PredictStream: Event Streams as the Foundation for Actionable Artificial Intelligence

> **Note: This white paper is currently in draft form and subject to ongoing revision.**

Title Page

- **Title:** PredictStream: Event Streams as the Foundation for Actionable Artificial Intelligence
- **Authors:** David L Remy, Aamar Hussain (prospectively)
- **Date:** May 21, 2025

Abstract

The accelerating proliferation of real-time data—across industries, infrastructures, and daily life—has rendered event streams the true pulse of the modern world [REF1] Yet, despite their ubiquity, most event streams remain underutilized: noisy, heterogeneous, and disconnected from the actionable intelligence that the world increasingly demands [Wrench et al., 2016] This white paper introduces PredictStream, an open source project and reference implementation, aimed at harnessing advances in large language models (LLMs) and modular, context-aware agent pipelines[REF3]—capable of leveraging a range of AI and statistical techniques for context enhancement—to transform raw, messy streams into semantically rich, actionable knowledge and to facilitate timely, intelligent action. By positioning event streams as the substrate for intelligence, PredictStream seeks to bridge the gap between the dynamism of event streams and the adaptability of modern AI.[REF4] We argue that event streams are not only natural fits for modern AI—mirroring the sequential, semantic structure of language—but are also foundational to a new era of context-aware, adaptive, and self-improving systems.[REF5] Through technical architecture, real-world use cases, and critical discussion, we invite feedback and collaboration on PredictStream as one open, extensible approach to unlocking the potential of event streams.

---

1. Introduction

The world is awash in streams. From the clickstreams of e-commerce, the sensor feeds of industrial IoT, and the transaction logs of finance, to the event-driven workflows of healthcare and IT operations, virtually every process can be described as a flow of events, interactions, and signals. These event streams represent the true heartbeat of modern enterprises and, by extension, the fabric of our digital society.

Despite their centrality, the vast majority of event streams are underleveraged. Traditional data architectures treat streams as transient, low-level artifacts—mere precursors to static datasets or batch-processed analytics. Artificial intelligence, meanwhile, has largely evolved in the context of static, curated corpora: images, tables, and, most notably, text. This disconnect has left a critical gap between the real-time, dynamic nature of the world and the actionable, adaptive intelligence that organizations and communities increasingly require.

PredictStream is founded on the premise that event streams are not only compatible with, but fundamentally suited to, the most advanced forms of AI. Like language, event streams are sequences of semantically meaningful units—tokens of action, change, and context. Modern LLMs have demonstrated the power of learning from sequential, context-rich data; PredictStream extends this paradigm to the world of events, enabling AI not only to interpret and predict, but to act upon the streams that define our reality in real time.

This white paper proposes PredictStream as an open source project and reference implementation, intended to serve as a catalyst for innovation and collaboration in the field of event-stream-driven AI. The core value of PredictStream lies in operationalizing real-time, AI-driven actions: transforming passive understanding of data into active, automated decision-making and intervention through agent technology. We recognize the urgent need for this type of infrastructure and invite feedback and contributions from researchers, practitioners, and organizations. By working together, we can unlock the full potential of event streams and create a new era of actionable, adaptive, and transparent systems that close the loop from insight to action.

---

2. Background & Related Work

In parallel, artificial intelligence—particularly deep learning and large language models (LLMs)—has advanced rapidly, but has largely operated on static, curated datasets. The sequential, context-rich nature of language has enabled LLMs to achieve remarkable results in understanding, summarizing, and generating text. Yet, the application of such models to real-time, heterogeneous event streams remains nascent. Most AI deployments in industry still rely on batch processing, periodic retraining, and post hoc analytics, leaving a critical gap between the potential of AI and the realities of streaming data.

Recent research has begun to explore the intersection of streaming data and AI, with efforts to apply machine learning to anomaly detection, predictive maintenance, and real-time recommendation systems.[REF6] However, these solutions are often domain-specific, inflexible, and challenged by the messiness and heterogeneity of real-world streams.[REF7] There remains a need for a universal, extensible platform that can bridge the gap between the dynamism of event streams and the adaptability of modern AI.[REF8]

Figure 1: Evolution from Traditional Data Pipelines to PredictStream's Event-Oriented, Agentic Architecture

```
flowchart LR
    A[Event Sources - Files APIs Sensors Queues] --> B[Ingestion and Normalization]
    B --> C[Semantic Enrichment and Knowledge Extraction]
    C --> D[Configurable Agent Pipeline]
    D --> E[Prediction and Evaluation]
    E --> F[Action and Feedback]
    F --> G[Operational Telemetry]
    G --> D
```

Figure 1. The evolution from traditional, batch-oriented data pipelines (top) to PredictStream's event-oriented, agentic, and feedback-driven architecture (bottom). PredictStream closes the loop from insight to action, enabling self-improvement and real-time operational intelligence.

---

3. Key Insights

Event Streams as Semantic Data

Event streams are often misconstrued as mere sequences of raw data points—timestamps, identifiers, and values flowing through digital infrastructure. In reality, each event carries semantic meaning: a click in an e-commerce application expresses intent; a sensor anomaly in an industrial setting signals potential failure; a medication administration event in healthcare marks a critical intervention. These events are not isolated, but are woven into broader narratives and contexts, much like tokens in natural language.

This semantic structure is foundational for effective AI. Just as language models derive meaning from the sequence and context of words, so too can AI systems extract actionable knowledge from the order, type, and context of events. Recognizing event streams as semantic data enables richer interpretation, more accurate prediction, and—crucially—the ability to trigger appropriate, timely actions in response to unfolding situations.

Traditional data processing pipelines often flatten or discard this semantic richness, treating events as undifferentiated records to be aggregated or batch-processed.[REF9] This approach misses the opportunity to understand intent, causality, and higher-order patterns that are only visible when events are interpreted in context.[REF10] PredictStream is designed to preserve and leverage this semantic structure, using modular agents to enrich, interpret, and act upon streams in real time. By doing so, it enables a shift from passive data collection to operational, context-aware intelligence that can drive meaningful intervention.

Example: In an e-commerce platform, a sequence of page views, searches, and cart additions is more than a list of actions—it is a story about a customer's intent, preferences, and journey. By interpreting these events semantically, PredictStream can anticipate needs, personalize experiences, and trigger timely offers or alerts, moving beyond analysis to real-time, AI-driven engagement.

Streams are Everywhere

Every organization, process, and system generates streams: from user interactions and sensor readings to financial transactions and log events. These streams are fundamental to actionable intelligence—whether for prediction, anomaly detection, or real-time response. Yet, most organizations treat streams as ephemeral, archiving them for later analysis rather than leveraging them for immediate, context-driven action. The underutilization of event streams represents a missed opportunity for both operational efficiency and strategic insight. Failing to act in real time can result in missed opportunities, increased risk, and an inability to respond to rapidly evolving situations—outcomes that can be costly or even catastrophic in domains such as finance, healthcare, and industrial operations.

Messy Data, Messy World

The real world is inherently messy. Event streams are often noisy, incomplete, and heterogeneous, reflecting the complexity and unpredictability of real processes.[REF11] Traditional data cleaning and transformation pipelines—especially those based on rigid rules or manual intervention—struggle to keep pace with this messiness, leading to lost value and missed signals.[REF12] These methods can be brittle, labor-intensive, and slow to adapt to new patterns or anomalies.[REF13] While AI offers promising new possibilities—such as learning directly from raw, unstructured streams to infer structure, semantics, and context—even the most advanced models have their limitations.[REF14] Not every challenge can be solved by AI alone, and many problems will require a blend of approaches.[REF15] PredictStream is designed to complement, not replace, established cleaning and transformation tools: its agent architecture can orchestrate and integrate traditional, rule-based, or manual methods alongside AI-driven techniques.[REF16] Zero-knowledge approaches, which require no prior assumptions about the data, are particularly promising for transforming messy streams into meaningful, actionable events, but are just one part of a broader toolkit.[REF17]

Figure 2: Transformation of Messy Event Streams into Semantic Events

```
flowchart LR
    A[Messy Event Stream - Raw Noisy Heterogeneous] --> B[Ingestion and Normalization]
    B --> C[Semantic Enrichment - Context Structure Meaning]
    C --> D[Semantic Events - Actionable Context Rich]
```

Figure 2. PredictStream transforms messy, heterogeneous event streams into structured, context-rich semantic events through ingestion, normalization, and enrichment. This enables actionable intelligence and real-time intervention.

Configurable Agent Pipelines

Real-time prediction and action require more than just batch-trained models; they demand flexible, composable pipelines of agents that can filter, enrich, predict, and react to events as they occur. PredictStream's architecture is built around modular agents, each responsible for a distinct phase of the pipeline. This modularity is not just a matter of software engineering elegance—it is essential for operationalizing real-time, AI-driven actions. Importantly, the agent-based approach is designed to be complementary and integrative: organizations can incorporate existing data cleaning, enrichment, or transformation tools as agents within the pipeline, alongside new AI-driven methods. By enabling rapid adaptation to new domains, seamless integration of emerging and established technologies, and continuous self-improvement through feedback and evaluation, PredictStream aims to help organizations move from static analysis to dynamic, automated intervention as situations unfold—while leveraging the best tools available, old and new.

Events All the Way Down

PredictStream is event-oriented at every level. Not only does it empower organizations to leverage external event streams, but its own internal operations, agent feedback, and learning cycles are themselves modeled as streams. This enables PredictStream to "eat its own dogfood": streaming its own telemetry, feedback, and operational events through its own pipelines, thereby enabling self-monitoring, self-improvement, and meta-learning. Over time, PredictStream will be able to recursively apply its own technology to optimize, adapt, and even redesign itself—creating a virtuous cycle of continuous improvement. By making streams the substrate for intelligence—rather than an afterthought—the platform enables self-improving, predictive, and agentic systems. Every agent action and pipeline phase is observable, auditable, and open to both automated and human-in-the-loop evaluation. This transparency not only supports trust and compliance, but also powers a rich ecosystem of visualization, introspection, and interactive exploration. This streaming foundation, both inward and outward, is what makes PredictStream uniquely adaptive, extensible, and future-proof.

From Insight to Action

Most data systems stop at analysis or prediction, leaving the crucial step of operationalizing those insights as a manual or disconnected process. PredictStream is designed to close this gap: its agent pipelines are not only capable of interpreting and predicting from event streams, but also of triggering timely, automated actions in response to unfolding events. This action orientation transforms streams from passive sources of information into engines of real-time intervention—enabling organizations to automate workflows, respond to anomalies, personalize experiences, and drive continuous improvement as events happen.

- **Example:** In an industrial setting, PredictStream can detect an anomaly in sensor data and immediately trigger an automated safety check or alert a human operator, minimizing downtime and risk.
- **Complementarity:** These actions can be fully automated, human-in-the-loop, or orchestrated with existing systems—ensuring flexibility and control.

---

4. PredictStream Architecture

PredictStream's architecture is designed to be both modular and extensible, enabling organizations to adapt to new domains, integrate emerging AI technologies, and continuously improve over time. At its core, the system is built around a flexible, agent-based pipeline, with each agent responsible for a distinct phase of event processing: ingestion, enrichment, prediction, evaluation, or action.

High-Level System Overview:

```
flowchart LR
    A[Event Sources - Files APIs Sensors Queues] --> B[Ingestion and
Normalization]
    B --> C[Semantic Enrichment and Knowledge Extraction]
    C --> D[Configurable Agent Pipeline]
    D --> E[Prediction and Evaluation]
    E --> F[Action and Feedback]
    F --> G[Operational Telemetry]
    G --> D
```

Figure 3: PredictStream High-Level Architecture (Mermaid diagram; view in a Mermaid-compatible markdown viewer)

- **Event Sources:** PredictStream ingests events from a wide variety of sources—files, APIs, message queues, sensors, or external services—abstracted via `StreamSource` interfaces. This abstraction allows seamless integration with both structured and unstructured data streams.
- **Ingestion and Normalization:** Incoming events are normalized and validated, ensuring consistent structure and metadata regardless of source. This phase also handles deduplication and initial filtering.
- **Semantic Enrichment and Knowledge Extraction:** Specialized agents (e.g., LLM-backed) enrich events with additional context, infer domains, extract entities, or perform summarization. Retrieval-augmented generation (RAG) agents can incorporate external knowledge for deeper semantic understanding.

- **Configurable Agent Pipelines:** The core of PredictStream is its agent pipeline: each phase is handled by modular, pluggable agents that can be independently developed, swapped, extended, or composed to filter, enrich, predict, evaluate, and act on events in real time. This pluggable architecture allows organizations to integrate custom logic, leverage best-in-class third-party tools, or rapidly adopt new AI models as the field evolves. Pipelines are fully configurable, supporting both linear and branched flows, and can be adapted to new domains or tasks with minimal effort.
- **Prediction and Learning Modules:** Prediction agents leverage LLMs or other models to forecast next events, classify anomalies, or generate recommendations. Learning is continuous, with agents able to adapt based on feedback and new data.
- **Feedback and Self-Improvement Loops:** Every agent emits structured feedback events to a central event log/bus, enabling auditability, introspection, and adaptive improvement. Evaluation agents compute metrics, assess prediction quality, and can trigger adaptation or alerting as needed.
- **Visualization and UI:** PredictStream provides real-time, interactive visualization and introspection tools, allowing users to explore event flows, agent actions, and prediction outcomes. The UI is powered directly by the structured feedback emitted at every stage.

This architecture ensures that PredictStream is not only powerful and flexible, but also transparent and auditable—key requirements for trust, compliance, and continuous improvement. The pluggable agent framework supports rapid experimentation, domain adaptation, and collaborative development, making it easy to tailor and evolve pipelines as needs change or new technologies emerge.

---

5. Handling Messy Data: From Noise to Knowledge

One of the most significant challenges in real-world event processing is the inherent messiness of data. Streams are often noisy, incomplete, ambiguous, and heterogeneous, reflecting the complexity and unpredictability of the environments from which they originate.

Challenges:

- **Heterogeneity:** Events may come from disparate sources, each with different formats, schemas, and semantics.
- **Missing Data:** Real-world streams frequently contain gaps, out-of-order events, or incomplete records.
- **Ambiguity:** The meaning of an event may depend on context, history, or external knowledge—often unavailable at ingestion time.

AI-Powered Solutions:

- **Stream Cleaning and Enrichment:** Agents equipped with AI models can impute missing values, resolve ambiguities, and harmonize heterogeneous data. LLMs can infer likely intent or context even from sparse or poorly structured events.
- **Entity and Event Recognition:** Semantic agents extract entities, relationships, and key attributes from raw streams, transforming low-level signals into meaningful, actionable events.[REF18]
- **Zero-Knowledge Learning:** PredictStream supports zero-knowledge approaches, in which agents learn the structure and semantics of streams without prior assumptions or labeled data. This enables rapid adaptation to new domains and the discovery of emergent patterns.[REF19]

Figure 4: Before and After—Transforming Messy Streams into Semantic Events

```
flowchart LR
    A[Messy Stream - Raw Noisy] --> B[Cleaning Agent]
    B --> C[Enrichment Agent]
    C --> D[Entity Recognition Agent]
    D --> E[LLM or ML Model]
    E --> F[Semantic Events - Structured Actionable]
```

Figure 4. Agents leveraging AI/ML technologies transform messy, heterogeneous event streams into structured, context-rich semantic events. Each agent specializes in a distinct phase, enabling deep semantic understanding and actionable intelligence.

By leveraging these AI-powered techniques, PredictStream turns the challenge of messy data into an opportunity for deeper insight, enabling organizations to extract value from even the most unruly streams.

---

6. Configurable Agent Pipelines for Real-Time Prediction

At the heart of PredictStream lies a modular, agent-based pipeline architecture, designed to serve as a reference implementation for event-stream-driven AI. Each pipeline is composed of pluggable agents, each specializing in a distinct phase of stream processing—such as cleaning, semantic enrichment, prediction, evaluation, or action. This modularity ensures that new agents can be introduced, replaced, or reconfigured with minimal disruption, enabling rapid adaptation to new domains, technologies, and organizational requirements.

The system is designed for discoverability and extensibility: agents are registered and easily integrated, allowing the platform to evolve alongside advances in AI and data engineering. Critically, PredictStream is intended to be open source and accessible—an infrastructure component that the industry, internet, and society at large can rely on as a foundation for real-time, actionable intelligence. Our vision is for PredictStream to become a standard, universally available platform: one that can be pointed at any event stream and, with minimal configuration, immediately begin delivering value—learning, adapting, and improving autonomously as it operates.

Real-time operation is a first-class concern. The pipeline processes events as they arrive, supporting low-latency response to emerging patterns and anomalies. Agents learn continuously from feedback, enabling the pipeline to improve over time. Structured feedback emitted at every stage supports auditability, introspection, and interactive visualization, making the system not only powerful but also transparent and trustworthy. For implementation details and usage patterns, we refer readers to the PredictStream design document, which provides concrete guidance for deploying, configuring, and extending the platform in real-world settings.

```
flowchart LR
    A[Event Ingestion] --> B[Cleaning Agent]
    B --> C[Semantic Enrichment Agent]
    C --> D[Prediction Agent]
    D --> E[Evaluation Agent]
    E -->|Feedback| C
    E --> F{Action Required?}
```

```
    F -- Yes --> G[Action Agent]
    F -- No --> H[End]
    G --> I[External System / Notification]
    G --> J[Human-in-the-Loop?]
    J -- Yes --> K[Human Review/Approval]
    K --> I
    J -- No --> I
    style G fill:#ffe082,stroke:#ffb300,stroke-width:2px
    style K fill:#c5cae9,stroke:#303f9f,stroke-width:2px
```

Figure 5: Example Agent Pipeline Flow (Mermaid diagram; view in a Mermaid-compatible markdown viewer)

This agent-based approach provides the flexibility and power needed for real-time, context-aware intelligence, supporting both automated and human-in-the-loop workflows. Crucially, PredictStream is designed to operationalize intelligence—enabling pipelines that not only analyze and understand streaming data, but also trigger timely, AI-driven actions and interventions. It is our belief that such a platform should be a public good—open, extensible, and easy to use—so that the benefits of event-stream-driven AI are available to all.

---

7. Use Cases & Value Propositions

PredictStream's agent-based architecture is uniquely suited not just for insight, but for automating action—especially the mundane, repetitive, or time-sensitive tasks that are often overlooked. Here are several action-oriented use cases and the distinctive value PredictStream brings:

1. Automated Incident Remediation (IT/DevOps)

- Streams logs and system events to detect common issues (e.g., service crashes, misconfigurations).
- Agents automatically restart services, roll back deployments, or apply fixes—before users notice a problem.
- Escalates to human operators only when automation is insufficient.

2. E-Commerce Workflow Automation

- Detects abandoned carts or failed payments in real time.
- Triggers personalized follow-up actions: sends reminders, applies discounts, or re-attempts transactions automatically.[REF20]
- Reduces manual intervention and increases conversion by acting immediately.[REF21]

3. Financial Operations Automation

- Monitors for routine, repetitive transaction patterns (e.g., daily reconciliations, compliance checks).
- Agents automatically perform reconciliations, flag exceptions, and generate required reports—freeing staff from tedious, error-prone tasks.

4. Healthcare Task Automation

- Identifies when routine tasks (e.g., medication reminders, appointment scheduling, supply restocking) are needed.
- Automatically sends reminders to patients, schedules appointments, or places supply orders.

- Allows clinical staff to focus on complex care, not administrative overhead.

5. Industrial IoT: Proactive Maintenance Actions

- Detects early warning signals from equipment.
- Automatically schedules maintenance, orders parts, or adjusts machine parameters.
- Prevents downtime with minimal human intervention.

Value Propositions (Action-Oriented):

- **Automates the Mundane:** Frees human talent for higher-value work by automating repetitive, routine actions.
- **Closes the Loop:** Goes beyond alerts—directly intervenes, remediates, and resolves issues in real time.
- **Customizable Workflows:** Action agents can be tailored to organizational needs, integrating with existing systems or triggering external processes.
- **Hybrid Automation:** Supports seamless handoff between automated actions and human oversight.
- **Open & Accessible:** Designed as open source infrastructure for broad adoption and community-driven innovation.

---

8. Discussion

Why Event Streams are the Future of AI—and AI is the Future of Event Streams

Event streams mirror the dynamism and complexity of the real world. As the fundamental substrate of digital processes, they offer a natural fit for modern AI—especially as models become more context-aware, adaptive, and capable of real-time reasoning. By treating streams as first-class citizens, PredictStream aligns with the way the world actually works: as a continuous flow of semantically rich events, not static tables or isolated records.

Universal, Real-Time Data Substrate: PredictStream's vision is for event streams to become the universal substrate for actionable intelligence—supporting not only prediction, but also understanding, evaluation, and action.[REF22] This shift enables organizations to move from reactive, batch-oriented processes to proactive, real-time operations.[REF23]

Barriers to Adoption: Despite these advantages, significant challenges remain. Many organizations are encumbered by legacy architectures, entrenched batch-processing pipelines, and siloed data.[REF24] The transition to event-driven AI requires not only technical innovation, but also cultural and organizational change.[REF25] Integration with existing systems, assurance of data quality, and the need for explainability and trust are all critical hurdles.[REF26]

How PredictStream Overcomes These: PredictStream addresses these barriers through its open, extensible, and agentic design.[REF27] Modular pipelines support incremental adoption, while structured feedback and auditability foster transparency and trust.[REF28] The platform's zero-knowledge learning capabilities enable rapid adaptation to new domains, even in the absence of labeled data or predefined schemas.[REF29]

Addressing Contrarian Challenges:

- **Event Streams ≠ Text Streams:** While event data is sparser and more heterogeneous than natural language, the agent-based approach and semantic enrichment capabilities of PredictStream allow for meaningful interpretation and prediction, even in challenging contexts.
- **Messiness and Label Scarcity:** Zero-knowledge and unsupervised learning agents can extract actionable semantics from messy, unlabeled streams, though further research is needed to maximize their effectiveness.
- **Latency and Scalability:** PredictStream is designed for low-latency, high-throughput operation, but scaling to massive, global streams will require ongoing optimization and research.
- **Explainability and Trust:** Structured feedback, audit trails, and interactive visualization support transparency and regulatory compliance, but explainability remains an open area for innovation.
- **Integration with Legacy Systems:** Incremental, modular deployment allows organizations to adopt PredictStream alongside existing infrastructure, easing the transition to event-driven intelligence.
- **Novelty:** While event-driven architectures exist, PredictStream's integration of modular AI agents, zero-knowledge learning, and end-to-end auditability represents a significant advance over existing solutions.

A Living Research Project: This white paper itself is an ongoing research effort. As the field evolves, new challenges and opportunities will emerge. PredictStream is designed to grow, adapt, and incorporate advances from both academia and industry.

---

9. Conclusion

Summary of Contributions and Differentiators

- PredictStream provides an open, modular, agent-based platform for transforming raw event streams into actionable, real-time intelligence.
- Its pluggable agent architecture enables flexible, incremental adoption and rapid integration of new AI models, custom logic, and third-party tools.
- The platform is designed for operationalization—closing the loop from prediction to action, automating both mundane and critical interventions, and supporting hybrid human/AI workflows.
- PredictStream is transparent and auditable, with structured feedback and visualization at every stage, supporting trust and compliance.

Limitations and Open Challenges

While PredictStream advances the state of the art, it is not a universal solution. Challenges remain in areas such as:

- Ensuring data quality and robust integration with legacy systems
- Scaling to massive, heterogeneous, and high-velocity streams
- Maintaining explainability, safety, and responsible AI deployment in sensitive domains
- Evolving standards and best practices for event-driven AI pipelines Ongoing research and community collaboration are needed to address these open questions.

Call to Action

We invite practitioners, researchers, and organizations to engage with PredictStream as an open source project and reference implementation. You can participate by:

- Contributing new agents, connectors, or visualization tools
- Sharing real-world use cases, benchmarks, and feedback
- Collaborating on best practices, standards, and governance for event-driven AI
- Helping to advance research on zero-knowledge learning, operational feedback, and safe automation

Looking Forward

PredictStream is designed to be a robust, production-grade platform that evolves as the field advances and real-world needs emerge. Our aim is to catalyze further innovation, collaboration, and widespread adoption of event-stream-driven AI—helping to establish a new, reliable foundation for operational, adaptive intelligence in the real world.

Operationalizing Real-Time Intelligence:

PredictStream's unique value lies in operationalizing real-time intelligence, enabling organizations to move from passive analysis to active, automated decision-making and intervention. By integrating AI-driven prediction and action into the core pipeline, PredictStream empowers organizations to respond to emerging patterns and anomalies in real time, unlocking new opportunities for operational efficiency, strategic insight, and competitive advantage.

References

References

- **[REF1]** Kreps, J., Narkhede, N., & Rao, J. (2011). "Kafka: a Distributed Messaging System for Log Processing." LinkedIn Engineering. https://engineering.linkedin.com/kafka/kafka-architecture
- **[REF2]** Wrench, C., Stahl, F., Di Fatta, G., Karthikeyan, V., & Nauck, D. D. (2016). "Data stream mining of event and complex event streams: a survey of existing and future technologies and applications in big data."
- **[REF3]** Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). "Attention is All You Need." Advances in Neural Information Processing Systems (NeurIPS 2017). https://arxiv.org/abs/1706.03762
- **[REF4]** Stonebraker, M., Çetintemel, U., & Zdonik, S. (2005). "The 8 requirements of real-time stream processing." SIGMOD Record, 34(4), 42–47. doi:10.1145/1107499.1107504

Further supporting materials and additional references will be included as the document is finalized.

Appendix