

## Real-Time Financial Risk Assessment Using Streaming Data and Adaptive Machine Learning

Inés Costa

Research Fellow, Iberian Centre for Adaptive Financial Intelligence, Spain

### Abstract

Financial risk can change rapidly as market volatility, liquidity, correlation structures, and trading activity evolve. Conventional risk models are commonly estimated from historical observations and updated periodically, which may be insufficient when statistical relationships change within high-frequency or continuously arriving data. Streaming data analytics and adaptive machine learning provide a methodological framework for updating predictive models incrementally and detecting changes in the relationship between incoming financial indicators and short-horizon risk conditions. This hypothetical methodological study examines real-time financial risk assessment using streaming data and adaptive machine learning. A fictional stream containing 50,000 sequential financial observations is divided chronologically into model-development, validation, and evaluation segments. Three conceptual approaches are compared: static logistic regression, an incremental Hoeffding-tree classifier, and an adaptive online ensemble supported by concept-drift detection. Predictor variables include recent return magnitude, realized volatility, liquidity spread, volume imbalance, downside movement, short-term drawdown, rolling correlation, and volatility change.

The illustrative analysis evaluates accuracy, precision, recall, and F1 score for identifying a predefined elevated-risk state. Under the assigned hypothetical values, the adaptive ensemble achieves the strongest performance, particularly in recall of elevated-risk observations. A second analysis evaluates the adaptive model during stable, gradually changing, abrupt high-volatility, and post-change environments. Performance decreases temporarily during simulated distributional changes and subsequently improves as the model adapts. The study concludes that adaptive learning provides a useful methodological approach for real-time financial risk monitoring when statistical relationships evolve over time. However, continuously updating models introduce additional risks related to concept-drift detection, unstable labels, delayed outcomes, model overreaction, false risk alerts, and validation under changing distributions. All results in this study are illustrative and require empirical verification using appropriately governed financial data before any operational risk-management application.

**Keywords:** Financial Risk Assessment, Streaming Data, Adaptive Machine Learning, Concept Drift, Online Learning, Real-Time Analytics, Financial Data Streams, Risk Prediction, Hoeffding Tree, Predictive Risk Monitoring.

### 1. Introduction

Financial risk assessment involves estimating the possibility and magnitude of unfavorable outcomes under uncertain market conditions. Changes in asset returns, volatility, liquidity, and dependence among

financial instruments can alter the level of risk over relatively short periods. Traditional risk models have therefore emphasized time-varying volatility and tail behavior rather than assuming constant market variability. McNeil and Frey demonstrated the importance of combining volatility models and extreme-value methods for conditional tail-risk estimation, while Engle and Manganelli modeled conditional Value-at-Risk directly as a dynamic quantile process.

Financial data increasingly arrive as sequential streams rather than isolated static datasets. Market prices, transaction summaries, spreads, volumes, and portfolio exposures can be updated continuously. This structure creates methodological challenges because many conventional machine-learning methods assume that training data are collected before a model is deployed. When new observations arrive continuously, repeated batch retraining can become computationally inefficient and may respond slowly to rapid changes in the underlying data-generating process.

Online machine learning addresses this problem by updating models incrementally as observations become available. Oza and Russell developed online versions of ensemble learning procedures, while Domingos and Hulten examined algorithms for mining high-speed data streams. Hulten et al. subsequently considered learning from time-changing data streams. These approaches established important principles for computationally efficient learning when complete datasets cannot be stored or repeatedly reprocessed after every new observation.

A more difficult problem arises when the relationship between predictors and outcomes changes through time. Gama et al. described this phenomenon as concept drift and reviewed strategies for detecting and adapting to changing statistical relationships. Bifet and Gavaldà developed adaptive windowing to alter the effective amount of historical data retained according to detected changes in the stream. Such methods are particularly relevant to financial environments because volatility regimes, correlations, liquidity conditions, and market behavior can vary substantially across periods.

The present hypothetical study examines adaptive machine learning for continuous financial risk assessment. Static logistic regression is compared with an incremental Hoeffding-tree model and an adaptive online ensemble that monitors distributional change. The objective is methodological rather than investment-oriented: the study examines whether adaptive learning can maintain predictive classification performance when a financial data stream changes over time. No actual security, stock exchange, portfolio, bank, or investor data are analyzed.

## 2. Objectives of the Study

The primary objective of this study is to examine whether adaptive machine learning can improve short-horizon financial risk classification when observations arrive sequentially and statistical relationships change over time. The investigation focuses on streaming-data processing, concept-drift adaptation, comparative model performance, and stability during changing market-like conditions.

### 2.1 Development of Streaming Risk Indicators

The first objective is to examine how sequential financial measurements can be converted into compact risk indicators suitable for online learning. The fictional feature stream includes recent absolute return, realized volatility, volatility change, liquidity-spread variation, order-flow imbalance, downside-return intensity, short-term drawdown, rolling cross-asset correlation, and normalized activity measures. Each feature is conceptually calculated using only information available before the corresponding prediction time.

The streaming design avoids unrestricted use of future information because financial prediction is especially vulnerable to temporal leakage. Rolling statistics, normalization parameters, and risk labels must be constructed chronologically. A feature that uses observations occurring after the prediction timestamp would provide artificially strong information about the future. The methodological framework therefore treats temporal integrity as a fundamental requirement rather than a secondary preprocessing consideration.

## 2.2 Comparison of Static and Adaptive Machine Learning

The second objective is to compare a fixed statistical classifier with models capable of incremental updating. Logistic regression represents a conventional static baseline trained during an initial historical period. A Hoeffding tree represents incremental learning from incoming observations, while the adaptive ensemble combines several online learners with a drift-detection mechanism. All configurations use the same conceptual features and final evaluation stream.

The comparison evaluates whether adaptive updating provides useful resilience when the data distribution changes. A static model can remain effective if relationships are stable but may deteriorate when market conditions depart substantially from its original training period. Adaptive models can respond to new observations, although excessive adaptation can also reduce stability. The study therefore evaluates predictive performance and behavior during different simulated stream regimes rather than assuming that continuous updating is automatically beneficial.

## 2.3 Assessment of Concept Drift and Risk-Monitoring Stability

The third objective is to examine how the adaptive model behaves during stable conditions, gradual distributional change, abrupt volatility increase, and subsequent stabilization. These regimes conceptually represent different forms of concept drift commonly discussed within streaming-data research. Performance is evaluated separately within each regime so that temporary model deterioration is not hidden by one aggregate score.

The study additionally considers the practical consequences of drift detection. A false drift signal could cause unnecessary model replacement or rapid reweighting, while delayed detection could leave the model using outdated patterns. The optimal response therefore depends on both statistical evidence and the operational meaning of the detected change. Adaptive risk systems should distinguish genuine structural change from temporary random fluctuations whenever possible.

## 3. Review of Literature

### 3.1 Financial Risk Measurement and Dynamic Forecasting

McNeil and Frey examined tail-related financial risk measures for heteroskedastic time series and combined GARCH filtering with extreme-value methods to estimate conditional Value-at-Risk and Expected Shortfall. Berkowitz and O'Brien evaluated the accuracy of commercial-bank Value-at-Risk models and emphasized the importance of testing risk forecasts against realized outcomes. These studies demonstrate that financial risk estimates should be evaluated dynamically rather than interpreted solely from in-sample model fit.

Engle and Manganelli developed Conditional Autoregressive Value-at-Risk models in which conditional quantiles evolve over time. Kuester et al. later compared alternative Value-at-Risk forecasting strategies and demonstrated that model performance can differ substantially according to

distributional assumptions and estimation method. Patton further examined volatility forecast comparison under imperfect volatility proxies. Collectively, this literature highlights the time-varying and evaluation-sensitive nature of financial risk measurement.

### 3.2 Online Machine Learning for Data Streams

Domingos and Hulten introduced efficient methods for mining high-speed data streams and developed the statistical principles underlying incremental decision-tree learning. Hulten et al. extended this work to time-changing streams, recognizing that older observations can become less representative. Oza and Russell developed online bagging and boosting, demonstrating that ensemble-learning principles can be adapted to sequential settings rather than requiring repeated batch estimation.

Bifet et al. developed the MOA environment for classification and clustering in evolving data streams. The framework supported online evaluation of decision trees, ensembles, and related methods while emphasizing scalability and evolving distributions. Aggarwal's edited volume on data-stream models and algorithms further consolidated methods for stream classification, clustering, summarization, and pattern analysis. These contributions provide a methodological basis for processing continuously arriving financial information.

### 3.3 Concept Drift and Adaptive Learning

Tsymbal examined concept drift as the problem of changing relationships in sequential data and reviewed strategies for updating models under changing conditions. Gama et al. developed drift-detection methods for online learning and subsequently provided a comprehensive survey of concept-drift adaptation. Their framework distinguished adaptation mechanisms, change detection, model updating, and evaluation under evolving data distributions.

Bifet and Gavalda proposed Adaptive Windowing, in which the effective historical window changes automatically according to statistically detected variation. Such mechanisms are relevant to financial applications because old observations may become less useful when volatility or market relationships change. However, adaptation must remain controlled: if the learner responds excessively to short-lived fluctuations, it may discard useful historical information and increase prediction instability.

## 4. Research Methodology

### 4.1 Research Design and Streaming Architecture

The present study adopts a hypothetical chronological streaming design. Incoming observations are processed sequentially, and predictions are generated before the corresponding risk-state labels become available. The static logistic-regression model is estimated during the initial development period and remains unchanged during final evaluation. The Hoeffding-tree classifier updates incrementally, while the adaptive ensemble updates its component learners and uses an ADWIN-style change detector to identify significant variation in prediction error.

The adaptive ensemble conceptually applies online bagging to several incremental tree learners. When sustained distributional change is detected, weaker components are replaced or reweighted using recent observations. The model therefore retains continuity while permitting adaptation. This architecture is illustrative and was not implemented experimentally. A genuine study would require explicit stream-processing infrastructure, delayed-label handling, model versioning, reproducibility controls, and computational latency measurements.

## 4.2 Hypothetical Financial Streaming Dataset

The fictional dataset contains 50,000 sequential observations derived conceptually from a diversified multi-asset risk-monitoring environment. Thirty-five thousand observations are assigned to model development, seven thousand to validation and parameter selection, and eight thousand to final chronological evaluation. Within the test stream, 2,000 observations are hypothetically labeled as elevated-risk states and 6,000 as ordinary-risk states.

Predictor variables include short-horizon absolute return, realized-volatility level, volatility acceleration, liquidity-spread change, trading-activity imbalance, downside movement, short-term drawdown, rolling cross-asset correlation, and normalized market-activity indicators. The elevated-risk label represents a fictional composite condition associated with unusually adverse short-horizon movement and increased variability. It does not correspond to any actual portfolio loss threshold, institution, security, or regulatory risk classification.

## 4.3 Evaluation Metrics and Prequential Assessment

The primary evaluation uses accuracy, precision, recall, and F1 score. True positives represent elevated-risk states correctly identified, while false negatives represent elevated-risk observations missed by the model. False positives represent ordinary-risk periods incorrectly flagged as elevated risk. Recall is particularly relevant because missed risk states reduce the usefulness of an early-warning system, whereas precision reflects how frequently alerts correspond to the predefined target condition.

Streaming evaluation additionally uses a prequential principle in which each observation is first predicted and then used for subsequent learning after its hypothetical label becomes available. This procedure reflects online operation more realistically than repeatedly training on the complete test stream. Performance is also reported separately across simulated stream regimes. No model selection is performed using final test outcomes, thereby preserving chronological separation between development and evaluation.

## 5. Analysis

### 5.1 Comparative Real-Time Risk Classification

The first hypothetical analysis compares the three models across 8,000 sequential evaluation observations. The static logistic-regression model correctly identifies 1,440 of 2,000 elevated-risk observations but generates 720 false-positive alerts. The incremental Hoeffding tree improves detection to 1,640 elevated-risk observations. The adaptive ensemble identifies 1,800 elevated-risk observations and reduces false positives to 480.

The adaptive ensemble therefore records an illustrative accuracy of 91.50%, precision of 78.95%, recall of 90.00%, and F1 score of 84.11%. The Hoeffding tree records an F1 score of 76.64%, while the static model reaches 69.23%. These differences are constructed examples and cannot establish empirical superiority. They demonstrate how adaptability could be evaluated when the distribution of streaming financial data changes during the test period.

**Table 1: Hypothetical Comparative Performance for Streaming Financial Risk Classification**

| Model                      | TP    | FP  | FN  | TN    | Accuracy (%) | Precision (%) | Recall (%) | F1 Score (%) |
|----------------------------|-------|-----|-----|-------|--------------|---------------|------------|--------------|
| Static Logistic Regression | 1,440 | 720 | 560 | 5,280 | 84.00        | 66.67         | 72.00      | 69.23        |
| Incremental Hoeffding Tree | 1,640 | 640 | 360 | 5,360 | 87.50        | 71.93         | 82.00      | 76.64        |
| Adaptive Online Ensemble   | 1,800 | 480 | 200 | 5,520 | 91.50        | 78.95         | 90.00      | 84.11        |

**Note:** The entire stream, labels, confusion-matrix counts, and performance metrics are fictional and do not represent actual financial-market forecasting.

## 5.2 Adaptive Performance Under Changing Stream Conditions

The second hypothetical analysis divides the evaluation stream into four conceptual regimes. During stable conditions, the adaptive model records an assigned accuracy of 94.2% and F1 score of 89.2%. Performance decreases under gradual drift and declines further during an abrupt high-volatility regime. After adaptation, performance recovers, although it remains slightly below the initial stable-period values.

The pattern demonstrates that adaptive learning does not eliminate the effect of structural change. Instead, the model may experience a temporary performance decline while accumulating sufficient evidence that the stream has changed. Faster adaptation can reduce this interval but may increase false drift detections during ordinary fluctuations.

**Table 2: Hypothetical Adaptive-Model Performance Across Streaming Regimes**

| Streaming Regime              | Accuracy (%) | Precision (%) | Recall (%) | F1 Score (%) | Illustrative Drift Response                     |
|-------------------------------|--------------|---------------|------------|--------------|-------------------------------------------------|
| Stable Statistical Conditions | 94.2         | 87.0          | 91.5       | 89.2         | No adaptation required                          |
| Gradual Distributional Drift  | 91.0         | 80.0          | 88.0       | 83.8         | Adaptation after approximately 140 observations |
| Abrupt High-Volatility Change | 86.8         | 75.0          | 85.0       | 79.7         | Adaptation after approximately 55 observations  |
| Post-Change Stabilization     | 92.6         | 84.0          | 90.0       | 86.9         | Updated ensemble stabilized                     |

**Note:** Regime definitions, drift timing, and all numerical results are illustrative. They do not correspond to an actual market event.

## 6. Discussion

### 6.1 Streaming Learning and Financial Risk Monitoring

The hypothetical findings illustrate why static models may become less effective when financial relationships evolve. A classifier estimated under one volatility and liquidity environment can assign inappropriate probabilities when return distributions, spreads, correlations, or activity patterns change. Concept-drift research provides a framework for interpreting this deterioration as a change in the statistical relationship between incoming predictors and the target variable rather than simply random prediction error.

Online learners address this issue by incorporating new observations without reconstructing the complete model from the beginning. Oza and Russell demonstrated how ensemble principles can be adapted to online settings, while stream-learning research established incremental decision trees and other scalable algorithms. Within financial risk monitoring, such approaches could conceptually support continuous updating when market conditions change faster than scheduled batch-model revisions.

### 6.2 Concept Drift and Adaptive Model Stability

The regime-specific analysis demonstrates the distinction between detecting change and adapting successfully to it. Drift detection does not guarantee immediate recovery because the model still requires informative new observations. Abrupt changes can therefore create temporary periods of weaker accuracy even when the change detector performs appropriately. Gradual drift creates a different challenge because small changes may accumulate before crossing a statistical detection threshold.

Gama et al. emphasized that concept-drift adaptation involves both change detection and model-management strategy. Bifet and Gavaldà's adaptive-window approach similarly adjusts the effective historical window as evidence of change accumulates. In financial streams, retaining too much old information can delay adaptation, while discarding historical data too aggressively can make models unstable. Successful adaptive learning therefore requires a controlled balance between memory and responsiveness.

### 6.3 Risk Prediction Versus Risk Measurement

The present model classifies a fictional elevated-risk state rather than replacing conventional financial risk measures. Value-at-Risk, Expected Shortfall, volatility forecasts, stress analysis, and scenario-based methods answer different questions from machine-learning classification. McNeil and Frey and Kuester et al. demonstrated the importance of tail behavior and careful out-of-sample evaluation in financial-risk forecasting. Adaptive machine learning should therefore be interpreted as a complementary predictive layer rather than a universal substitute for established risk models.

This distinction is particularly important because strong classification accuracy does not demonstrate accurate estimation of monetary losses or tail quantiles. A model can correctly identify many high-risk periods while misestimating the severity of losses. A complete empirical risk-management framework would therefore compare adaptive classification with statistical risk measures and evaluate calibration, exceedance behavior, economic relevance, and robustness under extreme conditions.

## 7. Challenges

### 7.1 Concept Drift Detection and False Adaptation

A central challenge is distinguishing genuine statistical change from short-lived random fluctuation. Financial streams are inherently noisy, and temporary volatility increases can occur without representing a persistent structural shift. A drift detector that reacts too quickly may repeatedly reset or alter models, reducing stability. Conversely, an insensitive detector may leave obsolete relationships active for too long after genuine change.

The appropriate sensitivity depends on the prediction task and error cost. Adaptive systems should therefore report change-detection thresholds, warning periods, model replacements, and performance immediately before and after adaptation. Evaluation should include both true structural changes and periods of high but temporary variation. Without such analysis, apparent adaptability may simply reflect aggressive response to noise.

### 7.2 Delayed Labels and Real-Time Evaluation

Online learning assumes that the outcome associated with a prediction eventually becomes available. In financial applications, the delay between prediction and confirmed risk outcome can vary according to the chosen horizon. A model predicting a future risk state cannot be updated with that label until the horizon has passed. This creates a delay between observation, prediction, verification, and model adaptation.

Real-time architectures must therefore distinguish event-arrival time from label-arrival time. Updating a model immediately using future outcomes would constitute temporal leakage. Genuine implementations require queues or delayed-supervision mechanisms that ensure only confirmed historical outcomes are used for adaptation. Performance metrics should also be calculated according to the same information constraints that would apply in live operation.

### 7.3 Model Governance and Extreme Conditions

Adaptive models change over time, creating governance challenges that are less prominent in fixed models. A financial institution or research environment must record which model version produced each risk score, what data triggered adaptation, and whether changes improved or degraded subsequent performance. Continuous learning without auditability can make model behavior difficult to reproduce or validate.

Extreme financial conditions present an additional problem because rare events may differ substantially from all previously observed data. Machine learning cannot reliably infer patterns that have no meaningful representation in its historical information. Stress testing, scenario analysis, and conservative risk controls therefore remain important even when adaptive prediction is available. Automated learning should supplement rather than eliminate broader risk-governance procedures.

## 8. Conclusion

Real-time financial risk assessment using streaming data requires models that can process observations sequentially while preserving strict temporal integrity. The present hypothetical study compared static logistic regression, incremental Hoeffding-tree classification, and an adaptive online ensemble incorporating concept-drift detection. The methodology demonstrates how continuous updating and regime-specific evaluation can be incorporated within a financial risk-classification framework.

The illustrative analysis assigned the strongest overall performance to the adaptive ensemble, with accuracy of 91.50% and recall of 90.00% under the fictional test stream. Performance temporarily declined during gradual and abrupt distributional changes before improving after adaptation. These values are completely hypothetical and do not demonstrate forecasting effectiveness in any real market, portfolio, or institution.

The study concludes that adaptive machine learning is methodologically promising for continuously changing financial data but introduces additional requirements involving drift detection, delayed supervision, temporal validation, model governance, and stability under extreme conditions. Future empirical research should compare adaptive learners with established financial-risk models, examine calibrated risk probabilities, evaluate multiple real market regimes, and document model changes comprehensively before operational deployment.

## References

1. Bahnsen, A. C., Aouada, D., Stojanovic, A., & Ottersten, B. (2015). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications*, 51, 134–142. <https://doi.org/10.1016/j.eswa.2015.12.030>
2. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602–613. <https://doi.org/10.1016/j.dss.2010.08.008>
3. Bifet, A., & Gavaldà, R. (2007). Learning from time-changing data with adaptive windowing. In *Proceedings of the 2007 SIAM International Conference on Data Mining* (pp. 443–448). Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611972771.42>
4. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–255. <https://doi.org/10.1214/ss/1042727940>
5. Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182–194. <https://doi.org/10.1016/j.inffus.2017.09.005>
6. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2015). Credit card fraud detection: A realistic modeling and a novel learning strategy. In *2015 IEEE International Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. <https://doi.org/10.1109/IJCNN.2015.7280527>
7. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications*, 41(10), 4915–4928. <https://doi.org/10.1016/j.eswa.2014.02.026>
8. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), Article 44. <https://doi.org/10.1145/2523813>
9. Gomes, H. M., Bifet, A., Read, J., Barddal, J. P., Enembreck, F., Pfharinger, B., Holmes, G., & Abdessalem, T. (2017). Adaptive random forests for evolving data stream classification. In *Proceedings of the 2017 IEEE 16th International Conference on Data Mining* (pp. 443–448). IEEE. <https://doi.org/10.1109/ICDM.2017.54>
10. Hand, D. J. (2009). Measuring classifier performance: A coherent alternative to the area under the ROC curve. *Machine Learning*, 77, 103–123. <https://doi.org/10.1007/s10994-009-5119-5>
11. Krawczyk, B. (2016). Learning from imbalanced data: Open challenges and future directions. *Progress in Artificial Intelligence*, 5, 221–232. <https://doi.org/10.1007/s13748-016-0094-0>

12. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1038/nature14539>
13. Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346–2363. <https://doi.org/10.1109/TKDE.2018.2876857>
14. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* 30 (pp. 4765–4774). Curran Associates.
15. Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *arXiv*. <https://arxiv.org/abs/1009.6119>
16. Sahin, Y., Bulkan, S., & Duman, E. (2013). A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications*, 40(15), 5916–5923. <https://doi.org/10.1016/j.eswa.2013.05.021>
17. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), Article e0118432. <https://doi.org/10.1371/journal.pone.0118432>
18. Su, R., Guo, H., & Wang, W. (2024). Elastic online deep learning for dynamic streaming data. *Information Sciences*, 676, Article 120799. <https://doi.org/10.1016/j.ins.2024.120799>
19. Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence* (pp. 159–166). IEEE. <https://doi.org/10.1109/SSCI.2015.33>
20. Zliobaite, I. (2010). Learning under concept drift: An overview. Vilnius University.