



TECHNOLOGY FEATURE

Research Communication / Data Science / Technology Report

Implementing AI in Decision making in the Age of Big Data: Evolution, Architecture and Challenges

AI decision systems are built today from batch processing statistical models to pipelines that optimize themselves in real-time and what still lies between them and reliable and trustworthy autonomy.

By **Muhammad Azeem Sabir** | Department of Computer Science | Published: October 2026

✉ 2023ag9978@uaf.edu.pk | Software Engineer

Introduction

All followed-by decisions – whether it's a credit approval, a treatment recommendation, a reroute in logistics, or a flag for fraud – are a result of two factors: the quality of data and the quality of the reasoning applied to the data. Both were limited by the human cognitive bandwidth for most of the history of humankind. There was no way to obtain and use data at speeds faster than the speed of human effort. The only way to reason was at the human rate. With the blending of AI and Big Data infrastructure, both limitations have been broken down.

Organizations today regularly deal with Petabytes of data and process it at milliseconds latency and feed decision systems that work without fatigue and drift. It's just 30 years since the early rule-based decision models to the current self-optimizing AI pipelines. It's only been about 30 years since the early days of rule-based decisions, to today's self-optimizing AI pipelines. Learning about the technical architecture of this journey, as well as what hasn't been solved, is crucial to any practitioner who is developing intelligent systems at scale.

“By 2026, the completely independent AI-driven data ecosystems are being created - self-healing, self-optimizing pipelines making more decisions than any human team could ever.”

Technology Overview: The Evolution of AI Decision Systems

The first automated decision systems were deterministic, meaning that they contained explicit rules, programmed by humans. If credit score is above threshold X, and debt-to-income ratio is below threshold Y, then accept the loan. The systems were easy to understand and track but were brittle; they could not handle data patterns that weren't expected or data sets that were larger than the ones that their rule makers had considered. The first wave of statistical machine learning built on logistic regression, decision trees, and support vector machines, focused on brittleness, meaning that models learned decision boundaries based on historical data, with decisions made based on yesterday's data and models trained weekly or monthly.

With the introduction of distributed computing frameworks, like Hadoop and later on Spark, the scope of what could be processed changed. Numbers that were previously considered processable and unusable were now usable. Numbers that were previously discarded as impossible to store and/or process became tractable. The size of the datasets on which gradient-boosted tree ensembles and early deep learning models could be trained would be larger than before. With the advent of deep learning — specifically convolutional networks, recurrent architectures, and in the end the transformers — feature engineering for high-dimensional data types such as



images, text, and time series was automated. At the same time, the streaming architectures like Apache Kafka and the new Apache Flink allowed to continuously ingest and compute model inference at production scale, filling for the first time the feedback loop of world state and decision output.

Today is the age of foundation models and agentic architectures. For the first time, LLM has been used to build generalized decision capability that works across domains without retraining on specific tasks using data collected from internet. Agentic systems brought AI from reactive inferences to proactive, multi-step actions: AI systems that are not just predictive but also planned and executed over a series of steps to achieve a goal. Today's AI decision pipeline is part of a decision system in its own right that can detect data quality problems on its own, redirect around data quality failures, and retrain components when it detects statistical drift.




Figure 1: The evolution of AI decision architectures — from rule-based expert systems and batch ML models through real-time streaming pipelines to autonomous agentic decision systems.

Modern AI Decision Pipeline – Architecture in Depth

In 2026, the production of AI decision system will not be a single model, it will be a network of models. The data ingestion layer is located at the bottom: Raw data comes in from a variety of sources and is heterogeneous, with varying latencies and quality. There are platforms like event streaming (Kafka, Pulsar) and batch ETL pipelines (ETL) for handling real-time feeds and historical data, respectively. It validates data quality, enforces schema and performs deduplication prior to any further downstream processing in the model layer.

The feature store is on top of ingestion: A centralized repository of features that are pre-computed and stored in a versioned format and that models use during inference. Feature stores address one of the most common engineering challenges: making sure a model doesn't see a skewed set of features in training compared to what it encounters in production. This eliminates one of the most common sources of performance degradation in production models: training-serving skew. The model layer sits on top of the trained algorithms that deliver decisions: tabular data—gradient-boosted trees, text/sequence data—transformer models, and relational data—



graph neural networks. Monte Carlo simulations capture a distribution of possible outcomes, and Markov decision processes deal with optimizing a sequence of decisions in an uncertain environment.

The decision layer converts the model output into decisions, takes into account business rules, confidence scores, and regulatory requirements. On top of this is the monitoring and governance layer: always checking model performance and data drift, fairness scores and audit logs, etc. With the advent of modern MLOps platforms, the feedback loop of retraining can become automated, allowing organizations to prevent drift from affecting business outcomes and move beyond firefighting to optimizing. The entire stack is increasingly referred to as an AI tech stack, with data ingestion and storage, model layer, retrieval and context through the use of RAG, orchestration and agents, integration through APIs, and a governance layer on security, monitoring, compliance and cost control.

Key Findings & Impact

AI-powered decision systems can make a significant measurable impact and it is increasing rapidly. By end 2026, Gartner expects 40% of enterprise apps to feature task-specific AI agents; potentially AI driven automation could account for up to 30% of enterprise application software revenue by 2035. Major financial institutions had already implemented AI factory architectures, which used technology platforms, pre-built algorithms, and data pipelines controlled by rules, several years before their rivals, and these had led to better credit decision making and fraud prevention rates and decision-making timeliness. In industries where milliseconds make the difference, real-time data lake architectures and ETL pipelines are now used to process data for real-time decisions.

Technically, too, the difficulties are great. Explainability is the inextricable paradox: powerful predictive models based on deep learning are often models that are not easily understood or explained by humans. The EU AI Act also includes provisions for transparency and human oversight for high-risk AI applications, which are directly opposed to the opacity of the most powerful models. The twin challenge is data quality and governance: AI decision systems are only as reliable as the pipelines that supply them, and organizations with underdeveloped data processes simply won't reap the benefits of the more advanced modelling. The issue of bias amplification, in which historical trends in training data embed and are passed on discriminatory outcomes, is an ongoing research and regulatory issue with a significant impact in areas such as lending, hiring, and criminal justice. In regulated fields, the dichotomy between most powerful and least explainable AI decision models is not an abstract, theoretical concept, it's a real barrier to deployment.

“The most capable AI decision models are often the least explainable — and in regulated domains, that tension is not theoretical. It is a deployment blocker.”

What's Next?

There will be three trends that will set the course for the next generation of AI-based decision-making systems. Firstly, there is a need for post-transformer architecture to address the problem of high computational costs and poor generalization capabilities of present-day large models. According to specialists, the direction will be towards smaller models enhanced by a particular domain knowledge base, which will allow developing decision-making algorithms in different sectors such as healthcare, manufacturing, law, and finance instead of creating one generic model capable of handling all tasks. Secondly, there will be the implementation of self-healing data pipelines, i.e., mechanisms for autonomous detection of potential errors in data quality, schema drift, and tracing data lineage without human involvement. Thirdly, explainable AI approaches will gain traction as a result of growing regulatory demands.



Organizations that will excel in using AI to make decisions are not those organizations that have the greatest volume of data or the most advanced AI algorithms. These organizations are those who lay the groundwork by building robust data pipelines, having policies regarding the need for human intervention, and having capabilities to evaluate, interpret, and continuously improve their decision systems. In the age of Big Data, intelligence is plentiful. Trust is still the scarcest of all resources.

About the Author

Muhammad Azeem Sabir a Software Engineering student at University of Agriculture Faisalabad. With keen interest in Digital Innovation Through Artificial Intelligence. He can be reached at 2023ag9978@uaf.edu.pk

Keywords: Artificial Intelligence; Big Data; Decision-Making; Machine Learning Pipelines; Data Architecture; Explainable AI; Real-Time Systems

References

- [1] Davenport, T.H. and Bean, R. (2026) Five trends in AI and data science for 2026, *MIT Sloan Management Review*. Available at: sloanreview.mit.edu.
- [2] IBM Institute for Business Value (2026) The trends that will shape AI and tech in 2026. Available at: ibm.com/think.
- [3] Gartner (2026) Top strategic technology trends for 2026. Available at: gartner.com.
- [4] Kellton (2026) AI tech stack 2026: Frameworks, MLOps and IDEs guide. Available at: kellton.com.
- [5] Vaswani, A. et al. (2017) Attention is all you need, *Advances in Neural Information Processing Systems*, 30, pp. 5998–6008.