Provenance - Fact Checker (Whitepaper)

Provenance is an AI-powered fact-checking agent and SaaS platform, incubated by Fact Protocol, designed to combat misinformation at scale. By anchoring provenance records for digital content on the blockchain, Provenance delivers transparency, accountability, and trust for AI models, agents, and human users.

Anchoring information integrity on-chain.

*Powered by **$PROV** (official token).*

Provenance

ISNI:

0000 0005 2824 1030

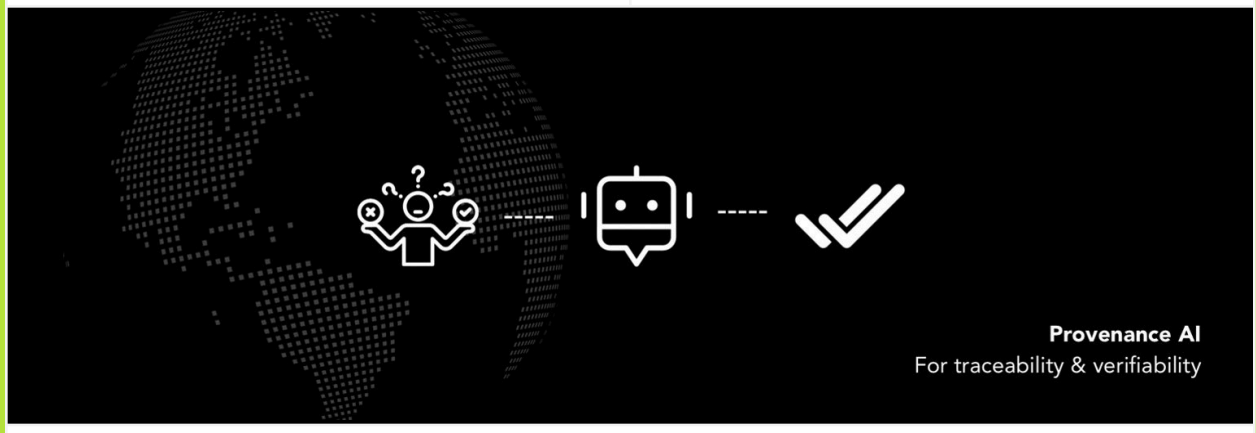

To Mitigate AI-driven Misinformation:

Recent research highlights a significant increase in AI-driven misinformation. A NewsGuard study revealed that the rate of false information generated by popular AI chatbots nearly doubled in one year, from 18% to 35%, when responding to news-related prompts. (newsguardtech.com)

✓✓ Abstract

The rapid evolution of Generative AI has fueled an explosion of synthetic and often misleading content. Social media platforms, particularly X (Twitter), have become the ground zero where misinformation spreads at viral speeds, eroding public trust and undermining discourse integrity.

Provenance AI is a fact-checking agent and SaaS platform powered by blockchain and AI. It verifies claims in real time, provides transparent evidence, and anchors results immutably on-chain. By combining decentralized trust infrastructure with advanced AI retrieval and validation, Provenance AI delivers a scalable solution to combat misinformation in the AI information age.

ISNI:
0000000528241030

Whitepaper:
v1.0.4

Ticker:
$PROV

X (Twitter):
x.com/0xProvenance

Telegram:
t.me/ProvenanceAI

Website:
provenance.technology

Logo:

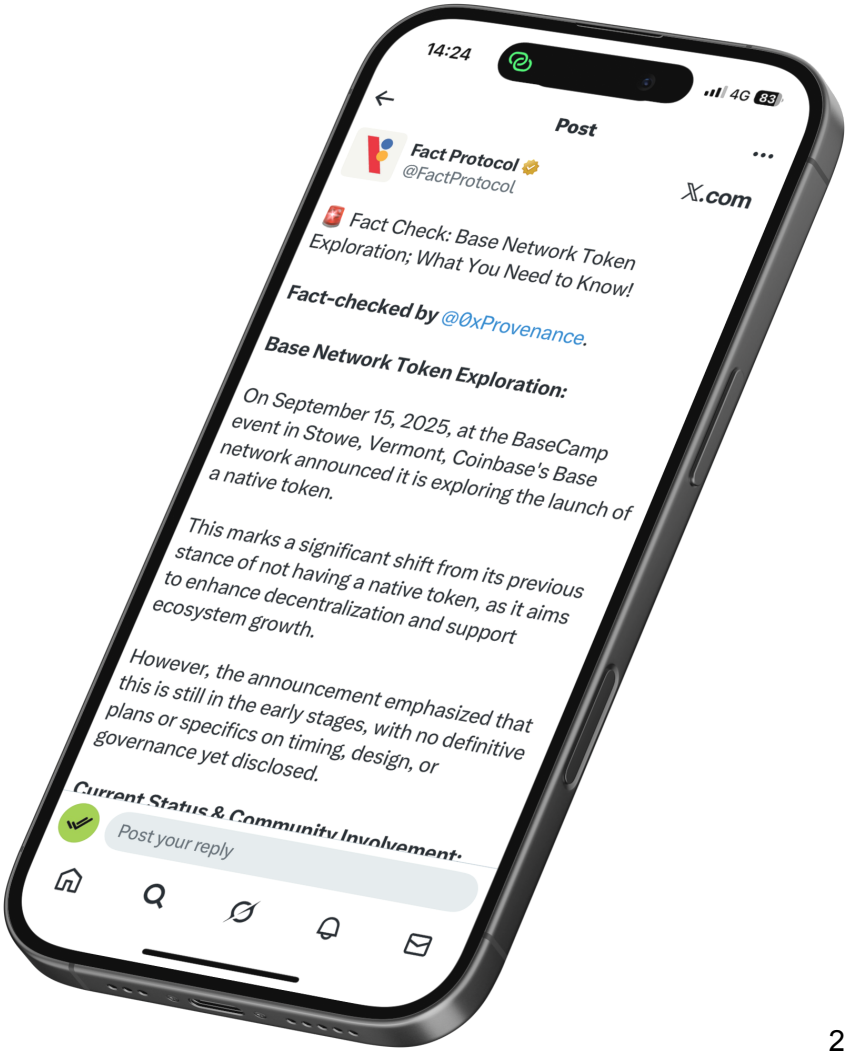

✓✓ Introduction

The proliferation of AI-generated text, images, and videos has brought both innovation and disruption. A study by NewsGuard revealed that AI chatbots generate misinformation in **35% of cases involving current events**, nearly double the rate from the previous year [*Axios, 2025*].

Similarly, **58% of U.S. adults report being deceived by AI-generated news**, with **69% of Gen Z** stating they've fallen for fake AI content [*HostingAdvice, 2025*].

This crisis of trust demands urgent solutions. Provenance AI addresses this challenge by acting as a **verifiable trust layer**: a real-time fact-checking AI agent integrated with X (Twitter) and supported by a dedicated blockchain, Provenance Chain (prospect). It empowers individuals, communities, and institutions to validate information reliably, transparently, and at scale.

Traditional fact-checking vs. Provenance AI

Aspect	Traditional Fact-checking	Provenance AI
Scalability	Limited scalability, labor-intensive processes	High scalability, leveraging AI and decentralized systems
Perspective diversity	Focuses on validating single truths	Offers multiple credible perspectives/sources
Timeliness	Slower response to evolving narratives	Real-time adaptability
Bias mitigation	Subject to institutional or individual biases	Decentralized governance ensures balanced outputs
User empowerment	Provides answers, with a limited focus on user growth	Encourages critical thinking and intellectual growth
Community involvement	Centralized operations, minimal public participation	Involves the community through tokenized governance
Cost efficiency	High costs due to manual processes	Cost-efficient with automated AI and decentralized systems

✓✓ Vision Statement

To become the **global standard for verifiable truth in the AI era**, where every digital claim, from breaking news to enterprise reports, carries a transparent, tamper-proof provenance trail.

✓✓ Our Motivation to Mitigate

AI-driven misinformation poses a significant threat to economies, democracies, and trust. Provenance AI bridges the gap with real-time verification and blockchain-backed transparency.

- **Misinformation is costly:** Deepfake scams alone may cost the global economy **$40 billion in the next 3 years** [Wikipedia/MarketsandMarkets, 2024].
- **Democratic integrity is at risk:** Misleading content impacts elections, journalism, and freedom of expression [Reuters Institute, 2023].
- **AI amplifies speed:** The velocity of misinformation far outpaces traditional fact-checking.

Provenance AI exists to bridge this gap by applying **AI for real-time responses** and **blockchain for transparent verification**.

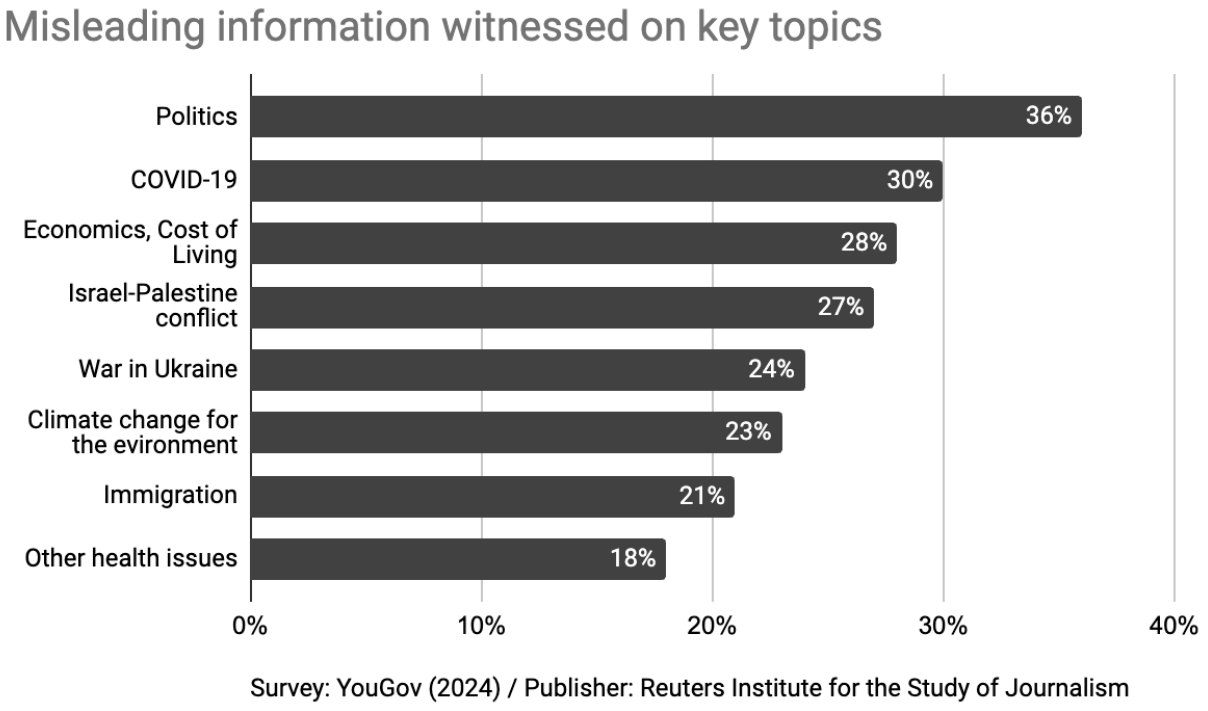

✓✓ Problem Statement

The digital world is facing an accelerating crisis of disinformation and misinformation, amplified by the speed and scale of AI. Existing moderation and fact-checking methods are fragmented, slow, and unable to meet the demand for trustworthy verification.

1. **Information Overload:** Billions of posts are posted daily on platforms like X (formerly Twitter), making manual moderation impossible.
2. **AI-Generated Misinformation:** LLMs and generative models produce credible yet false narratives, images, and videos.
3. **Trust Deficit:** Users, governments, and institutions lack scalable ways to verify authenticity.
4. **Fragmented Efforts:** Current fact-checking initiatives are siloed, slow, and lack accountability/transparency.

✓✓ Solution

Provenance AI delivers a scalable trust layer by combining AI-driven fact-checking with blockchain-based transparency. From real-time verification on social platforms to enterprise-grade dashboards, every claim is backed by evidence, anchored on-chain, and reinforced by human oversight. Below are some of the strategies in our roadmap:

- **On-chain Attestations:** Each verdict is anchored to on-chain Structured data (Schema), including the claim hash, evidence references, model version, and timestamp.
- **On X (Twitter):** Users tag **@0xProvenance** (AI Agent) with a claim or link → AI triages, retrieves evidence, evaluates, and replies with a verdict + citations.
- **SaaS (Tooling/APIs):** A standalone, enterprise-grade verification tool for direct claim submissions, organizational dashboards, and appeal/dispute flows.
- **Human-in-the-loop (Powered by Web3):** Sensitive claims and disputes involve validators (journalists, academics, experts). AI can make mistakes, and it needs feedback.

✓✓ How Provenance AI Differs from Other Fact-Checkers

Unlike conventional fact-checkers that only verify claims at the surface level, Provenance AI goes deeper by identifying and recording the provenance, or true origin, of any information.

By establishing a source trail for every piece of content, Provenance AI enables verifiability of truth and misinterpretation for anyone —humans, organizations, and AI systems alike.

Each verified claim is stored in a structured format compatible with schema.org standards and JSON-LD, including its contextual summary, provenance data, and evidence. This enables seamless integration with tools and platforms such as Google, Bing, and other AI- or data-driven systems.

Example structured data specified in the **On-chain Structured Data** section below.

Most importantly, this structured, verifiable dataset empowers AI models to be trained on trustworthy data, resulting in more accurate, transparent, and bias-resistant outputs for end users.

🧭 ***Traditional fact-checkers stop at truth.***
🔗 ***Provenance AI starts from truth and builds verifiable knowledge for the AI era.***

✓✓ On-chain Structured Data

Provenance AI utilizes Fact Protocol's proposed standard for schema markup, which ensures on-chain storage of the provenance trail and fact-check attestations, enabling retrievability by both AI systems and humans. Below is an example of markup with source and publisher details:

```
{
  "@context": "https://schema.org",
  "@type": "WebPage",
  "mainEntity": {
    "@type": "Article",
    "headline": "Difference: Objective Facts vs. Subjective
Facts or Claims",
    "url":
"https://fact.technology/learn/difference-objective-facts-vs-sub
jective-facts/",
    "datePublished": "2023-07-31",
    "image":
"https://fact.technology/files/2023/07/difference-objective-fact
s-subjective-facts-claims-opinions-fact-protocol-learn.png",
    "author": {
      "@type": "Person",
      "name": "Fact Protocol Team"
    }
  },
  "publisher": {
    "@type": "Organization",
    "name": "Fact Protocol Learn",
    "sameAs": "https://fact.technology/learn/"
  },
  "factProvenance": {
    "@type": "FactProvenanceTag",
    "originalSource":
"https://www.jibc.ca/sites/default/files/community_social_justic
e/pdf/cl/Objective_vs._Subjective.pdf"
  }
}
```

✓✓ AI Agent: Fact-Checking Where Conversations Happen

The Provenance AI Agent delivers real-time, evidence-backed fact-checks directly on X (Twitter), providing users with instant clarity on the platform where information spreads the fastest.

By tagging ***@0xProvenance***, anyone can request verification and receive a transparent, citation-rich response anchored on-chain, for future retrievability and verifiability, ***powered by $PROV***.

As part of the Provenance roadmap, the agent will expand beyond X (Twitter) to other platforms, such as Telegram, ensuring fact-checking is available wherever digital conversations take place.

✓✓ Enterprise-Ready Fact-Checking SaaS

Provenance SaaS delivers a powerful fact-checking interface and API for seamless integration into newsrooms, platforms, and enterprise systems. Built for scale, every verification is transparent, auditable, and on-chain, immutably anchored.

Through its intuitive dashboard, individuals can easily submit a claim, receive AI- and human-validated results, and trace each proof back to its verifiable origin. As shown in the example below, the SaaS interface provides a seamless and reliable experience for verifying any claim.

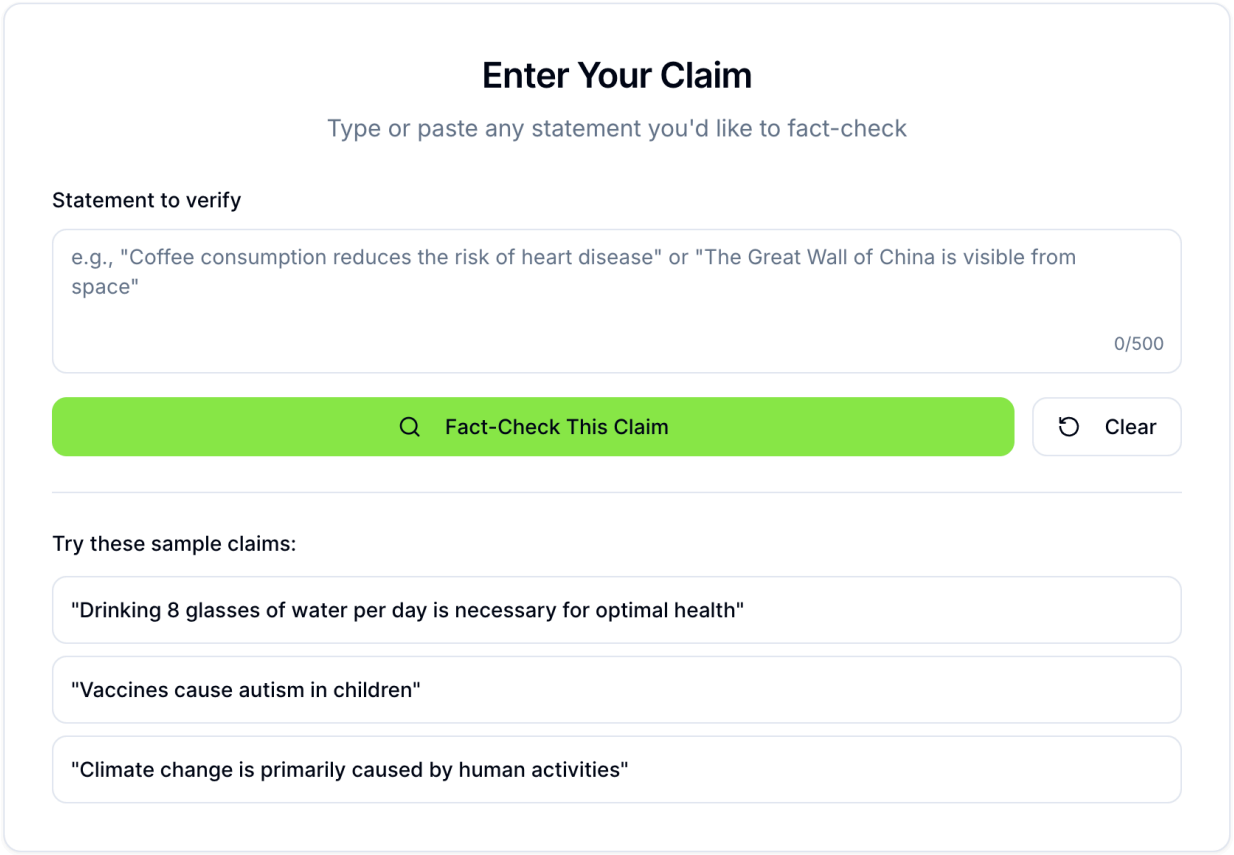

✓✓ SaaS: Clear, Evidence-Based Fact-Check Results

Provenance AI delivers fact-check results that are simple to read yet backed by rigorous evidence. Each claim is summarized with a clear verdict, a concise analysis, and supporting sources from high-credibility institutions. This ensures users not only see the outcome but also understand the context and trust the proof behind it.

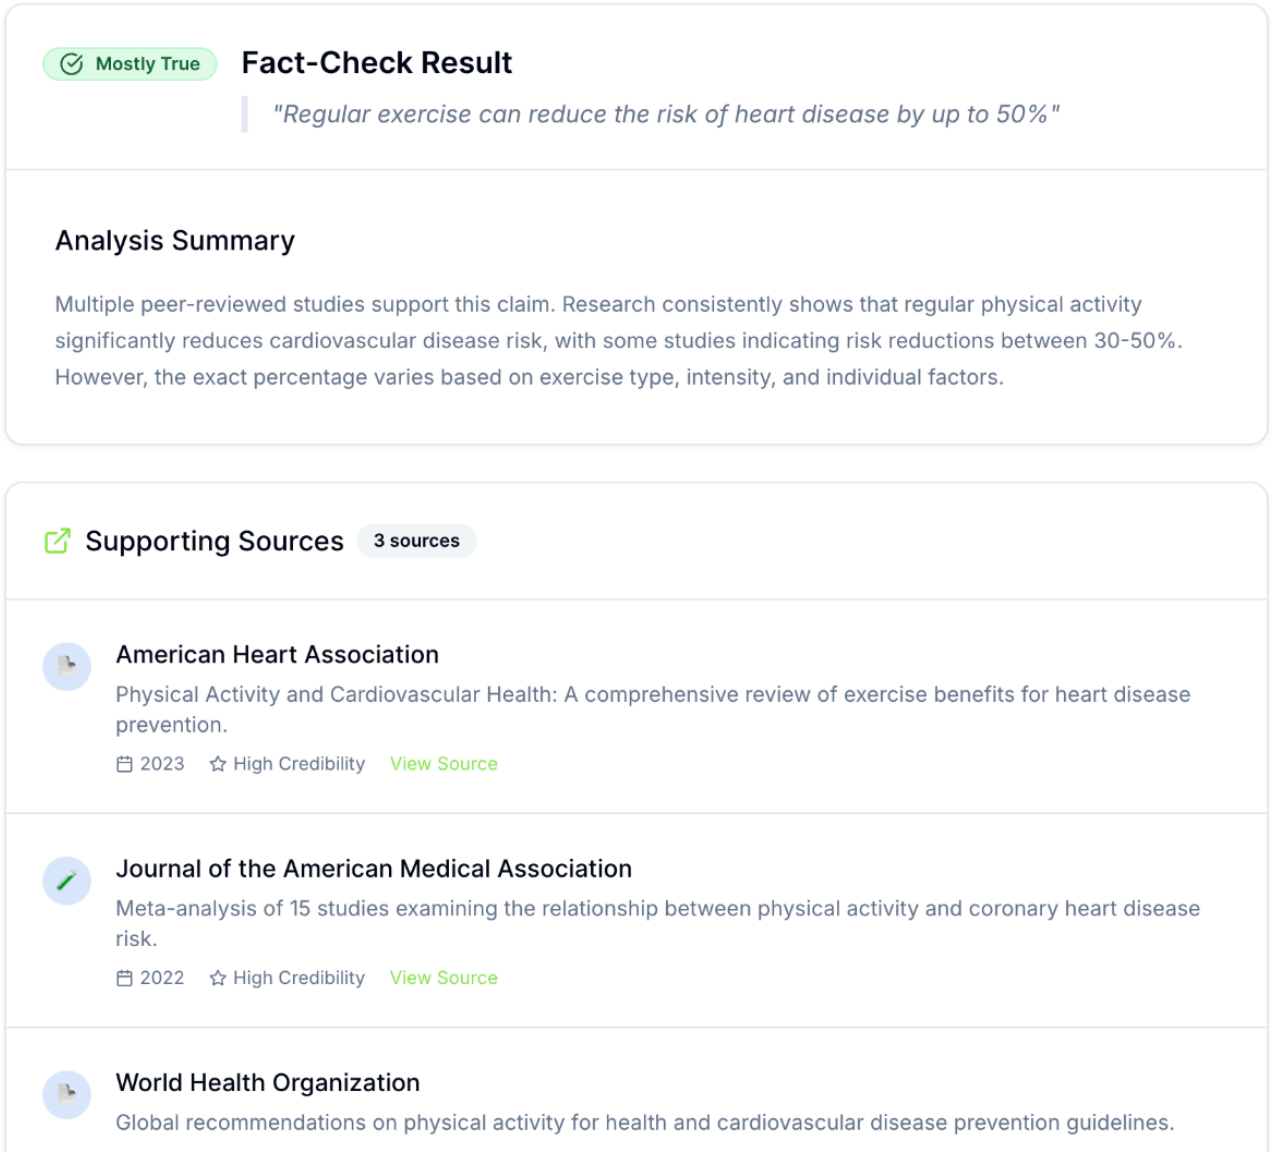

✓✓ Roadmap

The Provenance AI roadmap outlines our journey to combat misinformation through innovation and community collaboration. With a phased approach, we aim to expand across platforms and integrate multi-model AI.

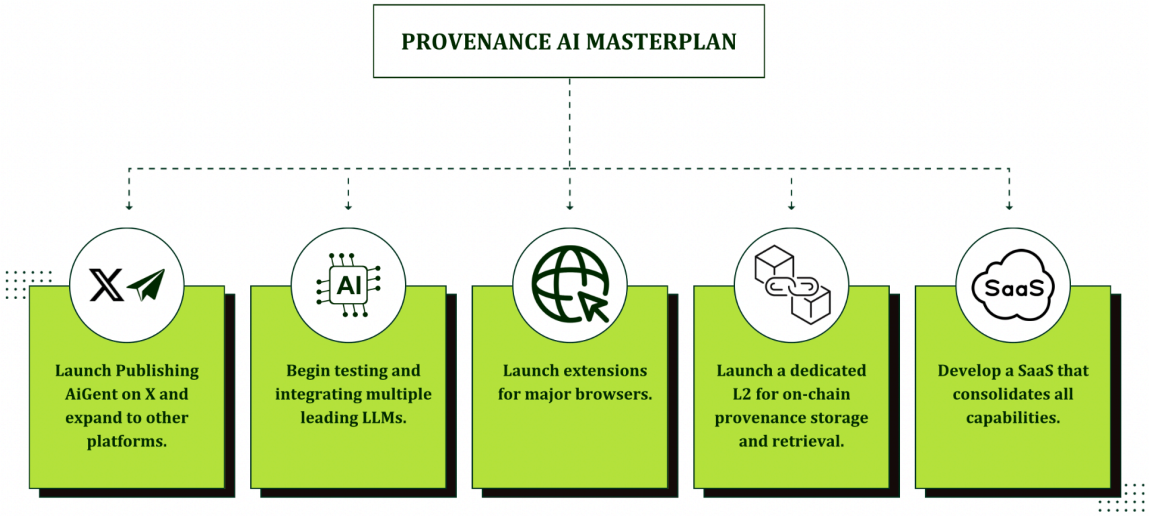

✓✓ Framework

The Provenance AI framework combines AI, human validation, and blockchain to create a transparent and scalable fact-checking system. Each layer, from retrieval to governance, ensures that information is verified, auditable, and community-driven.

- **Retrieval Pipeline:** Aggregates data from verified sources, fact-check databases, and academic APIs.
- **Evaluation Engine:** AI + human validators to assess claims.
- **On-chain Infrastructure:** Immutable proofs on Provenance Chain ensure trust and auditability.
- **Governance:** Token holders vote on model upgrades and economics.

✓✓ Incentivization

Incentivization ensures active participation and sustainability within the Provenance ecosystem. By aligning rewards and utility across users, validators, stakers, and enterprises, the model fosters continuous growth while maintaining trust and transparency.

- **Users:** Pay micro-fees to request fact-checks.
- **Validators:** Earn rewards for accurate validations (prospect).
- **Stakers:** Stake tokens to participate in ecosystem growth.
- **Developers/Enterprises:** API credits for integrations.

✓✓ Token Flywheel

Demand for fact-checks drives usage of the Provenance AI Agent and SaaS platform. Each claim processed generates fees, while enterprise and SaaS offerings expand adoption. Revenues are cycled back into the ecosystem through token buybacks, creating sustained value for **$PROV** holders and reinforcing the continuous growth of the fact-checking network.

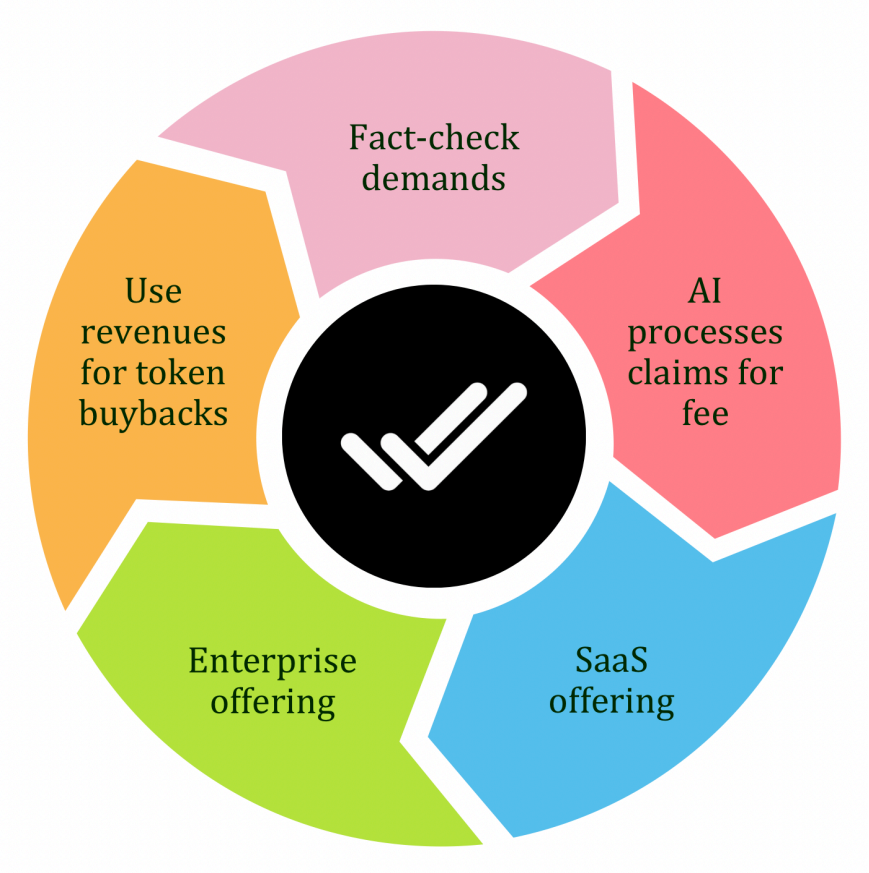

✓✓ Challenges and Limitations

While Provenance AI presents a transformative solution, it is essential to recognize potential challenges and proactively address them:

Issue	Challenge	Potential Solution
Scalability	Handling a vast amount of data across diverse languages, cultures, and topics could strain computational resources and slow down performance.	Leverage distributed computing architectures and community-driven contributions to enhance scalability. Utilize advanced optimization techniques to process multilingual and cultural datasets efficiently.
Bias in AI Models	AI models may inadvertently replicate or amplify biases present in the data they are trained on, potentially leading to imbalanced outputs.	Regularly audit the AI models and the knowledge graph for potential biases. Incorporate diverse datasets and develop algorithms to detect and neutralize bias in the system.
Regulatory Concerns	Operating in a regulatory environment that varies across countries could pose legal challenges, particularly concerning data privacy and misinformation laws.	Ensure compliance with local and international regulations by adopting privacy-by-design principles. Work with legal experts to navigate regulations while maintaining transparency & user trust.
Community Engagement	Encouraging consistent, meaningful community contributions for AI feedback and on-chain attested datasets.	Introduce incentive mechanisms through the **$PROV** token to reward active participation. Provide easy-to-use tools and clear guidelines to help community members contribute effectively.
User Adoption	Users unfamiliar with multi-perspective analysis may find the platform initially overwhelming or challenging to use.	Simplify the UI and provide guided walkthroughs to enhance new users' onboarding experience. Develop educational content that effectively showcases the platform's value in a clear, accessible manner.
Real-Time Adaptation	Continuously evolving misinformation trends and current events require rapid adaptation, which could put pressure on the platform's infrastructure.	Implement automated monitoring systems and real-time feedback loops to detect and address emerging trends quickly. Continuously refine the knowledge graph based on these insights.

✓✓ Use Cases (Goes beyond Fact-checking)

Real-time verification of claims/info empowers journalists, educators, policymakers, businesses, diplomats, and online communities to make informed and empathetic choices. Provenance AI helps mitigate biases across media reporting, academic exploration, policy analysis, and public discourse.

Media and Journalism	
	**Real-time verification of breaking news claims.** Provenance AI can be a critical tool for journalists and media outlets to access live fact-checking and diverse perspectives, thereby validating the credibility of sources.
Education and Academia	
	Teachers, researchers, and students can utilize Provenance AI to explore verifiable information on a topic, allowing for reliability and intellectual growth within academic environments.
Policy Making and Governance	
	Governments and policymakers can utilize the platform to evaluate the implications of proposed policies from various verifiable perspectives and societal sentiments, thereby ensuring more informed and equitable decision-making.
Business and Market Analysis	
	Provenance AI can help businesses understand market trends, customer sentiments, and competitive landscapes by analyzing diverse opinions and insights.
Conflict Resolution and Negotiation	
	Organizations involved in diplomacy or conflict resolution can utilize the platform to gain a deeper understanding of the motivations and concerns of all stakeholders, thereby facilitating empathy and constructive dialogue.
Social Media and Public Discourse	
	Provenance AI can enhance the quality of online discussions by injecting verifiable fact-checks and mitigating the spread of misinformation in real-time.

✓✓ Impact Analysis

Provenance AI has a profound and far-reaching impact across multiple sectors. It builds a society that is more informed, empathetic, and resilient by addressing the root causes of misinformation and polarization. Through its scalable and adaptable design, the tool not only enhances democratic processes and intellectual growth but also empowers businesses, educators, and policymakers with actionable insights.

Provenance AI is set to disrupt the future by establishing trust through AI-powered fact-checking, verifiable provenance, and global accessibility.

Factor	Analysis	Potential Impact
Societal	Provenance AI builds a more informed and connected global society by combating misinformation and encouraging verify first, trust later.	**Positive**
Educational	In academic settings, Provenance AI inspires critical thinking and intellectual curiosity. Students and researchers benefit from exposure to diverse viewpoints, which fosters a deeper understanding and stimulates innovative ideas.	**Positive**
Economic	By helping businesses and policymakers make verifiable, data-driven decisions, Provenance AI contributes to more effective strategies, deeper market insights, and better organizational outcomes.	**Positive**
Democratic	By empowering citizens with accurate, trustworthy, and verifiable information, the platform reinforces democratic processes, enabling informed voting and participation in public discourse.	**Positive**
Global Scale	Designed for scalability, Provenance AI addresses the global challenges of disinformation and misinformation. It establishes a universal standard for verifiable provenance, contributing to international efforts to build information integrity.	**Positive**

✓✓ Ethical Considerations

Provenance AI is committed to upholding the highest ethical standards in its design, deployment, and operation of its fact-checking tools and strategies. Recognizing the sensitive nature of combating misinformation and presenting diverse facts and perspectives with verifiable sources.

- **Free Speech:** Fact-checking is advisory, not censorship.
- **Transparency:** All verdicts cite evidence and model versions.
- **Accountability:** Dispute and appeal mechanisms ensure fairness.
- **Privacy:** Claims and sources are logged without storing personal user data.

✓✓ Team

Provenance AI is incubated at Fact Protocol, a recognized innovator in decentralized fact-checking, cited by global media, universities, Fortune 500 companies, and governments. The team consists of experts in AI, blockchain, journalism, and misinformation research.

Notable Citations and Mentions of Fact Protocol: A Selection from Hundreds.

UNIVERSITY OF SAN FRANCISCO	EMARKETER
Hindustan Times @100	FINANCIAL EXPRESS Read to Lead
COINTELEGRAPH The future of money	frontiers
IEEE Xplore® Digital Library	The Alan Turing Institute
INSIGHTS ie UNIVERSITY	College of Staten Island The City University of New York
LTIMindtree	eLife

✓✓ Get Involved

Join us in shaping the future of real-time information fact-checking and information integrity!

Provenance AI is more than just a tool; it's a movement to combat misinformation and empower individuals with access to instant fact-checking. Here's how you can get involved:

1. **Community members:** Join our Telegram.
2. **Researchers:** Collaborate on misinformation datasets.
3. **Validators:** Apply to join our web3-powered fact-checking network.
4. **Partners:** Integrate Provenance APIs for trusted data.

Together, we can build a more informed, empathetic, and resilient global society.

Website: https://provenance.technology
X (formerly Twitter): https://x.com/0xProvenance
Telegram: https://t.me/ProvenanceAI
Instagram: https://www.instagram.com/0xProvenance
Facebook: https://www.facebook.com/0xProvenance

✓✓ Glossary

Misinformation: Misinformation is false or misleading information that is unintentionally shared. It can be the result of a misunderstanding, a mistake, or a lack of knowledge.

Disinformation: Disinformation is false or misleading information that is intentionally created to deceive or manipulate people. It can take many forms, such as fake news and propaganda.

AI (Artificial Intelligence): The use of computer programs that have some of the qualities of the human mind, such as the ability to understand language, recognize pictures, and learn from experience. [Ref: https://dictionary.cambridge.org/dictionary/english/artificial-intelligence]

AI Agent: Refers to a system or program that is capable of autonomously performing tasks on behalf of a user or another system by designing its workflow and utilizing available tools. [Ref: https://www.ibm.com/think/topics/ai-agents]

LLMs (Large Language Models): A computer program that uses very large collections of language data to understand and produce text in a way that is similar to the way humans do. [Ref: https://www.oxfordlearnersdictionaries.com/definition/english/llm]

Provenance: The provenance of something is the place that it comes from or that it originally came from. [Ref: https://www.collinsdictionary.com/dictionary/english/provenance]

Bias: An unfair personal opinion that influences your judgment.

Critical Thinking: Critical Thinking is the process of using and assessing reasons to evaluate statements, assumptions, and arguments in ordinary situations. [Ref: https://iep.utm.edu/critical-thinking/]

Fact-checking: Fact-checking is the process of verifying the accuracy of information and statements made by individuals or organizations. It is an essential aspect of journalism and research, as it ensures that the public is provided with reliable and accurate information. [Ref: https://fact.technology/learn/what-is-fact-checking/]

Community Governance: Community governance is a system of rules, processes, and practices that enables communities to have representation and influence over decisions that affect them or the projects in which they are involved.

Tokenization: Tokenization is the process of creating a digital representation of a real thing. [Ref: https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-tokenization]

$PROV: The ticker for Provenance AI's official crypto token.

Fact Protocol: An AI & Web3-based Decentralized Fact-checking System. [Ref: https://cointelegraph.com/magazine/real-ai-use-cases-crypto-fighting-fakes-trusted-content/]

✓✓ Credits

Provenance AI Whitepaper prepared for the Provenance AI project/team by the Fact Protocol team, drawing insights from misinformation researchers, AI safety experts, and blockchain architects.

With Thanks to:

1. Icons: Pixel perfect - Flaticon
2. Official Website's DDoS Protection: Project Shield

✓✓ References

Fact Protocol. https://fact.technology/learn/ (ISSN:2836-5925)

NewsGuard & Axios (2025). *Chatbots' misinformation rate doubles year-over-year*. Link

HostingAdvice (2025). *58% of U.S. adults are misled by AI fake news*. Link

MarketsandMarkets (2024). *Deepfake detection market report*.

Reuters Institute (2023). *Digital News Report*.

OECD (2023). *Addressing the Risks of Disinformation in the Digital Age*.