



Establishing Fairness and Transparency Through AI-Driven Data Lineage

Ayaan Dilip Kadam

Department of Computer Engineering, SVIT, Nasik, India

ABSTRACT: Artificial intelligence (AI) systems are now integral to decision-making across industries, yet concerns around fairness, transparency, and accountability continue to undermine public trust. As models grow in complexity, understanding how data flows through these systems becomes essential. **AI-driven lineage** offers a solution by providing dynamic, real-time tracking of data and model transformations, enabling transparency throughout the AI lifecycle. This paper explores the critical role of lineage in establishing fair and accountable AI systems. We analyze existing tools, frameworks, and standards, propose a methodology for implementing AI-driven lineage, and present a layered framework that supports ethical governance and regulatory compliance.

KEYWORDS: AI Lineage, Data Governance, Explainability, Transparency, Responsible AI, Data Ethics, Fairness, Model Accountability, Provenance, AI Lifecycle

I. INTRODUCTION

The surge in AI adoption across domains such as healthcare, finance, and criminal justice has raised both optimism and concern. While AI promises efficiency and insight, it also risks reinforcing biases, making opaque decisions, and evading accountability. The question arises: how do we build AI systems that are not only intelligent but **fair and transparent**?

To answer this, we must start with **lineage**—the ability to track and understand how data and models evolve through the AI lifecycle. AI-driven lineage goes beyond traditional metadata tracking. It dynamically captures the flow of information, from data ingestion and feature engineering to model training, deployment, and monitoring. This visibility is critical for ensuring fairness (by tracing the origins of bias), transparency (by showing how decisions are made), and accountability (by identifying responsible agents and processes).

Despite its importance, lineage in AI systems remains underdeveloped. Existing tools often focus on data engineering pipelines, with limited integration into model behavior and ethical analysis. To address this gap, this paper proposes an **AI-driven lineage framework** that fuses technical lineage with ethical oversight. We examine the state of lineage in AI, explore its alignment with responsible AI principles, and provide a methodology for implementation that supports governance, auditability, and trust.

II. LITERATURE REVIEW

The foundation of data lineage lies in the field of data warehousing and database systems, where the need to trace data transformations and dependencies first emerged. **Moreau et al. (2011)** formalized provenance standards, while lineage tools such as **Apache Atlas**, **OpenLineage**, and **DataHub** emerged to support metadata collection in complex systems.

In the AI context, lineage must go further. According to **Gebru et al. (2018)** and **Holland et al. (2018)**, understanding data origins and transformations is essential for fairness and transparency. Their proposals, such as **datasheets for datasets** and **dataset nutrition labels**, laid the ethical groundwork for what lineage systems could support. However, these approaches are often static and lack automation, making them less practical for dynamic AI systems.

Schelter et al. (2018) and **Koshy et al. (2022)** argue that lineage must also include model-level artifacts: hyperparameters, training environments, and even interpretability tools. **NIST's AI RMF (2023)** and the **EU AI Act** emphasize traceability as a requirement for high-risk AI, showing growing regulatory support for lineage tracking.

Despite these developments, there remains a gap between ethical aspirations and technical execution. Most tools don't track lineage across the full AI lifecycle or align with governance standards. Our work addresses this by offering a comprehensive, AI-driven lineage methodology that links technical artifacts with ethical intent.



TABLE: Evaluation of AI Lineage Tools and Capabilities

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

Evaluation of AI Lineage Tools and Capabilities

Tool Platform	/ Lineage Granularity	ML-Specific Features	Integration (ML/Data)	Visualization	Versioning	Compliance Support	Notes
MLflow	Experiment- & model-level	Experiments, runs, models	Python, Spark, REST APIs	Basic (UI)	Models, runs	Partial	Strong for model tracking, limited full data lineage
Weights Biases	& Model training and tuning	Training, artifacts, hyperparameters	Pytorch, TensorFlow, etc.	Rich dashboards	Artifacts, runs	Partial	Visualization-rich, experiment lineage focus
Kubeflow Pipelines	Pipeline- and step-level	ML pipeline components, metadata	Native for K8s, TensorFlow	Interactive pipeline graphs	Via metadata	Pipeline auditability	Strong for reproducible workflows, Kubernetes-based
Apache Atlas	Table/column-level (data-focused)	△Basic support extensions	ML via Hive, Spark, HDFS, Kafka	Metadata UI	Via integrations	GDPR, governance features	Strong for enterprise data lineage, extensible
OpenLineage	Job-dataset-level and	ML support via orchestration (Airflow, dbt)	Airflow, dbt, Spark	Graph lineage view	& External tools	Partial	Open spec, ecosystem growing, good for pipeline metadata
DataHub	Dataset, schema, pipeline-level	ML metadata support growing	Kafka, Airflow, dbt, ML models	UI + graphs	Metadata snapshots	Role-based access, audit logs	General-purpose metadata engine, extensible
Neptune.ai	Experiment-level	Model training logs, artifacts	Python, Keras, XGBoost, etc.	Dashboards & traces	Experiment versions	⚠ Partial	Lightweight ML experiment tracking, great for teams
Pachyderm	File- and container-level	Full versioning	data Docker-based pipelines	CLI + UI	Full data lineage	Immutable data tracking	Strong provenance support, built for reproducibility



Tool Platform	/ Lineage Granularity	ML-Specific Features	Integration (ML/Data)	Visualization	Versioning	Compliance Support	Notes
LakeFS	Object-store (S3)-level	Not ML-specific	Git-like data control for lakes	CLI integrations	+ Branches, commits	Audit logs, access control	Best for data versioning at scale, Git-like semantics

Key Capabilities Evaluated:

- **Lineage Depth:** What level of data/model history is captured (job, dataset, parameter, etc.)
- **ML Awareness:** Does the tool understand and support ML-specific elements (models, hyperparameters)?
- **Integration Scope:** How well does it integrate with existing tools (Airflow, TensorFlow, etc.)?
- **Visualization:** Can users easily explore lineage?
- **Versioning:** Tracks versions of datasets, models, configs.
- **Compliance / Governance:** Support for audit, roles, traceability.

Summary:

- **Best for End-to-End ML Pipelines:** Kubeflow Pipelines, MLflow
- **Best for Data Lineage & Governance:** Apache Atlas, DataHub, OpenLineage
- **Best for Full Provenance & Versioning:** Pachyderm, LakeFS
- **Best for Experiment Tracking:** Weights & Biases, Neptune.ai

III. METHODOLOGY

To realize the benefits of AI-driven lineage, we propose a five-component methodology that integrates technical tracking with ethical oversight:

1. Lineage Instrumentation

Embed tracking capabilities into every stage of the AI pipeline—from data ingestion, feature selection, model training, to inference. Tools like **OpenLineage** or **Pachyderm** can automate this process.

2. Metadata and Versioning Layer

Use version control (e.g., DVC, Git) to record changes in data, code, and models. Include metadata for source attribution, transformation logs, and environment configurations.

3. Fairness and Bias Auditing

Incorporate fairness metrics and data distribution visualizations at each step. Flag potentially biased transformations or imbalanced training data during lineage capture.

4. Cryptographic Verification


Use hash functions or blockchain to make lineage records tamper-evident and verifiable, particularly important for audit and legal defense.

5. Governance Integration

Connect lineage data with dashboards for regulators, ethics boards, and internal reviewers. Map lineage artifacts to compliance frameworks like **NIST RMF** or **EU AI Act** standards.



FIGURE: AI-Driven Lineage Framework



A layered AI pipeline visual:

- **Data Sources** → **ETL/Preprocessing** → **Feature Engineering** → **Model Training** → **Validation** → **Deployment** → **Monitoring**
- Overlaid with:
- **Lineage Capture Nodes** at each stage
- **Audit Trails** and **Bias Indicators**
- **Cryptographic Checks**
- **Governance Dashboard Feed**

IV. CONCLUSION

In the age of AI, where systems increasingly influence society's most important decisions, **lineage is no longer optional—it is foundational**. AI-driven lineage systems offer a powerful approach to ensure transparency, fairness, and accountability by tracking the complete journey of data and models through the AI lifecycle. This paper has demonstrated that lineage is more than a technical feature—it is a strategic imperative. When aligned with ethical frameworks and regulatory standards, it becomes a tool for **building trust** in AI systems. From identifying sources of



bias to verifying compliance, AI-driven lineage bridges the gap between **ethical AI theory and operational reality**. While current tools offer partial solutions, our proposed framework pushes for full integration—combining data and model lineage, real-time tracking, governance mapping, and cryptographic verification. This holistic approach supports a future where AI is not only intelligent but **understandable, explainable, and just**. As regulatory landscapes evolve and public expectations grow, organizations that prioritize AI-driven lineage today will be best positioned to deliver **responsible and robust AI systems** tomorrow.

REFERENCES

1. Moreau, L., et al. The Open Provenance Model core specification. *Future Generation Computer Systems*, 27(6), 743–756.
2. Gebru, T., et al. Datasheets for datasets. *arXiv:1803.09010*.
3. Holland, S., et al. The dataset nutrition label. *arXiv:1805.03677*.
4. NIST. AI Risk Management Framework 1.0.
5. Apache Atlas. <https://atlas.apache.org>
6. OpenLineage. <https://openlineage.io>
7. DataHub. (<https://datahubproject.io>)
8. Schelter, S., et al. Automatically tracking metadata and provenance of machine learning experiments. *Data Engineering Bulletin*.
9. Pachyderm. <https://www.pachyderm.io>
10. MLflow. <https://mlflow.org>
11. EU Commission. EU Artificial Intelligence Act.
12. Comet ML. <https://www.comet.com>
13. Koshy, R., et al. Data governance and lineage in regulated AI systems. *Journal of Data and Information Quality*, 14(3).
14. Mittelstadt, B. D., et al. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2).
15. Microsoft. Responsible AI Standard.