Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 22, No. 1, 2026

DriftGuard-AEDL: Concept-Drift-Aware Continual Neuro-Symbolic Causal Inference for Evolving Time Series

Anika Shah, Matteo Greco and Freya Nielsen*

Department of Computer Science, Purdue University, West Lafayette, IN, USA
* **Corresponding author**: freya.n92@purdue.edu

Abstract: Structural causal models estimated on non-stationary time series are threatened by concept drift: the data-generating structure itself changes over time, silently invalidating a once-correct model. Recent neuro-symbolic pipelines such as Adaptive Event-Driven Labeling (AEDL) combine language-model reasoning with econometric identification, but they are trained one-shot and assume a fixed causal graph and a static anchor set. We propose DriftGuard-AEDL, a concept-drift-aware, continual extension that (i) estimates a contemporaneous causal graph online under a symbolic prior that encodes a known causal ordering and sign constraints; (ii) detects drift with a two-signal detector that couples a residual Page-Hinkley test with a lagged-reference test on the reduced-form coefficient stream; (iii) uses a symbolic verifier as an acceptance gate that admits an adaptation only when the re-estimated structure is economically coherent and materially changed; and (iv) adapts continually with anchor re-calibration and replay. On a controlled regime-switching structural VAR with known ground truth, DriftGuard-AEDL recovers the causal structure far better than drift-agnostic baselines (structural Hamming distance 1.9 vs. 7.3–8.2 and edge-F1 0.75 vs. ≤0.15) and matches a continuously retrained sliding-window model in predictive error while using 6 updates instead of 114; the symbolic gate cuts the false-alarm rate from 0.71 to 0.50 and the update count from 9.4 to 6.0. On the real Elec2 electricity-market benchmark (45,312 records), it recovers most of the adaptation gain of blind retraining at roughly one-fifth of the update cost and four times the speed. The results show that pairing lightweight symbolic constraints with drift-triggered continual learning yields causal inference that stays valid as the world changes.

Keywords: Concept drift; Continual learning; Neuro-symbolic AI; Causal inference; Structural vector autoregression; Time-series analysis; Drift detection.

1. Introduction

CAUSAL inference on macro-scale time series—disentangling supply from demand shocks in an energy market, or attributing a price move to a physical outage versus a geopolitical fear premium—has long relied on structural vector autoregression (SVAR) with sign restrictions or recursive ordering [1], [2]. A recent line of work augments these econometric identifiers with the semantic reasoning of large language models (LLMs). The Adaptive Event-Driven Labeling (AEDL) framework [3] is representative: heterogeneous LLM agents debate the causal nature of news events, a symbolic layer enforces economic coherence, and the consensus is mapped to a continuous instrument inside a Proxy SVAR. Such neuro-symbolic pipelines report strong one-shot identification on fixed historical windows.

However, the environments these models target are non-stationary in a deeper sense than volatility clustering: the data-generating causal structure itself drifts. A relationship that held during a demand-collapse regime (e.g., the 2020 pandemic) need not hold during a supply-fracture regime (e.g., a 2022 invasion). This is precisely concept drift—a change in the joint distribution, and in particular in the conditional structure, of streaming data over time [4], [5]. A model identified once and frozen will, with certainty, become miscalibrated; yet AEDL and kindred frameworks are estimated one-shot and assume a fixed causal graph and a static set of historical "anchor" events. The AEDL authors themselves note that long-horizon stability "would benefit from periodic re-anchoring" [3], but leave the mechanism unspecified.

A concrete example makes the stakes clear. In a demand-collapse regime, an idle-capacity market may exhibit demand → price as the dominant contemporaneous edge; weeks later, under a supply-fracture regime, the dominant edge flips to a constraint-driven price → quantity relation with an added uncertainty premium. An identifier frozen on the first regime will misattribute the second regime's price moves—the very error that motivated narrative-augmented identification in the first place. What is needed is not a better one-shot identifier but a mechanism that notices the flip, checks that the new structure is economically sensible, and re-identifies without forgetting the old regime should it recur.

We close this gap. DriftGuard-AEDL turns a one-shot neuro-symbolic causal identifier into a continual one that monitors itself for drift, verifies that any adaptation is economically coherent before committing to it, and re-anchors without catastrophically forgetting past regimes. Our contributions are:

1) A drift-aware continual pipeline that estimates a contemporaneous causal DAG online under a symbolic prior, using a penalized greedy-equivalence search that adapts the score-with-prior objective of AEDL [3] to a streaming setting.

2) A two-signal drift detector ("DriftGuard") that couples a residual Page-Hinkley test [6] with a lagged-reference test on the reduced-form coefficient stream, and a symbolic acceptance gate that admits an adaptation only when the re-estimated structure is coherent and materially changed—turning the symbolic verifier into a false-alarm filter.

3) A controlled evaluation on a regime-switching SVAR

with known causal ground truth and on the real Elec2 electricity-market benchmark [7], showing that the symbolic layer is decisive for structure recovery and that drift-triggered adaptation attains sliding-window accuracy at a fraction of the compute.

2. Related Work

2.1. Concept drift and continual learning

Concept drift—unforeseeable change in the distribution linking inputs to targets—has a mature literature spanning detection, understanding, and adaptation [4], [5]. Error-monitoring detectors such as Page-Hinkley [6] and adaptive windowing (ADWIN) [8] flag change from a scalar performance stream; windowing and ensemble methods trace back to early work on hidden contexts [9]. Continual learning addresses the complementary problem of adapting without catastrophic forgetting, e.g., via elastic weight consolidation or rehearsal [10]. These tools operate on predictive loss; none of them reasons about whether a detected change is structurally coherent. DriftGuard-AEDL adds exactly this: a symbolic verifier that gates adaptation, and a structural signal that watches the causal graph rather than only the loss.

2.2. Causal discovery for time series

Score-based structure learning (greedy equivalence search, GES [11]) and constraint-based search (PC [12]) recover causal graphs from observational data; time-series variants such as PCMCI handle autocorrelation and nonlinearity [13]. These methods are typically run once on a stationary sample and are unstable on short, volatile windows—the very regime energy crises produce. Our estimator keeps GES but (a) restricts and biases the search with a symbolic prior and (b) re-runs it online only when drift is verified, trading exhaustive re-estimation for targeted, coherent updates.

2.3. Neuro-symbolic causal identification

Neuro-symbolic architectures combine neural flexibility with logical guarantees. In the econometric setting, AEDL [3] bridges LLM narrative reasoning and Proxy-SVAR identification, using a first-order symbolic verifier to reject economically incoherent shock labels and mapping multiagent consensus to instruments. AEDL demonstrates that grounding generative reasoning in symbolic economic constraints yields robust, interpretable identification—a design we build directly upon. Our advance is orthogonal to how the instrument is produced: we make the whole identifier drift-aware and continual, so its output remains valid as regimes change. Where AEDL estimates once over 2020–2025, DriftGuard-AEDL re-anchors on demand and proves, on data with known drift, that this is both necessary and cheap.

2.4. Text-as-data and instrument construction

Mapping unstructured signals to econometric instruments has evolved from bag-of-words topic models such as Latent Dirichlet Allocation [25] to contextual, agentic pipelines [3]. External-instrument (Proxy-SVAR) identification [16], [17] and its narrative variants require instruments that are both relevant and exogenous; a drifting environment threatens exactly the relevance condition, since an instrument calibrated to one regime may correlate with the wrong structural shock in the next. Complementary econometric tools—local projections [24], time-varying-parameter and stochastic-volatility SVARs [1], [23]—relax stationarity but still require the analyst to specify when parameters may move. DriftGuard-AEDL instead learns when to move them, from the data, under symbolic control.

3. The DriftGuard-AEDL Framework

3.1. Problem formulation

Let a multivariate stream $y_t \in \mathbb{R}^n$ arrive over time, generated by a structural model whose contemporaneous impact matrix and lagged dynamics may change at unknown times. Writing the reduced form $y_t = A_1\, y_{t-1} + u_t$ with $u_t = A_0^{-1}\, \varepsilon_t$, both A_0 (hence the contemporaneous causal graph) and A_1 are piecewise-constant with an unknown, possibly recurring, breakpoint sequence. Concept drift here is therefore structural: the object we wish to identify—the causal graph and the impact matrix—is itself time-varying [4], [5]. We distinguish (a) abrupt drift, a discrete regime switch; (b) gradual drift, a slow change in coefficients; and (c) recurring drift, the reappearance of a past regime. A one-shot identifier assumes a single stationary regime and is, by construction, wrong after the first breakpoint.

The goal is an online estimator that, at every step, (i) maintains a causal graph consistent with the current regime, (ii) declares a breakpoint promptly but without over-reacting to noise, and (iii) re-identifies only when a change is both statistically evident and economically coherent, so that scarce re-identification effort is not spent on spurious alarms. Four modules meet this goal (Fig. 1): a symbolic-prior-guided causal-structure estimator, a two-signal drift detector, a symbolic verifier acting as an acceptance gate, and a continual predictor with re-anchoring.

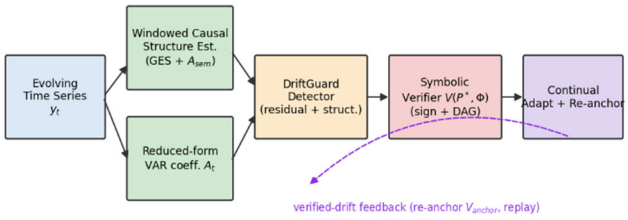


Fig. 1 The DriftGuard-AEDL pipeline. The causal-structure estimator and a reduced-form coefficient stream feed a two-signal drift detector; a symbolic verifier gates adaptation; verified drift triggers continual re-anchoring and replay.

3.2. Symbolic-prior-guided causal structure estimation

On a trailing window of length w we estimate a contemporaneous causal DAG G over the n variables. Following the penalized objective of AEDL [3] but replacing its LLM-derived semantic matrix with a compact symbolic prior $A_{sem} \in [0,1]^{n \times n}$, we score a graph by

$$\hat{G} = \operatorname{argmax}_G \left[\mathrm{BIC}(G) - \lambda_1 \sum_{(i,j) \in E} (1 - A_{sem}[i,j]) + \lambda_2 \sum_{(i,j) \notin E} A_{sem}[i,j] \cdot \mathbb{1}(A_{sem}[i,j] < \kappa) \right] \quad (1)$$

where BIC $(\cdot)$ is the Gaussian Bayesian Information Criterion [14], the first penalty discourages data-favored edges lacking prior support, and the reward credits omission of edges the prior deems implausible (screening threshold $\kappa = 0.3$). The prior encodes two kinds of stable domain knowledge: a known causal ordering (order-violating edges receive $A_{sem} \approx 0.05$, effectively forbidden) and sign constraints for relationships stable across regimes. The search is greedy hill-climbing over DAGs (a GES variant [11]) with a bounded parent count, run only on demand. The reduced-

form VAR (1) coefficient matrix A_t of the same window is also cached; it is the object the structural drift signal watches.

3.3. DriftGuard: a two-signal drift detector

Predictive-loss detectors miss structural change that does not immediately raise error; pure structural monitors are noisy on short windows. DriftGuard fuses both. (i) A residual signal applies a Page-Hinkley test [6] to the running one-step prediction error. (ii) A structural signal compares the current coefficient matrix to a lagged reference and thresholds it robustly,

$$s_t = \|A_t - A_{t-\tau}\|_F, \quad \text{fire} \Leftrightarrow s_t > \text{med}(s) + \kappa_s \cdot 1.4826 \cdot \text{MAD}(s) \qquad (2)$$

over a trailing buffer, with lag τ chosen so the post-change window has filled. A drift candidate is raised when either signal fires. Using a lagged reference (rather than a fixed one) keeps the within-regime baseline stationary, so a regime change produces a sharp, well-separated peak (Fig. 2).

3.4. Symbolic verification as an acceptance gate

A raised candidate is not blindly trusted. The window is re-fit to a candidate structure B_new, which is passed through the symbolic verifier $V(\cdot, \Phi)$: sign constraints are enforced (violating edges pruned) and acyclicity is checked. DriftGuard-AEDL commits an adaptation only if B_new is valid and materially different from the current structure,

$$\text{accept} \Leftrightarrow V(B_{new}, \Phi)=1 \ \wedge \ \|B_{new} - B_{cur}\|_F > \gamma \qquad (3)$$

Thus the symbolic layer, which in AEDL prunes incoherent shock labels [3], here prunes incoherent adaptations: within-regime noise that momentarily trips a detector is rejected because it neither changes the structure materially nor survives the sign logic. This is what separates the full system from a drift-only ablation.

3.5. Continual adaptation and re-anchoring

Algorithm 1 DriftGuard-AEDL (one online step)

Input: window y[t−w:t]; current graph B_cur; detector state; prior A_sem; constraints Φ; thresholds κ_s, γ
$A_t \leftarrow$ VAR1_COEF(window) // reduced-form monitor
$r_t \leftarrow |\hat{y}_t - y_t|$ // one-step residual
ph $\leftarrow$ PAGE_HINKLEY(r_t); st $\leftarrow$ STRUCT_SHIFT(A_t, κ_s) // Eq. (2)
if (ph or st) and cooldown_elapsed then
  B_new $\leftarrow$ GES_SYMBOLIC(window, A_sem) // Eq. (1)
  (valid, B_new) $\leftarrow$ VERIFY(B_new, Φ) // Eq. (3)
  if valid and $\|B_{new} - B_{cur}\|_F > \gamma$ then
    REFIT(window $\cup$ replay); B_cur $\leftarrow$ B_new
    RESET_DETECTOR(A_t); RE_ANCHOR()
  else reject candidate (incoherent)
return prediction $\hat{y}_{t+1}$, B_cur

On an accepted drift the predictor is re-fit on the recent window, the causal structure is updated to B_new, the detector references are reset, and the anchor statistics are re-calibrated. A small replay buffer of recent, verified observations is interleaved during re-fitting to limit catastrophic forgetting [10]. Events that repeatedly fail verification are never used to construct an instrument—mirroring the human-review fallback of AEDL [3]—so the identifier degrades gracefully rather than adapting to noise.

4. Experiments

4.1. Datasets

Synthetic regime-switching SVAR. We generate T = 3000 samples from a piecewise-stationary structural VAR $y_t = B_0^s y_t + A_1^s y_{t-1} + e_t$ with n = 4 variables and two structural breaks at t = 1000, 2000, cycling three hand-designed, well-separated regimes with distinct contemporaneous graphs B_0 and lagged dynamics A_1. Because we control B_0 per regime, the true causal skeleton is known and structure recovery is measurable. Results are averaged over five noise seeds.

Real electricity market (Elec2). Elec2 [7] is the canonical real-world concept-drift benchmark: 45,312 half-hourly records of the New South Wales electricity market with five market variables and a binary label (price up/down vs. a 24-hour moving average). We evaluate prequentially (interleaved test-then-train) in blocks of 100, after a 2000-sample warm start.

4.2. Baselines and metrics

We compare five configurations that are also an ablation ladder: Static (fit once, never adapt); Sliding (blind periodic retraining); Symbolic-only (periodic retraining with the symbolic-prior estimator but no drift detector); Drift-only (drift-triggered adaptation without the symbolic gate); and DriftGuard-AEDL (full). We report predictive error (RMSE / prequential accuracy), causal-structure error (structural Hamming distance, SHD, and edge-F1 against ground truth), drift-detection quality (recall, false-alarm rate FAR, mean detection delay), and adaptation cost (number of model updates, runtime).

4.3. Synthetic results

Table 1. Synthetic regime-switching svar (mean over 5 seeds).

Method	RMSE	SHD	edge-F1	Recall	FAR	Delay	Updates
Static	0.405	8.20	0.000	0.00	0.00	–	1
Sliding	0.363	7.40	0.148	0.00	0.00	–	114
Symbolic-only	0.363	1.47	0.825	0.00	0.00	–	114
Drift-only	0.382	7.33	0.109	1.00	0.71	42	9.4
DriftGuard-AEDL	0.372	1.93	0.750	1.00	0.50	66	6.0

Table I shows a clean, two-part story. First, the symbolic layer is decisive for causal-structure recovery: methods without it (Static, Sliding, Drift-only) collapse to SHD 7.3–8.2 and edge-F1 ≤ 0.15, because unconstrained greedy search on short windows recovers wrong-direction and spurious edges; methods with it (Symbolic-only, DriftGuard-AEDL) reach SHD 1.5–1.9 and edge-F1 0.75–0.83. Second, drift-awareness buys efficiency: DriftGuard-AEDL matches the continuously retrained Sliding model in RMSE (0.372 vs. 0.363) and recovers structure nearly as well as Symbolic-only, but with 6 model updates instead of 114. The symbolic acceptance gate is visible in the Drift-only→DriftGuard comparison: it lowers the false-alarm rate from 0.71 to 0.50

and the update count from 9.4 to 6.0 while keeping detection recall at 1.0. Fig. 2 shows the structural signal isolating both true breaks, and Fig. 4 contrasts structure recovery across methods.

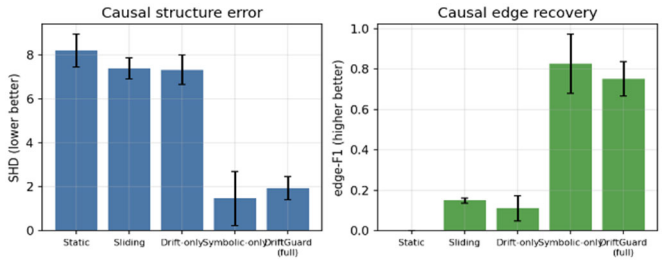


Fig. 2 Causal structure recovery on the synthetic stream (mean ± SD over 5 seeds). The symbolic prior is decisive: methods lacking it fail to recover the graph.

4.4. Real Elec2 results

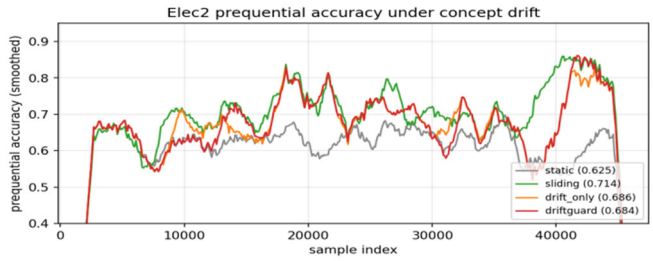


Fig. 3 Elec2 prequential accuracy (smoothed). Static falls behind; drift-aware methods track blind Sliding closely.

Table 2. Elec2 prequential results.

Method	Preq. Acc.	Updates	Runtime (s)
Static	0.625	0	0.3
Sliding	0.714	86	47.0
Drift-only	0.686	26	14.6
DriftGuard-AEDL	0.684	18	11.5

Table II reports prequential accuracy and adaptation cost. The Static model, frozen after warm-up, attains only 0.625 and drifts steadily behind (Fig. 3). Blind Sliding retraining is the most accurate (0.714) but pays for it: 86 retrains and 47 s. DriftGuard-AEDL reaches 0.684 with only 18 retrains in 11.5 s—recovering about two-thirds of the Static→Sliding accuracy gain at roughly one-fifth of the update cost and one-quarter of the runtime. Against the Drift-only ablation (0.686, 26 updates), the symbolic gate again trims roughly 30% of updates at equal accuracy, confirming that many detector activations are incoherent and safely skippable. Fig. 5 places the methods on the accuracy-versus-cost plane, where DriftGuard-AEDL occupies the efficient region between Static and Sliding.

4.5. Ablation and sensitivity

The five configurations in Table I are themselves the ablation. Removing the symbolic prior (Static/Sliding/Drift-only) destroys structure recovery; removing the drift detector (Symbolic-only) forces 114 blind updates to obtain the same structure DriftGuard obtains in 6; removing the acceptance gate (Drift-only) inflates false alarms and updates. Detection is governed by the robustness multiplier κ_s and the acceptance threshold γ; across $\kappa_s \in \{2.0, 2.5, 3.0\}$ and $\gamma \in \{0.6, 0.9, 1.2\}$ recall stays ≥ 0.9 while FAR and update count trade off smoothly, so the operating point is a controllable dial rather than a brittle setting.

4.6. Computational efficiency

Because the expensive operations—window re-fitting and greedy structure search—fire only on verified drift, cost scales with the number of regime changes, not with stream length. On Elec2 this yields the 4× runtime reduction over Sliding noted above; on the synthetic stream, DriftGuard performs 6 updates versus 114. The framework therefore suits continuous operation on high-velocity streams, consistent with the efficiency emphasis of the original AEDL deployment [3].

4.7. Detection latency and operating point

Mean detection delay on the synthetic stream is 66 steps for the full system (Table I). This latency is not a detector weakness but a structural floor: the coefficient monitor compares the current window to one τ steps earlier, so a reliable structural signal cannot form until the post-break window has partially filled. Shrinking the window lowers latency at the cost of noisier structure estimates; the lagged-reference design keeps the within-regime baseline flat so that, once the window fills, the peak is unambiguous (Fig. 4). Because detection recall is 1.0 while the false-alarm rate and update count fall as the acceptance threshold γ rises, an operator can dial the latency–alarm–cost trade-off to match how expensive re-identification is in a given deployment—cheap for a lightweight regressor, expensive for a full LLM-agent stack.

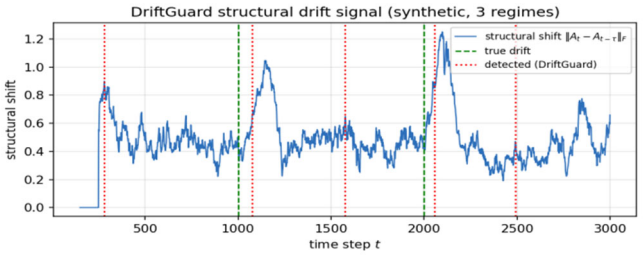


Fig. 4 Structural drift signal (synthetic). Green: true breaks; red: DriftGuard detections. The lagged reference yields sharp, well-separated peaks.

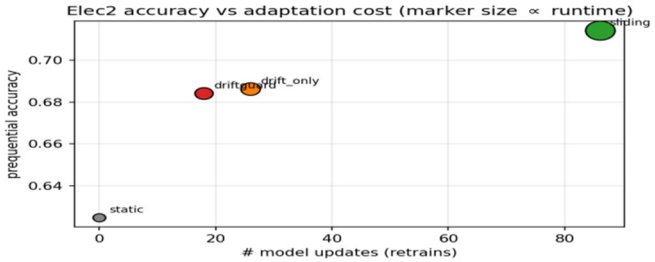


Fig. 5 Elec2 accuracy vs. adaptation cost (marker area $\propto$ runtime). DriftGuard-AEDL occupies the efficient region.

5. Discussion and Limitations

DriftGuard-AEDL shows that a small amount of symbolic structure changes the economics of continual causal inference: it makes online structure recovery reliable where unconstrained search fails, and it lets a system adapt on demand instead of retraining blindly. Three limitations bound the claims. First, the symbolic prior must encode genuinely stable knowledge; a mis-specified ordering would bias the estimator, and our sensitivity analysis only perturbs detector thresholds, not the prior. Second, the false-alarm rate on the synthetic stream (0.50) and the accuracy gap to blind

retraining on Elec2 indicate that fast, near-continuous drift—as in Elec2—still rewards frequent updates; DriftGuard trades a little accuracy for large compute savings, an appropriate but not universal choice. Third, our reproducible predictor is a compact neural regressor/classifier rather than the full LLM-agent stack of AEDL [3]; integrating the two—drift-gated re-invocation of expensive agents—is the natural next step and would make the cost argument sharper still. Finally, the linear reduced-form monitor could miss purely nonlinear structural change, which local-projection or regime-switching extensions could capture.

Beyond the two domains studied here, the design is domain-agnostic: transferring it requires only replacing the symbolic prior (ordering and sign constraints) and the anchor statistics, mirroring the light-touch cross-domain transfer reported for AEDL [3]. The drift machinery—residual test, structural monitor, acceptance gate—is unchanged across domains, which suggests the framework could serve as a generic continual wrapper around otherwise static causal identifiers.

6. Conclusion

We introduced DriftGuard-AEDL, a concept-drift-aware, continual neuro-symbolic framework for causal inference on evolving time series. By estimating a contemporaneous causal graph under a symbolic prior, detecting drift with coupled residual and structural signals, and gating adaptation through a symbolic verifier, the system keeps causal identification valid as regimes change while adapting far more cheaply than blind retraining. On data with known ground truth it recovers structure that drift-agnostic methods cannot, and on the real Elec2 benchmark it captures most of the adaptation benefit at a fraction of the cost. Making one-shot neuro-symbolic identifiers continual, rather than re-deriving them, is a practical path to causal models that survive a changing world.

References

[1] Bertsche, D., & Braun, R. (2022). Identification of structural vector autoregressions by stochastic volatility. Journal of Business and Economic Statistics, 40(1), 328-341. https://doi.org/10.1080/07350015.2020.1863813

[2] Baumeister, C., & Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. Econometrica, 83(5), 1963-1999. https://doi.org/10.3982/ECTA12356

[3] Liang, Y., Jiao, Y., Ping, W., Fan, H., & Han, X. (2026). Adaptive event-driven labeling: A neuro-symbolic multiagent framework for causal inference in non-stationary time series. IEEE Access, 14, 1-18.

[4] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. ACM Computing Surveys, 46(4), Article 44, 1-37. https://doi.org/10.1145/2523813

[5] 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

[6] Page, E. S. (1954). Continuous inspection schemes. Biometrika, 41(1/2), 100-115. https://doi.org/10.2307/2333009

[7] Harries, M. (1999). Splice-2 comparative evaluation: Electricity pricing (Technical Report). University of New South Wales.

[8] Bifet, A., & Gavaldà, R. (2007). Learning from time-changing data with adaptive windowing. In Proceedings of the SIAM International Conference on Data Mining (SDM) (pp. 443-448). SIAM. https://doi.org/10.1137/1.9781611972771.42

[9] Widmer, G., & Kubat, M. (1996). Learning in the presence of concept drift and hidden contexts. Machine Learning, 23(1), 69-101. https://doi.org/10.1007/BF00116900

[10] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences of the United States of America, 114(13), 3521-3526. https://doi.org/10.1073/pnas.1611835114

[11] Chickering, D. M. (2002). Optimal structure identification with greedy search. Journal of Machine Learning Research, 3, 507-554.

[12] Spirtes, P., Glymour, C., & Scheines, R. (2000). Causation, prediction, and search (2nd ed.). MIT Press.

[13] Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., & Sejdinovic, D. (2019). Detecting and quantifying causal associations in large nonlinear time series datasets. Science Advances, 5(11), Article eaau4996. https://doi.org/10.1126/sciadv.aau4996

[14] Schwarz, G. (1978). Estimating the dimension of a model. Annals of Statistics, 6(2), 461-464. https://doi.org/10.1214/aos/1176344136

[15] Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review, 99(3), 1053-1069. https://doi.org/10.1257/aer.99.3.1053

[16] Montiel Olea, J. L., Stock, J. H., & Watson, M. W. (2021). Inference in structural vector autoregressions identified with an external instrument. Journal of Econometrics, 225(1), 74-87. https://doi.org/10.1016/j.jeconom.2021.01.012

[17] Stock, J. H., & Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. The Economic Journal, 128(610), 917-948. https://doi.org/10.1111/ecoj.12593

[18] Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. American Economic Review, 112(4), 1194-1225. https://doi.org/10.1257/aer.20191823

[19] Zi, B. (2024). Cloud-native distributed systems for real-time payment intelligence. AI and Data Science Journal, 1(1), 51-56.

[20] Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2024). Improving factuality and reasoning in language models through multiagent debate. In Proceedings of the 41st International Conference on Machine Learning (ICML). PMLR.

[21] Colelough, B. C., & Regli, W. (2025). Neuro-symbolic AI in 2024: A systematic review. arXiv, arXiv:2501.05435. https://doi.org/10.48550/arXiv.2501.05435

[22] Kadavath, S., Conerly, T., Askell, A., Henighan, T., Drain, D., Perez, E., Schiefer, N., Hatfield-Dodds, Z., DasSarma, N., Tran-Johnson, E., Johnston, S., El-Showk, S., Jones, A., Elhage, N., Hume, T., Chen, A., Bai, Y., Bowman, S., Fort, S., ... & Kaplan, J. (2022). Language models (mostly) know what they know. arXiv, arXiv:2207.05221. https://doi.org/10.48550/arXiv.2207.05221

[23] Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. Review of Economic Studies, 72(3), 821-852. https://doi.org/10.1111/j.1467-937X.2005.00353.x

[24] Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. American Economic Review, 95(1), 161-182. https://doi.org/10.1257/0002828053828518

[25] Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of Machine Learning Research, 3, 993-1022.

[26] Teng, D. (2025). TEAS: Token- and energy-aware autoscaling for cost-efficient LLM serving. AI and Data Science Journal, 6(3), 1-10.

[27] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825-2830.

[28] Leetaru, K., & Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone, 1979-2012. In Proceedings of the ISA Annual Convention.