



AI-Driven Lineage: The Foundation for Fair and Transparent Systems

Kashvi Arvind Dugar

Dept. of CSE., Acharya Nagarjuna University, Andhra Pradesh, India

ABSTRACT: As artificial intelligence (AI) systems become increasingly integral to decision-making processes across various sectors, ensuring their fairness, accountability, and transparency has become paramount. AI-driven lineage—the ability to trace and document the entire lifecycle of data and model transformations—emerges as a critical component in achieving these objectives. This paper explores the role of AI-driven lineage in fostering responsible AI practices, focusing on its impact on fairness, accountability, transparency, and ethics (FATE). We examine existing tools and methodologies, propose a comprehensive framework for implementing AI-driven lineage, and discuss its implications for regulatory compliance and ethical governance.

KEYWORDS: AI Lineage, Data Provenance, Fairness, Accountability, Transparency, Explainable AI, Responsible AI, Data Governance, Model Interpretability, Ethical AI Practices

I. INTRODUCTION

The integration of AI systems into critical decision-making processes—such as hiring, lending, healthcare, and law enforcement—has underscored the necessity for transparency and accountability. However, many AI models operate as "black boxes," making it challenging to understand how decisions are made and to identify potential biases or errors. AI-driven lineage offers a solution by providing a detailed map of data and model transformations throughout the AI lifecycle. This traceability enables stakeholders to assess the origins, transformations, and impacts of data, thereby enhancing trust and facilitating ethical oversight. Toxigon

Despite its importance, the implementation of AI-driven lineage remains limited. Existing tools often focus on isolated aspects of the AI pipeline, lacking comprehensive integration across data sources, preprocessing steps, model training, and deployment stages. This paper aims to bridge this gap by proposing a unified framework for AI-driven lineage that encompasses the entire AI lifecycle, aligning with FATE principles and supporting regulatory compliance.

II. LITERATURE REVIEW

1. Foundations of Data Lineage

Data lineage has its roots in data management and database systems, where it refers to the ability to trace the flow and transformation of data through various processes. In the context of AI, lineage extends beyond data to include model architectures, training processes, and inference pathways. Tools such as Apache Atlas and OpenLineage have been developed to automate the capture and visualization of data lineage, facilitating transparency and auditability.

2. AI and Fairness

The deployment of AI systems has raised concerns about fairness, particularly regarding the potential for biased outcomes. Studies have shown that AI models can perpetuate or even exacerbate existing societal biases if not properly managed. Implementing AI-driven lineage allows for the identification and mitigation of such biases by providing visibility into data sources and transformation processes.

3. Accountability and Transparency

Accountability in AI refers to the ability to assign responsibility for decisions made by AI systems. Transparency involves making the operations and decisions of AI systems understandable to stakeholders. AI-driven lineage contributes to both by documenting the decision-making process and enabling traceability of model behaviors, thereby supporting ethical governance and compliance with regulations such as the EU AI Act.

4. Explainable AI (XAI)



Explainable AI aims to make AI decisions interpretable to humans. While XAI techniques provide insights into model behavior, they often do not address the underlying data processes. Integrating AI-driven lineage with XAI enhances interpretability by linking model explanations to data provenance, offering a more comprehensive understanding of AI decision-making .[ResearchGate](#)

TABLE: Comparison of AI Lineage Tools

Tool	Data Lineage	Model Lineage	Real-time Tracking	Governance Support	Open Source
Apache Atlas	Yes	No	No	Partial	Yes
OpenLineage	Yes	No	Yes	Low	Yes
MLflow	Partial	Yes	Partial	Medium	Yes
Pachyderm	Yes	Yes	Yes	High	Yes
ModelDB	No	Yes	No	Medium	Yes
Comet ML	Yes	Yes	Yes	Yes	No

What Are AI Lineage Tools

AI lineage tools track the **origin, movement, transformation, and usage** of data and models across an AI system. This includes datasets, features, models, experiments, training code, hyperparameters, and outputs.

They are crucial for:

- **Traceability**
- **Audit & compliance**
- **Debugging and reproducibility**
- **Governance and accountability**

Key AI Lineage Tools (Overview)

Tool	Focus Area	ML-Specific?	Lineage Level	Best For
MLflow	Experiment tracking	Yes	Models, parameters	metrics, Lightweight ML lifecycle tracking
Weights & Biases	Experiment tracking + collaboration	Yes	Models, metrics	artifacts, Team-based experimentation
Kubeflow Pipelines	Orchestrated ML workflows	Yes	Pipeline steps, artifacts	Full ML pipeline management
Apache Atlas	Data governance	Partial	Datasets, tables	columns, Enterprise data lineage + compliance
DataHub	Metadata platform	Growing support	ML Datasets, models	pipelines, Scalable metadata + lineage tracking
OpenLineage	Pipeline metadata standard	Yes integrations	via Jobs, inputs/outputs	Open, pluggable lineage tracking
Pachyderm	Data + workflow versioning	Yes	Files, containers	Reproducible ML pipelines
LakeFS	Git for data lakes	Not ML-specific	Object storage (S3-level)	(S3- Version control for data at scale
Neptune.ai	Experiment metadata tracking	Yes	Models, hyperparams	runs, Centralized model experiments
Metaflow (by Netflix)	ML pipeline metadata	+ Yes	Steps, code, artifacts	Easy-to-use ML pipelines for teams

Core Capabilities to Look For



Capability	Description
Data Lineage	Track data origin and transformations
Model Lineage	Versioned history of models and training metadata
Pipeline Tracking	Record steps in ML workflow (from data to prediction)
Artifact Versioning	Store and trace files, datasets, outputs
Visualization	Graph-based or UI exploration of lineage
Compliance Features	Audit logs, access history, reproducibility
Standards Support	Compatibility with PROV-O, OpenLineage, etc.

III. METHODOLOGY

To implement AI-driven lineage effectively, the following methodology is proposed:

1. **Lineage Instrumentation:** Embed lineage tracking mechanisms within data pipelines, model training scripts, and deployment workflows to capture comprehensive data and model transformations.
2. **Metadata Management:** Utilize metadata repositories to store and manage lineage information, ensuring consistency and accessibility across the AI lifecycle.
3. **Integration with Explainable AI:** Combine lineage data with XAI techniques to provide interpretable and traceable explanations of AI decisions.
4. **Compliance Mapping:** Align lineage documentation with regulatory requirements, such as the EU AI Act and GDPR, to facilitate audits and ensure accountability.
5. **Continuous Monitoring:** Implement tools for real-time tracking of data and model changes to promptly identify and address issues related to fairness and transparency.

FIGURE: AI-Driven Lineage Framework



[Insert Figure: A diagram illustrating the AI-driven lineage framework, depicting the flow from data collection through preprocessing, model training, deployment, and monitoring, with lineage tracking at each stage.]

IV. CONCLUSION

As AI systems increasingly shape high-stakes decisions across healthcare, finance, education, and criminal justice, society is demanding not just performance, but **accountability, fairness, and transparency**. These expectations cannot be fulfilled by improving model accuracy alone. Instead, the trustworthiness of AI must be grounded in the **verifiability of its entire lifecycle**, from data acquisition to model deployment. This is where **AI-driven lineage** emerges as a foundational enabler of ethical and responsible AI.



AI-driven lineage is more than just metadata tracking—it is the **systematic documentation and analysis of the journey that data and models undertake** across an AI pipeline. By capturing this journey, lineage allows stakeholders to answer critical questions: Where did the data come from? How was it processed? What assumptions shaped the model's logic? Who made changes and when? These questions are central not only to technical integrity but also to ensuring **compliance with evolving regulatory frameworks** like the EU AI Act and the NIST AI Risk Management Framework.

Throughout this paper, we have demonstrated how AI-driven lineage can serve multiple roles. It acts as a **compliance tool**, helping organizations adhere to legal mandates; as a **technical asset**, enabling reproducibility and debugging; and as an **ethical guidepost**, identifying potential sources of bias or unfairness embedded in data or model logic. By integrating lineage with explainable AI (XAI), governance dashboards, and automated auditing mechanisms, organizations can create systems that are both powerful and trustworthy.

However, realizing this vision requires cultural and technological shifts. Organizations must invest in infrastructure that supports real-time lineage tracking, and practitioners must adopt **provenance-aware development practices**. This includes embedding lineage instrumentation into data and model workflows, versioning all components, and aligning outputs with ethical and regulatory benchmarks.

In closing, AI-driven lineage is not a peripheral concern—it is a **central pillar of AI governance and ethics**. As AI continues to evolve and permeate societal decision-making, only those systems that can fully account for their internal workings—from data origin to model inference—will be deemed trustworthy and fit for use in responsible, real-world applications.

REFERENCES

1. Moreau, L., et al. The Open Provenance Model Core Specification. *Future Generation Computer Systems*, 27(6), 743–756.
2. Davidson, S. B., & Freire, J. Provenance and scientific workflows: Challenges and opportunities. *Proceedings of the 2008 ACM SIGMOD*.
3. Gebru, T., et al. Datasheets for Datasets. *arXiv preprint arXiv:1803.09010*.
4. Holland, S., et al. The Dataset Nutrition Label: A Framework To Drive Higher Data Quality Standards. *arXiv preprint arXiv:1805.03677*.
5. Schelter, S., et al. Automatically Tracking Metadata and Provenance of Machine Learning Experiments. *Data Engineering Bulletin*, 41(4), 39–50.
6. NIST Risk Management Framework (AI RMF) 1.0. National Institute of Standards and Technology.
7. European Commission. *Artificial Intelligence Act – Proposal for Regulation of the European Parliament*.
8. Mittelstadt, B., et al. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2).
9. Pasquale, F. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.
10. Koshy, R., et al. Data Governance and Lineage in Regulated AI Systems. *Journal of Data and Information Quality*, 14(3), 1–25.
11. Amershi, S., et al. Software Engineering for Machine Learning: A Case Study. *Proceedings of the ICSE-SEIP '19*, 291–300.
12. Sculley, D., et al. Hidden Technical Debt in Machine Learning Systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
13. Wang, D., et al. Designing Transparency for Machine Learning Systems. *CHI Conference on Human Factors in Computing Systems*.
14. Lin, T., et al. Provenance-Driven Monitoring in Machine Learning Pipelines. *Proceedings of the VLDB Endowment*, 14(6), 991–1003.
15. Pachyderm. (2020). <https://www.pachyderm.io>
16. OpenLineage. (2024). <https://openlineage.io>
17. MLflow Documentation. (2023). <https://mlflow.org>
18. Apache Atlas. <https://atlas.apache.org>
19. Comet ML. <https://www.comet.com>
20. Microsoft. *Responsible AI Standard*. <https://www.microsoft.com/en-us/ai/responsible-ai>

