Europe Withdrew Its AI Liability Directive, and the Exposure It Left Behind Is Larger

Basil C. Puglisi, MPA
A Human-AI Collaboration

For executives shipping AI-enabled products into the European market, the lawyers defending them, and the insurers pricing the risk. Verified as of June 2026.

Two Laws Changed, and the Safer-Sounding One Is the Trap

In 2022 the European Union proposed two liability laws as a matched pair. One, the Artificial Intelligence Liability Directive, would have harmonized fault-based claims for harm caused by AI, easing the burden of proof for a person suing over an AI system that injured them. The other, a revision of the decades-old Product Liability Directive, would update strict liability for defective products to reach the digital age. The plan was for the two to work together, one covering fault, the other covering defect.

Only one survived, and the survivor is the one that may raise exposure most directly for AI-enabled products. The Artificial Intelligence Liability Directive is gone. The European Commission listed it for withdrawal in its 2025 Work Programme on February 11, 2025, citing no foreseeable agreement among Member States, and the withdrawal became formal when the notice appeared in the Official Journal on October 6, 2025. Many companies read that news as a regulatory burden lifting. The revised Product Liability Directive tells a different story, and the practical lesson is the first fact worth pairing with an action. Because the fault directive is gone while the product directive expanded, a company that treats the withdrawal as relief and pauses its documentation work increases its exposure, so the tactic is to keep building the evidence file on the assumption that the regime to prepare for is strict product liability rather than the easier fault standard, and the measure is whether the company can produce a complete technical and decision record on demand before December 2026.

Directive (EU) 2024/2853 was adopted on October 23, 2024, entered into force on December 8, 2024, and must be transposed into national law in every Member State by December 9, 2026. It repeals the 1985 Product Liability Directive and applies to products placed on the market or put into service after the transposition date. The change that matters for anyone building with AI is the definition. The new directive redefines a product to include software, including AI systems, and digital manufacturing files. Strict liability, the regime that holds a maker responsible for a defective product without the victim having to prove negligence, now reaches the software itself. The action that follows is an inventory: a company maps which of its AI-enabled offerings will be placed on the EU market after the deadline, since those carry the new strict-liability exposure, and the measure is a dated product list tied to release schedules that cross December 9, 2026.

The reach extends beyond Europe's borders the way product rules tend to. A company outside the European Union can be drawn into the regime through an EU importer, authorized representative, or fulfilment service provider when it places an AI-enabled product on the European market, so a firm in New York selling AI-driven diagnostic software or an autonomous system into Europe is exposed. This article covers the liability directives alone. The EU AI Act sets separate obligations and penalties that apply alongside them, the General Data Protection Regulation governs the personal data these systems process, and sector rules add further layers, each addressed on its own terms.

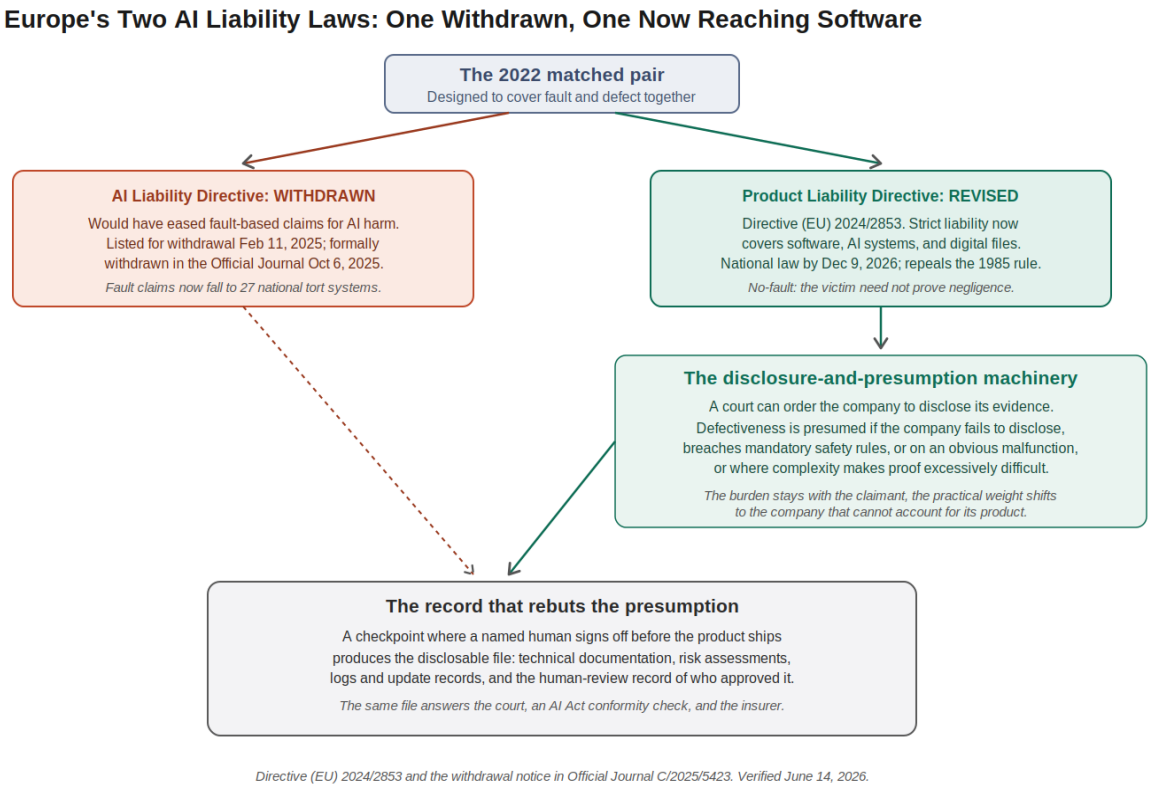

Figure 1: The two-law split, where the AI Liability Directive was withdrawn and the revised Product Liability Directive extended strict liability to software and AI.

Strict Liability Now Reaches Software as a Product

The revised directive keeps strict liability as its core and adapts the procedure around it to the realities of complex technology. A claimant still bears the burden of proving the defect, the damage, and the causal link between them. What changed is the machinery that surrounds that burden, and the machinery is where the exposure lives.

The first piece is disclosure. When an injured person presents a plausible claim, a national court can order the company to disclose the evidence in its control, subject to protections for genuine confidentiality. The second piece is a set of rebuttable presumptions that lower the claimant's

burden when certain conditions hold. A court presumes the product was defective if the company fails to disclose the evidence it was ordered to produce, if the product breached mandatory safety requirements, or if the claimant shows an obvious malfunction. The directive treats cybersecurity as a safety question explicitly, listing relevant cybersecurity requirements among the factors that decide whether a product is defective, and it limits a maker's ability to escape liability where the defect comes from a failure to supply the security updates needed to keep the product safe. A court presumes the causal link when the product is found defective and the damage is of a kind typically consistent with that defect. And where a claimant faces excessive difficulty proving defectiveness or causation because of the technical or scientific complexity of the product, the court may presume the element the claimant struggled to prove. The directive's own recitals name complex software and medical devices as the cases this provision is built for.

Read together, these provisions point at a single instruction, and it is the fact that drives the whole governance argument. The presumption of defectiveness for non-disclosure exists to make disclosure worthwhile, so the company that cannot produce its records faces the presumption it might otherwise have rebutted. The burden of proof formally stays with the claimant, yet the practical weight shifts to the company that arrives in court unable to show how its product was built, tested, and maintained. The tactic this fact demands is to treat the disclosable record as a deliverable of the build process rather than a document assembled after a claim, and the measure is the time it takes the company to produce, from a standing archive, the design, test, and decision records for any shipped AI product. For an AI system, whose behavior is hard to reconstruct after the fact, that production time is the difference between rebutting the presumption and inheriting it.

The withdrawal of the fault-based directive sharpens the point rather than softening it. Strict product liability covers defective products, and a great deal of AI harm sits outside that category. Discriminatory outcomes in hiring or lending, privacy violations by automated systems, and pure economic loss from an algorithmic error may not fit neatly as product-defect claims, and with the harmonizing fault directive gone, those harms fall to the national tort law of each Member State. Twenty-seven systems now apply, with different rules on fault, causation, and evidence, so a person harmed by the same AI system in the same way may have a strong claim in one country and no realistic path in another. The withdrawn directive left the exposure in place and fragmented it across those systems, which makes it harder to plan around.

Why Autonomous Agents Need a Named Human Risk Owner

The structure of the directive bears hardest on one deployment pattern in particular, and naming it plainly matters because the market is moving toward it fast. An agentic AI system, one that acts across tools and takes consequential actions without a human approving each output, is automation: the machine checks the machine, and no named human is bound to what it ships.

For a high-risk system the EU AI Act requires, in Article 14, that the system be designed so a natural person can effectively oversee it in use, with the ability to monitor it, interpret its output,

and disregard, override, or stop it. A company can read that as a design obligation satisfied by building in those controls, and a defense will argue that a monitoring dashboard and a working stop function meet it even when no person is watching.

The reading argued here is that the oversight the Act calls effective is oversight a person actually exercises, because a control no one uses is latent capability rather than exercised oversight, and a reviewer who only ratifies what the system already produced is the rubber stamp that regulators have said does not pass. The Act assigns that work to two different parties, and keeping them apart is what makes the liability argument precise. Article 14 is a duty on the provider, the party that builds the system, to design it so a natural person can effectively oversee it. Article 26 is a duty on the deployer, the party that puts it to use, to assign a named person with the competence and authority to do the overseeing. The Product Liability Directive holds the manufacturer strictly liable for a defective product, and its recitals point toward treating the provider of an AI system as that manufacturer where that actor places the product on the market or otherwise fits the directive's economic-operator role, so under this argument the strongest version runs through a system whose design itself prevents oversight, which is a provider failure under Article 14. A breach of a mandatory safety requirement laid down in Union law is one of the conditions under which a court presumes a product defective, so a claimant could argue that a design defeating effective oversight is that kind of breach and carries into the liability directive, and whether a court accepts the bridge would turn on the system, the harm, the national implementation, and whether the obligation reads as a safety requirement rather than a procedural step. The harder case is the deployer that runs a perfectly governable system without governing it. The product is sound, so the exposure sits with the deployer as an Article 26 failure, and it travels through the deployer's own duties and national law rather than the producer-centric presumption, until the deployer puts its own name on the system, makes a substantial modification that keeps it high-risk, or repurposes it into a high-risk use, at which point Article 25 of the AI Act can bring the deployer within the provider's obligations, and the manufacturer's strict liability can follow. The boundary matters and is worth stating plainly: the argument is cleanest for high-risk systems whose design defeats oversight, and for agentic products outside the high-risk category the relevant hook is the general product-safety rules rather than Article 14. The named-human question does not disappear outside that category, because where an automated decision carries legal or similarly significant effects, the GDPR supplies a separate right to human intervention under Article 22, which the SCHUFA ruling read broadly. Across these regimes the exposure follows from whether a person held binding authority over the output, and the pattern that emerges is that liability tends to move toward whoever held control of the substance, whatever the system's level of autonomy. The operating response follows directly: an organization classifies every AI deployment by whether a named human holds binding authority to approve, modify, or halt its consequential outputs, and the deployments that lack that authority are not placed on the EU market in a strict-liability posture without it. The measure is the share of consequential AI outputs that carry a named, authorized human sign-off, and the target for anything inside the product-liability regime is complete coverage.

Article 25 is, on this reading, where the EU comes closest to writing that principle into binding law, and it appears to sit ahead of other jurisdictions in doing so for the AI value chain. Accountability tracks control, and the party that takes command of the substance inherits the maker's exposure, which is the line the governance argument has drawn throughout. Courts are starting to draw it too. In May 2026 the Regional Court of Munich I held Google directly liable for false statements its AI Overviews generated about two publishers, treating the company as a direct source of the output rather than a neutral intermediary because it builds and controls the system. That ruling is a preliminary injunction under German personality and competition law, it is not a product-liability decision, and Google has said it will appeal. It is offered here as a sign of where courts are heading on control and accountability, with no suggestion that it settles how the product-liability directive applies to AI. Lawmakers are moving in the same direction. Italy's June 2026 implementing decrees, still in preliminary approval and not yet final law, would tie a presumption of causation to a proven AI Act breach and let an injured person reach the technical documentation and the operator's insurer directly, which is the same logic this analysis draws, that the disclosable record decides who carries the presumption.

The standard for what makes that sign-off meaningful is drawn here from GDPR case law, where the courts have already tested it. In the SCHUFA ruling the Court of Justice of the European Union treated an automated credit score as the decision itself where it played a determining role in the outcome, so a downstream human who merely ratified it did not bring the decision outside the automated-processing prohibition, and the Article 29 Working Party guidance requires that the human review be meaningful, carried out by someone with the authority and competence to change the decision. Read into a product setting, a person placed only at the final gate is not a real checkpoint when the system held the determining role over the substance, because the points where a person exercises genuine control are the checkpoints, and a finding that the system controlled the substance is a finding that those checkpoints were missing.

The market data shows how wide this gap runs. A 2026 Grant Thornton survey of roughly 950 senior business leaders found that 78 percent lack strong confidence that their organization could pass an independent AI governance audit within 90 days, and only 5 percent allow agents to execute high-stakes decisions without human review. The 5 percent figure is the encouraging one, because it means the overwhelming majority already keep a human in the loop for high-stakes calls, which is the posture the liability regime favors when a claim arrives. The 78 percent figure is the warning, because keeping a human in the loop is worth little if the organization cannot produce the record proving it did. The tactic is to convert informal human review into a documented checkpoint, and the measure is whether the organization could, today, pass the 90-day audit it currently doubts.

There is more than one way to place that human, and the right one depends on how fast the outputs arrive. Where a person can read or watch the consequential outputs as they happen, holding the authority to disregard, override, or halt them, the checkpoint is continuous, a live position held in real time. Where outputs occur faster than any person could follow them one by one, an algorithmic trading system or an automated cyber response, the checkpoint moves

upstream instead, and a named officer signs the policy that authorizes a defined class of automated actions so that accountability attaches to that signature. Where a product ships in discrete releases, the checkpoint is the sign-off before each release. All three are real checkpoints, because in each one a named human exercises binding authority over the substance. The pattern that fails is the same across all three, a human who is present without being engaged, the signature or the dashboard that ratifies whatever the system produced. The Act itself draws only one bright line at the level of the individual output, requiring for remote biometric identification, and in most uses outside law enforcement, that no decision be taken unless two competent persons separately confirm it, and it leaves the degree of engagement open everywhere else, which is the space an exercised checkpoint is built to fill. The deciding line is whether a named human answers for what the system ships, whatever its level of autonomy, and that binding can sit at the policy level for a class of actions, at a live position during operation, or at the output level for a single consequential decision, as long as it exists and is documented.

Why Documentation Beats the Presumption, and Why It Is Not Enough on Its Own

The same machinery that creates the exposure also names the evidence that answers it, and that evidence is the upside for a company that holds it before a court asks. Because the presumption of defectiveness is triggered by a failure to disclose, the records a company keeps are the difference between rebutting the presumption and inheriting it. The governance-pays case is now measurable rather than asserted. The same 2026 Grant Thornton survey found that organizations with fully integrated, governed AI are ten times more likely to pass an independent governance audit and nearly four times more likely to report revenue growth than those still piloting without controls. The fact pairs cleanly: governance correlates with both the defense and the growth, so the tactic is to build the control set once and let it serve audit-readiness, litigation defense, and the revenue case together, and the measure is the audit-pass rate alongside the revenue attributable to governed AI.

Several artifacts carry that weight, and they are the documents a regulator or an opposing lawyer expects to see. Technical documentation showing how the system was designed and tested. Risk assessments completed before the product shipped. For a high-risk AI system, the file the EU AI Act already requires, the event logs, the data-governance records, the post-market monitoring. Records of updates and security patches, since a product that was safe at release can become defective through a later change the maker controlled. And the human-review record, the contemporaneous proof that a named person checked a consequential output, held the authority to change it, and approved it.

A distinction inside this evidence is where most organizations go wrong, and it is the reason documentation alone does not settle the question. A system can produce logs, traceability, and a full record of its own decision tree, machine checking machine, and still have no named human bound to the output. That posture, often labeled Responsible AI, is automation that documents

itself. The logs and the decision tree satisfy the technical-documentation and logging obligations, which the EU AI Act sets out separately in Articles 11 and 12. They do not satisfy the oversight obligation in Article 14, because a log is the machine accounting for what it did, not a person accounting for the decision to ship it. The record proves the automation ran. It does not produce a named human who can answer for the output, and that gap is exactly the line between Responsible AI and AI Governance: a record is not a governor. An opposing lawyer will press on that gap, because a system that only ever checked itself, however completely it logged the checking, has no human answer to a question about why a defective product reached the market. The tactic is to add the named-human checkpoint on top of the automated controls rather than treating the logs as a substitute for it, and the measure is whether, for any consequential release, the disclosure file names the person who approved it and shows the authority they held to stop it.

These artifacts share a quality worth holding onto. They are recognized, auditable, and credited by the people who decide exposure, the court weighing a disclosure order, the regulator checking conformity, and increasingly the insurer pricing a product-liability or technology policy. A company that can produce them holds evidence that travels, because the same technical file that rebuts a presumption of defectiveness may also satisfy an AI Act conformity check and answer an underwriter's question about controls. Meeting the AI Act helps build that file, though it does not by itself defeat a product-liability claim.

There is a second reward worth naming without overselling it. A company that builds a real human checkpoint into the path a product takes to market, rather than a final rubber stamp, tends to catch the defects and edge cases that an ungoverned pipeline would have shipped. The benefit there reaches past liability. It is the quality and the trust that come from a person who can stop a flawed release, and it shows people kept in the loop rather than removed from it. That claim is modest and depends on the organization, and it is offered as an observation rather than a guarantee.

Governance That Produces the Evidence

The mechanism that turns this exposure into those artifacts is a checkpoint, a defined point in the workflow where a named human must sign off before a product ships or an automated output takes effect, exercising binding authority that cannot be delegated to the AI being governed. The recognized baselines for what good practice looks like are the established standards, the NIST AI Risk Management Framework and ISO/IEC 42001, recognized voluntary baselines that increasingly inform standard-of-care and underwriting discussions. Checkpoint-Based Governance, a method proposed by the author, is one approach to producing the specific evidence the liability regime calls for, offered as one way to do the work rather than as a standard in its own right.

The connection to the directive is direct. The presumption of defectiveness rewards the company that can disclose a clean record and penalizes the one that cannot. A checkpoint built into the release path produces, as a byproduct, exactly that record, who reviewed the product, against what

tests and data, with the authority to halt the release. The governance and the documentation are the same act. The person governs the decision to ship, and the trail of that governance is the artifact a court later reads when it decides whether the presumption applies. This is the move that separates a defensible deployment from an exposed one, and that line turns on whether a person with authority to stop the release actually stood behind it, whatever the system's level of autonomy.

The pattern is visible in a single pair of directives. The liability a company must avoid is a presumption of defectiveness flipping against it because it cannot show how its AI product was built and maintained. The growth it wants to claim is a defensible, insurable position in one of the largest consumer markets in the world, and the survey evidence ties that growth to the same governance that produces the defense. Both come from the same move, putting a named human with real authority at the point where the product ships, and keeping the record.

The Limits Worth Naming

A documented checkpoint is persuasive evidence rather than a guarantee, and the honest limits deserve naming. A disclosure file cuts both ways, since a thin, contradictory, or self-incriminating record handed to a court becomes evidence of the defect rather than a defense against it, so the governance has to be real to help. The directive is not yet uniformly in force as national law, because the transposition deadline is December 9, 2026, and implementation is uneven, with Hungary having completed its transposition, roughly a third of Member States having published draft bills, and neither France nor Italy having yet transposed the directive itself, though Italy is moving early on adjacent ground, with implementing decrees under its national AI law given preliminary approval by its Council of Ministers in June 2026. The rules apply only to products placed on the market after that date, so products already in service stay under the 1985 regime, and the national implementations will vary in their detail.

The timing of the AI Act itself is in motion, and it cuts in a direction worth naming. Through the Digital Omnibus, the institutions reached a provisional agreement in May 2026 to defer the high-risk obligations, moving the Annex III systems from August 2026 to December 2027 and the Annex I embedded systems to August 2028, with formal adoption anticipated in July 2026. That deferral does not touch the product-liability timeline, since the transposition deadline stays at December 9, 2026, so the liability regime arrives on schedule while the governance obligations it leans on slip more than a year. It also complicates the cleanest version of the Article 14 argument, because a defendant can point out that the high-risk design duty is not yet in force for many systems when the strict-liability regime begins to bite. The harm, though, does not wait for the governance clock. Product liability, the GDPR, and national sectoral law all reach AI-caused harm now, which is the reading argued throughout, that the exposure runs ahead of the obligations meant to manage it.

The fault gap is real and unresolved, since the withdrawn directive would have addressed the discrimination, privacy, and pure-economic-loss harms that strict product liability does not reach,

and those claims now depend on whichever Member State's tort law applies. And adopting a standard and then ignoring it can deepen exposure, because a documented process a company failed to follow supports an argument that it should have known better. None of this removes the value of the record. It means the record must reflect real oversight, honestly kept.

The shape of European AI liability for the rest of this decade is now set by the law that survived. Strict liability reaches the software, the burden shifts in practice to the company that cannot account for its product, and the evidence that rebuts the presumption is the same evidence good governance produces as a matter of course. The methods to produce that record are published and available, and the deadline to have them in place arrives in December 2026.

Sources

Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC, Official Journal L, 2024/2853, 18 November 2024 (entry into force December 8, 2024; transposition by December 9, 2026; product definition extended to software, AI systems, and digital manufacturing files; disclosure obligations and rebuttable presumptions of defectiveness and causation).

European Commission, withdrawal of Commission proposals, Official Journal C/2025/5423, 6 October 2025 (formal withdrawal of the proposed Artificial Intelligence Liability Directive, 2022/0303(COD), originally proposed as COM(2022) 496).

European Commission, Commission Work Programme 2025, COM(2025) 45, 11 February 2025 (listing the Artificial Intelligence Liability Directive for withdrawal, citing no foreseeable agreement).

Council Directive 85/374/EEC of 25 July 1985 on liability for defective products (the prior strict-liability regime, repealed with effect from December 9, 2026).

Regulation (EU) 2024/1689 (the EU AI Act), high-risk obligations including risk management, data governance, logging, technical documentation, and post-market monitoring, which apply alongside the liability regime. Article 14 (human oversight, a provider design obligation), Article 25 (responsibilities along the AI value chain, reclassifying a deployer as a provider on rebranding, substantial modification, or repurposing into a high-risk use), and Article 26 (deployer obligation to assign and exercise human oversight).

Landgericht München I (Regional Court of Munich I), Case 26 O 869/26, judgment of 28 May 2026 (preliminary injunction holding Google directly liable for false statements generated by its AI Overviews, under German personality and competition law, classifying Google as a direct source rather than a neutral intermediary; under appeal). Cited as evolving case-law direction, not as product-liability precedent.

Italy, Law No. 132/2025 and the implementing decrees given preliminary approval by the Council of Ministers on 10 June 2026 (aligning national law with the AI Act; introducing a rebuttable presumption of causation tied to an established AI Act breach, access to technical documentation, and a direct claim against the operator's insurer). Cited as draft national law in preliminary approval, not yet final, as evolving implementation direction.

Digital Omnibus on AI, provisional political agreement of 7 May 2026 (deferring the AI Act high-risk obligations, Annex III standalone systems to 2 December 2027 and Annex I embedded systems to 2 August 2028; formal adoption and Official Journal publication pending, anticipated July 2026). The product-liability transposition deadline of 9 December 2026 is unaffected.

Grant Thornton, 2026 AI Impact Survey (approximately 950 senior US business leaders, fielded early 2026): 78 percent lack strong confidence they could pass an independent AI governance

audit within 90 days; 5 percent allow agents to execute high-stakes decisions without human review; organizations with fully integrated, governed AI are ten times more likely to pass an independent governance audit and nearly four times more likely to report revenue growth.

NIST AI Risk Management Framework (AI RMF 1.0) and ISO/IEC 42001:2023, recognized governance baselines.

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

> The Other AI: Audio Briefings on Augmented Intelligence and AI Governance
> Spotify: https://open.spotify.com/show/033dvhzMIcWLdY7IUgsu7F
> Apple Podcasts: https://podcasts.apple.com/us/podcast/id1896506152
> Amazon Music: https://music.amazon.com/podcasts/923d1a79-533f-4623-bae3-e2ba83453dfb
> YouTube Playlist: https://www.youtube.com/playlist?list=PLchpU2bIYoEEBh2hdY-BVP9ckyTPiHOAQ

#AIassisted using HAIA Ecosystem