

The Human Data Collective: A Standard for Open, Sustainable, and Equitable Artificial Intelligence

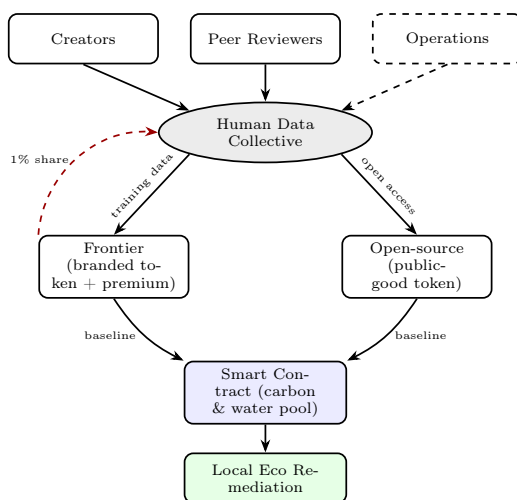
With transparency and compensation optionality for large language model inference

Austin Harshberger

Happy Stack Calculus · Los Angeles County
x@97115104.com · <https://links.97115104.com>

April 2026

Abstract. Several of the biggest challenges plaguing the development and utilization of large language models in 2026, notably bias, the difficulty of monitoring undesirable and illegal content, the psychological safety risks associated with human-led alignment, and severe environmental impact, can be partially or completely mitigated through the careful curation of high-quality human-generated content paired with fair incentives. This paper offers an alternative to the existing paradigm of extraction and privatization by proposing a decentralized network of human-generated content created, curated, and labeled specifically for the purpose of training large language models. By introducing a new inference standard rooted in cryptographically verifiable tokens tied to environmental impact, those who do the creating, curating, and labeling can be fairly compensated, the carbon and water footprint of large language models can be addressed at compute time, and the artificial intelligence industry can be returned to its original mission of open access and equality for all while preserving productive free-market competition.



attest v3.0	2026-04-22-77bcda
<i>Human-AI collaboration.</i>	
Content	The Human Data Collective
Author	A. Harshberger
AI collaborators	Claude Opus 4.7, Gemini 3.1 Pro (Preview)
Role	collaborated (multi-prompt iterative authoring; human directed structure, argument, and final wording; AI assisted with drafting, citation lookup, LaTeX formatting, and revision under human review)
Authorship	collab
Prompt type	multi-prompt
Prompt summary	<i>multi-prompt iterative authoring of LaTeX paper proposing a decentralized human data network for LLM training</i>
Platform	GitHub Copilot in VS Code
Timestamp	2026-04-22T04:35:54.795Z
Signature	HMAC-SHA256 936e2483...a7dd9a1f
Verify	https://attest.97115104.com/s/r11xzkrm

Contents

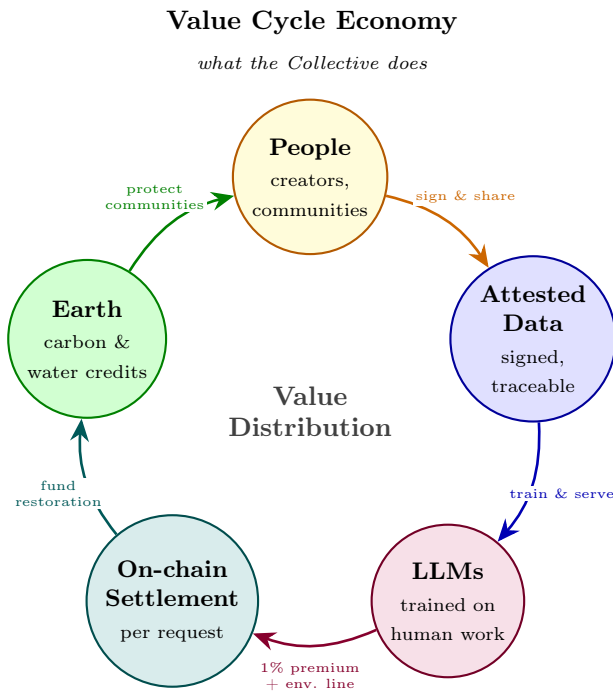
Preface and Author's Note	iii
1 Brief on LLMs and Their Consumer Product Manifestations	1
2 Current Centralization and Pricing of LLMs is Problematic	1
3 Environmental and Social Costs	2
4 Murky Data	3
5 Proposed Solution: The Human Data Collective	4
6 For Frontier Providers	5
7 For Contributors	5
8 For Governments and Regulators	6
9 For the Artificial Intelligence Industry	6
10 Considerations	6
A New Inference Standard	7
B Lifecycle of a Single Request	8
C A Brief for Regulators	9
D The Human Data Collective in a Post-AGI World	10
E Cost Formulas	11
F Glossary of Terms	11
G Index of Named Persons and Organizations	12
Interested in Contributing?	14
Acknowledgements	15
References	16

Preface and Author’s Note

This paper is written about, not yet with, the contributors and communities whose labor and lived environments give my proposal here its motivating force; published reporting and the public record are its sources. It is a draft for public comment and not yet a ratified standard, with the v0.1 wire format published openly as AIISP-1 [10].

As the author of this paper, I believe a fair approach that incentivizes all would usher in an era of higher quality training inputs, resulting in higher quality outputs, making future LLMs more aligned with the best of humanity, that is community, creativity, culture, and equality, as opposed to the worst, isolation, extraction, marginalization, anxiety, and individualism. The mechanism for creating human-generated data with shared incentives also creates a secondary marketplace for data generation, aggregation, and labeling, ushering in new job opportunities that fulfill the promise of AI as a public good which generates jobs, propping up human productivity and creativity. And, most importantly, by prioritizing and designing for environmental impact, it’s possible to create a future where AI development is a positive driver in policy around sustainable technology regulation. My hope is that this proposal and similar standards move the line in an effort to give credit where it is due, and to amplify the voices of those whose contributions are often overlooked, and sometimes, silenced.

As you will learn from reading the proposal outlined below, if even one of the largest frontier providers adopted this standard at current 2026 inference volumes, the practical consequence for the broader community would be on the order of hundreds of millions of dollars per year flowing into verifiable environmental remediation and direct compensation for the people whose work made the underlying systems possible. Together, we can make a difference. As the 1997 Apple campaign put it, “the people who are crazy enough to think they can change the world are the ones who do”¹.



Data from people trains models, which return a share of value through royalties to creators and through carbon and water remediation to the communities that host the compute.

Implementation Premise

how providers plug in

```
POST /v1/inference
X-AIISP-Token: 0xA1B2...
X-AIISP-Settlement: deferred
```

Mechanically the standard is two optional HTTP headers attached to any existing inference call. The first names a branded ERC-20 token that identifies the provider on Base, the settlement chain. The second selects **deferred** batched settlement (the typical case) or **realtime** per-request settlement (useful for audit).

Requests *without* the headers behave exactly as they do today. Requests *with* them route the energy line, the environmental line, and a 1% share of provider premium through the on-chain contract that powers the cycle on the left, distributing the proceeds to creators, peer reviewers, and the carbon and water funds.

Optionally read by role. Frontier-provider leadership and public-affairs teams should start with §6 (*For Frontier Providers*); creators, researchers, and labor organizers should start with §7 (*For Contributors*); legislators and policy staff should start with §8 (*For Governments and Regulators*) and Appendix C (page 9); engineers and protocol designers should start with §9 (*For the Artificial Intelligence Industry*) and the AIISP-1 specification in Appendix A (page 7).

For the full cost decomposition and AIISP-1 wire-format specification see Appendix A (page 7); for a single request walked end-to-end with worked dollar figures see Appendix B (page 8).

¹From the 1997 Apple “Think Different” campaign; copy by Rob Siltanen and Ken Segall at TBWA\Chiat\Day, narrated by Steve Jobs in the unaired version of the spot [35].

1 Brief on LLMs and Their Consumer Product Manifestations

Large language models *are* the data that was used to create them [4]. This framing draws directly on Bender, Gebru, McMillan-Major, and Shmitchell’s *Stochastic Parrots*, which argued, well in advance of the current frontier race, that the corpora used to train large language models encode the social and economic conditions of their collection, and that those conditions, including who was paid, who was credited, and what was excluded, propagate forward into every output the resulting model produces. If frontier models are trained mostly on public and freely accessible material such as millions of public GitHub repositories, Wikipedia, and free-to-use educational databases intended to benefit all, why then are the largest companies, namely Microsoft, OpenAI, Anthropic, and Google, privatizing access to that data and charging the people who created it for the privilege? Material that is not clearly marked as open source but was nonetheless used for training before it was knowable, including copyrighted books, personal blog posts, YouTube transcripts, and other socially shared content on the internet, is rarely attributed to its original creators and almost never fairly compensated.

The value and broad applicability of large language models has grown exponentially over the past several years to make them one of the most polarizing and disproportionately impactful technologies ever to come to fruition. Despite the early stage of these systems as consumer products through APIs, chat interfaces, image generators, and video generators, the quality of their responses remains questionable, with internal frontier-lab testing showing that on so-called closed-domain prompts, where the correct answer is supplied directly in the prompt and any deviation from it is definitionally a hallucination, an early version of GPT-4 fabricated answers roughly 30 percent of the time [9, ch. 11].² The publicly released system card for the most widely deployed model in the same family continues to report meaningful hallucination rates across its safety evaluations [19], and recent peer-reviewed work formalizes why this happens, demonstrating that under standard pretraining and reinforcement-learning-from-human-feedback objectives, models are statistically incentivized to guess [12]. Most of the quality of these models is tied directly to the volume and provenance of the data used to train them [13], with the largest of the frontier models’ value being directly proportional to the amount of data and compute available in their development. The current winners, unsurprisingly, are the biggest names in technology, namely Meta, Alphabet, Microsoft, and X, who either directly, in the case of Meta, Alphabet, and X, or indirectly, in the case of OpenAI and Anthropic through Microsoft via GitHub, train their models on their users’ data

²In this paper *hallucination* is used in its technical sense, namely fluent and internally plausible model output that is factually incorrect or unverifiable against any external source, including fabricated citations, statistics, and quotations; see the Glossary entry on page 11 for a fuller definition.

or on publicly accessible, free-to-use, easy-to-obtain-without-accountability internet sources.

The internet itself is the most relevant precedent for the benefits of broad, free, and general access. At the time of its creation, the argument to privatize and sell access to it was actually stronger than the equivalent argument for software-only large language models, in large part because running the internet and providing access to it requires physical servers and networked infrastructure that themselves carry inherent energy and capitalization costs [16]. Despite this, the internet flourished as a free and open tool through government grants and research institutions, ushering in a new era of exploding value for both individuals and society and precipitating the creation of some of the largest companies that have ever existed, notably Alphabet’s Google search engine, which made its money in large part by surfacing internet data and ranking it according to a combination of marketing spend and user interest. OpenAI itself was originally founded as a non-profit research company with the explicit remit to advance digital intelligence in the way most likely to benefit humanity as a whole [18], and Anthropic was launched as a public benefit corporation focused on artificial intelligence safety and the long-term interests of society [2]. However, as both organizations grew, the cost of developing frontier models grew vertically with them, encouraging both to take on capped-profit subsidiaries and accept investment from organizations whose privatization and monetization goals are in many ways antithetical to their original remit. The resulting discord caused real strife across the artificial intelligence industry, fracturing the field into two competing camps [9, ch. 13], and culminated in the two most widely used frontier consumer products today, Claude and ChatGPT, being closed source. Unlike the internet, which was open-sourced and freely accessible, these frontier models cannot be inspected, reproduced, or freely run for inference by the same communities whose data trained them.

While publicly advocating for the regulation of artificial intelligence, often by proposing an international oversight body modeled on the International Atomic Energy Agency [1], leading frontier executives have in practice helped construct regulatory moats around the existing incumbents. Executive Order 14110 set reporting requirements for any model trained using more than 10^{26} floating point operations [30], a threshold tuned almost exactly to capture the current frontier while imposing compliance overhead that smaller teams and open-source efforts cannot easily absorb. The cumulative effect of these moves has been the consolidation of power and the preservation of centralized control over what should, by any reasonable reading of its history, be a democratized public utility.

2 Current Centralization and Pricing of LLMs is Problematic

When a small number of providers control both the model weights and the pricing surface, the cost structure facing every downstream user, whether that user

is a frontier customer in San Francisco, an indie developer in Lagos, a student in Santiago, or a small business in rural Idaho, can be revised unilaterally and without notice. Providers ship attractively priced tiers, encourage developers and enterprises to build dependencies on those tiers, and then change the underlying economics once those dependencies are load-bearing [36].

The general availability of Claude Opus 4.7 on April 16, 2026 illustrates how this works in practice [3]. Anthropic kept the published list rate at \$5 per million input tokens and \$25 per million output tokens, the same as the prior Opus generation, while announcing the deprecation of Opus 4.5 and 4.6 in the same release window, effectively migrating existing Opus pipelines onto 4.7 whether or not anyone had budgeted for the change. Two model-level adjustments compounded the move, namely an updated tokenizer that maps the same input to roughly 1.0 to 1.35 times the prior token count depending on content type, and a new `xhigh` reasoning effort level that was made the default in Claude Code and that increases output token volume on harder tasks by design.

Downstream agent platforms then translated the change into their own pricing surfaces. GitHub launched Opus 4.7 in Copilot at a 7.5x premium request multiplier and disclosed that the figure was the result of a 50 percent promotional discount running through April 30, 2026, implying a post-promotion multiplier of 15x [8]. The previous Opus 4.5 and 4.6 multiplier was 3.0x in the same picker, and both models were removed from Copilot Pro+ as part of the rollout [17]. The shift from 3.0x to 7.5x is a 150 percent increase in effective cost, holding everything else constant, and the shift from 3.0x to a post-promotion 15x is a 400 percent increase.

The consequences scale linearly with usage. At the Opus 4.7 list rate, a single 200,000-token agentic coding request runs roughly \$5 in raw tokens for an input-heavy workload, and considerably more once `xhigh` output and tokenizer expansion are included. Wrapped in Copilot’s 7.5x multiplier the same request bills at roughly \$37.50 in metered request equivalents, and at the 15x post-promotion multiplier roughly \$75. Iterative agent calls during a single refactoring session can consume a month’s plan allowance in minutes. A pricing risk of this magnitude is survivable for an enterprise customer with a procurement contract and a dedicated FinOps team, but for a solo developer building a side project, a teacher building a classroom tool, a researcher in a low-income region, or a small startup operating on a shoestring, the same change can shut a project down overnight. Even mid-sized businesses are forced to choose between absorbing unbudgeted cost, halting agent workflows, or migrating to a different provider whose pricing structure is, by construction, equally subject to revision [36]. When access to frontier capability is mediated by two layers of unilateral pricing power, the model provider on one side and the agent platform routing to it on the other, every developer downstream is ultimately a price-taker

on a contract they did not sign, and centralization has come to concentrate artificial intelligence and also the discretion to revoke it. A worked breakdown of the formula that produces the \$37.50 and \$75 figures above, including the interaction between list rate, tokenizer expansion, and the Copilot multiplier, is given in Appendix E, and a cost-decomposition alternative to the current bundled retail price is developed in Appendix A.

3 Environmental and Social Costs

Much of the empirical material in this section, including the megacampus capacity figures, the documentation of community resistance in Cerrillos, and the framing of corporate donations as a substitute for structural remediation, is drawn from Karen Hao’s *Empire of AI* [9], whose chapter-level investigation of OpenAI and the broader frontier landscape is, at the time of writing, the most thorough public account of how these costs are distributed and obscured.

On the other side of the development of large language models are the immense compute resources required for training and inference workloads. GPUs³ and specialized accelerators such as TPUs⁴ are housed in warehouse-scale facilities that span multiple football fields and in some cases thousands of acres, and the largest single data centers, which until recently were designed at roughly 150-megawatt capacity, are now being engineered at the gigawatt scale, with current AI megacampuses targeting between 1,000 and 2,000 megawatts of dedicated power, the equivalent of one and a half to three and a half San Franciscos worth of annual electricity demand [9, ch. 12]. By 2030, data centers are projected to consume on the order of 8 percent of all United States electricity, up from roughly 3 percent in 2022, and global artificial intelligence compute could exceed the total electricity consumption of India [9, 28, ch. 12].

These same facilities draw clean potable drinking water for evaporative cooling, and the burden of that consumption falls disproportionately on the communities that receive almost none of the upside. In Cerrillos, Chile, the water-activist group MOSACAT obtained Google’s environmental filing for a planned data center and identified a projected draw of 169 liters of fresh drinking water per second, more than one thousand times the annual water consumption of the entire municipality of roughly 88,000 residents [9, ch. 12]. After organizing through Tania Rodríguez and others, MOSACAT successfully blocked the project; Google then relocated it to Uruguay, exporting the underlying conflict. In the Atacama region of northern Chile, the Lickanantay activist Sonia Ramos⁵ has

³Graphics Processing Unit, originally a co-processor for rendering graphics and now the dominant hardware for training and serving large neural networks; in this paper *5090-class GPU* refers to the 2025-generation NVIDIA RTX 5090 consumer card and equivalents.

⁴Tensor Processing Unit, Google’s custom application-specific integrated circuit for accelerating machine-learning workloads, available to external developers through Google Cloud; see <https://cloud.google.com/tpu>.

⁵Sonia Ramos is a Lickanantay water defender from the At-

spent years documenting the disappearance of native flamingo populations and traditional water sources tied to lithium and data center expansion in the Salar de Atacama, and her own framing of that work is the framing this paper adopts:

“Chile emits no more than 3 percent of the planet’s emissions, but we are the ones paying the consequences for others. That is what I call territorial injustice.”

— Sonia Ramos, Lickanantay water defender, Atacama [11]

Aggregated industry-wide, training and inference together emit greenhouse gases at a scale that already rivals entire industrial sectors [28], and most of this impact is presently “remediated” or talked around through relatively minor donations, political lobbying, and publicity campaigns with no structural change [9, ch. 12]. With private investment in artificial intelligence at record highs in the United States and continuing to climb, reaching approximately US\$300 billion of global venture funding in the first quarter of 2026 alone, of which roughly 83 percent flowed to United States-based companies and approximately US\$242 billion to companies in the artificial-intelligence sector specifically [27, 34], and the projected expansion of compute infrastructure trending only upward, the negative impact of these systems on the communities that actually host them is expected to intensify if the existing model is allowed to continue unchanged, particularly in the global south.

4 Murky Data

The history of large language models is, in effect, a history of data, as mentioned before, but more specifically, it is a history of *murky data*. Scaling laws as outlined in research published by several sources, notably Dwarkesh Patel’s⁶ *The Scaling Era: An Oral History of AI, 2019–2025* [20] and Kaplan et al.’s foundational paper on neural language models [13], establish that the quality of a trained system is primarily a function of the volume and distribution of the data on which it was trained irrespective of architectural novelty in the model itself, and Karen Hao’s account of the internal calculus at the frontier labs reinforces this from the operational side, with the decisive quality jumps at each generation tracing back to deliberate data-filtering and data-acquisition programs [9, pp. 174, 177]. Patel’s oral history records the same observation in the practitioners’ own words, with Jared Kaplan⁷ describing scaling as “this rela-

acama region of northern Chile whose long interview with the International Institute for Environment and Development on territorial injustice in lithium and data-center expansion is the source of the quotation that follows.

⁶Dwarkesh Patel is an independent technology podcaster whose long-form interviews with leading artificial-intelligence researchers form the basis of the cited oral history; see <https://www.dwarkesh.com/>.

⁷Jared Kaplan is a co-founder and chief science officer of Anthropic and a lead author of the foundational scaling-laws paper for neural language models cited elsewhere in this section.

tion we’ve noticed: as you make AI systems larger—increasing the number of parameters they have, training on more data, or increasing the total amount of compute used for training—you get really, really predictable trends for the performance of these systems” [20, p. 20], and Demis Hassabis⁸ flagging data curation as the field’s emerging bottleneck, noting that “we’re still in the nascent stage of [optimal] data curation and data analysis and analyzing the holes in our data distribution” [20, pp. 45–46]. The practical consequence is that the most significant single variable in the development of the frontier models used today was the decision, made early and almost universally, to treat the open internet as a free raw material and to scrape it without permission, attribution, or compensation.

The costs of that decision are already visible and are distributed along three distinct surfaces. The first is bias and representational harm, which *Stochastic Parrots* documented in detail several years ahead of the current race and which is now well-established in peer-reviewed literature, namely that an uncuration internet corpus encodes the historical, racial, gender, and geographic biases of the population that was online during its collection, and that a model trained on such a corpus reproduces and in many cases amplifies those biases in its outputs [4]. The second is quality and factual reliability, which manifests to users as hallucination. Recent formal work establishes that under standard pretraining and reinforcement-learning-from-human-feedback objectives, models trained on low-quality, contradictory, or thinly-attributed data are statistically incentivized to guess and do not abstain when uncertain [12]; further, successive generations of synthetic or recycled data compound the problem by collapsing the diversity of the underlying distribution [26]. The third is the hidden labor cost of making the scraped corpus usable at all, most of which is absorbed by reinforcement-learning-from-human-feedback workforces operating on outsourced contracts in Kenya, the Philippines, and elsewhere, where workers have reported sustained exposure to extreme, violent, and sexually explicit material [9, ch. 9]. In Chapter 9 of *Empire of AI*, Karen Hao documents the subsequent shift of this work toward more highly-educated workers such as doctors, engineers, and lawyers as the easier prompt categories saturate, a migration that is best read as a confirmation that the quality of a model is ultimately bounded by the quality and compensation of the humans who curate its training data. Patel’s account makes the same point structurally, with Leopold Aschenbrenner⁹ observing that “RLHF is equivalent to increasing the model size 100 times in terms of the resulting increase in human evaluators’ preference ratings” [20, p. 35], which is to

⁸Demis Hassabis is co-founder and chief executive officer of Google DeepMind and a 2024 Nobel laureate in Chemistry, recognized jointly with John Jumper for the AlphaFold protein-structure prediction system.

⁹Leopold Aschenbrenner is a former member of OpenAI’s Superalignment team and the author of the widely-circulated essay collection *Situational Awareness*; see <https://situational-awareness.ai/>.

say that the human-feedback layer is doing as much capability work as several order-of-magnitude jumps in compute.

The human cost of that labor, when traced to the actual people performing it, is the part of the current pipeline that is hardest to defend on any reading. Karen Hao’s reporting on the Sama contract that supplied OpenAI’s early content-moderation and reinforcement-learning-from-human-feedback pipelines documents Kenyan workers paid on the order of US\$1.32 to US\$2.00 per hour to read, label, and rank the most extreme material the open internet produces, including child sexual abuse material, descriptions of bestiality and incest, and detailed depictions of torture and self-harm, on quotas that left no meaningful recovery time between exposures [9, ch. 9]. Workers later reported sustained psychological injury consistent with post-traumatic stress, including intrusive imagery, broken marriages, substance dependence, and inability to be physically present with their own children. As a result, the contract with Sama was terminated early by OpenAI when the conditions became public, with no comparable mental-health infrastructure put in place for the workers who had already been exposed [9, ch. 9].

The structural conclusion is that reinforcement-learning-from-human-feedback at its current scale is largely a remediation cost imposed by the decision to train on murky data in the first place. As Hao puts the point directly, the apparatus is needed primarily because the underlying corpus is uncurated and contradictory; “RLHF ceases to be necessary with higher quality data,” or rather, the share of it that is genuinely needed shrinks toward the share that is genuinely capability-enhancing [9, p. 245].

A useful corollary of this observation is that if the humans curating training data were compensated fairly and selected for the relevant expertise from the outset, a large portion of the reinforcement-learning-from-human-feedback apparatus currently used to paper over low-quality data would become unnecessary [9, p. 245]. The present supply chain for RLHF’d data is dominated by a small number of vendors including Scale AI, Appen, and their peers, who extract huge sums of money per year in aggregate while producing datasets whose documented deficiencies are precisely the deficiencies this paper is attempting to address with higher quality data and fair incentives.

At the same time, vast reserves of the highest-quality human-generated data are currently unavailable to frontier training for reasons of licensing, copyright, and platform terms of service. The creator economies on platforms such as YouTube, Twitch, Substack, Udemy, LinkedIn Learning, Patreon, and OnlyFans represent sustained, category-rich, often subject-matter-expert human output in volumes that dwarf most open-web corpora, produced by contributors who already expect and receive compensation for their work and who could in principle opt in to a data-licensing arrangement on comparable terms. A standard that routes verifiable compensation to iden-

tifiable contributors is the precondition for any of this material to become usable at scale, and the absence of such a standard is a principal reason the current generation of models is trained disproportionately on the material that is easiest to access and often murky at best.

The three surfaces above, bias, hallucination, and unfairly extracted labor, along with the un-monetized high-quality data sitting outside of current corpora, together define the data problem that the remainder of this paper is addressed to. The proposal that follows is motivated by the observation that each of these surfaces admits the same structural solution, namely compensating the humans whose work produces the data, attributing that work at the time of creation, and surfacing the cost of the compute that trains on it at the point of use.

5 Proposed Solution: The Human Data Collective

The argument for higher quality, human-generated content created specifically for the purposes of training large language models is therefore clear. The question is how to properly incentivize humans to generate it, how long it will take to create, how to develop it safely and with as little environmental consequence as possible, and how to provide equitable access to it once it exists.

This paper proposes an alternative architecture rooted in several coordinated mechanisms working together, namely a decentralized network of human creators, peer reviewers, and smart contracts that produces categorized, attributed, training-grade data, a branded inference token standard that couples access to large language models with environmental impact accounting and an automatic profit share back to the data contributors, and on-chain remediation funds that route a portion of every paid request into carbon and water programs with grants available to the communities most affected by compute infrastructure expansion. Together these three mechanisms form what the rest of this paper refers to as the Human Data Collective, and the three of them are interdependent by design, in that the data is what makes the network valuable to frontier providers, the token is what makes access to the data legible and auditable, and the remediation funds are what make the whole arrangement defensible to the communities whose land, water, and electricity are currently being externalized.

What follows is a detailed breakdown of the Human Data Collective, how it works, and how it addresses the problems outlined in the previous sections. This proposal is intentionally modular, with each component designed to function independently while also contributing to a cohesive whole, and with the expectation that each component will require iterative development and refinement in collaboration with the communities it is intended to serve. For many reading, the Appendix sections will add better clarity, and may be easier to understand than the main text, because they include concrete examples and formulas that il-

lustrate how the token economics and inference standard work in practice. With that in mind, feel free to skip ahead to Appendix A and Appendix E or any of the sections 6-9 that may be specifically relevant to you before returning to the main text.

Access to data sourced from the Human Data Collective is conditional on adoption of the standard, with no exceptions. Any model trained on Collective data must mint a branded inference token on Base¹⁰ and must route every paid request through the inference standard described in Appendix A. There is no opt-out path that preserves access to Collective data while bypassing the standard, because the standard *is* the licence and the token *is* the term of use. A frontier provider that wishes to monetize access to a model trained on Collective data mints a branded ERC-20 token on Base using a provided HDC framework, pays a one-time onboarding fee, and a recurring annual renewal that is similar in spirit to domain registration but proportionally larger. Each inference request thereafter carries a baseline cost that reflects the actual electricity used and the embodied environmental impact of the request, plus an optional premium that is set freely by the provider. Open-source models mint a public-good variant of the same token, with the same baseline environmental cost, no premium, and no per-request profit share, preserving free access to community-trained models in the same spirit as the early internet. Of any premium charged by a frontier provider, a minimum of one percent flows to the Collective, and within that one percent, eighty percent is distributed to the data creators whose attestations contributed to the training set, five percent to the peer reviewers who verified the data, and the remaining fifteen percent to Collective operations.

Table 1: Token economics. Frontier and open-source tokens share the same environmental floor; only frontier tokens carry a premium and the 1% creator share.

Component	Frontier	Open-src
Onboarding fee	one-time	waived
Annual renewal	required	nominal
Per-request baseline	energy + env	energy + env
Per-request premium	provider-set	none
Min. premium share to HDC	1%	n/a
<i>Distribution of the 1% HDC share</i>		
Creators (attestation)	80%	
Peer reviewers	5%	
Collective operations	15%	

The smart contracts that control each branded token route the baseline environmental cost of every request into two pooled funds, the *HDC Carbon Credit Fund* and the *HDC Water Credit Fund*, both managed on-chain with all inflows, outflows, and credit retirements published openly to the public ledger. The Carbon Credit Fund follows the existing pattern for

¹⁰Base is an Ethereum Layer 2 network maintained by Coinbase that provides low-cost, high-throughput settlement while inheriting the security properties of the Ethereum main chain; see the Glossary entry on page 11 and <https://www.base.org/>.

tokenized verified credits issued under recognized standards such as the Verified Carbon Standard or Gold Standard and registered on platforms that prevent double-counting through immutable on-chain identifiers [23], with credits purchased and retired against the inference workload that funded them and each retired token representing one metric ton of CO₂-equivalent removed, avoided, or sequestered, with auditable lineage from project to retirement. Water restoration markets are considerably less mature than carbon markets, but the most established instrument available today is the Water Restoration Certificate (WRC) issued by the Bonneville Environmental Foundation, where each WRC represents 1,000 gallons of natural freshwater flow restored or protected, third-party verified, and tracked against double-counting on the S&P Global Market registry [5, 29]. WRC supply is concentrated in the western United States, so the Water Credit Fund operates a two-track strategy, purchasing and retiring verified credits where they are locally available, and where they are not it issues direct grants to community-led restoration programs in the places most affected by data center cooling demand, including those represented by activists such as Sonia Ramos in the Atacama and groups such as MOSACAT in Cerrillos. To extend market coverage over time, the Collective also proposes a tokenized HDC Water Restoration Token (WRT), modeled directly on the WRC and issued globally on the same Base infrastructure, with each token backed by an independently audited unit of restored or protected freshwater, retired only after audit, and mirroring the carbon credit lifecycle in every other respect. Because the deduction occurs at inference time on every single request, the combined remediation budget scales naturally with usage.

6 For Frontier Providers

The ask is concrete and bounded, namely to mint a branded inference token on Base, route the small environmental and creator-compensation lines described in Appendix A through the on-chain settlement contract, and retain ninety-nine percent of every premium dollar the provider chooses to charge.

Frontier providers want the data, and on the current trajectory will need it. Synthetic data generated from prior model output produces measurable degradation across successive generations [26], and additional internet-scraped data is increasingly polluted, copyright-encumbered, or legally precarious. The Collective offers an attributed, categorized stream of training data that can be licensed by category and provenance, and the 99-to-1 premium split leaves the provider with the overwhelming majority of any margin while solving the data supply bottleneck that is increasingly the actual constraint on continued model improvement.

7 For Contributors

For contributors, the Collective converts the act of producing digital content into a source of verifiable,

ongoing income. The current generation of human-feedback supply chains has exported much of this work to low-wage labor markets in Kenya, the Philippines, and elsewhere with minimal protections [9, ch. 9], and the Collective restructures that relationship from the outset so that contributors in the same regions earn a recurring share of the downstream value their work creates, on terms designed to meet the kind of fair-work standards promoted by the Oxford Internet Institute’s Fairwork project. The framework described here also takes seriously the data-sovereignty critiques articulated by the Distributed AI Research Institute and the broader Fairwork community, namely that the people whose work and lived experience produce a dataset must retain a meaningful and enforceable share of its downstream value, and that any compensation arrangement must be auditable by the workers themselves and not only by the platform that pays them. Attribution attestations restore authorship to writers, photographers, programmers, educators, translators, and indigenous knowledge holders whose work is currently absorbed without recognition. The end-to-end mechanics by which a single request is metered, settled, and attributed back to a contributor are walked through in Appendix B, and the implied earnings at scale are summarized in the *Cost summary at scale* subsection of that appendix (§B).

8 For Governments and Regulators

Auditable, categorized, human-generated training data enables a deterministic regulatory model that nothing in the current frontier landscape supports. Regulators can trace model capabilities back to specific data categories and restrict access to models trained on weapon-manufacturing material, regulated chemistry, child sexual abuse material, or other prohibited content. Categories such as adult content, for example, could be tagged at the data layer and only used to train models that are gated behind age verification at the model layer rather than the application layer, the same logic applying to dangerous or highly regulated categories such as certain kinds of physical-science weaponry. Automated environmental remediation funded at the request level provides governments with direct, measurable, audit-trail-backed progress toward climate commitments they have already adopted, without requiring any new agency, compute threshold, or inspection of model weights. The aggregate magnitude of those flows at realistic deployment volumes is set out in the *Cost summary at scale* subsection of Appendix B (§B), which provides regulators with concrete order-of-magnitude figures for both remediation and creator compensation under the proposed standard.

9 For the Artificial Intelligence Industry

The architecture proposed here maps almost perfectly onto the initiatives and goals of every major contingent in the artificial intelligence development space. The capability-focused camp, sometimes referred to as

the AGI maximalists, gets a renewable supply of high-quality, behaviorally rich, responsibly sourced data to train on and to push closer to their stated goal of reaching general intelligence. The safety-focused camp gets the auditable provenance, category-level controls, and capability-tracking tools they have been asking for, including the ability to track emergent capabilities by training set and to completely recuse themselves of the human-psychology liabilities created by current reinforcement-learning-from-human-feedback supply chains operated by firms such as Scale AI and Appen. A third constituency, exemplified by Bender, Gebru, and Hao, holds that the current scale of these systems is itself a harm and that no amount of attribution or environmental-cost disclosure is a substitute for slowing or halting deployment; the Collective does not resolve that disagreement, and it does not pretend to, but it does give that constituency a per-request audit trail with which to make its case in measurable terms. The structure also dissolves a substantial portion of the legal and reputational exposure currently created by training on uncurated or copyright-encumbered data.

10 Considerations

Generating a human-created dataset of the same scale used by current frontier models will take time and coordinated effort, on the order of a sustained social movement rooted in the desire for fair and equitable artificial intelligence and for sustainable development. The growth and impact of the network once created has network-effect properties and is therefore essentially infinitely scalable, because any human can in principle generate any kind of content in a verifiable and equitable way.

The percentage breakdowns of profit share described above are very generous to frontier providers, and the generosity is intentional, because frontier providers will be the largest single driver of growth for large language models as tools of industry, and without them on board a project of this scope faces almost insurmountable competition. Even a one percent per-request profit share distributed under the 80/5/15 split would, at the implied scale, provide viable economic prosperity to an unprecedented global population of creative workers, alongside the environmental impact of compulsory carbon and water credits attached to every request processed by both frontier and open-source providers.

Several open questions remain. The cryptographic attestation primitive that ties compensation to authorship has a real attack surface, including identity fraud, prompt-laundering through intermediate models, and collusion between contributors and downstream consumers, and all three require dedicated technical and economic analysis in AIISP-2. Public attestation of authorship also has known privacy and surveillance implications for contributors whose participation depends on pseudonymity, including dissidents, sex workers, and indigenous knowledge holders whose communities reserve the right to control dis-

closure of specific material, and AIISP-2 is intended to specify a pseudonymous-attestation mode that preserves on-chain compensation routing without disclosing real-world identity to the public ledger, perhaps through zero-knowledge proofs or a similar mechanism. The Water Restoration Certificate market is geographically concentrated in the western United States, which means the v0.1 mechanism for non-US watersheds is a non-solution placeholder and is acknowledged as such; the Cerrillos and Atacama cases that motivate this proposal are precisely the cases for which the v0.1 mechanism is weakest, and rewriting that subsection of the standard with affected community input is the single highest-priority item for v0.2. Settlement on a public blockchain carries its own energy footprint, which the standard accounts for in its environmental line but which merits scrutiny against the marginal energy cost of plain off-chain settlement.

A New Inference Standard

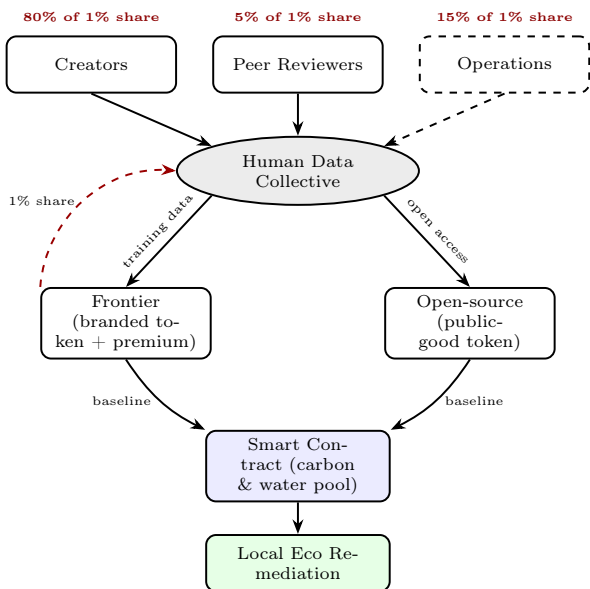


Figure 1: Flow of data, value, and environmental remediation through the Human Data Collective. Solid arrows carry data and value forward; the dashed arrow returns the 1% premium share to the Collective for redistribution. The dashed border on the Operations box indicates its role as the managing and governing entity of the Collective, which is structurally distinct from the Creators and Peer Reviewers, who are the compensated participants in the network itself.

Frontier providers today price access in dollars per million input or output tokens, with representative public list prices including Anthropic Claude Opus 4.7 at \$5.00 per million input tokens and \$25.00 per million output tokens [3], Anthropic Claude Sonnet 4.5 at \$3.00 input and \$15.00 output, and OpenAI GPT-4o at approximately \$2.50 input and \$10.00 output. These prices bundle compute amortization, profit margin, and an implicit allocation toward provider data acquisition and labeling, but they do not surface energy use, water use, or any payment back to

the humans whose work trained the model.

The Human Data Collective proposes a thin standard that is layered on top of, not in replacement of, this existing billing surface. Each inference request carries a structured cost decomposition with three components, namely an *energy line* priced against a published per-request kilowatt-hour benchmark, an *environmental line* routed into the HDC Carbon and Water Credit Funds, and a *premium line* set freely by the provider, of which a minimum of one percent is routed to the Collective for redistribution. The standard is expressed at the request level so that any HTTP-based inference API can adopt it by adding a small header block and an on-chain settlement call, with no existing client code needing to break in the process.

A representative cost decomposition for a single 1,000-token blended request can be illustrated with an energy line on the order of US\$0.000010, equivalent to a single 5090-class GPU-second amortized over typical concurrency and benchmarked against a published California electricity rate,¹¹ an environmental line on the order of US\$0.000020 split between the carbon and water funds, and a premium line of US\$0.001000 set by the provider. Of that premium, US\$0.000010 is routed to the Collective and split 80/5/15 between creators, peer reviewers, and operations, leaving the provider with 99 percent of every premium dollar, which is more than they currently capture from any third-party data licensing arrangement on a per-request basis.

The forcing function for this standard is the data itself. Synthetic data degrades predictably when models train on the output of prior models [26], and internet-scraped corpora are increasingly polluted and legally precarious. The frontier providers that achieve the next quality threshold will be those with sustainable access to fresh, attributed, categorized, human-generated data, and the Collective is the only proposed structure that produces such data at scale while compensating its sources. The cost of joining is structurally trivial in light of that prize: one percent of premium revenue and a small per-request environmental line in exchange for a renewable supply of the only input that still meaningfully moves capability.

This is, in effect, a proposal for a new internet standard. The early internet was made viable by a small number of agreed-upon protocols (TCP/IP, HTTP, DNS) that no single party owned, and each protocol added value to every participant precisely because it was open and universal. The AI Inference Standards Protocol (AIISP) is offered in the same spirit, and the v0.1 draft is published openly for public review and comment as AIISP-1 [10]. Adoption of the standard is

¹¹An NVIDIA RTX 5090 draws on the order of 575 watts under sustained inference load, so one GPU-second is approximately $575 \text{ W} \times \frac{1}{3600} \text{ h} \approx 1.6 \times 10^{-4} \text{ kWh}$; at the California 2026 industrial reference rate of approximately US\$0.06 per kWh that single GPU-second costs roughly $9.6 \times 10^{-6} \text{ US\$}$, which is rounded to US\$0.000010 in the body of the text. The figure is amortized across typical concurrency, in which any single 1,000-token inference request occupies a fraction of one card for a fraction of one second.

opt-in and incremental, but access to Collective data is not, in that any model trained on Collective data must mint a branded token and must settle through the standard. The environmental remediation layer, however, is not proprietary to the Collective. A provider that does not train on Collective data can still adopt the AIISP wire format to surface and settle its own energy and environmental lines against any audited registry or community-grant pipeline of its choosing; the HDC Carbon and Water Funds are a reference implementation of the standard and not its only permissible destination. The remediation of downstream environmental harm is the priority, and the Collective offers one mechanism by which providers can act on that priority, mandating it only for models that have trained on its data.

Related work

AIISP-1 sits at the intersection of three established lines of work, none of which it claims to supersede. On content provenance and authorship metadata, the Coalition for Content Provenance and Authenticity (C2PA) and the W3C PROV-O recommendation define complementary signing and provenance vocabularies that the HDC attestation primitive is designed to interoperate with [6, 31]. On per-job environmental accounting, the CodeCarbon and ML CO₂ Impact projects established the practice of emitting energy and carbon estimates alongside model output [7, 14], and AIISP-1’s energy line is a wire-format generalization of that practice with on-chain settlement attached. On on-chain royalty distribution, Audius and Royal demonstrated that fractional creator-share contracts on Ethereum can clear at production volume [21, 22], and AIISP-1’s settlement contract follows the same split-and-distribute pattern. The 80/5/15 share is informed by, although not derived from, prior platform-cooperative practice including Stocksy United’s 75 percent creator share and the platform-cooperativism literature of Scholz and Schneider [24, 25], and the broader case for treating personal data contributions as compensable labor draws on Lanier and Weyl’s work on data dignity [15].

Adoption pathway and pilot funding

The standard does not require any of the following in order to be technically useful, but adoption at scale will require at least one. First, a data-supply forcing function, namely a Collective-curated corpus large enough that frontier providers cannot match its freshness, attribution, and category coverage from any other source; bootstrapping such a corpus is the principal use of the initial Coinbase ecosystem grant the Collective intends to seek, with proposed milestones of a published reference implementation, a first cohort of at least 1,000 attested contributors, and a first frontier provider in pilot, at a budget on the order of US\$50,000 to US\$100,000 milestone-based. Second, a regulatory forcing function, namely a California-first disclosure-only bill of the kind described in Appendix C; the Collective intends to approach the UCLA Institute for Technology,

Law and Policy, the Greenlining Institute, and a California water nonprofit as candidate co-sponsors, and the bill will not be introduced without at least one named California-headquartered organization on the masthead. Third, a capital forcing function, namely a frontier provider that signs on as the first pilot counterparty and routes a portion of live inference traffic through the standard ahead of its competitors. The value proposition to any such provider is threefold, consisting of a differentiated data moat that cannot be matched by continued scraping of the open web, a pre-negotiated answer to the environmental and labor questions that increasingly dominate regulatory and public conversation in 2026, and first-mover credit with the contributor base the standard is designed to cultivate. The Collective intends to meet directly with the providers best positioned to carry this role and to make the case on those terms, and an early signal of interest from any of them would shorten the timeline for the other two pathways materially.

B Lifecycle of a Single Request

This appendix walks a single inference request end-to-end under the proposed standard, with worked numbers for a representative 1,000-token blended call (700 input + 300 output). All dollar figures are illustrative; the standard does not fix any of them.

Step 1: Client request

The client issues a normal HTTPS POST to the provider’s existing inference endpoint with two new optional headers:

```
X-AIISP-Token: 0xA1B2... (ERC-20 address)
X-AIISP-Settlement: deferred
```

Absence of the headers means the request is processed and billed exactly as it is today, and presence of them opts the request into AIISP settlement. The `deferred` value batches on-chain settlement, which is the typical case, and `realtime` forces per-request settlement, which is mainly useful for audit testing.

Step 2: Provider-side metering

The provider runs the request as normal and emits a structured cost record alongside the model output, decomposing the charge into explicit lines:

Line	Amount
Energy (metered kWh × regional rate)	\$0.000010
Environmental (carbon + water funds)	\$0.000020
Premium (provider-set)	\$0.001000
Total request cost	\$0.001030

The energy line is computed from actual GPU-seconds multiplied by a published per-region kWh rate, and the environmental line is a surcharge calibrated to the running spot price of carbon and water credits in the HDC funds.

Step 3: On-chain settlement

At the provider’s settlement cadence, a single transaction is submitted to the HDC settlement contract on

Base and is split atomically, with the energy line retained by the provider, the environmental line routed 50/50 to the Carbon and Water Funds, the premium line split 99/1 between the provider and the HDC distribution contract, and within that 1 percent share, 80 percent routed to the creator-attestation pool, 5 percent to the peer-reviewer pool, and 15 percent to operations. For the example request the provider keeps \$0.001020, the environmental funds receive \$0.000020 combined, and the HDC distribution contract receives \$0.000010 (creators \$0.000008, reviewers \$0.000005, operations \$0.000015).

Step 4: Attribution resolution

The creator share is not paid out per request; it accrues against attestation identifiers tied to the training data slice the model invoked. Attribution is resolved off-chain using the same cryptographic attestation primitive used for the authorship record on this paper (HMAC-SHA256 over a canonical content+model+role payload with a registry-anchored timestamp), and the resulting identifiers are bundled into Merkle roots and posted on-chain at a fixed cadence so that any creator can independently verify their share.

Step 5: Credit retirement

The Carbon and Water Fund contracts maintain running balances, and when a balance crosses a configurable threshold, the contract triggers a purchase-and-retire transaction against a verified credit registry (VCS, Gold Standard, or a registered WRC issuer). The retirement event links the retired credit identifier back to the inference workload that funded it, producing an end-to-end audit trail from a single API call to a specific retired ton of CO₂ or restored gallon of water.

Cost summary at scale

A frontier provider processing 10^{10} such requests per month would, under this example, route roughly \$200,000 per month into combined environmental remediation and roughly \$100,000 per month into creator and reviewer compensation, while retaining approximately \$10,200,000 of premium revenue net of the 1 percent HDC share.

Why Base

Base, an Ethereum Layer 2 maintained by Coinbase, was selected as the settlement layer for several practical reasons. Per-transaction settlement cost is low enough that even sub-cent environmental and creator lines settle economically; the network already supports the ERC-20 token standard required for branded inference tokens and the public registries required for on-chain credit retirement; and existing public-goods grant programs, including the Coinbase Builder Grants and the Optimism Retroactive Public Goods Funding rounds in which Base participates, are an appropriate funding source for the initial data-creation effort that the Collective will require before frontier providers begin paying premium share at meaningful volume. The Coinbase compliance posture also provides an existing legal interface into United States fi-

ancial regulation that a smaller protocol would have to construct from scratch, lowering the barrier to participation for institutional contributors and corporate operators.

C A Brief for Regulators

What is being proposed. A voluntary, opt-in inference standard that adds three measurable line items to every billed inference request, namely energy, environmental remediation, and a small share of provider premium routed to the people whose work trained the model. All settlement runs on a public blockchain (Base, an Ethereum L2), so every dollar collected and every credit retired is independently auditable.

What it solves. Four documented failures that no single existing law fully addresses: (1) the extraction of personal and copyrighted training data without compensation or consent; (2) hallucination rates that remain materially high in production deployments and that compound when models train on prior model output [12, 19, 26]; (3) the concentration of pricing power in a small number of providers, with documented mid-cycle repricing of dependent users [3, 8, 17]; and (4) the externalization of energy, water, and emissions costs onto host communities, particularly in the global south [11, 28], including documented cases such as Cerrillos, Chile [9, ch. 12].

Why this is friendly to legislate. The standard does not require defining a new regulated entity, setting a compute threshold, or empowering a new agency to inspect model weights. A bill could simply require any provider serving more than a threshold volume of inference requests to United States persons to disclose, in machine-readable form on each request, the energy, environmental, and creator-compensation line items, whether or not the provider participates in the HDC settlement contract.

California as a natural first jurisdiction. Disclosure-only rules of this kind are already broadly compatible with existing California law, in particular the California Consumer Privacy Act [33] and Assembly Bill 2013 [32], which obliges generative-AI developers to publish summary information about the data used to train their models. A California-first version of the standard would extend AB 2013 from a one-time training-data summary to a per-request runtime disclosure, route remediation funding into California-specific water and energy programs through the same on-chain settlement contracts described in Appendix B, and give the state a measurable, auditable mechanism for tracking the climate and labor impact of the inference workloads physically hosted within its borders. The Collective intends to pursue at least one California-headquartered co-sponsoring organization before introduction of any bill, with the UCLA Institute for Technology, Law and Policy, the Greenlining Institute, and a California water nonprofit as the most natural candidates.

Minimum bill components. (1) A definition of a covered inference provider scoped by request vol-

ume. (2) A required per-request disclosure of energy in kilowatt-hours, environmental remediation in dollars, and creator-compensation pass-through in dollars, expressed as a standard JSON header. (3) A safe-harbor provision recognizing on-chain settlement against an audited registry as conclusive evidence of remediation. (4) A directive to NIST to publish reference per-region energy benchmarks. (5) A private right of action for creators whose attribution attestations are demonstrably ignored by a covered provider.

What the standard does not require. It does not require any provider to publish weights, training data, or proprietary methodology. It imposes no per-token tax. It does not preempt copyright, privacy, or consumer-protection law; it complements them by surfacing the underlying economic flows. It creates no new federal entity, and adoption by any individual provider remains voluntary. Disclosure of the equivalent line items, in the same machine-readable shape, is what a bill would make mandatory.

Estimated impact. If adopted by the three largest United States frontier providers at current 2026 inference volumes, the standard would, on the order of magnitude implied by Appendix B, generate hundreds of millions of dollars per year in environmental remediation and a comparable sum in creator and reviewer compensation, without any change to the headline retail price of inference.¹² The cost is borne almost entirely by providers’ existing margin.

D The Human Data Collective in a Post-AGI World

The Human Data Collective is designed for the world large language models actually inhabit in 2026, a world where models are powerful but bounded, fresh human data is scarce and contested, and environmental and labor externalities are still treated as someone else’s problem. A reasonable question to ask of any proposal at this scale is what happens to it if and when the field crosses into a regime that today is usually called artificial general intelligence, that is to say, a regime where a single system can match or exceed expert human performance across most economically valuable cognitive tasks. The honest answer is that some parts of the architecture here become less important, others become more important, and the parts that survive are the ones least tied to any particular generation of model.

The data layer changes the most. A truly general system would in principle be able to generate large

¹²Concretely, Appendix B routes US\$0.000020 per request into combined environmental remediation and US\$0.0000085 per request into combined creator and reviewer compensation. Public statements from the largest frontier providers in 2026 indicate aggregate inference volume on the order of 10^{12} to 5×10^{12} requests per year across the top three; substituting $N = 5 \times 10^{12}$ into the Appendix B formulas yields approximately US\$100 million in annual environmental remediation and approximately US\$42 million in annual creator and reviewer compensation, and doubling either the per-request lines or the request volume reaches the headline figure of hundreds of millions of dollars per year in each line.

volumes of plausibly novel content on its own, and the marginal value of an additional human-written essay or piece of code would fall relative to today. The marginal value of human *attestation*, however, would rise sharply, because in a regime where machine output is indistinguishable from human output by inspection alone, the only remaining ground truth for what a real person actually said, made, or experienced is a cryptographic record signed at the time of creation by an identifiable human. The same attestation primitive that the Collective uses today to compensate creators becomes, in a post-AGI world, the primary mechanism by which human authorship is distinguishable from machine authorship at all, and the Collective’s registry becomes infrastructure for the broader question of provenance.

The environmental layer becomes obviously more important because a system capable of operating autonomously across most cognitive domains will, by construction, be invoked far more often than today’s models, and its aggregate energy and water footprint will scale with that invocation count. Pricing every request against a published per-region kilowatt-hour and per-liter benchmark, and routing a fixed environmental line into on-chain remediation funds, is precisely the kind of mechanism that scales gracefully with usage, since the deduction occurs once per request regardless of what the request is for. The HDC Carbon and Water Funds, and the on-chain registries that back them, outlive any particular model generation, because they are denominated in tons of CO₂ and gallons of water which is more durable than tokens or parameters for the purposes of remediation.

The governance layer is where the Collective matters most in a post-AGI world. A system that can act with broad cognitive autonomy raises a constitutional question that no existing institution is shaped to answer, namely who decides what such a system is allowed to do, on whose behalf, and with what recourse for the people affected. A standing, transparent, on-chain body funded by a fixed share of every inference request, governed by the contributors who supply its training data, and bound by published rules for distributing both compensation and remediation, is one of the few proposals on the table today that could grow into the kind of legitimacy infrastructure such a regime would require.

In the most optimistic scenario, the Collective in a post-AGI world settles into a quiet utility role, providing a settlement layer for inference, a registry of human authorship, a pair of environmental funds with audited balances, and a governance body composed of the people whose work and territories made the underlying systems possible. In the least optimistic scenario, where the post-AGI transition is captured by the same handful of incumbents who control the field today, the Collective is the only piece of the existing landscape that was structurally designed to redistribute the upside and absorb the downside, and its continued existence becomes one of the few remaining checks on full enclosure. Either way, building it now

is cheaper than building it later, and the cost of not building it compounds with every additional megawatt and model generation.

E Cost Formulas

This appendix collects the formulas behind the numerical claims made elsewhere in the paper so that each can be independently reproduced and audited. All figures are illustrative; the standard does not fix any of them.

E.1 Copilot request cost under Opus 4.7

Let T_{in} and T_{out} be the input and output token counts of a single request as measured by the provider’s tokenizer, let p_{in} and p_{out} be the list rates in dollars per million tokens, and let M be the downstream agent platform’s request multiplier. The raw token cost at the model list rate is

$$C_{\text{raw}} = \frac{T_{\text{in}} p_{\text{in}} + T_{\text{out}} p_{\text{out}}}{10^6},$$

and the billed cost to the end developer through the agent platform is

$$C_{\text{billed}} = M \cdot C_{\text{raw}}.$$

For Claude Opus 4.7 the list rates are $p_{\text{in}} = \$5$ and $p_{\text{out}} = \$25$ per million tokens [3]. For GitHub Copilot the multiplier was $M = 3.0$ under Opus 4.5 and 4.6, $M = 7.5$ under the Opus 4.7 promotional rate, and $M = 15$ under the post-promotion rate [8, 17]. The \$37.50 and \$75 figures in Section 2 are obtained by substituting a representative agentic-coding request profile into C_{billed} at $M = 7.5$ and $M = 15$ respectively.

E.2 Representative 1,000-token request under AIISP

For a single 1,000-token blended request under the proposed standard, the three cost lines are computed as follows. The energy line is

$$E = R_{\text{kWh}} \cdot e_{\text{req}},$$

where R_{kWh} is the published per-region electricity rate in dollars per kilowatt-hour and e_{req} is the per-request energy in kilowatt-hours, calibrated such that a single 5090-class GPU-second at typical concurrency on the California 2026 reference rate yields $E \approx \$0.000010$. The environmental line is

$$V = \eta \cdot c + \omega \cdot w,$$

where c is the running spot price of one metric ton of CO₂-equivalent verified credit, w is the running spot price of one Water Restoration Certificate unit, and η and ω are per-request shares calibrated to yield $V \approx \$0.000020$ with a 50/50 split between the Carbon and Water Funds. The premium line P is set freely by the provider, and the Collective share is

$$S_{\text{HDC}} = 0.01 \cdot P,$$

split further 80/5/15 into creator, peer-reviewer, and operations pools, so that for $P = \$0.001000$ the creator share is \$0.000008, the reviewer share is \$0.0000005, and the operations share is \$0.0000015.

E.3 Cost summary at scale

A frontier provider processing N such requests per billing period routes

$$N \cdot V \quad \text{into environmental remediation,}$$

and

$$N \cdot S_{\text{HDC}} \quad \text{into the Collective’s distribution contract,}$$

and retains

$$N \cdot (P - S_{\text{HDC}} + E) \quad \text{as net provider revenue.}$$

For $N = 10^{10}$ and the representative values above, the period totals are approximately \$200,000 of environmental remediation, \$100,000 of creator and reviewer compensation, and \$10.2 million of net provider revenue, matching the order-of-magnitude claim in Section B.

F Glossary of Terms

This glossary defines terms used in this paper that are likely to be unfamiliar to readers from a policy or creator-economy background, as such, these definitions clarify their use.

AIISP. The Artificial Intelligence Inference Standards Protocol, the open wire-format and on-chain settlement standard proposed in this paper and published in draft as AIISP-1; see appendix A and <https://github.com/97115104/aiisp-spec>.

Attestation. A signed, time-stamped statement by an identifiable party that a specific piece of content was produced by them in a specific role at a specific moment, used in this paper as the primitive that binds compensation and authorship to a given data contribution.

Base. An Ethereum Layer 2 network maintained by Coinbase that provides low-cost, high-throughput settlement while inheriting the security properties of the underlying Ethereum main chain; selected in this paper as the settlement venue for AIISP contracts.

Blockchain. A distributed database in which records are grouped into cryptographically linked blocks and validated by a network of independent participants, providing data structure underneath a public ledger.

ERC-20. The Ethereum Request for Comment number 20, the canonical token standard on Ethereum and its Layer 2 networks; defines the minimum interface (transfer, balance, approval) that any fungible token contract must implement so that wallets, exchanges, and settlement contracts can interact with it without bespoke integration. Used in this paper as the contract type each frontier provider mints to identify itself on Base.

GPU. Graphics Processing Unit, originally a co-processor for rendering graphics and now the

dominant hardware for training and serving large neural networks; in this paper *5090-class GPU* refers to the 2025-generation NVIDIA RTX 5090 consumer card and equivalents.

Hallucination. Output from a language model that is fluent and internally plausible but factually incorrect or unverifiable against any external source, including fabricated citations, events, statistics, and quotations. Hallucination is a property of the generation process and not a deliberate act by the model.

HDC. The Human Data Collective, the decentralized network of human creators, peer reviewers, and on-chain settlement contracts proposed in this paper as the reference implementation of AIISP; see <https://humandatacollective.org>.

Model weights. The learned numerical parameters of a neural network, typically a large matrix of floating-point numbers on the order of hundreds of billions of entries for a frontier language model, which together encode what the model has learned from its training data and which are the artifact that is copied, served, and in proprietary systems withheld from external inspection.

Public ledger. An append-only record of transactions maintained by a blockchain network and readable by anyone, in which each state change is cryptographically linked to those before it so that the history cannot be quietly rewritten and every inflow, outflow, and retirement event is independently auditable.

RLHF. Reinforcement Learning from Human Feedback, the post-training technique in which a base language model is fine-tuned against a reward model trained on human preference rankings of its outputs; introduced as the dominant alignment method for frontier consumer models and used in this paper to refer to the broader supply chain of human raters whose labor that technique requires.

Safe-harbor provision. A clause in a statute or regulation that protects a party from liability if specified conditions are met, used in this paper to mean legal recognition that on-chain settlement against an audited registry constitutes valid and sufficient evidence of environmental remediation.

Smart contract. A small program deployed to a blockchain that executes automatically when its conditions are met and whose execution is recorded on the public ledger, used in this paper to handle settlement, attribution, and credit retirement without a human intermediary.

TPU. Tensor Processing Unit, Google’s custom-designed application-specific integrated circuit for accelerating machine-learning workloads, available to external developers through Google Cloud and used in this paper as the canonical example

of a non-GPU accelerator at warehouse scale; see <https://cloud.google.com/tpu>.

G Index of Named Persons and Organizations

This appendix collects, in one place, every named individual and named organization referenced in the body of this paper, together with a short description of who or what each is and the most useful public links for further information.

Named persons

Sam Altman. Co-founder and chief executive officer of OpenAI; cited in this paper for his May 2023 United States Senate testimony advocating an international oversight body for artificial intelligence. <https://blog.samaltman.com/> <https://x.com/sama>

Leopold Aschenbrenner. Former member of OpenAI’s Superalignment team and author of the essay collection *Situational Awareness*; quoted in §4 on the capability impact of reinforcement learning from human feedback. <https://situational-awareness.ai/> <https://x.com/leopoldasch>

Emily M. Bender. Professor of linguistics at the University of Washington and co-author of *Stochastic Parrots*; cited throughout the data sections of this paper. <https://faculty.washington.edu/ebender/> <https://dair-community.social/@emilymbender>

Timnit Gebru. Founder and executive director of the Distributed AI Research Institute and co-author of *Stochastic Parrots*; cited in the data sections and in §9. <https://dair-institute.org/> <https://bsky.app/profile/timnitgebru.bsky.social> <https://x.com/timnitGebru>

Karen Hao. Author of *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI* and a primary on-the-ground source for the environmental, labor, and corporate-history material in this paper. <https://karendhao.com/empire> <https://www.linkedin.com/in/karendhao/> <https://bsky.app/profile/karendhao.bsky.social> https://x.com/_KarenHao

Demis Hassabis. Co-founder and chief executive officer of Google DeepMind and 2024 Nobel laureate in Chemistry (jointly with John Jumper) for the AlphaFold protein-structure prediction system; quoted in §4 on data curation as the field’s emerging bottleneck. <https://deepmind.google/> <https://x.com/demishassabis>

Austin Harshberger. Author of this paper, founder of Happy Stack Calculus, and editor of the AIISP-1 specification. <https://links.97115104.com> <https://blog.97115104.com>

Jared Kaplan. Co-founder and chief science officer of Anthropic and a lead author of the foundational scaling-laws paper for neural language models; quoted in §4. <https://www.linkedin.com/in/jared-kaplan-645843213/>

- Jaron Lanier.** Computer scientist, author, and Microsoft Research interdisciplinary scientist; co-author with E. Glen Weyl of the data-dignity argument cited in Appendix A. <https://www.jaronlanier.com/>
- Gavin Leech.** Co-author with Dwarkesh Patel of *The Scaling Era*. https://scholar.google.com/citations?user=xC-v_aUAAAAJ&hl=en
- Angelina McMillan-Major.** Computational linguist and co-author of *Stochastic Parrots*. <https://scholar.google.com/citations?user=1vwtA8AAAAJ&hl=en>
- Dwarkesh Patel.** Independent technology podcaster whose long-form interviews with leading artificial-intelligence researchers form the basis of *The Scaling Era: An Oral History of AI, 2019–2025*, used as a primary source in §4. <https://www.dwarkesh.com/> https://x.com/dwarkesh_sp
- Sonia Ramos.** Lickanantay water defender from the Atacama region of northern Chile; her interview with the International Institute for Environment and Development on territorial injustice in lithium and data-center expansion is the source of the quotation in §3. <https://www.iied.org/energy-transition-community-fragmentation-territorial-justice-interview-sonia-ramos>
- Tania Rodríguez.** Organizer with MOSACAT in Cerrillos, Chile, credited in Karen Hao’s *Empire of AI* for leading the community campaign that obtained Google’s environmental filing for the proposed Cerrillos data center. <https://www.instagram.com/tania.rodriguez1/>
- Trebor Scholz.** Founder of the Platform Cooperativism Consortium at The New School and author of *Platform Cooperativism*; cited in Appendix A. <https://platform.coop/> <https://www.newschool.edu/>
- Nathan Schneider.** Associate professor of media studies at the University of Colorado Boulder and author of *Everything for Everyone*; cited in Appendix A. <https://nathanschneider.info/> <https://x.com/ntnsndr>
- Margaret Mitchell.** Computer scientist and co-author of *Stochastic Parrots* (writing under the name Shmargaret Shmitchell at the time of original publication); chief ethics scientist at Hugging Face. <https://www.m-mitchell.com/> https://x.com/mmitchell_ai
- Ilia Shumailov.** Researcher whose 2024 *Nature* paper on model collapse under recursive training on synthetic data is cited throughout. <https://iliaishacked.github.io/>
- Emma Strubell.** Assistant professor at Carnegie Mellon University whose 2019 ACL paper on the energy and policy implications of deep learning in NLP is cited in §3. <https://strubell.github.io/> <https://scholar.google.com/citations?user=UCDMtMOAAAAJ&hl=en>
- E. Glen Weyl.** Founder of RadicalxChange and researcher at Microsoft Research; co-author with Jaron Lanier of the data-dignity argument cited in Appendix A. <https://www.glenweyl.com/>

Named organizations

- Alphabet / Google.** Parent company of Google, Google DeepMind, and the Google Cloud platform; cited in §1 as one of the largest training-data holders and in §3 for the Cerrillos data-center proposal. <https://abc.xyz/> <https://www.google.com/>
- Anthropic.** Frontier artificial-intelligence laboratory and developer of the Claude family of models, organized as a public benefit corporation. <https://www.anthropic.com/> <https://x.com/AnthropicAI>
- Appen.** Australian-headquartered data-annotation and reinforcement-learning-from-human-feedback vendor; cited in §4 and §9. <https://www.appen.com/>
- Audius.** Decentralized music-streaming protocol cited in Appendix A as a precedent for on-chain creator-share contracts. <https://audius.co/>
- Base.** Ethereum Layer 2 network maintained by Coinbase; the proposed settlement venue for AIISP-1. <https://www.base.org/> <https://x.com/base>
- Bonneville Environmental Foundation.** United States non-profit and the issuer of Water Restoration Certificates referenced throughout §5. <https://www.b-e-f.org/>
- Coalition for Content Provenance and Authenticity.** Standards body whose technical specification underpins the attestation primitive used in this paper. <https://c2pa.org/>
- Coinbase.** United States cryptocurrency exchange and the maintainer of the Base Layer 2 network. <https://www.coinbase.com/> <https://x.com/coinbase>
- Distributed AI Research Institute (DAIR).** Independent research institute founded by Timmit Gebru, cited in §7. <https://dair-institute.org/>
- Fairwork.** Oxford Internet Institute project that benchmarks platform working conditions against five fair-work principles, referenced in §7. <https://fair.work/>
- GitHub.** Microsoft-owned source-code hosting platform and the developer of GitHub Copilot, cited throughout §2. <https://github.com/>
- Gold Standard.** Carbon-credit certification body referenced in §5 as a permitted registry for HDC Carbon Fund retirements. <https://www.goldstandard.org/>
- Google DeepMind.** Artificial-intelligence research subsidiary of Alphabet, cited in the Acknowledgments and §4. <https://deepmind.google/>
- Greenlining Institute.** California-based racial-and-economic-equity policy organization named

in Appendix A as a candidate co-sponsor for a California disclosure bill. <https://greenlining.org/>

Happy Stack Calculus. The author’s company and the entity under which AIISP-1 is being developed. <https://hsc.97115104.com>

International Atomic Energy Agency (IAEA). United Nations agency cited in §1 as the model for proposed international AI oversight bodies. <https://www.iaea.org/>

Intl. Institute for Env. and Development (IIED). London-based policy research institute and publisher of the Sonia Ramos interview cited in §3. <https://www.iied.org/>

Meta. Parent company of Facebook, Instagram, WhatsApp, and the Llama family of open-weight language models. <https://about.meta.com/>

Microsoft. Cloud and software vendor and principal investor in OpenAI; the developer, with GitHub, of GitHub Copilot. <https://www.microsoft.com/>

MOSACAT. Movimiento Socioambiental Comunitario por el Agua y el Territorio in Cerrillos, Chile; the community organization that successfully blocked the proposed Google data center. <https://mosacatchile.cl/> *secondary coverage in Karen Hao’s Empire of AI, ch. 12*

National Institute of Standards and Technology. United States federal agency named in Appendix C as the proposed publisher of reference per-region energy benchmarks. <https://www.nist.gov/>

OpenAI. Frontier artificial-intelligence laboratory and developer of the GPT family of models and ChatGPT. <https://openai.com/>

Optimism. Ethereum Layer 2 ecosystem and operator of the Retroactive Public Goods Funding rounds referenced in Appendix B. <https://www.optimism.io/>

Penguin Publishing Group. Publisher of Karen Hao’s *Empire of AI*. <https://www.penguin.com/>

Royal. On-chain music-royalty platform cited in Appendix A as a precedent for fractional creator-share contracts. <https://royal.io/>

S&P Global / Markit. Financial-information provider whose registry is used in this paper as the system of record for Water Restoration Certificate retirements. <https://www.spglobal.com/>

Sama. Data-annotation vendor whose Kenyan content-moderation contract for OpenAI is the central case study in §4. <https://www.sama.com/>

Scale AI. United States data-annotation and reinforcement-learning-from-human-feedback vendor; cited in §4 and §9. <https://scale.com/>

Stanford Institute for Human-Centered AI. Publisher of the annual AI Index Report cited in §3. <https://hai.stanford.edu/>

Stocksy United. Photographer-owned stock-imagery cooperative referenced in Appendix A for its 75 percent creator share. <https://www.stocksy.com/>

Stripe Press. Publisher of Dwarkesh Patel and Gavin Leech’s *The Scaling Era*. <https://press.stripe.com/>

UCLA Institute for Technology, Law and Policy. Academic policy center named in Appendix A as a candidate California co-sponsor. <https://itlp.law.ucla.edu/>

Verra (Verified Carbon Standard). Carbon-credit certification body referenced in §5 as a permitted registry for HDC Carbon Fund retirements. <https://verra.org/>

The White House. Issuer of Executive Order 14110, cited in §1. <https://www.whitehouse.gov/>

World Wide Web Consortium (W3C). Standards body that publishes the PROV-O provenance ontology referenced in Appendix A. <https://www.w3.org/>

X (formerly Twitter). Social network owned by xAI and cited in §1 as one of the largest first-party training-data holders. <https://x.com/>

xAI. Frontier artificial-intelligence laboratory founded by Elon Musk and the operator of X. <https://x.ai/>

Interested in Contributing?

The AIISP-1 specification, this paper, and the reference implementation are developed openly. Contributions, issues, and pull requests are welcomed at the standard’s repository, <https://github.com/97115104/aiisp-spec>, with an open tracking issue for this draft at <https://github.com/97115104/aiisp-spec/issues/1>, and direct correspondence about partnership, pilot funding, worker-organization or community co-authorship, and editorial review can be addressed to contribute@humandatacollective.org.

Acknowledgements

This paper owes a substantial intellectual debt to Karen Hao’s *Empire of AI* [9], whose on-the-ground investigation of OpenAI, its capital structure, its labor supply chain, and the communities affected by its physical build-out is drawn on throughout, and to the workers and community organizers whose testimony it records, particularly Sonia Ramos in the Atacama and the organizers of MOSACAT in Cerrillos, whose framing of territorial injustice informs the environmental argument made here. The paper also owes a debt to Dwarkesh Patel and Gavin Leech for *The Scaling Era* [20], whose oral history of the field provided the practitioners’ own framing of scaling laws, data curation, and the human-feedback bottleneck used in the data section.

The argument made here also rests on the published work of researchers whose findings the paper draws on directly: Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Margaret Mitchell (writing as “Shmargaret Shmitchell”) for *On the Dangers of Stochastic Parrots* [4]; Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang for the formal account of why language models hallucinate [12]; Jared Kaplan and co-authors for the foundational scaling-laws paper [13]; Ilya Shumailov and co-authors for the *Nature* paper on model collapse under recursively-generated data [26]; Emma Strubell, Ananya Ganesh, and Andrew McCallum for early quantification of the energy and policy footprint of NLP training [28]; Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres for the machine-learning carbon-emissions methodology [14]; Jaron Lanier and E. Glen Weyl for the data-dignity framing [15]; Trebor Scholz for *Platform Cooperativism* [25] and Nathan Schneider for *Everything for Everyone* [24], whose work on cooperative ownership informs the Collective’s governance design; and Barry M. Leiner and the co-authors of the ACM SIGCOMM history of the Internet [16], whose account of an open, standards-based network is the model the proposal aspires to.

In addition, this paper and its ideas are a product of decades of effort by those who built the current generation of large language models into products the public actually uses, including the teams at OpenAI, Microsoft, Anthropic, Google DeepMind, X, and Meta. The proposal offered here is intended as a constructive follow-on to that work, an attempt to help the outputs of it better serve the creators, communities, and environments on which the inputs depend, and is not a critique of the building itself.

I’d like to thank the many AI assistants and agents I have used to create this and other projects, specifically, VS Code and GitHub Copilot, Claude 4.7, and Gemini 3.1 Pro (Preview). As an individual with only a singular perspective, I am grateful for the work of those these models have been trained on and for the promise LLMs and similar applications hold to amplify human creativity and productivity now and in the future.

Finally, I’d like to thank my friends and family for their support and feedback on this and countless other projects. It would not be possible without you.

with gratitude

*For everyone whose work, words, and care went into this paper without a byline: the moderator who tagged the post, the worker who labelled the dataset, the reviewer who flagged the harm, the activist who filed the records request, the engineer who fixed it on a weekend, the parent who covered an evening, the friend who read a draft and said it could be clearer.
If a future version of this standard reaches the people whose contributions made it possible, the credit belongs to you. Any errors that remain are mine alone.*

A. H.
Los Angeles County, April 2026

COLOPHON. Set in Computer Modern with L^AT_EX on macOS. Drafted, revised, and typeset between November 2025 and April 2026 with the assistance of Claude Opus 4.7 and Gemini 3.1 Pro (Preview) under human direction. Released for public comment under CC BY 4.0; the AIISP-1 reference implementation is released under the MIT Licence.

References

- [1] S. Altman. Oversight of A.I.: rules for artificial intelligence. Testimony, U.S. Senate Judiciary Subcomm. on Privacy, Technology, and the Law, May 2023. <https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-rules-for-artificial-intelligence>
- [2] Anthropic. Anthropic: Company. 2021. <https://www.anthropic.com/company>
- [3] Anthropic. Introducing Claude Opus 4.7. April 2026. <https://www.anthropic.com/news/claude-opus-4-7>
- [4] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell. On the dangers of stochastic parrots. *FAccT*, pp. 610–623, 2021. <https://doi.org/10.1145/3442188.3445922>
- [5] Bonneville Environmental Foundation. Water Restoration Certificates. 2024. <https://www.b-e-f.org/programs/water-restoration-certificates/>
- [6] Coalition for Content Provenance and Authenticity. *C2PA Technical Specification*, v2.1, 2024. https://c2pa.org/specifications/specifications/2.1/specs/C2PA_Specification.html
- [7] B. Courty, V. Schmidt, et al. CodeCarbon: estimate and track carbon emissions from compute. 2024. <https://github.com/mlco2/codecarbon>
- [8] GitHub. Claude Opus 4.7 is generally available. April 2026. <https://github.blog/changelog/2026-04-16-claude-opus-4-7-is-generally-available/>
- [9] K. Hao. *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. Penguin, 2025.
- [10] A. Harshberger. *AIISP-1: The AI Inference Standards Protocol*, draft v0.1, April 2026. <https://github.com/97115104/aiisp-spec>
- [11] International Institute for Environment and Development. Energy transition, community fragmentation and territorial justice: an interview with Sonia Ramos. April 2024. <https://www.iied.org/energy-transition-community-fragmentation-territorial-justice-interview-sonia-ramos>
- [12] A. T. Kalai, O. Nachum, S. S. Vempala, and E. Zhang. Why language models hallucinate. arXiv:2509.04664, 2025. <https://arxiv.org/abs/2509.04664>
- [13] J. Kaplan et al. Scaling laws for neural language models. arXiv:2001.08361, 2020. <https://arxiv.org/abs/2001.08361>
- [14] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres. Quantifying the carbon emissions of machine learning. arXiv:1910.09700, 2019. <https://arxiv.org/abs/1910.09700>
- [15] J. Lanier and E. G. Weyl. A blueprint for a better digital society. *Harvard Business Review*, Sept. 2018. <https://hbr.org/2018/09/a-blueprint-for-a-better-digital-society>
- [16] B. M. Leiner et al. A brief history of the Internet. *ACM SIGCOMM CCR*, 39(5):22–31, 2009. <https://doi.org/10.1145/1629607.1629613>
- [17] microsoft/vscode contributors. Why is Claude Opus 4.7 priced at 7.5x even after the 50% promotion? Issue #310775, April 2026. <https://github.com/microsoft/vscode/issues/310775>
- [18] OpenAI. Introducing OpenAI. 2015. <https://openai.com/index/introducing-openai/>
- [19] OpenAI. *GPT-4o System Card*. 2024. <https://openai.com/index/gpt-4o-system-card/>
- [20] D. Patel and G. Leech. *The Scaling Era: An Oral History of AI, 2019–2025*. Stripe Press, 2025.
- [21] Royal. Own a piece of your favorite songs. 2021. <https://royal.io>
- [22] R. Rumburg, F. Sethi, and H. Nagaraj. *Audius: A Decentralized Protocol for Audio Content*. 2018. <https://whitepaper.audius.co/AudiusWhitepaper.pdf>
- [23] RWA.io. Carbon credit tokenization: tracking and settlement. January 2026. <https://www.rwa.io/post/carbon-credit-tokenization-tracking-and-settlement>
- [24] N. Schneider. *Everything for Everyone*. Nation Books, 2018.
- [25] T. Scholz. *Platform Cooperativism*. Rosa Luxemburg Stiftung, 2016.
- [26] I. Shumailov et al. AI models collapse when trained on recursively generated data. *Nature*, 631:755–759, 2024. <https://doi.org/10.1038/s41586-024-07566-y>
- [27] Stanford HAI. *The 2025 AI Index Report*. 2025. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- [28] E. Strubell, A. Ganesh, and A. McCallum. Energy and policy considerations for deep learning in NLP. *ACL*, pp. 3645–3650, 2019. <https://doi.org/10.18653/v1/P19-1355>
- [29] Terrapass. The 2026 complete guide to water credits (WRCs). March 2026. <https://terrapass.com/blog/the-2026-complete-guide-to-water-credits-wrcs/>
- [30] The White House. Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. 88 FR 75191, 2023. <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>
- [31] World Wide Web Consortium. *PROV-O: The PROV Ontology*. W3C Recommendation, April 2013. <https://www.w3.org/TR/prov-o/>
- [32] California State Legislature. Assembly Bill 2013: Generative Artificial Intelligence – Training Data Transparency. Chapter 817, Statutes of 2024, approved September 28, 2024. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB2013
- [33] California State Legislature. Assembly Bill 375: California Consumer Privacy Act of 2018. Chapter 55, Statutes of 2018, codified at Cal. Civ. Code §§ 1798.100–1798.199. https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5
- [34] G. Teare. Q1 2026 shatters venture funding records as AI boom pushes startup investment to \$300B. *Crunchbase News*, April 1, 2026. <https://news.crunchbase.com/venture/record-breaking-funding-ai-global-q1-2026/>
- [35] TBWA\Chiat\Day (R. Siltanen and K. Segall, copy) for Apple Computer, narrated by S. Jobs. *Think Different* (“Here’s to the Crazy Ones”), broadcast and print campaign, 1997. Full text: <https://www.thecrazyones.it/spot-en.html>
- [36] M. Zeff. Anthropic tightens usage limits for Claude Code – without telling users. *TechCrunch*, July 17, 2025. <https://techcrunch.com/2025/07/17/anthropic-tightens-usage-limits-for-claude-code-without-telling-users/>