The Fundamental Dilemma of the Foundation Model Ecosystem: A New Proposal for a Sustainable AI Era

The Dawn of a New Paradigm

Artificial intelligence (AI) technology is heralding the dawn of a new era, with "Foundation Models" at its core. Coined by the Stanford Institute for Human-Centered AI (HAI), a foundation model is a large-scale, general-purpose AI model trained on vast data that can be applied to a wide range of tasks.[1] This represents a paradigm shift from traditional AI, which was limited to specific tasks. Now, once built, these models can be reused as the foundational infrastructure for multiple applications.[3]

The emergence of these models has driven explosive growth in the AI ecosystem. In 2023 alone, 149 new foundation models were released, more than double the number from 2022.[4] Their application scope is expanding infinitely, moving beyond text and images to include video, audio, and even world foundation models that simulate the physical world.[4]

However, behind this dazzling progress lies a structural characteristic known as "homogenization".[2] The phenomenon where a few massive foundation models become the basis for countless applications increases efficiency but is also a double-edged sword that concentrates the entire system's risk in one place. This goes beyond mere technical glitches. If the "foundation" model itself harbors legal and ethical flaws, it creates a systemic vulnerability where the thousands, or even tens of thousands, of applications and services built upon it could collapse in an instant. The current AI ecosystem stands precariously on this unstable ground, and

AHNIST aims to solve this structural problem and present a future of sustainable technology [5, 5].

The Invisible Architecture of AI

To understand the problems within the AI ecosystem, we must first grasp its technical structure. The AI Tech Stack is generally composed of four core layers [5]:

1. **Infrastructure Layer:** The physical hardware foundation, including semiconductor chips like GPUs, servers, and the cloud that process AI model computations.
2. **Data Layer:** The domain where vast amounts of data for training models are collected, stored, and managed.
3. **Model Layer:** The core of AI's intelligence, where trained foundation models, LLMs, and other models reside.
4. **Application Layer:** The services and software, like ChatGPT, that end-users directly interact with.

The AI workflow—"Data Collection → Model Training → Content Generation → Distribution → Monetization"—operates across these layers.[5] However, the current tech stack is missing a crucial element: the "Meta-Infrastructure" layer that governs the fundamental rules of the ecosystem.[8]

The Meta-Infrastructure should be positioned between the data and model layers, serving as the ecosystem's registry and trust system. It must verify data provenance, prove content authenticity, track intellectual property (IP), and automate value exchange. The absence of this layer is the root cause of all the problems plaguing the current AI ecosystem. AHNIST provides the proprietary solution to build this essential, missing layer.[8]

Table 1: The AI Tech Stack and the Missing Meta-Infrastructure Layer

Layer	Description	Key Components	Governing Principle
Application	AI services that end-users interact with	ChatGPT, AI Editors, Chatbots	User Experience
Model	The core intelligence of AI, learned from data	Foundation Models, LLMs, Fine-tuned Models	Prediction & Generation
Meta-Infra (AHNIST)	**** Governs the rules and trust of the ecosystem	**Content Wallet, Absolute Identifier (AID), Content Chain**	**Provenance, Authenticity, Ownership Proof**
Data	Collection and management of data for AI model training	Databases, Data Lakes, Web Scraping	Scale & Quality
Infrastructure	The physical foundation for AI computation	GPUs, Cloud Servers, Data Centers	Computing Power

The Trilemma at the Heart of the AI Ecosystem

The absence of a meta-infrastructure creates three fundamental, seemingly unsolvable dilemmas—a "trilemma." These are not separate issues but a single, complex crisis with interconnected parts.[8]

3.1 The Provenance Paradox: A World Built on an Unknowable Foundation

To achieve their performance, foundation models learn from vast, indiscriminately collected data from the internet, the origins of which are often unknown.[1] This inevitably leads to a crisis of "Data Provenance," or the verification of data sources.[8]

This problem is no longer theoretical. In 2023, the New York Times (NYT) filed a massive copyright infringement lawsuit against OpenAI and Microsoft for the use of millions of its articles.[9] The NYT presented over 100 examples of ChatGPT "memorizing" and "regurgitating" its articles nearly verbatim, arguing this was not mere learning but clear copyright infringement.[9] Potential damages in this lawsuit could range from at least billions to hundreds of billions of dollars, posing an enormous legal risk that could shake the foundations of the AI industry.[9]

Such lawsuits are not limited to the NYT. Countless creators, including writers, artists, and computer programmers, have filed suits claiming their works were used for AI training without permission, highlighting a systemic problem across the entire AI industry.[12] The lawsuit filed by Getty Images against Stability AI demonstrates how difficult it is to prove copyright infringement after training is complete, underscoring the need for a proactive provenance tracking system rather than a reactive one.[14]

3.2 The Authenticity Abyss: The Collapse of Digital Trust

As the quality of AI-generated content improves dramatically, it is becoming nearly impossible to distinguish between human creations and AI-generated ones.[8] This is leading to the collapse of "Content Authenticity," eroding the foundation of trust in the entire digital world.[16]

In 2023 alone, at least 500,000 deepfake videos and audio clips were shared on social media [16], and public trust in media has plummeted to an all-time low.[17] The severity of this problem extends beyond the spread of fake news. It gives rise to a

secondary effect known as the "Liar's Dividend".[18] This phenomenon occurs when the mere existence of deepfake technology gives malicious actors a pretext to dismiss unfavorable "real" evidence as "fake".[19] Ultimately, as the line between truth and falsehood blurs, the very basis for social consensus and public discourse is eroded, entrenching a "post-truth" environment.[17]

3.3 The Value Vacuum: The Black Hole of Intellectual Property and Creator Earnings

The current AI ecosystem treats the intellectual labor of millions of creators not as a valuable asset, but like a free raw material to be mined.[8] There is no mechanism to track which creations an AI model has learned value from, nor to distribute fair compensation for that contribution.[8]

This creates a paradox in the "Creator Economy." While the market size of the creator economy is projected to reach approximately $500 billion by 2030 [20], that wealth rarely flows back to the creators. Statistics show that only 4% to 12% of full-time creators earn a sustainable income (over $50,000 annually).[21]

This unfair distribution structure has already materialized in the music streaming market. While platform companies generate billions in revenue, the revenue per stream returned to artists is a fraction of a cent, a figure that has only stabilized after years of decline.[23] The structure where platforms monopolize distribution and take the majority of the revenue (in the Korean music market, distributors/producers take 84%) is being replicated in the AI ecosystem.[8]

In conclusion, these three dilemmas stem from a single root. Because the **provenance** of data is unknown, the **authenticity** of content cannot be guaranteed, and the contributed **value** cannot be fairly distributed. This is a singular crisis born from the absence of a reliable "registry system" to record and certify the entire history of a digital asset from its birth. Solving this problem is the prerequisite for

building a sustainable AI ecosystem.

The Industry's Response: A Necessary but Flawed First Step

In response to these issues, the industry has begun to seek solutions voluntarily. The most prominent example is the "Coalition for Content Provenance and Authenticity (C2PA)," led by Adobe and joined by major companies like Microsoft and Google.[25]

C2PA's core technology is a standard called "Content Credentials".[27] This method involves attaching a "Manifest"—a set of provenance information like the creator and edit history (metadata)—to a digital file and signing it with cryptographic technology to verify if it has been tampered with.[28] This technology is a significant and meaningful attempt to increase content transparency.

However, C2PA has a fundamental Achilles' heel: its security relies on metadata, akin to a "label attached" to a file, rather than the essence of the content itself. This approach reveals critical vulnerabilities in the real-world internet environment where content is constantly copied and transformed.

- **Metadata Stripping Vulnerability:** C2PA's Content Credentials can be easily and permanently removed.
    - **File Format Conversion:** One of the most common actions, such as converting a JPG file to a PNG or compressing and re-encoding it for platform upload, can cause the loss of C2PA metadata.[30]
    - **Screen Capture:** The most common way to share content on the internet, taking a screenshot, creates a completely new image file with none of the original's metadata. This is a fatal flaw for a system that aims to track content across the internet.[32]
    - **Platform Incompatibility:** Many social media platforms have historically stripped unnecessary metadata during the upload process, which leads to

the loss of C2PA information.[35]

C2PA itself acknowledges this problem. Its official security documentation lists "Stripping C2PA Manifests" as a threat that cannot be prevented.[28] As a solution, it proposes complex "Soft Binding" technology that links to watermarks or external databases, but this is merely a reactive supplement, not a complete solution.[36]

Ultimately, C2PA's fundamental limitation is not a technical bug but a philosophical and structural problem. It is an attempt to impose order on a chaotic digital environment by attaching external "labels." But in an environment where content duplication and modification are routine, a security system that relies on "labels" is bound to fail. A truly robust solution must be inherent to the digital asset's essence—its "digital DNA"—not an external label.

A New Foundation: The AHNIST Meta-Infrastructure Solution

AHNIST presents a new paradigm that fundamentally overcomes the limitations of metadata-based approaches like C2PA. It is a "Meta-Infrastructure" solution that embeds security and ownership into the essence of the content, not an external label.

5.1 From Unstable Metadata to Immutable Hardware: The Content Wallet

The core of AHNIST's technology is the "Content Wallet." This is an innovative concept of a hardware-based "'physical wallet for digital data'".[8]

The core principle is simple and powerful. The original data and the certificate proving its authenticity are inseparably generated and stored inside a physically

secure, dedicated device. Security becomes an intrinsic physical property embedded at the asset's birth, not a detachable "label".[8] This resolves the C2PA's metadata stripping problem at its source. Even if an asset managed by a Content Wallet is captured via screenshot, it is merely an image without the value of the original. The unique, tradable, and legally valid original remains securely inside the owner's Content Wallet.

5.2 An Identity Rooted in Reality: Absolute Identifier (AID) and Democratic OTP (D-OTP)

To prevent copyright infringement and evasion of responsibility based on anonymity, AHNIST proposes a new identity authentication system that remedies the shortcomings of existing Decentralized Identifiers (DID).[8] This system is comprised of the "Absolute Identifier (AID)" and "Democratic OTP (D-OTP)."

This system verifies identity through encrypted communication between two physical devices owned by the user: a "Content Wallet" and a "node" such as a smartphone.[8] This inseparably links a digital identity (AID) to an individual's physical ownership. This prevents the creation of fake accounts at the source and assigns clear legal accountability for all digital activities, laying the groundwork for trustworthy transactions.

5.3 The Human Chain: A Scenario of Technology in Action

Let's explore how AHNIST's technology works in practice through a hypothetical scenario.

- **Step 1 (Creation & Proof):** An independent musician, 'Hana,' finishes a new song. She 'mints' this master audio file through her 'Content Wallet.' At this moment, a unique hash value for the audio is generated, and it is registered as an absolute original. Simultaneously, Hana's AID is recorded as the creator, and the song's 'digital DNA' is born. All this information is recorded on the immutable 'Content Chain' [5, 5].
- **Step 2 (Distribution & Transaction):** Hana decides to sell this song as a limited edition of 1,000 copies, setting the price and usage conditions directly from her wallet's interface. A fan named 'Jin' initiates a P2P transaction to purchase the song.
- **Step 3 (Authentication & Duplication):** When Jin's device requests the purchase, the network verifies Jin's identity (AID) and payment. Once everything is confirmed, Hana's Content Wallet authorizes the creation of a legitimate licensed copy for Jin and transmits it. The transaction record—'Hana sold to Jin, at a specific time, for a specific price'—is also recorded on the Content Chain.[8]
- **Step 4 (Value Distribution):** The payment is instantly and automatically deposited into Hana's account, without complex intermediary distribution processes or high platform fees. She directly owns the majority of the value she created.

This scenario demonstrates how AHNIST's system simultaneously solves the trilemma. **Provenance** is embedded in hardware at the moment of creation, **authenticity** is guaranteed by a physical device, and **value** is transparently tracked and fairly distributed on the chain.

Table 2: Comparative Analysis of Content Authenticity Solutions

Feature	AHNIST (Patent-based)	C2PA / CAI	Numbers Protocol	Other DID Solutions
Core Principle	Hardware-based original proof	File metadata attachment	NFT-based registration	Digital identity verification
Original Proof Method	Immutable original exists in wallet	In-file/external metadata	NFT issuance on blockchain	N/A
Metadata Stripping Vulnerability	Immune	High - Removed by screenshot, format conversion	Medium - NFT can be separated from asset	N/A
Identity Accountability	AID linked to physical device	Relies on signer's reputation	NFT owner address	Lacks physical entity, allows anonymous creation
IP Tracking & Licensing	Chain-based automation	Lacks feature (manual/off-chain)	Limited tracking via NFT transaction history	Lacks feature
Creator Revenue Distribution	P2P direct transaction	Platform-intermediated	NFT transaction fees	N/A
Legal Enforceability	Strong - Clear chain of evidence	Weak - Evidence chain easily broken	Medium - Utilizes blockchain logs	Weak

Source: Reconstructed based on AHNIST IR materials [8] and C2PA technical analysis

Conclusion: A Proposal for a Sustainable and Fair AI Future

The path the current AI ecosystem is on is not sustainable. Built on data of unknown origin, with the lines between truth and falsehood blurring, and exploiting the labor of creators without fair compensation, the current paradigm is fraught with massive legal, social, and ethical crises.

We are now at a crossroads. This is not merely a choice between different technologies or products. It is a choice between a future riddled with chaos and exploitation, and one based on order and fairness.

What AHNIST proposes is not just an app or a service. It is a 'social innovation movement' born from the firm philosophy that the distorted structure of the digital economy must be corrected and that the benefits of technology should be for all, not just a few.[8] AHNIST's technology provides the new social infrastructure essential for the digital age: a true "Meta-Infrastructure" that enables a trustworthy and fair AI ecosystem. This technology will be the first step toward realizing the decentralized, democratic internet ideal dreamed of by the web's founders and ushering in a new era where the rights of all creators are respected.

참고 자료

1. Foundation model - Wikipedia, 7 월 14, 2025 에 액세스, https://en.wikipedia.org/wiki/Foundation_model
2. Reflections on Foundation Models | Stanford HAI, 7 월 14, 2025 에 액세스, https://hai.stanford.edu/news/reflections-foundation-models
3. What Are Foundation Models? - IBM, 7 월 14, 2025 에 액세스, https://www.ibm.com/think/topics/foundation-models
4. What Are Foundation Models? | NVIDIA Blogs, 7 월 14, 2025 에 액세스, https://blogs.nvidia.com/blog/what-are-foundation-models/
5. AI Tech Stack Solutions - Intel, 7 월 14, 2025 에 액세스, https://www.intel.com/content/www/us/en/learn/ai-tech-stack.html
6. The AI Tech Stack: A Business Guide to the Moving Parts of the AI Ecosystem,

7 월 14, 2025 에 액세스, https://www.pymnts.com/news/artificial-intelligence/2025/ai-tech-stack-business-guide/
7. AI Tech Stack: A Guide to Frameworks & Best Practices, 7 월 14, 2025 에 액세스, https://www.spaceo.ai/blog/ai-tech-stack/
8. 1 월 1, 1970 에 액세스,
9. The New York Times v. OpenAI: The Biggest IP Case Ever - Sunstein LLP, 7 월 14, 2025 에 액세스, https://www.sunsteinlaw.com/publications/the-new-york-times-v-openai-the-biggest-ip-case-ever
10. The New York Times Case against OpenAI is Different. Here's Why. - Patent Docs, 7 월 14, 2025 에 액세스, https://www.patentdocs.org/2024/02/the-new-york-times-case-against-openai-is-different-heres-why.html
11. NYT v. OpenAI: The Times's About-Face - Harvard Law Review, 7 월 14, 2025 에 액세스, https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/
12. AI and intellectual property rights - Dentons, 7 월 14, 2025 에 액세스, https://www.dentons.com/ru/insights/articles/2025/january/28/ai-and-intellectual-property-rights
13. AI, Copyright, and the Law: The Ongoing Battle Over Intellectual Property Rights, 7 월 14, 2025 에 액세스, https://sites.usc.edu/iptls/2025/02/04/ai-copyright-and-the-law-the-ongoing-battle-over-intellectual-property-rights/
14. Getty Images v Stability AI: main copyright claims dropped - Pinsent Masons, 7 월 14, 2025 에 액세스, https://www.pinsentmasons.com/out-law/news/getty-images-stability-ai-copyright-claims-dropped
15. Navigating representative actions: takeaways from Getty Images v Stability AI, 7 월 14, 2025 에 액세스, https://www.hsfkramer.com/notes/ip/2025-01/navigating-representative-actions-takeaways-from-getty-images-v-stability-ai
16. How Do Deepfakes Affect Media Authenticity? - Identity.com, 7 월 14, 2025 에 액세스, https://www.identity.com/deepfake-ai-how-verified-credentials-enhance-medi

a-authenticity/
17. How Deepfakes Are Impacting Public Trust in Media - Pindrop Security, 7 월 14, 2025 에 액세스, https://www.pindrop.com/article/deepfakes-impacting-trust-media/
18. Deepfakes and Democracy (Theory): How Synthetic Audio-Visual Media for Disinformation and Hate Speech Threaten Core Democratic Functions - PubMed Central, 7 월 14, 2025 에 액세스, https://pmc.ncbi.nlm.nih.gov/articles/PMC9453721/
19. Understanding the Impact of AI-Generated Deepfakes on Public Opinion, Political Discourse, and Personal Security in Social Media - IEEE Computer Society, 7 월 14, 2025 에 액세스, https://www.computer.org/csdl/magazine/sp/2024/04/10552098/1XApkaTs5I6
20. 32+ Creator Economy Statistics of 2025 (Market Size Data) - Demand Sage, 7 월 14, 2025 에 액세스, https://www.demandsage.com/creator-economy-statistics/
21. Creator Economy Statistics 2024: Everything You Need to Know - Shanna Lindinger, 7 월 14, 2025 에 액세스, https://www.shannalindinger.com/articles/creator-economy-statistics
22. 22 Eye-Opening Creator Economy Statistics (2025 Data) - Blogging Wizard, 7 월 14, 2025 에 액세스, https://bloggingwizard.com/creator-economy-statistics/
23. Duetti's 2024 Music Economics Report Finds Industry-Wide Per Stream Rates for Independent Artists Are Finally Stabilizing Following Years of Decline - PR Newswire, 7 월 14, 2025 에 액세스, https://www.prnewswire.com/news-releases/duettis-2024-music-economics-report-finds-industry-wide-per-stream-rates-for-independent-artists-are-finally-stabilizing-following-years-of-decline-302358273.html
24. Streaming Music Revenue Shows Sluggish Growth in 2024 as Industry Faces New Challenges - Vinyl Me, Please, 7 월 14, 2025 에 액세스, https://www.vinylmeplease.com/blogs/music-industry-news/streaming-music-revenue-shows-sluggish-growth-in-2024-as-industry-faces-new-challenges
25. C2PA | Verifying Media Content Sources, 7 월 14, 2025 에 액세스, https://c2pa.org/
26. The Content Authenticity Initiative - Transparency in the Age of AI - YouTube, 7 월 14, 2025 에 액세스, https://www.youtube.com/watch?v=q2VpOuctSF4

27. Content Credentials, 7 월 14, 2025 에 액세스, https://contentcredentials.org/
28. C2PA Specifications :: C2PA Specifications, 7 월 14, 2025 에 액세스, https://c2pa.org/specifications/specifications/2.2/index.html
29. How it works - Content Authenticity Initiative, 7 월 14, 2025 에 액세스, https://contentauthenticity.org/how-it-works
30. The Role of Blockchain in Securing Content Without C2PA Labels - Numbers Protocol, 7 월 14, 2025 에 액세스, https://www.numbersprotocol.io/blog/blockchain-secures-content-without-c2pa
31. Unable to apply Content Credentials - Adobe Support, 7 월 14, 2025 에 액세스, https://helpx.adobe.com/creative-cloud/help/cai/unable-to-apply-content-credentials.html
32. How C2PA can safeguard the truth from digital manipulation - SC Media, 7 월 14, 2025 에 액세스, https://www.scworld.com/perspective/how-c2pa-can-safeguard-the-truth-from-digital-manipulation
33. Research Integrity, Image Manipulation, Content Provenance and the C2PA, 7 월 14, 2025 에 액세스, https://scholarlykitchen.sspnet.org/2025/03/13/research-integrity-content-provenance-and-c2pa/
34. Frequently-asked questions | Open-source tools for content authenticity and provenance, 7 월 14, 2025 에 액세스, https://opensource.contentauthenticity.org/docs/faqs/
35. Do not strip C2PA manifests from EXIF of uploaded photos · Issue #35100 - GitHub, 7 월 14, 2025 에 액세스, https://github.com/mastodon/mastodon/issues/35100
36. C2PA Security Considerations, 7 월 14, 2025 에 액세스, https://c2pa.org/specifications/specifications/1.0/security/Security_Considerations.html
37. C2PA Security Considerations, 7 월 14, 2025 에 액세스, https://c2pa.org/specifications/specifications/2.0/security/Security_Considerations.html