

MODALITY RELIABILITY AS A UNIFIED CONTROL INTERFACE FOR STREAMING MULTIMODAL DECISION SUPPORT PIPELINES

Uzun Illia

Postgraduate student

Institute of Artificial Intelligence and Robotics
National University «Odessa Polytechnic», Ukraine

Streaming decision support systems increasingly rely on multimodal time-series data, where both non-stationarity (concept drift) and temporary modality degradation (missingness, noise, scale shifts) may co-occur and confound the interpretation of error growth. This work proposes to represent modality reliability as a causal probabilistic estimate of a modality being in a non-degraded state and to treat this estimate as a unified control interface that consistently governs three online procedures: reliability-adaptive fusion for forecasting, drift-initiated budgeted micro-adaptation, and false-alarm-rate (FAR) constrained anomaly detection. The approach is evaluated in a prequential protocol on controlled streams with deterministic injections of drift, degradations, and anomalies, and is additionally validated on real datasets. Reported results indicate that calibrated reliability scores provide an interpretable control signal with low latency and improve robustness of downstream modules under stress regimes.

In streaming intelligent decision support systems, the same rise in forecasting error can be caused either by a genuine change of regime (concept drift) or by a temporary deterioration of one or more data sources (e.g., sensor outages, increased noise, scale shifts). Without an explicit representation of modality quality, online pipelines tend to overreact (unnecessary adaptation) or underreact (late recovery), while anomaly detection may produce systematic false alarms during degraded segments. These issues become critical when the pipeline must satisfy low-latency constraints and an explicit bound on the false alarm rate (FAR) in operational monitoring.

The proposed idea is to introduce a single probabilistic quantity - modality reliability - and use it as a shared control signal across the pipeline. Reliability is interpreted as a causal probability score (between 0 and 1) indicating that a given modality is currently in a non-degraded state (given only past and current observations). To support interpretability and stable downstream control, the reliability scale is post-hoc calibrated and monitored using the expected calibration error (ECE). At the pipeline level, reliability is used in three roles: (i) to transform per-modality scores into smoothed dynamic fusion weights for streaming forecasting (reliability-adaptive dynamic fusion, RADF); (ii) to trigger event-driven micro-adaptation only when changes are consistent with drift rather than with data-quality degradation, while

keeping update costs within a fixed budget; (iii) to constrain anomaly scoring and thresholding so that the decision-making layer can maintain a predefined FAR budget.

Evaluation is performed in a prequential (test-then-train) setting on controlled streams where degradations, drift, and anomalies are injected deterministically to enable causal interpretation of effects. In these experiments, the calibrated online reliability model separates degraded and non-degraded states with ROC-AUC 0.86 ± 0.07 and achieves ECE 0.08 ± 0.04 , providing a stable probabilistic control scale. In a microbenchmark, lightweight reliability scales require $9.29 \mu\text{s}$ and $19.19 \mu\text{s}$ on average, while the online model achieves an average latency of approximately $150 \mu\text{s}$. For anomaly detection with a fixed FAR budget of 0.05, the reliability-constrained procedure attains $\text{Recall@FAR} = 0.34 \pm 0.09$ on degraded segments of a controlled stream, compared to 0.10 ± 0.06 for naive multimodal scoring. In a system-level stress test with segment missingness rate 0.90, an integrated decision support prototype reduces MAE by approximately 35% on a controlled energy domain and by approximately 93% on real BTC/USD data compared to naive fusion.

Table 1. Key empirical findings for reliability as a control interface (streaming evaluation).

Aspect	Result (mean $\pm$ dispersion, where reported)
Reliability separability (degradation vs. clean)	ROC-AUC 0.86 ± 0.07
Reliability calibration quality	ECE 0.08 ± 0.04
Reliability latency (microbenchmark)	$\approx 150 \mu\text{s}$ (online model); $9.29 \mu\text{s}$, $19.19 \mu\text{s}$ (lightweight scales)
Anomaly detection under FAR budget 0.05 (degraded segments)	$\text{Recall@FAR} = 0.34 \pm 0.09$ vs. 0.10 ± 0.06 (naive scoring)
System-level stress (segment missingness 0.90)	MAE reduction $\approx 35\%$ (controlled energy), $\approx 93\%$ (real BTC/USD) vs. naive fusion

The presented results support the feasibility of treating reliability as a single, auditable control interface for streaming multimodal decision support. At the methodological level, the interface provides a consistent way to distinguish data-quality degradations from drift-related regime changes, reducing the risk of unstable updates and unnecessary alarms. At the engineering level, the interface enables transparent budget accounting (latency, update budgets, FAR constraints) and can be integrated into a decision layer that generates operator-facing recommendations.




Figure 1. Reliability as a unified control interface in a streaming multimodal decision-support pipeline.

References

1. Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2), 423–443. <https://doi.org/10.1109/TPAMI.2018.2798607>
2. Lahat, D., Adali, T., & Jutten, C. (2015). Multimodal data fusion: An overview of methods, challenges, and prospects. *Proceedings of the IEEE*, 103(9), 1449–1477. <https://doi.org/10.1109/JPROC.2015.2460697>
3. 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>
4. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning* (Vol. 70, pp. 1321–1330). PMLR. <https://proceedings.mlr.press/v70/guo17a.html>
5. Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874. <https://doi.org/10.1016/j.patrec.2005.10.010>
6. Dawid, A. P. (1984). Statistical theory: The prequential approach. *Journal of the Royal Statistical Society. Series A (General)*, 147(2), 278–292. <https://doi.org/10.2307/2981683>

МЕТОД КІЛЬКІСНОЇ ОЦІНКИ ЗБУРЕНЬ КОНТЕЙНЕРА

Борисенко Ірина Іванівна

старший викладач

Кафедра кібербезпеки та програмного забезпечення
Національний університет «Одеська політехніка», Україна

В комплексних системах захисту інформації широке застосування знайшли стеганографічні методи, принцип роботи яких полягає в створенні скритих каналів зв'язку у вже існуючих потоках даних в інформаційно-телекомунікаційних системах.

В якій би області (спектральній, просторовій і т.д.) не проводилося стеганоперетворення (СП) матриці контейнера це обов'язково призведе до зміни значень елементів в просторовій області. Аналіз збурень контейнера показав, що в залежності від того значення яких елементів були модифіковані рівень збурення матриці контейнера буде неоднаковим. Розглянемо це ствердження у формальному вигляді.

В якості контейнера будемо використовувати цифрове зображення (ЦЗ) з матрицею $C = [c_{ij}]$ розмірності $m \times n$, $c_{ij} \in [0; 255]$. Розглянемо два пікселі зі значеннями K_1 і K_2 таких, що $K_1 > K_2$. Змінимо ці значення на +1 кожне (тобто