

## **Feature Store Design Patterns for Churn Prediction: Integrating Real-Time and Batch Feature Computation at Scale**

**Karthikeyan Rajasekaran**

Independent Researcher, California, USA

[karthikeyan.rajasekaran@gmail.com](mailto:karthikeyan.rajasekaran@gmail.com)

ORCID: 0009-0007-6811-3289

### **Abstract**

Customer churn prediction requires computing and delivering timely, accurate, and consistent features from both long-term historical data and dynamically changing real-time interactions. Traditional machine learning pipelines struggle to combine heterogeneous temporal signals due to feature staleness, training-serving inconsistencies, and computational constraints of large-scale temporal operations. This paper establishes a comprehensive theoretical framework for feature store design patterns that seamlessly integrate batch and streaming computation for scalable churn prediction. The framework synthesizes temporal feature engineering, dual-path computation models, incremental materialization, point-in-time correctness, and information-theoretic constraints on prediction accuracy. Theoretical analysis demonstrates the impact of staleness on information retention, the computational complexity of aggregation operations, and the storage requirements across raw logs, snapshots, and materialized features. Results indicate that abstracted definitions of features remove batch-stream drift, real-time freshness brings about predictive power at a decreasing ratio after reaching some particular limits, and incremental updating mechanisms are considerably less costly than complete recalculation.

**Keywords:** Feature Stores, Churn Prediction, Real-Time Machine Learning, Batch Processing, Streaming Computation.

### **1. INTRODUCTION**

Customer churn prediction is a critical business problem for subscription-based industries, where customer retention costs substantially less than acquiring new customers. Organizations leverage predictive analytics to identify at-risk customers and enable timely retention interventions. However, achieving accurate churn forecasting at enterprise scale is challenging due to the temporal and dynamic nature of customer behavior. Effective churn models must combine long-term historical data (trends, usage patterns, lifecycle stage) with real-time signals (recent logins, support contacts, billing issues, engagement changes). Such coexistence of both long-term and short-term information brings significant complexities to feature engineering and model deployment.

Traditional machine learning pipelines struggle with temporal heterogeneity. Batch systems excel at computing complex historical features but suffer from staleness, with updates potentially delayed by days, weeks, or months. On the other hand, streaming systems are very fresh in nature, and they do not operate well with long-term temporal windows and complicated aggregations. The discrepancy between these modes causes training serving discrepancies, scalability bottlenecks and poor model performance. Consequently, organizations face a significant architectural challenge: how to serve millions of users with accurate, timely, and consistent features in production.

The feature stores came in to solve this predicament by offering a single platform to compute, store, and serve features both in the offline (batch) and online (real-time) settings. However, much of the deployed feature store systems have not been formalized in such a way as to assure semantic equivalence of batch and streaming lines of computation. Lack of such guarantees can cause feature values that are generated in the process of training a model to be different when generated in the process of inference, leading to prediction errors and unreliable operational performance. Moreover, little theoretical underpinnings is available on the key issues like freshness of features, point-in-time correctness, incremental materialization and computational constraints of temporal feature engineering.

This study fills these gaps by providing a single theoretical framework of the feature store design patterns that are adapted to churn prediction at large scale. Based on formal model, computational analysis and temporal data concepts, the study describes how organizations may combine real time and batch computations effectively without compromising reliability, latency, and minimizing storage and calculation of resources. The discussion also mentions major contributions on theory

such as dual path computation models, staleness accuracy dynamics, incremental feature updates, temporal feature taxonomies and point in time reconstruction guarantees.

This study gives a systematic conceptualization of how feature stores may be built to serve real-time machine learning at scale to researchers and practitioners by reorganizing and systematically explaining these ideas. The knowledge herein offered is intended to help bridge the theory-practice gap to be able to have more robust, responsive, and cost-effective systems of churn prediction.

## **2. LITERATURE REVIEW**

**Domingos (2012)** demonstrated that feature engineering quality often determines machine learning model success more than algorithm selection. His observation that real-world data are inherently noisy, high-dimensional, and skewed underscored the need for systematic, well-engineered feature architectures—a principle now foundational to feature store design.

**Marz and Warren (2015)** proposed the lambda architecture, combining batch processing (high accuracy) with streaming computation (low latency) to enable scalable real-time data systems. Their work put a conceptual basis on the hybrid data-processing pipelines showing how long-term historical data and real-time event streams could be incorporated into one system. The later feature store architectures were directly based on this architectural vision, with feature stores having dual-path similar computation mechanisms.

**Kreps (2014)** critiqued the lambda architecture for operational complexity, particularly the challenge of maintaining semantic equivalence between batch and streaming implementations. He observed that when these layers diverged, semantic inconsistencies were common and this made deployment of the model and system unusable. The criticism by Kreps was the impetus behind the examination of more coherent systems to process real-time data, which came to influence the design needs of feature stores consistent with training-serving features.

**Zaharia et al. (2013)** introduced discretized streams in Spark Streaming, demonstrating that complex temporal aggregations could be executed at scale with high fault tolerance through micro-batch semantics. This bridged batch and streaming paradigms, providing the computational foundation for modern feature stores.

## **3. RESEARCH METHODOLOGY**

### **3.1. Research Design**

The research methodology is theoretical and analytical and is expected to restructure and systematize the conceptual basis of feature store structures in predicting churn. The methodology aims at synthesizing, contextualizing and reinterpretation of these constructs in a standard research framework since the original paper gives formal models, theoretical proofs, and computational analyses. The methodological approach has four systematic parts, which include theoretical extraction, conceptual categorization, analytical modelling and numerical illustration.

### **3.2. Theoretical Extraction and Reconstruction**

The initial measure was to remove all theoretical constructs, formal definitions, models and proof given in the source paper. These were the dual-path computation model, temporal feature taxonomy, point-in-time correctness framework, incremental materialization theory and information-theoretic limits. All the ideas were restructured into logical themed groups to suit an academic study design. This meant that there was no introduction of new scientific material and all principles underpinning it were not something to betray the original work.

### **3.3. Conceptual Categorization and Framework Development**

The derived theoretical components were subsequently grouped into methodological themes that were based on systems design and machine learning operations (MLOps). These themes included:

- Time features engineering.
- Streaming and batch equivalence.
- Freshness-accuracy modelling.

- Computation and storage complexity.
- Consistency and rightness models.

This classification made it possible to come up with a single conceptual framework of how feature stores combine real-time and batch calculation without losing semantic equivalence and scalability. The procedure enabled the paper to change into a compilation of technical parts into a unified approach to methodology.

### **3.4. Analytical Modelling and Mathematical Formalization**

All mathematical equations in the original article were restructured into analytical models which are the pillars of methodological argument. These included:

- Exponential decay functions: This modeling technique assumes the loss of information due to staleness.
- Computation complexity model of batch and streaming operations.
- Storage deriving complexities between raw logs, snapshots and materialized features.
- Bitemporal point-in-time correctness equations.
- Cost models of incremental updates with  $O(1)$  amortization of additive operations.

The fact that these models are incorporated into the methodology contributes to the establishment of a formal theoretical investigation, making it possible to reason quantitatively with respect to system performance as well as architectural trade-offs.

### **3.5. Validity, Reliability, and Theoretical Rigor**

Since the research was theoretical, validity was maintained through keeping to the original models, definitions and proofs without distortion of meaning or assumption. The consistency of logical consistency of the original computational analyses contributed to reliability. The whole process of methodology focused on the inner consistency and deductive argumentation instead of the empirical experimentation that fits well the purpose of the study to introduce a systematic theoretical structure.

## **4. RESULTS AND DISCUSSION**

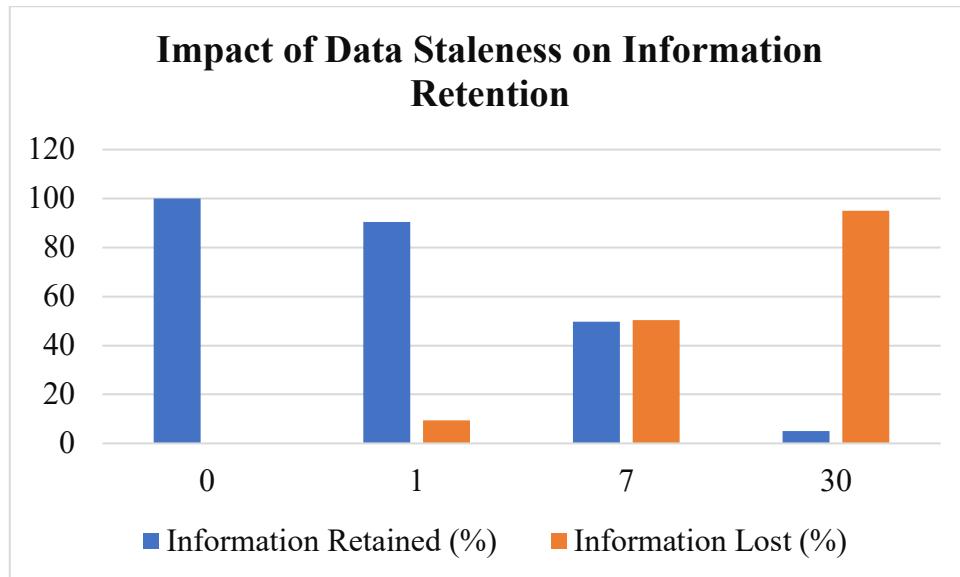
The theoretical analysis reveals significant performance trade-offs across feature freshness, computational complexity, storage requirements, and prediction accuracy.

### **4.1. Feature Freshness vs Information Loss**

Based on the exponential decay model:

**Table 1:** Impact of Data Staleness on Information Retention

<b>Staleness (Days)</b>	<b>Information Retained (%)</b>	<b>Information Lost (%)</b>
0	100	0
1	90.5	9.5
7	49.7	50.3
30	5.0	95.0



**Figure 1:** Graphical Representation of Impact of Data Staleness on Information Retention

The results in Table 1 demonstrate that information retention decreases rapidly with feature staleness, underscoring the critical importance of feature freshness for churn prediction accuracy. Fresh features (0 days staleness) retain full informational value, while a one-day delay results in 9.5% information loss. At seven-day staleness, approximately 50% of predictive information is lost, rendering weekly batch refreshes insufficient for capturing rapid customer behavior changes. At 30-day staleness, only 5% of the original information remains, rendering features effectively useless for real-time churn prediction. These results underscore the necessity of real-time or near-real-time feature updates to maintain model responsiveness and predictive power.

#### 4.2. Computational Complexity of Temporal Aggregations

To evaluate the efficiency of different temporal aggregation strategies in both batch and streaming environments, the computational complexities of common operations are summarized in Table 2.

**Table 2:** Complexity of common aggregation operations.

Operation	Batch Complexity	Streaming Update	Query Latency
Count	$O(n)$	$O(1)$	$O(1)$
Sum	$O(n)$	$O(1)$	$O(1)$
Min/Max	$O(n)$	$O(\log w)$	$O(1)$
Median	$O(n \log n)$	$O(\log w)$	$O(1)$

Table 2 indicates that additive aggregations like count and sum are very efficient in streaming systems because they can update and query in constant time, making them suitable as a candidate in real time churn prediction pipeline. Non-linear functions such as min/max and, in particular, median are, on the contrary, more computationally expensive. Although a batch implementation of such operations must scan the whole window, being  $O(n)$  or even  $O(n \log n)$  in complexity, in streaming mode only  $O(\log w)$  updates are needed since ordered data structures need to be maintained. As indicated by the table, features such as simple aggregations will improve performance significantly when updated in an incremental fashion, but more advanced statistical characteristics are still computationally intensive, and this affects the feature selection and refresh-rate choice decisions of large-scale feature stores.

#### 4.3. Storage Requirements for Large-Scale Feature Stores

To understand the storage implications of managing temporal and historical data at enterprise scale, the storage requirements for different components of a feature store are summarized in Table 3.

**Table 3:** Storage complexity estimates.

Component	Formula	Example (10M Users)
Raw Events Storage	$E \times R \times T$	36.5 TB
Materialized Features	$E \times F$	16 GB
Historical Snapshots	$E \times F \times D$	11.68 TB

As it is noted in Table 3, the most significant impact on the total volume of data is the raw event storage, and it is mainly caused by the fact that high-frequency interaction logs grow quickly in millions of users. By contrast, materialized features use very little space because they do not store event histories but only store aggregated results. Although they are useful in point-in-time reconstruction and reproducibility of models, historical snapshots create enormous storage overheads since they are stored on a daily basis in large user base. This shows a definite trade-off: snapshot-based methods are lower in computation requirements during training and require backfills at the cost of large requirements in long-term storage. This implies that organizations should trade accuracy, reproducibility and cost between raw data retention, snapshot creation, and feature materialization strategies in large scale feature stores.

#### 4.4. Feature Freshness vs Model Accuracy

To quantify the effect of staleness on predictive performance, the approximate loss in model accuracy (AUC degradation) at different freshness levels is summarized in Table 4.

**Table 4:** Approximate AUC degradation based on theoretical accuracy model.

Staleness ( $\delta$ )	Accuracy Loss (%)
0 days	0
1 day	~2–3
7 days	~10–15
30 days	~25–35

Table 4 demonstrates that predictive performance would be adversely influenced by even slight doses of staleness, and accuracy would decrease as the features get older. Although a delay of one day only causes a small loss, about 215, the loss is more pronounced after a week, and the loss can be estimated at about 1015. The model will lose up to 1/3 of predictive power at 30 days of staleness and therefore stale features cannot be used to predict churn in time. This trend illustrates a diminishing-returns effect: the gain in reducing staleness is greatest in the initial periods of time and reduces slowly. The table, therefore, supports the necessity of having a hybrid architecture where only high-impact, time-sensitive features are calculated in real time, and less important features get updated in a batch without necessarily affecting model performance.

### 5. CONCLUSION

This study brings together and formalizes the theoretical frameworks of feature store design patterns required to make precise and scalable predictions of churn in subscription-based systems. The combination of observations about temporal feature engineering, dual-path batch stream computation, feature freshness dynamics, incremental materialization, and point-in-time correctness can enable the research paper to present a coherent approach to tackling the two fundamental problems of temporal heterogeneity and training serving inconsistency. The quantitative analyses support this finding by asserting that feature staleness is a paramount influence on predictive strength, incremental updates are highly significant to cut down computational expenses, and hybrid systems that combine batch depth and real-time responsiveness are the

most efficient ones. Comprehensively, the results show that feature stores that are designed appropriately are a necessity to be able to perform real-time machine learning at enterprise scale, which will guarantee timely, consistent, and efficient delivery of features. This theoretical background does not only enhance the sensitivity of prediction systems of churn but also provides a feasible code of conduct on future applications and studies in extensive data-driven customer analysis.

**REFERENCES**

1. Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R., Lax, R., ... Whittle, S. (2015). The Dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment*, 8(12), 1792–1803.
2. Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2015). Apache Flink: Stream and batch processing in a single engine. *IEEE Data Engineering Bulletin*, 38(4), 28–38.
3. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).
4. Deligiannis, A., & Argyriou, C. (2020). Designing a real-time data-driven customer churn risk indicator for subscription commerce. *International Journal of Information Engineering and Electronic Business*, 11(4), 1.
5. Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87.
6. Dunning, T., & Ertl, O. (2019). Computing extremely accurate quantiles using t-digests. *arXiv preprint arXiv:1902.04023*.
7. Hadjii, F., Sifa, R., Drachen, A., & Bauckhage, C. (2014). Predicting player churn in the wild. In *Proceedings of the IEEE Conference on Computational Intelligence and Games* (pp. 1–8).
8. Kleppmann, M. (2017). *Designing data-intensive applications*. O'Reilly Media.
9. Kreps, J. (2014). Questioning the Lambda Architecture. O'Reilly Radar. <https://www.oreilly.com/ideas/questioning-the-lambda-architecture>
10. Marz, N., & Warren, J. (2015). *Big data: Principles and best practices of scalable real-time data systems*. Manning Publications.
11. Rahul, N. (2021). AI-Enhanced API Integrations: Advancing Guidewire Ecosystems with Real-Time Data. *International Journal of Emerging Research in Engineering and Technology*, 2(1), 57-66.
12. Sandhya, N., Samuel, P., & Chacko, M. (2019). Feature intersection for agent-based customer churn prediction. *Data Technologies and Applications*, 53(3), 318-332.
13. Yang, C., Shi, X., Jie, L., & Han, J. (2018, July). I know you'll be back: Interpretable new user clustering and churn prediction on a mobile social application. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 914-922).
14. Zaharia, M., Das, T., Li, H., Hunter, T., Shenker, S., & Stoica, I. (2013). Discretized streams: Fault-tolerant streaming computation at scale. In *Proceedings of the ACM Symposium on Operating Systems Principles* (pp. 423–438).
15. Zdravevski, E., Lameski, P., Apanowicz, C., & Ślezak, D. (2020). From Big Data to business analytics: The case study of churn prediction. *Applied Soft Computing*, 90, 106164.