

Graph-Driven Portfolio Optimization: EI-Minimizing GAT, Pattern Recognition with FFT, and HMM Volatility Regimes



Anika Dhuri, Valentina Montesano, Miao Wang, Professor Ali Hirsra
Columbia University

Introduction and Background

Financial markets are **non-stationary** and exhibit regime shifts, volatility clustering, and time-varying correlations—motivating tools that can summarize performance comprehensively and extract robust signals from structure, spectra, and states.

The following components constitute the prior foundation on which our study is built.

Explainability Index (EI). Prior work introduces EI as a composite, bounded $[0, 1]$ score that aggregates multiple performance measures. Raw metrics are first mapped using quantile-based linear scaling and a sigmoid transformation, then combined via a weighted distance metric; lower EI indicates more desirable outcomes.

Graph Attention Network: We used a previously-implemented Graph Attention Network for portfolio construction that **maximized Sharpe ratio**: node/edge features are normalized, a learned adjacency supports message passing, and a portfolio head outputs long-only weights used to compute portfolio returns.

Volatility-regime detection: A Bayesian supervised HMM was trained on labeled regimes; **forward-backward** inference produced state posteriors and point estimates on both train and validation sets.

Objective

Goal: align training and allocation with an **interpretable, multi-metric objective**, while making weights **pattern-aware** and **regime-adaptive**.

1

Whole-Graph GAT → EI-minimizing portfolio head

Objective: replace the max-Sharpe loss with **Explainability Index (EI)** minimization so the model learns weights that balance performance and for smoother, more interpretable allocations and lower drawdowns/turnover sensitivity.

2

FFT Pattern Recognition → spectral features to weights

Objective: convert windowed FFT features into **allocation scores** and normalize to a portfolio, targeting patterns that persist across windows rather than short-horizon noise.

3

Supervised HMM Volatility Regimes → probabilistic blending


Objective: use **posterior regime probabilities** from the supervised HMM to **blend regime-conditional portfolios** to improve stability in turbulent periods while remaining responsive to volatility shifts.

Data

Our original dataset contained daily data from 15 benchmark identifiers. Data was converted to **log-returns**. Data was cleaned by replacing $\pm\infty$ by NaN, dropping all-NaN columns, and removing columns with $>40\%$ missing. The models were evaluated on rolling windows of a **5-year train period** and **1-year test period**, sliding every **30 days**. The windows used span from **2009 to 2016**.

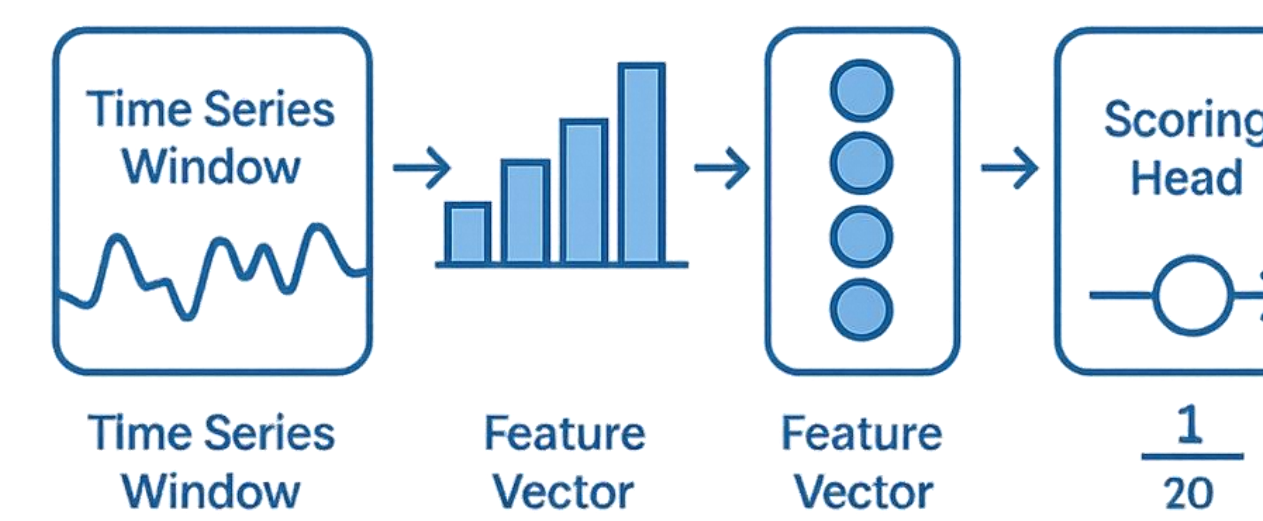
Overview of Model Architecture

GAT (EI-Minimization)



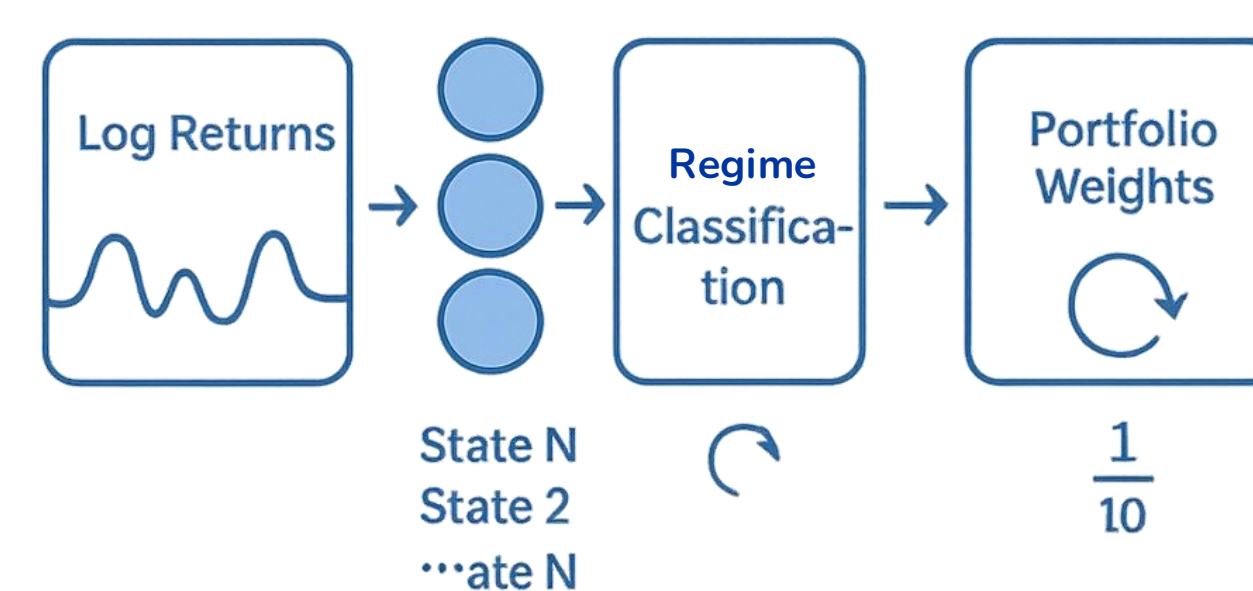
- Data prep.** Compute log-returns R_{train}, R_{test} . Build feature matrix X_{train} . Construct learned adjacency A_{train} . Initialize equal base-weights $w_0 = \frac{1}{N} \mathbf{1}$.
- GAT model.** Instantiate.
- EI metric selection.** From the metric set, select $\{maxDD, ann. mean, ann. vol\}$.
- Training objective (EI-min).** Train the GAT with `train_model_ei` on $(R_{train}, X_{train}, A_{train})$, to minimize a composite *Explainability Index (EI)* over chosen metrics.
- Final weights and returns.** Extract portfolio weights \hat{w} from the trained GAT. Compute portfolio returns: $r^{train} = R_{train} \hat{w}$, $r^{test} = R_{test} \hat{w}$.
- Rolling EI (post-hoc).** For each rolling window of length L , aggregate portfolio EI by weighted average $EI_L = \hat{w}^T EI_L^{(assets)}$.

FFT Pattern Recognition



- Preprocess.** Log-returns $R_{train}, R_{test} \in \mathbb{R}^{T \times N}$.
- Rolling FFT features.** For each window end $t = L, \dots$, average across assets, apply FFT, keep first K real/imag parts to form $f_t \in \mathbb{R}^{2K}$.
- Regime discovery.** Fit a C -component GMM on $\{f_t\}$ (train) to obtain hard labels z_t and soft posteriors $\pi_t \in \Delta^{C-1}$.
- Per-regime portfolios.** For each regime c , collect train rows with $z_t = c$, compute per-asset std $\sigma_{c,i}$, set inverse-vol weights $\tilde{w}_{c,i} = 1/\sigma_{c,i}$, then normalize $w_c = \tilde{w}_c / (\mathbf{1}^T \tilde{w}_c)$.
- Returns & blending.** Train: $r_t^{train} = R_{train}[t,:]\tilde{w}_{z_t}$. Test: get w_c^{test} from the trained GMM on test FFT features, blend $W_u = \sum_{c=1}^C \pi_{c,t}^{test} w_c$ (normalize), and compute $r_t^{test} = R_{test}[t,:]\tilde{w}_u$.
- Final train weights.** $w_{final}^{train} = \sum_c \pi_{c,t}^{train} w_c$ (normalize).
Outputs: $w_{final}^{train} \in \mathbb{R}^N$, r^{train} , r^{test} , train labels $\{z_t\}$, and test posteriors $\{\pi_t^{test}\}$.

Hidden Markov Model



- Normalize driver.** Compute S&P 500 TR log-returns $r_t = \log P_t - \log P_{t-1}$ and z-score; split into train/test.
- Label regimes (train).** Extract volatility, change-points, and attributes; cluster (≤ 5) to obtain train regime labels z_t .
- Bayesian fit.** Fit a supervised HMM to estimate transition matrix P and Gaussian emissions $\{(\mu_k, \sigma_k^2)\}_{k=1}^C$.
- Inference.** Forward-backward produces hard train states z_t and test posteriors $\pi_t = p(s_t = k | data_t)$.
- Per-regime portfolios.** For each regime k , compute ridge MV weights $w_k \propto (\Sigma_k + \lambda I)^{-1} \mu_k$, $\sum_k w_{k,i} = 1$.
- Test-time blending.** Blend regime portfolios by posteriors:
$$W_t = \sum_{k=1}^C \pi_{k,t} w_k$$
- Returns.** Train: $r_t^{train} = \mathbf{x}_t^T w_{z_t}$, Test: $r_t^{test} = \mathbf{x}_t^T W_t$.

Average Portfolio Weighting Per Strategy

Benchmark Identifier Name	equal_weights	mean_var_opt	gat_model_sharpe_obj	gat_model_ei_obj	pattern_recognition_fft	vol_regime_supervised_hmm
A001	0.0909	0.1000	0.0026	0.4383	0.0120	0.0305
A002	0.0909	0.1000	0.0514	0.0010	0.0741	0.0276
A003	0.0909	0.1000	0.8043	0.0010	0.5285	0.6433
A004	0.0909	0.1000	0.0008	0.5491	0.0463	-0.0416
A006	0.0909	0.1000	0.1089	0.0010	0.2188	0.3499
A008	0.0909	0.1000	0.0248	0.0016	0.0493	0.0271
A009	0.0909	0.1000	0.0043	0.0038	0.0312	-0.0016
A011	0.0909	0.1000	0.0008	0.0013	0.0110	-0.0325
A012	0.0909	0.1000	0.0008	0.0010	0.0091	-0.0045
A013	0.0909	0.0000	0.0008	0.0010	0.0099	-0.0130
A015	0.0909	0.1000	0.0008	0.0010	0.0099	0.0147

Results

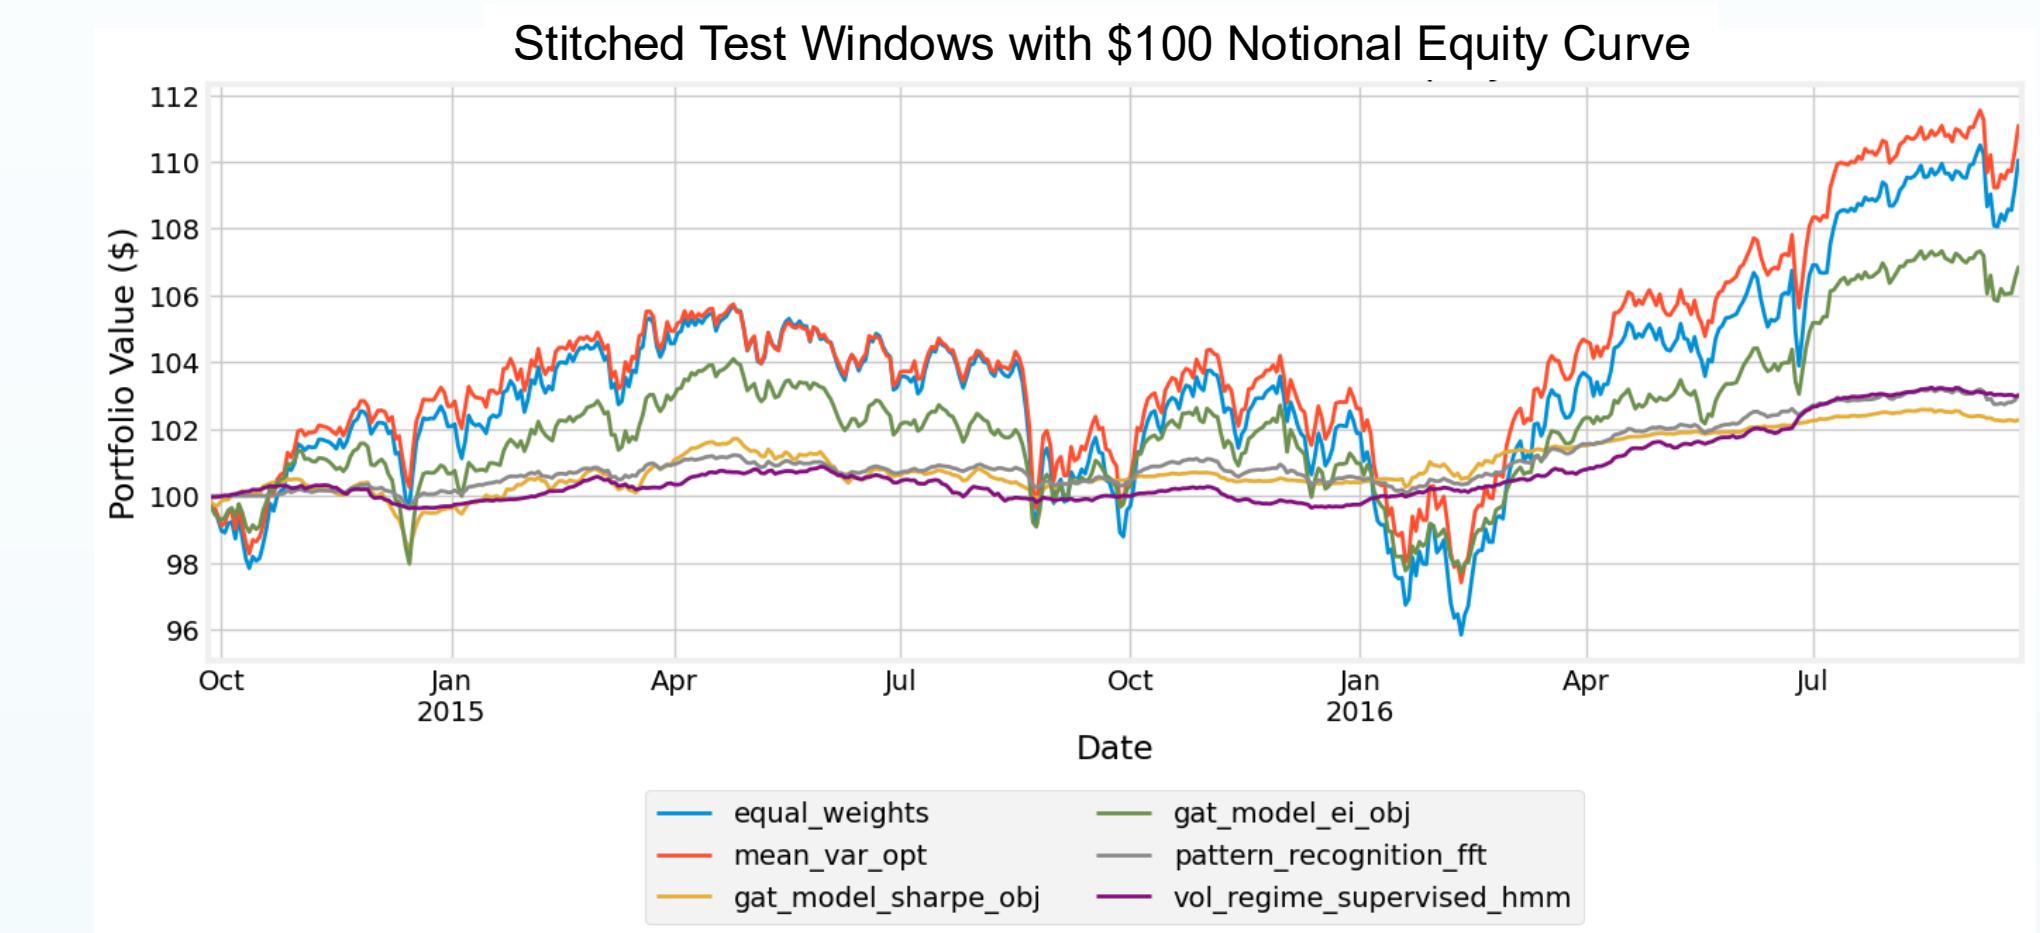
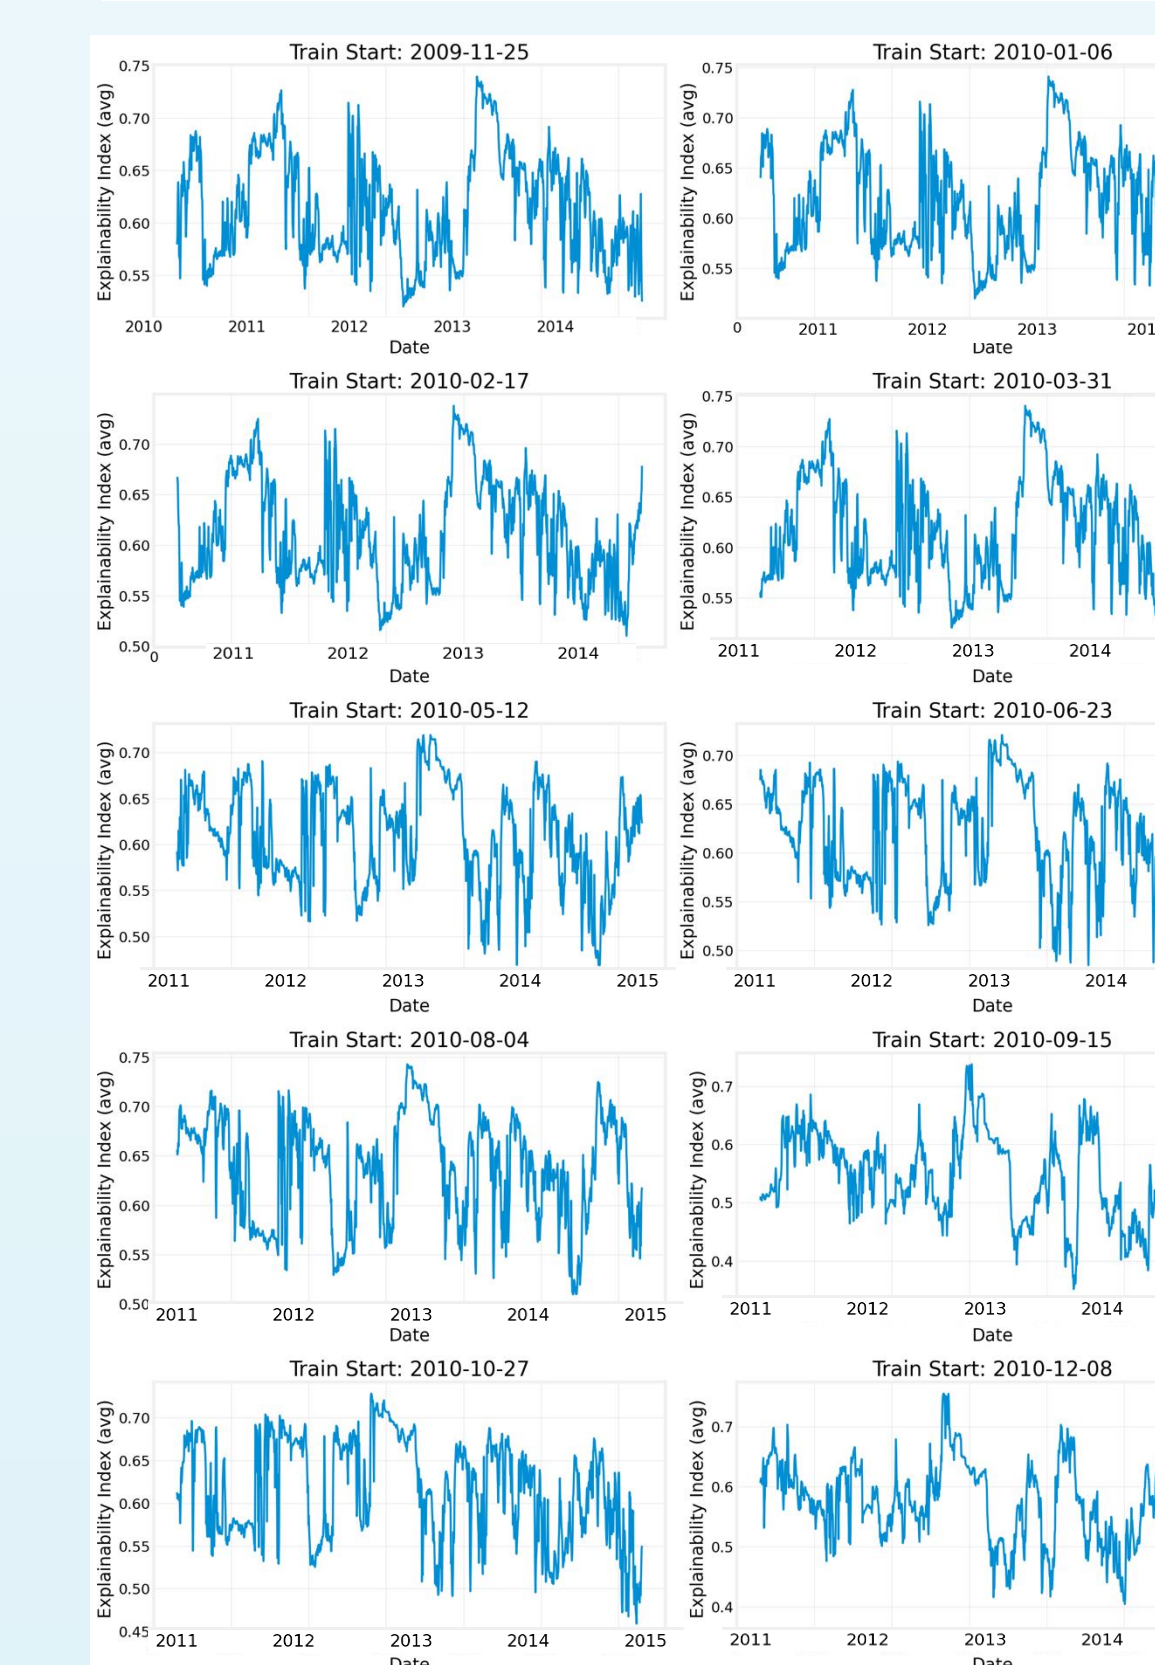


Figure 1. Mean-variance finishes highest but shows the deep drawdowns around Aug-2015 and Jan-2016. Equal-weight is close behind with similar volatility. GAT (EI objective) delivers competitive gains with noticeably shallower troughs. FFT and supervised-HMM are lower-return but smoother.

Figure 2. Across windows, the top performers are HMM and GAT (Sharpe objective). FFT and GAT (EI objective) form a middle tier. Mean-variance and equal-weight are weakest, including negative periods. Regime- and graph-aware methods outperform simple allocations on risk-adjusted returns.

TestStart	Annualized Sharpe Ratio Per Test Period					
	equal_weights	mean_var_opt	gat_model_sharpe_obj	gat_model_ei_obj	pattern_recognition_fft	vol_regime_supervised_hmm
2009-11-25	0.1304	0.2519	1.9796	0.1321	0.5017	-0.1612
2010-01-06	0.2868	0.3188	2.4228	0.2488	0.9628	0.2140
2010-02-17	0.3244	0.3403	2.1191	0.5471	0.4766	1.1087
2010-03-31	-0.8040	-0.7450	0.9241	-0.8608	-0.7122	0.7572
2010-05-12	-0.4807	-0.3709	1.4380	-0.0915	-0.4621	0.2675
2010-06-23	-0.2710	-0.1321	1.4924	0.0908	0.6108	1.8582
2010-08-04	-0.0880	0.0018	2.3447	0.5212	1.2641	3.0788
2010-09-15	0.3243	0.5053	3.1494	0.4608	1.7183	2.3495
2010-10-27	1.2793	1.4644	3.9595	2.0242	2.7771	4.0634
2010-12-08	1.0984	1.2596	3.9476	1.2121	2.2086	5.1432

Explainability Index in Rolling Train Windows



Portfolio Turnover Rate Heatmap Across Windows

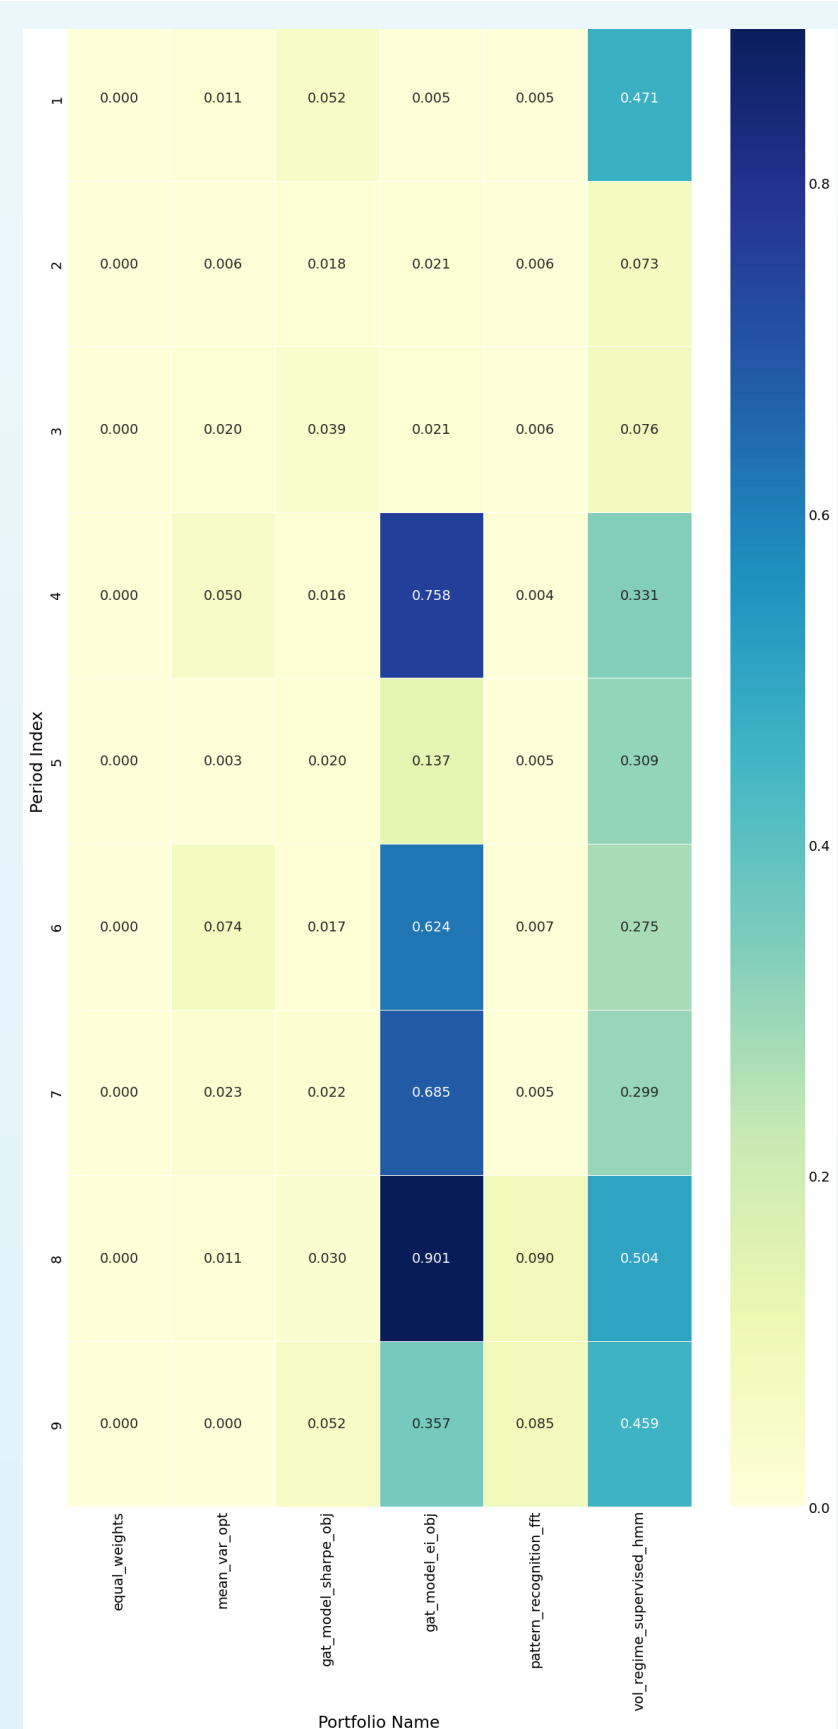


Figure 3. (left) Each panel shows the rolling portfolio EI over each training window for the GAT (EI objective). The stability of the EI shows **consistent interpretability** of the model across regimes, with variation indicating **sensitivity to market conditions**.

Figure 4. (right) Portfolio turnover measures how much a portfolio's holdings change over time. (Note: 0=no change, 1=full reallocation) EI-GAT shows the highest turnover; HMM is moderate; FFT low-moderate; MVO and Sharpe-GAT are very low; equal-weights remains 0.

Next Steps

- Extend the volatility-regime model by replacing the supervised HMM with a **Supervised Markov Switching Model**, a **semi-supervised HMM** (partial labels), or a **neural HMM** to learn state durations.
- For pattern recognition beyond FFT, explore **wavelets** (CWT/DWT), **scattering transforms**, **EMD-Hilbert spectra**, **spectrogram features**, or **singular spectrum analysis**.
- Hybridize by feeding regime posteriors and time-frequency features into the **GAT** as node/edge attributes or using a **mixture-of-experts** that gates among sub-portfolios.