Graph Anomaly Detection as Finite-Horizon Control: Training-Free Scoring via Empirical Bayes

Fred Xu[1,*], Thomas Markovich[2], Florence Regol[3], Yizhou Sun[1]
[1]UCLA [2]Block, Inc. [3]Mila – Quebec AI Institute [*]Work done during an internship at Block, Inc.

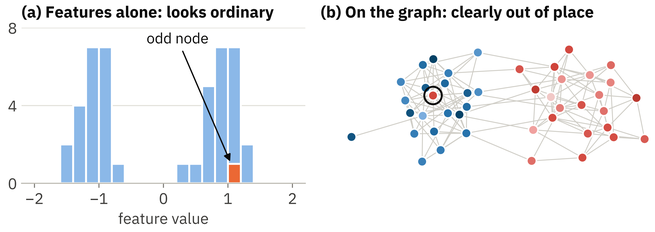
UCLA Samueli Computer Science

block

Motivation

Node-Level Graph Anomaly Detection

- Rank nodes whose features *and* neighborhoods deviate from the graph's regularities: fraud, fake reviews, bots
- **No labels**: anomalies are rare, and confirmed cases are scarce and arrive late



Choices Hidden in Prior Detectors

- **Graph trust**: how strongly should edges define normal? They reveal communities in homophilic graphs, but hide fraud under camouflage and heterophily
- Trust, spectral weighting and score choice are buried in architectures, losses and sweeps: scores are costly, opaque and unstable, and several baselines fall *below chance* on Elliptic and Facebook

Graph-Aware Prior

- **Template**: low-pass smoothing of the features; the residual is what the neighborhood does not explain

$$\mathbf{m} = (\gamma^2 I + L)^{-1} X, \qquad \Delta = X - \mathbf{m}$$

- **Precision**: graph Matérn blend; on Laplacian mode j (eigenvalue λ_j)

$$Q_\rho = \rho\,(\kappa^2 I + L) + (1-\rho)\,I, \qquad q_j = \rho\,(\kappa^2 + \lambda_j) + 1 - \rho$$

- **Three interpretable scalars**: graph trust ρ ($\rho = 0$ ignores the graph, $\rho = 1$ prices rough residuals most), inverse length-scale κ, template bandwidth γ

Empirical Bayes, No Training

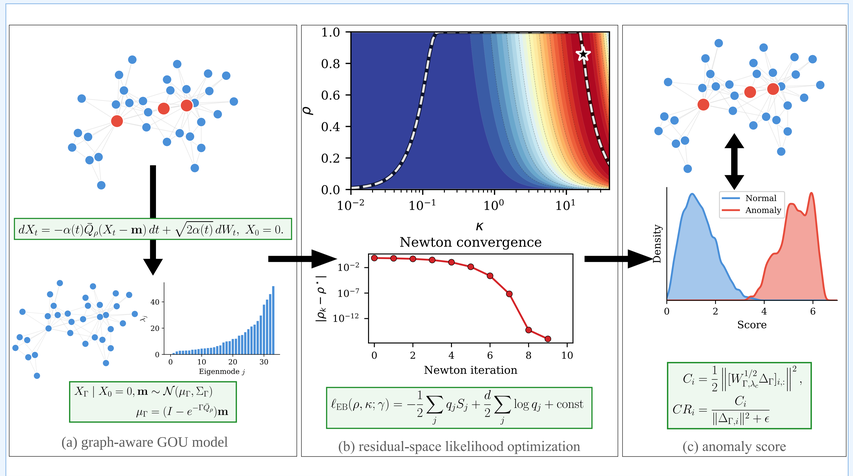
- Fit (ρ, κ, γ) on the unlabeled graph by the closed-form residual likelihood, with $S_j = \|\Delta^\top v_j\|^2$:

$$\log p(\mathcal{D} \mid \rho, \kappa, \gamma) = -\tfrac{1}{2}\textstyle\sum_j q_j S_j + \tfrac{D}{2}\sum_j \log q_j$$

- **Occam trade-off**: data fit against the Gaussian normalizer; anomalies are rare, so normal nodes dominate the fit
- **Cheap**: concave in ρ (clamped Newton), small grids over κ and γ
- **Training-free**: no neural parameters, gradient descent, labels or validation set

EB-GAD: Anomaly Scores as Finite-Horizon Control Energy

TL;DR: EB-GAD fits a graph Gaussian prior by empirical Bayes, reads it as the resting state of a graph-aware Ornstein–Uhlenbeck (GOU) process, and scores each node by the least effort needed to steer a neutral start to its features. No labels or training at any step

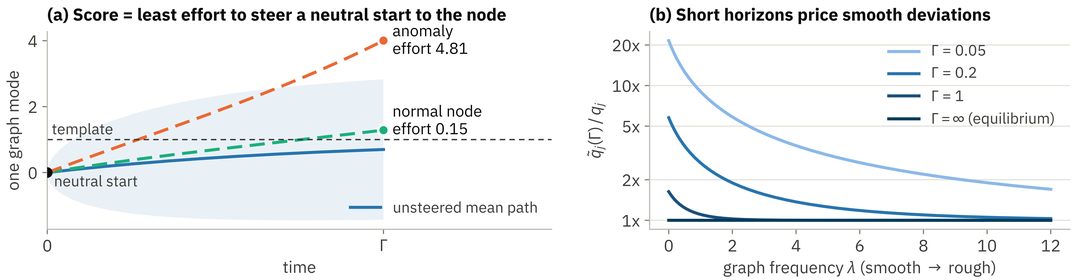


GOU dynamics: $dX_t = -\alpha(t)\,\bar{Q}_\rho\,(X_t - \mathbf{m})\,dt + \sqrt{2\alpha(t)}\,dW_t$ has stationary law $\mathcal{N}(\mathbf{m}, \bar{Q}_\rho^{-1})$; mode j relaxes at rate q_j (rough modes fast, smooth modes slowly) and from $X_0 = 0$ has variance $\sigma_j^2(\Gamma) = (1 - e^{-2\Gamma q_j})/q_j$ at horizon Γ

Control energy: started at $X_0 = 0$, the least steering effort to reach the observed x at horizon Γ is the negative log-density of the GOU endpoint

$$\mathcal{E}(0 \to x; \Gamma) = \tfrac{1}{2}(x - \mu_\Gamma)^\top \Sigma_\Gamma^{-1}(x - \mu_\Gamma) = -\log p(x \mid X_0 = 0, \Gamma) + \mathrm{const}$$

How Hard Is It to Reach a Node?



- **Per-mode price**: $\tilde{q}_j(\Gamma) = 1/\sigma_j^2(\Gamma) \geq q_j$; short horizons ask whether a node is surprising before its neighborhood has relaxed, and $\Gamma \to \infty$ recovers the equilibrium score $J_i^\star = \frac{1}{2}\|[Q_\rho^{1/2}\Delta]_{i,:}\|^2$
- **Endpoint tolerance**: $c_j = (\lambda_c^{-1} + \sigma_j^2(\Gamma))^{-1}$ tests whether the evidence survives soft matching

Score Bank and Label-Free Selector

- Sweeping (Γ, λ_c) gives a bank of scores that share one fitted prior, hence one plug-in null
- **Four families**: equilibrium anchor ($J^\star$, R), control energy (C), control ratio, and profile aggregation of per-node p-values across horizons
- **Control ratio**: $CR_i = C_i/(\|\Delta_{\Gamma,i}\|^2 + \epsilon)$ separates graph roughness from residual magnitude
- **Selector**: feature homophily, edge density and feature dimension pick the family; within it, fitted-null deviation plus neighboring-rank stability pick the score
- **Stable**: ±25% on each of its 21 cutoffs leaves the choice unchanged in 99.4% of evaluations

Main Results (AUROC, %)

	Dataset	Best of 9 baselines		EB-GAD	Δ
Standard GAD	Weibo	ANOMALOUS	94.5	**95.1**	+0.6
	Reddit	CoLA	60.3	**60.6**	+0.3
	Amazon	DIF	77.1	**78.0**	+0.9
	YelpChi	LOF	57.2	**71.7**	+14.5
	BlogCatalog	TAM	**82.5**	78.6	−3.9
	Facebook	TAM	**91.4**	**91.4**	0.0
	ACM	TAM	**88.8**	86.0	−2.8
Finance	Elliptic	CoLA	65.3	**74.5**	+9.2
	Elliptic++	CoLA	62.9	**72.5**	+9.6
	DGraph	DIF	64.2	**66.8**	+2.6
	T-Finance	CONAD	63.1	**84.8**	+21.7

- **Best or tied-best on 9 of 11**, with no label at any step; leads on all four financial networks (up to 3.7M nodes) by up to 21.7 points
- **Trails TAM** on BlogCatalog and ACM
- **Anchor bandwidth**: Weibo and Reddit use a rule adopted after the first evaluation; without it they are 92.5 and 48.5 (7 of 11)

What the Finite Horizon Adds

Dataset	Family	Equil.	Selected	Δ
T-Finance	profile	80.1	84.8	+4.7
ACM	ratio	81.3	86.0	+4.7
Facebook	profile	87.8	91.4	+3.6
BlogCatalog	energy	78.4	78.6	+0.2
Amazon	energy	77.9	78.0	+0.1

- Same fitted model, equilibrium score vs. selected score; the other six datasets are routed to the anchor
- **Scalable**: 8 s (Facebook) to 33 min (DGraph) on 16 CPU threads; the trained baselines that fit in memory on DGraph need 65–69 min on a GPU