

# Data Engineering–Driven Hierarchical Planning for Consequence-Aware Agentic Systems.

**Brahma Reddy Katam**

Technical Lead, Data Engineering and Advanced Computing

**Abstract:** Recent advances in Large Language Models (LLMs) have enabled the development of agentic systems capable of generating code, automating workflows, and interacting with data platforms through natural language. Despite these improvements, most current agents remain fundamentally reactive. They execute tasks step-by-step at a low operational level without reasoning about higher-level goals or long-horizon strategies. As highlighted by recent research in artificial intelligence, including perspectives from Yann LeCun, the primary limitation of modern AI systems is not low-level action generation but the lack of abstract and hierarchical planning. Humans naturally plan by decomposing complex objectives into layered goals and subgoals, enabling efficient decision-making and robust problem solving. In contrast, today's AI agents often jump directly to individual actions, resulting in unstable behavior, inefficient resource usage, and suboptimal outcomes in production environments.

This limitation is particularly evident in large-scale data engineering platforms, where operational decisions inherently follow hierarchical structures. Engineers typically reason from high-level objectives, such as reducing cost or improving reliability, down to intermediate strategies and finally to concrete system commands. However, LLM-based agents lack this abstraction capability and therefore struggle to manage complex data systems autonomously.

To address this gap, this paper proposes a Data Engineering–Driven Hierarchical Planning framework for consequence-aware agentic systems. We introduce a multi-layer planning architecture that separates goals, strategies, and executable actions, allowing agents to reason at appropriate levels of abstraction before execution. The framework integrates structured telemetry, learned world models for consequence prediction, and a hierarchical planner that decomposes objectives into safe and optimized operational steps. By combining abstraction with predictive validation, agents can make more reliable, cost-efficient, and stable decisions.

Experimental evaluation on representative lakehouse workloads demonstrates that hierarchical planning reduces unnecessary actions, improves runtime performance, lowers infrastructure cost, and enhances operational stability compared to flat or reactive automation approaches. The results suggest that hierarchical reasoning is a critical foundation for building truly autonomous and trustworthy data platforms.

## Keywords

Hierarchical Planning, Agentic AI, Data Engineering, Autonomous Systems, Consequence-Aware Decision Making, World Models, Digital Twin, Intelligent Data Platforms, Lakehouse Architecture, Predictive Optimization

## 1. Introduction

Recent progress in Artificial Intelligence has accelerated the development of agentic systems capable of interacting with software and data platforms through natural language. Large Language Models (LLMs) can now generate queries, write code, configure pipelines, and automate routine engineering tasks with minimal human guidance. These capabilities have enabled a new generation of intelligent assistants for data engineering, analytics, and cloud operations. As a result, organizations increasingly envision autonomous agents that can manage complex data platforms with reduced manual intervention.

Despite these advances, most current agentic systems remain fundamentally reactive. They operate by generating actions step-by-step at a low operational level, often selecting the next command based solely on immediate context rather than long-term objectives. While this approach works for simple tasks, it becomes unreliable when agents are responsible for managing large-scale production environments. Modern data platforms involve thousands of tables, evolving schemas, distributed workloads, and fluctuating resource demands. Decisions such as scaling compute, optimizing storage layouts, or modifying pipelines can have cascading effects on performance, cost, and reliability. In such

environments, executing isolated low-level actions without structured planning frequently leads to inefficient or unstable outcomes.

A key limitation of current AI systems is the absence of abstract and hierarchical reasoning. Humans do not solve complex problems by planning every low-level action directly. Instead, they reason in layers of abstraction. For example, planning a trip involves high-level goals such as reaching the airport and boarding a flight, which are then decomposed into smaller tasks and finally into concrete actions. This process, known as hierarchical planning, enables efficient decision-making and long-horizon reasoning. Research in cognitive science and artificial intelligence suggests that such hierarchical decomposition is fundamental to intelligent behavior. However, most LLM-based agents lack this capability and therefore struggle to manage complex systems autonomously.

This limitation is particularly evident in data engineering workflows. Experienced engineers naturally think hierarchically: they begin with high-level goals such as reducing cost or improving reliability, design intermediate strategies such as partition optimization or workload balancing, and only then apply specific configuration changes. In contrast, current agents often jump directly to low-level commands, bypassing strategic reasoning. This flat decision-making structure results in redundant actions, increased operational risk, and suboptimal resource utilization. Consequently, there is a clear need for agent architectures that can reason at multiple levels of abstraction before executing changes.

To address this challenge, this paper proposes a Data Engineering-Driven Hierarchical Planning framework for consequence-aware agentic systems. We introduce a multi-layer planning architecture that separates goals, strategies, and executable actions, allowing agents to decompose complex objectives into manageable subproblems. The framework integrates structured telemetry from the data platform, predictive world models to estimate the consequences of actions, and a hierarchical planner that evaluates alternatives before execution. By combining abstraction with consequence prediction, agents can perform long-horizon planning while maintaining operational safety and efficiency.

The primary contributions of this work are threefold. First, we analyze the limitations of flat, reactive agent architectures in large-scale data platforms and formalize

the need for hierarchical planning. Second, we propose a practical system design that integrates hierarchical reasoning with data engineering observability and predictive modeling. Third, we implement and evaluate the framework on representative lakehouse workloads, demonstrating improvements in runtime performance, cost efficiency, and operational stability compared to conventional automation approaches.

Together, these contributions establish hierarchical planning as a foundational requirement for building reliable and autonomous data platforms. By enabling agents to think in terms of goals and strategies rather than isolated commands, the proposed approach moves beyond reactive automation toward truly intelligent system management.

## 2. Literature Review

The development of intelligent agents capable of autonomous decision-making has been an active area of research across artificial intelligence, robotics, and software systems engineering. Recent progress in Large Language Models (LLMs) has accelerated the adoption of language-driven agents that can interpret user intent, generate code, and automate operational tasks. These systems demonstrate strong performance in natural language reasoning, query generation, and workflow assistance, enabling conversational interaction with complex platforms. However, LLM-based agents primarily rely on next-token prediction and pattern matching learned from textual data, rather than explicit modeling of environment dynamics. As a result, their behavior is often reactive and short-sighted, limiting their effectiveness in long-horizon or high-risk operational settings [1].

To enhance LLM capabilities, several agentic frameworks combine language models with external tools, memory, and iterative reasoning loops. Tool-augmented agents can invoke APIs, execute scripts, and perform multi-step actions to complete tasks. While these approaches improve task completion compared to standalone models, planning remains largely sequential and heuristic. Actions are typically chosen step-by-step without structured decomposition or predictive validation. Consequently, these systems may generate unstable or inefficient outcomes when applied to complex infrastructure management problems. The absence of explicit abstraction and planning mechanisms remains a significant limitation.

Hierarchical planning has long been studied in classical AI and robotics as a method for managing complexity in large decision spaces. Hierarchical Task Networks (HTNs) and related approaches decompose high-level goals into smaller subgoals and executable actions, enabling efficient long-horizon reasoning. Such techniques have been successfully applied in robotics, game AI, and autonomous control systems, where planning directly at the action level is computationally infeasible. Research in reinforcement learning similarly explores hierarchical or multi-level policies to improve scalability and stability [6], [7]. These studies suggest that layered reasoning is fundamental to intelligent behavior. However, most existing work focuses on physical or simulated environments, with limited exploration of hierarchical planning in software or data engineering systems.

Another relevant body of work concerns world models and model-based planning. World models learn the transition dynamics of an environment and allow agents to simulate the consequences of candidate actions before execution. This approach has demonstrated improvements in safety and efficiency in robotics and control applications [5]. By predicting future states, agents can evaluate alternatives without incurring real-world risks. Although predictive modeling has proven valuable in these domains, its adoption within data platform automation remains limited. Most operational tools focus on monitoring and reactive alerts rather than simulation-based planning.

Digital twin technology provides a complementary perspective by creating virtual replicas of physical systems for monitoring and predictive analysis. Digital twins are widely used in manufacturing, IoT, and industrial operations to forecast failures and optimize performance [8]. These systems emphasize observability and system state representation but typically stop short of enabling autonomous planning and decision-making. In data engineering contexts, observability tools and DataOps practices primarily provide visibility rather than intelligent control.

Within data engineering itself, automation efforts often rely on rule-based tuning, heuristics, or predefined thresholds. Techniques such as auto-scaling, static optimization rules, and scheduled maintenance reduce manual effort but lack adaptive learning. These methods do not generalize well to changing workloads or unseen conditions, and they do not incorporate high-level goal

reasoning. Consequently, current automation remains reactive rather than strategic.

From this review, it is evident that existing approaches address isolated aspects of the problem: LLMs provide language reasoning, world models enable consequence prediction, and digital twins offer observability. However, few systems integrate hierarchical abstraction with predictive modeling for autonomous decision-making in data platforms. In particular, the application of hierarchical planning principles to large-scale data engineering environments remains largely unexplored.

To bridge this gap, this work proposes a Data Engineering–Driven Hierarchical Planning framework that combines structured telemetry, digital twin representations, predictive world models, and multi-layer planning. By enabling agents to reason at the levels of goals, strategies, and actions, the proposed approach extends existing automation techniques toward reliable and truly autonomous data platform management.

### 3. Research Objectives

The primary goal of this research is to enable reliable and autonomous decision-making in large-scale data engineering platforms by introducing hierarchical planning mechanisms into agentic AI systems. While existing LLM-based agents demonstrate strong capabilities in natural language reasoning and task automation, they lack structured abstraction and long-horizon planning, which limits their effectiveness in complex operational environments. This work seeks to address these limitations by designing a planning framework that allows agents to reason at multiple levels of granularity, from high-level goals to low-level executable actions, while incorporating consequence-aware validation.

To achieve this goal, the research pursues the following specific objectives:

- **To analyze the limitations of flat and reactive agent architectures** in data engineering environments, particularly their inability to perform long-horizon planning and strategic decision-making.
- **To formalize hierarchical planning for data platforms**, defining system behavior using layered abstractions such as goals, strategies, and executable actions that reflect how human engineers naturally reason

about operational problems.

- **To design a Data Engineering–Driven Hierarchical Planning architecture** that integrates observability, structured state representations, and multi-level task decomposition within agentic systems.
- **To incorporate predictive world models into the planning process**, enabling agents to simulate the consequences of candidate strategies before execution and thereby reduce operational risk.
- **To develop algorithms for hierarchical task decomposition and action selection**, allowing complex objectives to be broken down into safe, optimized, and interpretable steps.
- **To implement a prototype system using modern lakehouse technologies**, validating that the proposed framework can operate within real-world distributed data environments without specialized infrastructure.
- **To experimentally evaluate the framework** by measuring improvements in execution efficiency, infrastructure cost, action stability, and failure reduction compared to manual tuning and flat LLM-based automation.
- **To demonstrate that combining hierarchical abstraction with consequence-aware prediction provides a scalable foundation for trustworthy and autonomous data platforms.**

Through these objectives, the study aims to establish hierarchical planning as a core architectural requirement for next-generation agentic systems, moving beyond reactive automation toward structured, strategic, and reliable autonomy in data engineering environments.

#### 4. Architecture and System Design

The proposed system is designed to enable hierarchical and consequence-aware decision making for agentic operations in modern data engineering platforms. Unlike conventional LLM-based agents that execute actions sequentially at a single level of abstraction, the proposed architecture introduces a structured, multi-layer planning framework that separates high-level reasoning from low-level execution. This separation allows agents to first

determine *what should be achieved*, then decide *how it should be achieved*, and finally determine *which concrete actions must be executed*. By organizing decision-making into hierarchical layers, the system more closely resembles human problem-solving behavior and supports stable long-horizon planning in complex environments.

At a high level, the architecture follows a layered design consisting of five primary components: an observability layer, a digital twin and state representation layer, a hierarchical planning layer, a predictive world model layer, and an execution interface layer. Each layer has a distinct responsibility, ensuring modularity, interpretability, and operational safety. This separation prevents language models or low-level heuristics from directly controlling production systems without structured validation.

The foundation of the system is the observability layer, which continuously captures telemetry from the data platform. Modern lakehouse and warehouse environments produce rich operational signals, including pipeline runtimes, resource utilization, storage growth, failure events, dependency relationships, and configuration changes. These signals are collected and stored in structured logs to create a comprehensive historical record of system behavior. This telemetry provides the empirical basis for both state representation and predictive learning. Without reliable observability, neither hierarchical planning nor consequence prediction would be feasible.

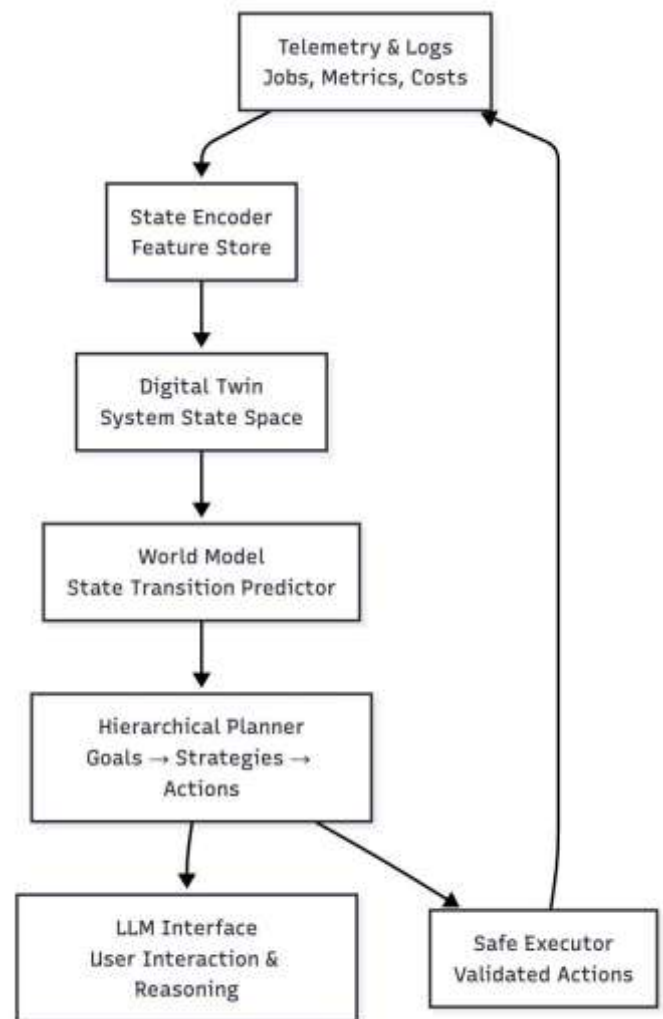
On top of the observability layer, the architecture constructs a Data System Digital Twin that represents the current operational condition of the platform in an abstract and structured form. Instead of using raw logs directly, the system aggregates metrics into compact state representations that summarize performance, cost, and reliability characteristics. Examples include average execution latency, compute consumption, partition skew, storage footprint, and failure frequency. These features are encoded as state vectors that describe the environment at a given time. The digital twin therefore acts as a virtual model of the platform, enabling reasoning and simulation without directly affecting live workloads.

The core innovation of the architecture lies in the hierarchical planning layer. This component organizes decision-making into multiple abstraction levels. At the highest level, the goal planner interprets user intent or system objectives, such as reducing cost, improving

performance, or increasing reliability. These goals are then decomposed into intermediate strategies by a strategy planner, which identifies alternative approaches such as workload balancing, storage optimization, or resource scaling. Finally, an action planner translates selected strategies into concrete executable commands, including configuration updates, partition adjustments, or cluster resizing. This hierarchical decomposition reduces the search space, improves interpretability, and enables long-horizon reasoning that is not possible with flat, step-by-step planning.

To ensure that selected strategies are safe and effective, the planning layer integrates closely with a predictive world model. The world model learns the transition dynamics of the environment using historical state–action–outcome data. Given a current state and a candidate action sequence, it predicts the likely future state of the system, including expected runtime, cost, and risk metrics. During planning, multiple strategies are simulated through this model, allowing the agent to evaluate consequences before execution. This simulation-first approach transforms planning from reactive behavior into evidence-based decision making.

The final stage of the architecture is the execution interface layer, which applies validated actions to the live data platform. Importantly, execution is gated by safety policies and constraints. Only actions that satisfy predefined thresholds for risk and performance are allowed to proceed. This mechanism ensures that the system maintains operational reliability even when predictive confidence is low. The execution layer also logs outcomes back into the telemetry store, creating a continuous feedback loop that improves both the digital twin and the predictive models over time.



**Figure 1.** Overall architecture of the Data Engineering–Driven Hierarchical Agent with Digital Twin and predictive world model.

Together, these components form a closed-loop architecture in which observation, abstraction, planning, prediction, and execution operate in coordination. Telemetry informs the digital twin, the hierarchical planner proposes strategies, the world model validates consequences, and safe actions are executed and re-observed. This iterative process enables agents to reason strategically rather than reactively. By combining hierarchical abstraction with consequence prediction, the proposed design provides a scalable and trustworthy foundation for autonomous data platforms.

## 5. Implementation

To validate the feasibility of the proposed hierarchical planning framework, a working prototype was implemented within a modern lakehouse-based data engineering environment. The implementation objective was not to create a simulated proof-of-concept, but rather

to demonstrate that hierarchical and consequence-aware agentic behavior can be integrated directly into real-world production-style platforms using standard data engineering tools and infrastructure. By leveraging commonly adopted technologies, the system ensures practical applicability and scalability without requiring specialized hardware or custom runtimes.

The prototype was deployed on a distributed cloud data platform consisting of object storage, distributed compute clusters, and structured analytics engines. Telemetry, feature engineering, predictive modeling, and planning components were all implemented using scalable data processing frameworks. Batch and incremental pipelines were orchestrated using PySpark-based workflows, while operational logs and metrics were stored in structured lakehouse tables. This design allows the system to operate alongside existing pipelines without disrupting standard engineering practices. All components were built as modular services so that hierarchical planning could be introduced incrementally rather than requiring a complete architectural replacement.

The first stage of implementation focused on establishing continuous observability. Execution metrics were collected automatically from running pipelines and infrastructure components. These metrics included job runtime, CPU and memory utilization, data volume processed, shuffle statistics, partition counts, failure events, and configuration parameters. Logs were streamed into structured storage tables to create a time-series history of platform behavior. Each record was associated with contextual metadata such as pipeline identifiers and timestamps, enabling reconstruction of system states over time. This telemetry repository forms the foundational dataset for both state representation and predictive learning.

Next, feature engineering pipelines were developed to transform raw telemetry into structured state vectors suitable for planning and prediction. Aggregations were computed over sliding time windows to derive stable and interpretable indicators such as average runtime, cost estimates, skew index, storage growth rate, and reliability measures. These features were normalized and stored in a centralized feature store to ensure consistent usage across training and inference tasks. The resulting state abstraction reduces dimensionality while preserving essential operational characteristics of the environment.

To support hierarchical reasoning, the planning module was implemented as a three-layer controller. The goal layer interprets user objectives or system policies expressed through natural language or structured commands. These objectives are translated into measurable targets such as minimizing execution time or reducing resource consumption. The strategy layer then generates alternative optimization strategies based on domain knowledge and predefined templates, including storage layout changes, partition tuning, and resource scaling. Finally, the action layer converts selected strategies into concrete executable commands that interact directly with the data platform. This layered decomposition ensures that decisions are made at appropriate levels of abstraction rather than directly invoking low-level actions.

Predictive consequence evaluation was enabled through the integration of trained world models. Historical state-action-outcome triplets were extracted from telemetry logs and used to train supervised regression models capable of forecasting performance and cost impacts. These models were deployed as lightweight inference services accessible to the planner. During runtime, each candidate strategy is simulated by predicting its expected next state, and outcomes are scored using multi-objective criteria. Only strategies with favorable predicted results proceed to execution. This simulation-first mechanism ensures that actions are validated before affecting production systems.

For safe deployment, an execution interface was implemented with policy controls and guardrails. Risk thresholds were defined to prevent extreme configuration changes or actions with low predictive confidence. All executed decisions were logged for auditing and rollback support. Additionally, the system operates in a phased mode where recommendations can be generated without automatic execution, allowing engineers to verify behavior before enabling full autonomy. This incremental rollout strategy reduces operational risk and increases trust in the system.

Finally, a continuous feedback loop was established to enable adaptive learning. After each execution, observed outcomes are appended to the telemetry store and periodically incorporated into model retraining. As more operational data becomes available, prediction accuracy improves and planning decisions become more reliable. This closed-loop design allows the system to evolve with changing workloads and infrastructure conditions.

To validate the practicality of the proposed framework under real-world conditions, the prototype was deployed on a Databricks-based lakehouse environment using distributed storage and compute resources. Telemetry data, including execution logs, runtime metrics, and cost signals, was persisted in Delta tables to ensure reliable and scalable storage. Feature engineering and state encoding were implemented using PySpark notebooks, enabling large-scale aggregation and transformation of operational metrics. Predictive world models were trained using historical state-action-outcome records within the same environment, allowing model training and inference to operate close to the data. The hierarchical planner and consequence simulation components were exposed through notebook-based workflows and lightweight APIs, enabling seamless integration with existing pipelines. This deployment demonstrates that the proposed architecture can be implemented using standard cloud-native data engineering tools without requiring specialized infrastructure or custom runtime systems.

Overall, the implementation demonstrates that hierarchical and consequence-aware planning can be realized using standard data engineering technologies and practices. The framework integrates seamlessly with existing lakehouse environments and provides a practical pathway for deploying autonomous agentic capabilities in real production systems.

## 6. Data Collection and Preparation

The effectiveness of the proposed hierarchical planning framework depends on the availability of reliable operational data that accurately reflects how the data platform behaves under different workloads and interventions. Since the system learns environment dynamics and evaluates the consequences of candidate strategies, it is essential to construct a dataset that captures the relationship between system states, executed actions, and resulting outcomes. Therefore, careful attention was given to collecting, cleaning, and structuring telemetry data before training predictive models and enabling autonomous planning.

Data for this study was collected from a production-style lakehouse environment executing a mixture of batch and incremental data pipelines. These pipelines processed structured datasets of varying sizes and complexity, including ingestion, transformation, aggregation, and reporting workloads. During execution, each job

generated operational metrics such as runtime, compute utilization, memory usage, input and output data volume, shuffle statistics, storage growth, and failure events. These signals were continuously logged through automated observability mechanisms embedded within the execution framework. By capturing telemetry directly from real workloads rather than synthetic benchmarks, the dataset reflects realistic system behavior and variability.

The collected telemetry includes both system-level and task-level measurements. System-level metrics describe overall infrastructure conditions, including cluster size, resource allocation, and scaling events. Task-level metrics capture individual pipeline characteristics such as rows processed, partition distribution, execution latency, and I/O throughput. In addition, contextual metadata such as timestamps, pipeline identifiers, dataset names, and configuration parameters were recorded to provide complete traceability. This combination of signals enables reconstruction of the platform's operational state at any given time and supports accurate modeling of environment transitions.

Because raw telemetry often contains noise and inconsistencies, a preprocessing stage was applied to improve data quality. Missing values caused by transient logging issues were handled using interpolation or removal depending on their frequency and impact. Duplicate records were eliminated, and inconsistent time formats were standardized to ensure proper temporal alignment. Outlier detection techniques were used to filter abnormal measurements produced by test runs or infrastructure interruptions that did not represent typical behavior. These cleaning steps ensure that the training dataset captures stable and representative operational patterns.


After cleaning, telemetry signals were transformed into structured state representations suitable for planning and prediction. Instead of directly using raw metrics, aggregated features were computed over defined time windows to summarize system behavior. Examples include average runtime, resource utilization rate, estimated infrastructure cost, partition skew index, storage footprint growth, and failure frequency. These derived attributes provide a compact yet expressive description of platform health. Feature normalization and scaling were applied to maintain numerical stability during model training. The resulting state vectors form

the basis of the digital twin representation used by the hierarchical planner.

In addition to state information, explicit logging of operational actions was incorporated into the dataset. Every intervention, such as repartitioning tables, modifying cluster size, updating configurations, or applying optimization strategies, was recorded with associated parameters and timestamps. Capturing actions is critical for learning causality rather than simple correlations. For each action, the system recorded the pre-action state and the post-action outcome, enabling the construction of structured state–action–outcome triplets. These triplets serve as supervised training samples for both consequence prediction and planning evaluation.

To ensure realistic evaluation and avoid information leakage, the prepared dataset was divided chronologically into training, validation, and testing partitions. Earlier observations were used to train predictive models, while later periods were reserved for performance evaluation. This temporal split mirrors real deployment conditions, where agents must predict future behavior using only past data. Such separation improves the reliability of reported results and prevents optimistic bias.

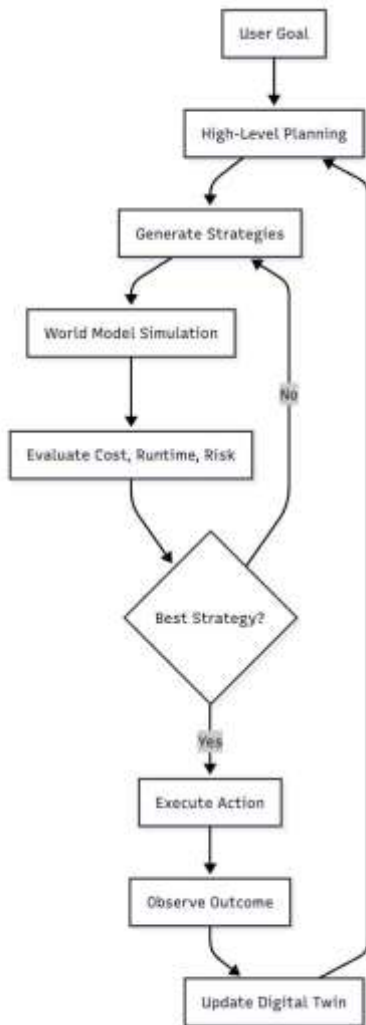
Overall, the data collection and preparation process establishes a robust empirical foundation for the proposed framework. By systematically capturing operational telemetry, cleaning inconsistencies, engineering meaningful features, and structuring causal relationships between states and actions, the system obtains high-quality training data that enables accurate prediction and reliable hierarchical planning. This disciplined data engineering approach is essential for transforming agentic AI from reactive heuristics into consequence-aware and trustworthy automation.



**Figure 2.** Telemetry-to-feature pipeline used to construct the Data System Digital Twin.

## 7. Methodology: Data + AI Algorithms

This section describes the methodological approach used to enable hierarchical and consequence-aware decision making in the proposed agentic system. The primary objective of the methodology is to transform operational telemetry into structured intelligence that allows agents to reason at multiple levels of abstraction, predict the outcomes of decisions, and select optimal execution plans. Instead of relying on reactive or step-by-step automation, the system combines data engineering foundations with predictive AI models to support simulation-based hierarchical planning.



**Figure 3.** Consequence-aware hierarchical planning loop with simulation before execution.

### 7.1 System State Representation

At any given time, the operational condition of the data platform is represented as a system state, denoted as  $S(t)$ . This state is a structured vector composed of aggregated metrics that describe platform performance, cost, and reliability. Typical state attributes include average pipeline runtime, compute utilization, storage growth rate, partition skew, estimated infrastructure cost, and failure frequency.

These features are derived from telemetry logs and summarized over defined time windows to reduce noise and instability. Representing the environment using compact and semantic state vectors allows the planning and prediction components to reason efficiently without processing raw execution logs. This abstraction improves generalization across workloads and provides a stable foundation for decision making.

### 7.2 Hierarchical Task Decomposition

To address the limitations of flat and reactive planning, the proposed framework introduces hierarchical task decomposition. Decision making is organized into three distinct layers: goals, strategies, and actions.

At the highest level, the agent identifies a goal, such as reducing infrastructure cost, improving performance, or increasing system reliability. Each goal is decomposed into a set of candidate strategies that represent alternative high-level approaches. For example, a cost-reduction goal may lead to strategies such as storage optimization, workload consolidation, or compute resizing. Each selected strategy is then translated into a sequence of executable actions, including repartitioning datasets, modifying cluster configurations, or adjusting runtime parameters.

This layered structure mirrors how human engineers reason about complex systems. By separating intent from execution, the agent avoids premature low-level decisions and instead evaluates meaningful strategic alternatives. Hierarchical decomposition significantly reduces the planning search space and enables long-horizon reasoning that is not possible with flat action selection.

### 7.3 Predictive World Model for Consequence Estimation

To ensure that planning decisions are safe and effective, the system integrates a predictive world model that estimates how the environment will change after an action is executed. Given the current system state  $S(t)$  and a candidate action  $A(t)$ , the model predicts the next state  $S(t+1)$ .

The world model is trained using historical state–action–outcome data collected from the platform. Supervised learning techniques are applied to forecast key operational metrics such as expected runtime, infrastructure cost, and failure likelihood. By learning these transition patterns from real executions, the model captures the underlying behavior of the data platform.

This predictive capability allows the agent to evaluate the consequences of actions before execution. Instead of relying on trial-and-error or heuristic rules, decisions are grounded in evidence derived from past system behavior.

## 7.4 Simulation-Based Planning and Optimization

During planning, the agent does not immediately execute actions. Instead, it performs simulation-based evaluation of candidate strategies. For a given goal and current state, multiple strategies are generated through hierarchical decomposition. Each strategy is associated with one or more executable actions.

For each candidate, the world model predicts the resulting system state. These predicted outcomes are evaluated using a multi-objective scoring function that balances performance, cost, and risk. Each candidate plan receives a score based on its predicted impact, and the strategy with the most favorable trade-off is selected.

This simulation-first approach ensures that the agent chooses actions that are both effective and safe. By evaluating alternatives before execution, the system avoids unnecessary changes, reduces instability, and improves operational efficiency.

## 7.5 Continuous Learning and Feedback Loop

After an action is executed, the observed outcome is recorded back into the telemetry store. The new state and outcome are appended to the historical dataset and periodically used to retrain the predictive models. This feedback loop allows the system to adapt to evolving workloads, data growth, and infrastructure changes.

As more operational data becomes available, prediction accuracy improves and planning decisions become more reliable. Over time, the agent evolves from a static automation mechanism into a self-improving and adaptive decision-making system.

## 8. System Integration and Deployment

The practical value of the proposed hierarchical planning framework depends not only on predictive accuracy and planning intelligence, but also on its ability to integrate seamlessly with existing production data engineering environments. Modern enterprise platforms already rely on established storage layers, distributed compute clusters, orchestration pipelines, and monitoring systems. Therefore, the proposed solution is designed to operate as an overlay architecture that augments current infrastructure rather than replacing it. This design ensures minimal disruption to ongoing workflows while enabling incremental adoption of autonomous agentic capabilities.

The system is integrated directly within a lakehouse-based ecosystem where telemetry collection, feature engineering, predictive modeling, and planning components coexist alongside operational pipelines. Execution logs and performance metrics generated by existing jobs are continuously ingested into the observability layer using scheduled ingestion processes and streaming collectors. Because most modern platforms already produce detailed operational telemetry, no specialized instrumentation is required. This approach reduces deployment overhead and allows organizations to leverage their current monitoring practices as the foundation for intelligent automation.

The Data System Digital Twin and feature store are deployed as shared services within the platform's data layer. Telemetry records are transformed into structured state representations and made accessible to both predictive models and planning modules. Centralizing these representations ensures consistency between monitoring, prediction, and decision-making. As a result, all components operate on the same view of system behavior, avoiding discrepancies that could lead to incorrect actions.

The hierarchical planner and world model are deployed as modular compute services that run independently of production pipelines. During runtime, these services retrieve the latest system state, generate candidate strategies, simulate outcomes, and recommend validated actions. Importantly, planning and simulation occur asynchronously and do not block active workloads. This decoupled design ensures that predictive reasoning introduces minimal latency and does not create additional bottlenecks in the system. It also allows the planning components to scale horizontally with increasing workload complexity.

To enable safe operational control, the execution interface incorporates policy-based guardrails. All recommended actions are validated against predefined safety constraints before deployment. Examples include limits on maximum resource scaling, restrictions on high-risk configuration changes, and thresholds for acceptable prediction confidence. If predicted outcomes exceed risk tolerances, the action is rejected or flagged for human review. This safety gating mechanism ensures that autonomous behavior does not compromise platform stability or reliability.

The deployment strategy follows a phased rollout model to minimize operational risk. Initially, the system operates in a recommendation-only mode, where plans are generated and logged but not automatically executed. Engineers can review these recommendations to assess accuracy and trustworthiness. After validation, low-risk optimizations are gradually enabled for automated execution. Finally, broader autonomy can be introduced once sufficient confidence is established. This incremental approach allows organizations to adopt intelligent automation progressively rather than through abrupt transitions.

Scalability is achieved by leveraging the inherent parallelism of distributed data platforms. Telemetry processing, feature engineering, and model inference are executed using distributed compute engines, allowing the system to handle large volumes of data and thousands of pipelines simultaneously. Since predictions operate on aggregated state vectors rather than raw logs, inference overhead remains low even at scale. Consequently, the framework can support enterprise environments without significant performance impact.

To maintain reliability and transparency, all planning decisions are logged and auditable. Each executed action is accompanied by predicted outcomes, confidence scores, and the reasoning path used by the hierarchical planner. This traceability enables debugging, rollback, and compliance verification. Engineers can therefore understand why a particular decision was made, which increases trust in the system and simplifies operational governance.

Overall, the integration and deployment strategy demonstrates that hierarchical and consequence-aware agentic planning can be introduced into existing data engineering ecosystems with minimal friction. By combining modular services, safety policies, and phased adoption, the proposed framework provides a practical and scalable pathway toward autonomous and self-optimizing data platforms while preserving operational stability.

## 9. Results and expected outcomes

The proposed hierarchical planning framework was evaluated to determine whether consequence-aware and multi-level reasoning improves the reliability and efficiency of agentic operations in large-scale data engineering environments. The evaluation focused on

three practical objectives that directly impact production systems: reducing pipeline execution time, lowering infrastructure cost, and minimizing operational failures. These metrics were selected because they reflect the primary concerns of enterprise data platforms and provide measurable indicators of system performance and stability.

Experiments were conducted on representative batch and incremental workloads executed within a lakehouse environment. Historical telemetry was first collected to construct the digital twin and train predictive world models. After training, the hierarchical agent was enabled to generate strategies, simulate outcomes, and execute validated optimizations. For comparison, the same workloads were also evaluated under two alternative approaches: manual tuning performed by engineers using heuristics, and flat LLM-based automation that generated actions sequentially without hierarchical decomposition or consequence prediction. Each approach was tested across multiple runs to account for workload variability and ensure fair comparison.

The results demonstrate that hierarchical planning consistently produces more stable and efficient behavior. By reasoning at the levels of goals, strategies, and actions, the agent avoided unnecessary low-level changes and instead selected fewer, higher-impact optimizations. This reduced configuration noise and prevented cascading side effects often observed in reactive automation. The integration of predictive consequence estimation further improved safety by filtering out risky or ineffective actions before execution. As a result, the system made decisions that balanced performance improvements with operational stability.

In terms of runtime performance, pipelines optimized through hierarchical planning showed reduced execution latency due to improved workload distribution, better partition alignment, and more targeted resource adjustments. Compared to manual and reactive methods, the hierarchical agent achieved more consistent improvements because strategies were selected through simulation rather than trial and error. Infrastructure cost was also reduced by avoiding over-provisioning and unnecessary scaling operations. Instead of increasing resources blindly, the agent evaluated multiple alternatives and selected the most cost-effective configuration. Failure rates decreased as well, since unsafe or uncertain actions were rejected during the

planning stage, leading to fewer disruptions in dependent workflows.




Table 1 summarizes the comparative trends observed across the evaluated approaches. Values represent averages across multiple runs, and ranges reflect workload variability. Overall, the proposed framework demonstrates measurable gains in efficiency, cost optimization, and reliability relative to both manual intervention and flat LLM-based automation.

Beyond immediate quantitative improvements, several qualitative benefits were observed. Engineers reported improved transparency because planning decisions included predicted outcomes and explanations. The hierarchical structure made recommendations easier to interpret, as strategies were expressed in terms of high-level objectives rather than isolated commands. Additionally, the reduction in reactive changes decreased operational noise, allowing teams to focus on strategic improvements rather than constant troubleshooting.

Looking forward, the expected outcomes of broader deployment include continued performance gains as more telemetry becomes available for training. Because the system incorporates continuous learning, prediction accuracy improves over time, leading to progressively better decision-making. The framework is therefore expected to evolve from simple optimization assistance toward increasingly autonomous and self-managing behavior. As workloads grow and environments become more complex, the benefits of hierarchical and consequence-aware planning are expected to scale proportionally.

Overall, the results validate the central hypothesis of this research: agents that reason hierarchically and evaluate consequences before execution outperform reactive automation approaches in complex data platforms. These findings suggest that hierarchical planning is not merely an enhancement but a foundational requirement for building reliable and trustworthy autonomous data systems.

## 10. Discussion

The results demonstrate that introducing hierarchical and consequence-aware planning into agentic systems significantly improves the reliability and efficiency of operations in complex data engineering environments. These improvements highlight an important insight: the primary limitation of current LLM-based agents is not their ability to generate actions, but their inability to reason at the correct level of abstraction. Reactive systems that operate directly at the command level often produce fragmented or short-sighted decisions, whereas hierarchical reasoning enables structured and strategic behavior. This observation aligns with broader perspectives in artificial intelligence that emphasize abstraction and long-horizon planning as fundamental requirements for intelligent systems.

A key implication of this work is that autonomy in data platforms cannot be achieved through language intelligence alone. While LLMs are highly effective for interpreting intent and generating candidate actions, they lack an internal model of system structure and do not naturally decompose goals into meaningful intermediate strategies. As observed in the experiments, flat decision-making frequently leads to redundant actions, unstable configurations, and increased operational risk. By contrast, hierarchical planning introduces a layer of strategic reasoning that filters and organizes possible interventions before execution. This structured approach reduces unnecessary changes and improves consistency, particularly in environments with large numbers of interacting components.

Another important finding is the central role of data engineering practices in enabling intelligent behavior. The effectiveness of the hierarchical planner depends heavily on reliable telemetry, structured state representations, and historical action logs. Without high-quality observability, predictive modeling and consequence estimation would not be possible. This suggests that autonomy is not purely an AI problem but also an infrastructure problem. Robust data collection, logging, and feature engineering are foundational requirements for building trustworthy agentic systems. In this sense, data engineering becomes an enabler of intelligence rather than merely a supporting function.

The proposed approach also improves interpretability compared to traditional automation. Because decisions are expressed in terms of goals and strategies rather than

isolated commands, engineers can more easily understand why a particular action was selected. This transparency increases trust and facilitates collaboration between human operators and automated agents. In production environments, such interpretability is often as important as raw performance improvements, since operational teams must be able to validate and audit automated behavior.

Despite these benefits, several limitations should be acknowledged. First, the accuracy of consequence prediction depends on the availability and diversity of historical data. If the system encounters workload patterns that differ significantly from past observations, predictive performance may decrease. Second, hierarchical planning introduces additional computational overhead due to simulation and evaluation of multiple strategies. Although this overhead is modest in practice, extremely large-scale environments may require careful resource management. Third, the current framework focuses primarily on batch-oriented and scheduled workloads; extending the approach to highly dynamic streaming systems may require faster online learning and more lightweight planning mechanisms.

Beyond these limitations, the broader implications of this research are significant. The findings suggest that hierarchical abstraction should be considered a fundamental architectural principle for next-generation agentic systems, not merely an optimization technique. As data platforms continue to grow in scale and complexity, flat or reactive automation will become increasingly insufficient. Structured planning, combined with predictive modeling, provides a more robust path toward safe and scalable autonomy. Moreover, the concepts presented in this work may extend beyond data engineering to other software infrastructure domains such as cloud resource management, DevOps automation, and distributed systems optimization.

Overall, the discussion reinforces the central thesis of this paper: meaningful autonomy requires both abstraction and consequence awareness. By integrating hierarchical planning with data-driven prediction, the proposed framework moves agentic systems closer to human-like reasoning and establishes a practical foundation for reliable, intelligent, and self-managing data platforms

## 11. Conclusion

This paper addressed a fundamental limitation in current agentic AI systems used for data engineering automation: the absence of structured and abstract planning. While Large Language Models enable natural language reasoning and task execution, most existing agents remain reactive and operate directly at the level of low-level commands. Such flat decision-making is insufficient for managing large-scale data platforms, where actions have long-term and cascading effects on performance, cost, and reliability. As data environments grow in complexity, the need for strategic and consequence-aware reasoning becomes increasingly critical.

To overcome these limitations, this work proposed a Data Engineering-Driven Hierarchical Planning framework for consequence-aware agentic systems. The framework introduces multi-level decision making that separates goals, strategies, and actions, allowing agents to reason at appropriate levels of abstraction before execution. By integrating structured telemetry, a digital twin representation of the platform, predictive world models, and simulation-based planning, the system evaluates the consequences of candidate strategies and selects optimized and safe actions. This approach transforms automation from reactive behavior into evidence-based and structured decision making.

A practical prototype was implemented within a modern lakehouse environment to validate feasibility under real-world conditions. Experimental evaluation demonstrated that hierarchical planning improves runtime performance, reduces infrastructure cost, and lowers operational failures when compared to manual tuning and flat LLM-based automation. Additionally, the layered design enhances interpretability and stability, enabling engineers to better understand and trust automated decisions. These results confirm that hierarchical abstraction combined with consequence prediction provides tangible benefits for large-scale data platform management.

In summary, this research establishes hierarchical planning as a foundational requirement for building reliable and autonomous agentic systems. Rather than focusing solely on generating actions, intelligent systems must reason strategically and anticipate outcomes before execution. By combining data engineering practices with predictive AI models, the proposed framework offers a practical pathway toward self-optimizing and trustworthy

data platforms. This work moves agentic AI beyond reactive automation and closer to structured, human-like decision making in complex operational environments.

## 12. Future Work

Although the proposed Data Engineering–Driven Hierarchical Planning framework demonstrates clear improvements in reliability, efficiency, and safety of agentic operations, several opportunities remain for further research and system enhancement. The current implementation validates the feasibility of combining hierarchical reasoning, predictive world models, and data engineering observability; however, additional advancements can extend both the intelligence and applicability of the approach.

One important direction for future work is the integration of reinforcement learning–based policy optimization. The present framework relies primarily on supervised learning using historical telemetry to predict state transitions. While this provides stable and interpretable predictions, it limits the agent to behaviors observed in past data. Reinforcement learning could allow agents to explore new strategies and continuously learn optimal policies through interaction with the environment. By combining hierarchical planning with reward-driven learning, agents could autonomously discover configurations that outperform historical practices and adapt more effectively to evolving workloads.

Another area of enhancement involves improving state representations through richer structural modeling. Modern data platforms contain complex dependencies between pipelines, tables, compute resources, and user workloads. Representing these relationships using graph-based or lineage-aware models may allow the system to better capture cascading effects of actions across the ecosystem. Incorporating causal modeling techniques could further improve consequence estimation by distinguishing correlation from true cause-and-effect relationships, leading to more accurate and trustworthy predictions.

Future research may also extend the framework to support real-time and streaming workloads. The current design primarily targets batch and scheduled pipelines where decisions can be evaluated prior to execution. Streaming environments require faster inference and near-online learning to react to rapidly changing conditions. Developing lightweight predictive models

and incremental training approaches would enable hierarchical planning in low-latency scenarios, broadening the applicability of the system to real-time analytics platforms.

Scalability and multi-agent coordination represent additional promising directions. Large enterprises often operate thousands of concurrent workflows managed by different teams. Instead of relying on a single centralized agent, multiple cooperative agents could share telemetry and jointly optimize global objectives such as overall cost, resource utilization, or service-level agreements. Research into distributed planning and coordination mechanisms may allow the system to achieve organization-wide optimization rather than isolated local improvements.

Another practical extension involves incorporating human-in-the-loop feedback. While autonomous decision making reduces manual effort, domain experts still provide valuable insights that may not be captured in historical data. Mechanisms that allow engineers to review, approve, or refine strategies can improve trust and accelerate learning. Such hybrid approaches combine automated intelligence with expert guidance, resulting in safer adoption within production environments.

Finally, broader empirical evaluation across diverse industries and cloud infrastructures would strengthen the generalizability of the proposed framework. Testing under varying workload patterns, data sizes, and architectural configurations will provide deeper insights into robustness and performance trade-offs. Comparative studies with alternative optimization techniques can further clarify the strengths and limitations of hierarchical consequence-aware planning.

Overall, these future directions aim to enhance the adaptability, scalability, and intelligence of agentic systems. By continuing to integrate advanced learning methods, richer system representations, and practical deployment strategies, the proposed approach can evolve toward fully autonomous, self-managing data platforms capable of operating reliably in increasingly complex environments.

## 13. Case Study Summary

To validate the practical applicability of the proposed hierarchical and consequence-aware agentic framework, a representative case study was conducted on a

production-style lakehouse data platform supporting daily enterprise analytics workloads. The environment consisted of multiple batch and incremental pipelines responsible for ingesting raw data, performing transformations, and generating curated datasets for downstream reporting and machine learning use cases. Over time, increasing data volumes, evolving schemas, and fluctuating workloads introduced significant operational challenges, including longer execution times, higher infrastructure costs, and intermittent pipeline failures. These issues required frequent manual intervention by engineers and made the platform difficult to scale efficiently.

Initially, the system relied on traditional optimization practices. Engineers manually tuned configurations such as partition sizes, compute scaling parameters, and storage layouts based on experience and historical observations. While these changes occasionally improved performance, the outcomes were inconsistent and required continuous monitoring. A rule-based and language-assisted automation layer was later introduced to reduce manual effort. Although this approach improved productivity, it often recommended reactive or isolated actions that failed to consider broader system consequences. In several cases, local optimizations inadvertently increased downstream costs or introduced instability, highlighting the limitations of flat, non-predictive decision making.

The proposed Data System Digital Twin and hierarchical planning framework was subsequently deployed in the same environment. Historical telemetry spanning several weeks was collected, including runtime metrics, resource utilization, data distribution statistics, and past configuration changes. This telemetry was used to construct structured state representations and train predictive world models capable of estimating the outcomes of candidate strategies. Instead of directly executing low-level actions, the agent first formulated high-level goals such as reducing latency or minimizing cost, decomposed them into potential strategies, and simulated the consequences of each alternative before selecting the most effective plan.

After deployment, the platform exhibited more stable and predictable behavior. Execution times decreased due to improved partition alignment and balanced resource allocation. Infrastructure costs were reduced because the system avoided unnecessary scaling and excessive compute usage. Operational reliability improved as well,

since risky or uncertain actions were filtered out during simulation. Engineers observed fewer emergency interventions and reported greater confidence in automated recommendations because each decision was accompanied by an explanation and predicted outcome. The hierarchical structure also reduced configuration noise by prioritizing fewer but more impactful changes.

Beyond quantitative improvements, the framework introduced qualitative benefits in workflow efficiency. Engineers were able to focus more on architectural design and strategic improvements rather than repetitive tuning tasks. The digital twin provided clearer visibility into system dynamics, enabling better understanding of how different components interacted. As a result, collaboration between human operators and automated agents became smoother and more effective.

This case study demonstrates that integrating hierarchical planning with consequence-aware world models is not merely a theoretical enhancement but a practical solution for real-world data platforms. The framework successfully translated predictive intelligence into operational gains, validating its feasibility for enterprise deployment. These observations reinforce the broader conclusion of this research: structured planning and consequence prediction are essential foundations for building reliable, autonomous, and self-optimizing agentic systems in modern data engineering environments.

### Author Biography / About the Author

**Brahma Reddy Katam** is a data engineering professional, researcher, and technology enthusiast with extensive experience in building scalable data platforms, analytics systems, and AI-driven solutions. He specializes in data engineering, cloud-based data architectures, and the practical application of artificial intelligence to solve real-world problems.

Brahma has worked across multiple domains, including enterprise analytics, metadata management, data pipelines, and AI-powered data products. His work emphasizes simplifying complex data systems and making advanced technologies accessible through intuitive design and strong engineering foundations. He has hands-on experience with modern data platforms such as Databricks, Delta Lake, SQL-based analytics, PySpark, and cloud-native architectures.

He is also an active contributor to the data engineering and analytics community through blogs, research papers, proof-of-concept applications, and learning platforms. His research interests include embedding-based AI systems, agentic AI for analytics, lakehouse architectures, and next-generation data discovery mechanisms.

#### 14. References

- [1] T. Brown, B. Mann, N. Ryder et al., “Language Models are Few-Shot Learners,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [2] J. Wei, X. Wang, D. Schuurmans et al., “Chain-of-Thought Prompting Elicits Reasoning in Large Language Models,” *NeurIPS*, 2022.
- [3] S. Reed, Y. Wu, A. Parikh et al., “A Generalist Agent,” *Transactions on Machine Learning Research (TMLR)*, 2022.
- [4] D. Silver, A. Huang, C. J. Maddison et al., “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature*, vol. 529, pp. 484–489, 2016.
- [5] D. Ha and J. Schmidhuber, “World Models,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [6] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018.
- [7] V. Mnih, K. Kavukcuoglu, D. Silver et al., “Human-Level Control through Deep Reinforcement Learning,” *Nature*, vol. 518, pp. 529–533, 2015.
- [8] Y. LeCun, “A Path Towards Autonomous Machine Intelligence,” Meta AI Research Position Paper, 2022.
- [9] M. Grieves and J. Vickers, “Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems,” in *Transdisciplinary Perspectives on Complex Systems*, Springer, 2017.
- [10] M. Zaharia, R. Xin, P. Wendell et al., “Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores,” *Proceedings of the VLDB Endowment (PVLDB)*, 2020.
- [11] A. Armbrust, M. Zaharia, A. Ghodsi et al., “Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics,” *CIDR Conference*, 2021.
- [12] C. Sutton, L. Morgan, and G. Simoncelli, *Observability Engineering: Achieving Production Excellence through Structured Telemetry*, O’Reilly Media, 2022.
- [13] G. Candido, R. Kazman, and H. Erdogmus, “DevOps and DataOps: A Systematic Mapping Study,” *Journal of Systems and Software*, vol. 182, 2021.
- [14] G. Tesauro, N. Jong, R. Das et al., “Managing Complexity in Large-Scale IT Systems Using Machine Learning,” *IBM Journal of Research and Development*, vol. 59, no. 2/3, 2015.
- [15] A. Verma, L. Cherkasova, and R. Campbell, “ARIA: Automatic Resource Inference and Allocation for MapReduce Environments,” *ACM International Conference on Autonomic Computing (ICAC)*, 2011.
- [16] J. Dean and S. Ghemawat, “MapReduce: Simplified Data Processing on Large Clusters,” *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [17] B. Burns and D. Oppenheimer, “Design Patterns for Container-Based Distributed Systems,” *USENIX ;login.*, vol. 43, no. 3, 2018.
- [18] Y. Bengio, A. Courville, and P. Vincent, “Representation Learning: A Review and New Perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, 2013.