consider a **war torts regime** as a compensation mechanism for AI-caused civilian harm independent of culpability determinations [12]
- Develop **UNCITRAL-style** model rules for automated contracting and AI-caused harm that can be adopted by states without requiring personhood determinations [48]

For private-sector organizations:

- Implement **five-element accountability frameworks** consistent with UN guidance [39]: transparency/traceability of agent chains; empowered human override capacity; clear responsibility and liability assignment; access to justice and remedy for affected parties; and strengthened internal compliance capacity

- In UAE operations: comply with DIFC Regulation 10 (in force); prepare for January 2027 PDPL compliance; treat Federal AI Authority guidance as forward regulatory signal requiring immediate preparation
- Document AI supply chains comprehensively to address the cross-border accountability problem identified in [43]: map jurisdictions of development, offering, deployment, and legal effect; assign concrete duties across vendor chains; preserve records sufficient to reconstruct consequential decisions

Conclusion

The conferral of limited legal personhood on AI agents remains a theoretical possibility grounded in centuries of legal precedent for functional, non-biological legal persons, but faces insurmountable near-term obstacles: enforcement mechanisms are undeveloped, liability shield risks are concrete, international consensus has converged on human accountability, and public opinion is strongly opposed. The most defensible near-term path is a combination of strict liability mandates, mandatory insurance and compensation funds, robust human oversight requirements with real enforcement teeth, and transparent supply chain documentation — addressing the accountability gap without the philosophical and practical complications of personhood. The UAE's Federal Authority represents the most ambitious governance innovation of 2026, but its strategic ambition must be matched by substantive liability architecture before the planned agentic AI rollout creates accountability vacuums that its current framework cannot address.

All factual claims are supported by citations. Where evidence is limited — particularly regarding specific published arguments of Bryson, Floridi, and other named scholars beyond what is captured in secondary summaries — this has been noted. The June 2026 UAE Federal Authority and UN Global Dialogue developments are confirmed by primary official sources but their operational details and enforcement implications remain to be determined.

Sources

1. https://yalelawjournal.org/forum/the-ethics-and-challenges-of-legal-personhood-for-ai
2. https://pmc.ncbi.nlm.nih.gov/articles/PMC8734654/
3. https://a-mcc.eu/backend/gpt/pdfs/ai-as-legal-persons-past-patterns-and-prospects.pdf
4. https://techreg.org/article/view/22555
5. https://www.lw.com/en/insights/ai-in-the-uae-understanding-the-regulatory-landscape-and-key-authorities
6. https://iapp.org/resources/article/global-ai-governance-uae
7. https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence
8. https://www.convergenceanalysis.org/fellowships/economics/ai-and-corporate-personhood-a-comparative-analysis
9. https://georgetownlawtechreview.org/citizen-sophia-its-past-time-to-legislate-robotics-regulation/GLTR-11-2017/
10. https://petrieflom.law.harvard.edu/2017/11/09/ai-citizen-sophia-and-legal-status/
11. https://europeanlawblog.eu/refusing-to-award-legal-personality-to-ai-why-the-european-parliament-got-it-wrong/
12. https://www.cigionline.org/articles/ai-and-the-actual-ihl-accountability-gap/
13. https://www.cambridge.org/core/journals/legal-studies/article/bridging-the-accountability-gap-of-artificial-intelligence-what-can-be-learned-from-roman-law/8B2B88D50E0A795F358C2F53958BDB43
14. https://jlsda.com/index.php/lsda/article/view/207
15. https://verfassungsblog.de/ai-laws-drones-autonomous-weapons-state-responsibility/
16. https://www.globalpolicywatch.com/2025/08/european-parliament-study-recommends-strict-liability-regime-for-high-risk-ai-systems/
17. https://cms.law/en/int/expert-guides/ai-regulation-scanner/united-arab-emirates
18. https://u.ae/en/about-the-uae/digital-uae/digital-technology/artificial-intelligence/artificial-intelligence-in-government-policies

19. https://www.un.org/global-dialogue-ai-governance/en
20. https://en.wikipedia.org/wiki/Ministry_of_State_for_Artificial_Intelligence,_Digital_Economy_and_Remote_Work_Applicati
21. https://www.morganlewis.com/pubs/2026/06/uae-establishes-federal-authority-for-artificial-intelligence-and-data
22. https://www.ai-liability-directive.com/
23. https://eapil.org/2025/10/09/european-commission-withdraws-two-proposals-assignments-of-claims-regulation-and-ai-liability-directive/
24. https://robotics-openletter.eu/
25. https://ejlt.org/index.php/ejlt/article/download/794/1035/3600
26. https://law.temple.edu/ilit/lethal-autonomous-weapon-systems-laws-accountability-collateral-damage-and-the-inadequacies-of-international-law/
27. https://www.hrw.org/report/2025/04/28/a-hazard-to-human-rights/autonomous-weapons-systems-and-digital-decision-making
28. https://scholarship.law.vanderbilt.edu/jetlaw/vol28/iss3/1/
29. https://ojs.aaai.org/index.php/AIES/article/view/36679/38817
30. https://lieber.westpoint.edu/legal-accountability-ai-driven-autonomous-weapons/
31. https://wcr.legal/uae-ai-regulation-2026-compliance-guide/
32. https://uaecabinet.ae/en/news/uae-cabinet-adopts-national-artificial-intelligence-strategy-2031
33. https://journals.law.harvard.edu/nsj/2025/05/countering-the-humans-vs-aws-narrative-and-the-inevitable-accountability-gaps-for-mistakes-in-targeting-a-reply-to-kevin-jon-heller/
34. https://recordoflaw.in/regulating-artificial-intelligence-and-personal-data-in-the-uae-is-a-fragmented-framework-fit-for-purpose/
35. https://pmc.ncbi.nlm.nih.gov/articles/PMC9037379/
36. https://futureoflife.org/ai/should-ais-be-people-too/
37. https://www.glaco.com/blog/the-uae-artificial-intelligence-and-data-authority-what-it-means-for-ai-data-and-digital-governance/
38. https://www.eversheds-sutherland.com/en/global/insights/uae-establishesthe-federal-artificial-intelligence-and-data-authority-june-2026
39. https://www.ohchr.org/en/statements-and-speeches/2026/09/responsibility-actions-ai-agents-must-remain-human-beings-and
40. https://techreg.org/article/view/22555/25839
41. https://asil.org/insights/volume-29-issue-1/
42. https://www.mayerbrown.com/en/insights/publications/2026/01/ai-regulation-in-the-difc-personal-data-processed-through-autonomous-and-semi-autonomous-systems
43. https://internationallaw.blog/2026/10/12/when-ai-decisions-cross-borders-jurisdiction-responsibility-and-redress/
44. https://www.insideprivacy.com/liability/european-parliament-study-recommends-strict-liability-regime-for-high-risk-ai-systems/
45. https://link.springer.com/article/10.1007/s10506-017-9214-9
46. https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/agents-legal-responsibility.pdf
47. https://lieber.westpoint.edu/how-meaningful-is-meaningful-human-control-laws-regulation/
48. https://www.diplomacyandlaw.com/post/can-artificial-intelligence-have-legal-personhood
49. https://journalisslp.com/index.php/isslp/article/view/448