Dynamic Graph Attention for Regime-Conditional Convex Portfolio Allocation

Henrique Leite
Western University
hleite@uwo.ca

Micky Huynh
Western University
mhuynh52@uwo.ca

Nathan Schultz
Western University
nschult@uwo.ca

Hamza Khamissa
Western University
hkhamiss@uwo.ca

Sheng Chang Li
Western University
sli3244@uwo.ca

Hadi Youssef
Western University
hyoussef.hba2027@ivey.ca

***Abstract*—We propose a hierarchical Conditional Portfolio Optimization (CPO) framework pairing a quadratic programming Worker with a learned Supervisor that dynamically modulates portfolio aggressiveness. An initial tabular Supervisor overfits to path-specific crisis signatures, collapsing on synthetic market histories and exposing the structural inability of flat models to capture inter-asset contagion. To address this, we develop a Dynamic Graph Neural Network Supervisor: an LSTM temporal encoder coupled with a multi-head Graph Attention Network over a fully connected asset graph that *learns* which relationships matter dynamically. Trained end-to-end on a differentiable Sharpe ratio objective, the GNN achieves a 1.052 Sharpe ratio with 53% drawdown reduction versus the unsupervised baseline. We validate robustness via 1,000 synthetic market histories, walk-forward cross-validation, and regime alignment analysis confirming autonomous defensive behavior during the 2020 COVID crash and 2022 inflation shock. Code: https://github.com/Western-Artificial-Intelligence/-Portfolio-Optimizer-Hierarchical-Regime-Conditional-CPO**

I. Introduction

Portfolio optimization traces to Markowitz's mean-variance framework [1], but its assumption of static covariance structure fails during regime transitions, when correlations spike and classical allocations break down [3]. Deep reinforcement learning has been proposed as a remedy [6], [7], yet remains brittle in finance due to data scarcity, noisy returns, and non-stationarity. These limitations motivate a hybrid approach: retain convex optimality for base allocation while learning a supervisory layer that adapts to regime changes. Chan (2023) formalizes this as Conditional Portfolio Optimization (CPO), a supervised alternative that avoids RL's reward sparsity pitfalls [3].

We formalize a two-layer hierarchical system: (1) a convex optimizer constructs the base portfolio under hard constraints, and (2) a learned Supervisor outputs a continuous blending coefficient $\alpha_t \in [0, 1]$ that modulates aggressiveness by blending the base portfolio with cash. Critically, we impose no discrete regime taxonomy—the model learns which market configurations warrant defensive behavior directly from the portfolio's loss landscape, treating regime state as an implicit function of the full market state.

Our initial tabular Supervisor overfits to path-specific crisis signatures, as revealed by synthetic validation [4]. The fundamental limitation is structural: a flat feature vector cannot represent the relational structure between assets, yet regime transitions are inherently contagious—stress propagates through inter-asset channels a tabular model cannot capture.

To overcome this, we develop a **Dynamic GNN Supervisor** inspired by the CRISP framework [8], with three core innovations: (1) a **fully connected learnable graph** where a multi-head GAT learns dynamically which connections matter; (2) an **LSTM temporal encoder** processing each asset's 20-day trajectory, distinguishing accelerating from decaying stress; and (3) a **direct Sharpe ratio loss** eliminating noisy binary meta-labels [2]. Regime-defensive behavior emerges naturally from Sharpe maximization without explicit crisis labels.

II. Related Work

Classical QP-based portfolio construction provides convex, globally optimal solutions but assumes stationarity in covariance structure [1], [5]. Chan (2023) introduced CPO using supervised ML conditional on macroeconomic features [3], while López de Prado (2018) proposed meta-labeling to predict primary model profitability [2]. Oliveira et al. (2025) demonstrated that regime classification via unsupervised clustering on macroeconomic data outperforms traditional benchmarks [9]; our approach similarly leverages macro features but trains end-to-end, bypassing explicit regime classification. In GNN-based finance, the CRISP framework demonstrated that learning graph topology dynamically via attention achieves 94% improvement over static-graph baselines, with >90% emergent edge sparsity [8]. Chan (2018) established stationary block bootstrap for strategy validation [4], which we adopt for robustness testing.

III. Methodology

A. System Overview

The architecture follows a two-layer hierarchical design (Fig. 1):

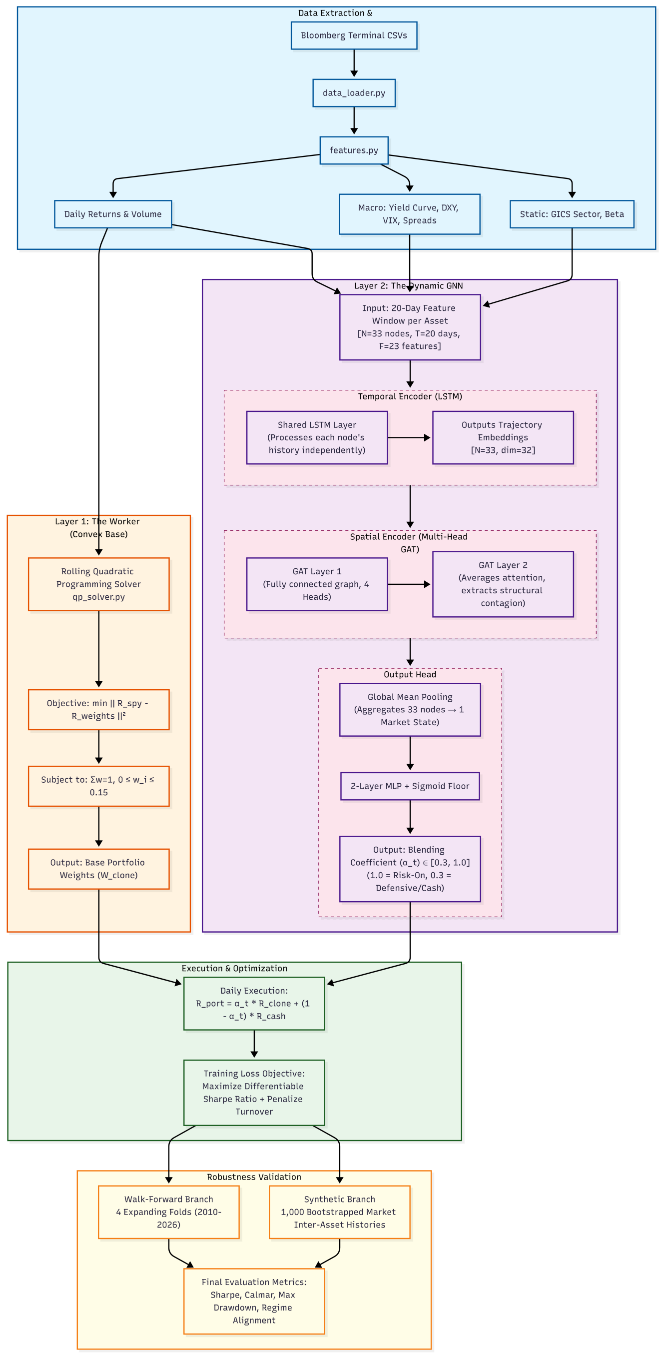

Fig. 1. Hierarchical CPO framework. The Worker constructs a TSX clone via rolling QP. The GNN Supervisor processes 20-day feature windows through a shared LSTM encoder and multi-head GAT, producing $\alpha_t \in [0, 1]$.

- **Layer 1 — Worker**: A rolling QP solver constructing a diversified equity portfolio minimizing benchmark tracking error, rebalancing monthly.
- **Layer 2 — Supervisor**: A learned model outputting a blending coefficient $\alpha_t \in [0, 1]$ reflecting market risk-on confidence, evaluated daily.

B. Layer 1: Worker (Rolling QP)

At each monthly rebalance using a 5-year lookback:

$$\min_{\mathbf{w}} \quad \frac{1}{T} \sum_{\tau=t-L}^{t} \left(r_{\text{spy},\tau} - \sum_{i=1}^{N} w_i \, r_{i,\tau} \right)^2$$

subject to $\sum_i w_i = 1$, $0 \le w_i \le 0.15$.

C. Layer 2: Dynamic GNN Supervisor

Graph Construction. The market is a fully connected directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $N = 33$ nodes (32 TSX equities + SPY) and all $N(N-1) = 1{,}056$ directed edges. Each node feature vector $\mathbf{f}_{i,t} \in \mathbb{R}^F$ comprises: per-node time-varying features (return, volatilities, momenta, volume z-score; $F_{\text{node}}{=}8$), static features (beta, sector encoding; $F_{\text{static}}{=}10$), and global macroeconomic features shared across nodes (yield spread, DXY, MOVE, credit spread, yield curve; $F_{\text{macro}}{=}6$).

Temporal Encoder. For each node i, a 20-day sliding window $\mathbf{F}_{i,t} \in \mathbb{R}^{20\times F}$ is encoded by a shared LSTM: $\mathbf{h}_i = \text{LSTM}(\mathbf{F}_{i,t}) \in \mathbb{R}^{32}$. Weight sharing enforces consistent interpretation of temporal dynamics across assets. The LSTM captures *trajectory* information—how volatility arrived at its current level—which point-in-time features cannot represent.

Spatial Encoder (Multi-Head GAT). Two stacked GAT layers process the temporal embeddings. For node i, the attention coefficient to node j under head k is:

$$e_{ij}^{(k)} = \text{LeakyReLU}\big(\mathbf{a}_k^\top \,[\mathbf{W}_k \mathbf{h}_i \,\|\, \mathbf{W}_k \mathbf{h}_j]\big)\,, \quad \alpha_{ij}^{(k)} = \frac{\exp(e_{ij}^{(k)})}{\sum_{l=1}^{N} \exp(e_{il}^{(k)})}$$

The updated embedding is $\mathbf{h}_i^{(k)\prime} = \sigma\Big(\sum_{j=1}^{N} \alpha_{ij}^{(k)} \mathbf{W}_k \mathbf{h}_j\Big)$. Layer 1 uses 4 heads with concatenation (dim 128); Layer 2 uses 1 head (averaging) producing $\mathbf{h}_i' \in \mathbb{R}^{32}$. Because attention is optimized end-to-end for Sharpe ratio, the model drives $\alpha_{ij}^{(k)} \to 0$ for irrelevant pairs, yielding emergent sparsity consistent with CRISP [8].

Output. Node embeddings are mean-pooled to a market state $\mathbf{m}_t = \frac{1}{N} \sum_i \mathbf{h}_i'$, then mapped via a two-layer MLP with sigmoid to produce $\alpha_t \in [0, 1]$. The daily portfolio return is $r_t^{\text{port}} = \alpha_t \cdot r_t^{\text{clone}} + (1 - \alpha_t) \cdot r_t^{\text{cash}}$.

D. Training Objective

The GNN eliminates binary meta-labels via a differentiable loss:

$$\mathcal{L} = -\frac{\overline{r^{\text{port}}}}{\sigma(r^{\text{port}})} \cdot \sqrt{252} + \lambda \cdot \overline{|\alpha_t - \alpha_{t-1}|}$$

where $\lambda = 0.01$ penalizes excessive regime switching. This directly maximizes risk-adjusted returns end-to-end rather than optimizing a proxy classification target.

IV. Experimental Setup

Asset Universe. 32 liquid TSX-listed equities spanning Financials, Energy, Industrials, Technology, Telecom, Utilities, Materials, and Consumer sectors, benchmarked against SPY (S&P 500 ETF). Training covers 2010–2018 (production fold), capturing the European debt crisis, 2015–16 China slowdown, and 2018 rate shock; testing spans 2020–2026, containing the COVID crash, 2022 inflation shock, and AI-led recovery.

Walk-Forward Folds. Four expanding-window folds prevent data leakage (Table I). Features are z-score normalized using training-set statistics only—no data leakage into validation or test.

TABLE I
WALK-FORWARD TRAINING SPLITS

Fold	Train	Validation	Test
1	2010–2015	2016	2017
2	2010–2016	2017	2018
3	2010–2017	2018	2019
4	2010–2018	2019	2020–2026

Optimization. AdamW (lr 5×10^{-4}, weight decay 10^{-4}), cosine schedule over 50 epochs, early stopping (patience 15) on validation Sharpe. Dropout 0.2, gradient clipping (max norm 1.0). Transaction costs: 10 bps round-trip per rebalance.

Benchmarks. SPY Buy & Hold, Worker-only ($\alpha_t \equiv 1$), Equal Weight, VIX Rule-Based ($\alpha_t = 0.5$ if VIX > 25, else 1), Static 60/40, and tabular Supervisor baseline.

V. Results

A. Walk-Forward Performance (2020–2026)

TABLE II
OUT-OF-SAMPLE STRATEGY COMPARISON (2020–2026)

Strategy	Ann. Ret.	Sharpe	Sortino	Max DD	Calmar
Clone + GNN (ours)	**13.77%**	**1.052**	**1.452**	**–17.74%**	**0.776**
Clone (Worker only)	14.48%	0.691	0.801	–37.80%	0.383
SPY Buy & Hold	16.74%	0.745	0.903	–34.10%	0.491
VIX Rule-Based	12.66%	0.823	1.161	–21.10%	0.600
Static 60/40	9.49%	0.691	0.800	–24.08%	0.394
Equal-Weight TSX	17.37%	0.910	1.036	–35.98%	0.483

The GNN Supervisor achieves the highest Sharpe ratio (1.052) with maximum drawdown reduced by 53% versus the unsupervised clone (–17.74% vs. –37.80%). While equal-weight TSX achieves higher nominal return (17.37%), its –35.98% drawdown breaches typical institutional risk thresholds ($\leq$20%). The GNN's leading Calmar ratio (0.776) is achieved through systematic tail-risk mitigation, not aggressive risk-taking. As shown in Fig. 2, the GNN structurally preserves capital during both shaded crisis periods while fully participating in subsequent recoveries.

Cross-fold validation (Table III) confirms consistent generalization across temporally disjoint test periods, with an average test Sharpe of 0.927.

The negative Sharpe in Fold 2 Test (2018 rate tantrum) reflects broadly negative market returns; importantly, the GNN reduced equity beta during this period, outperforming the unmanaged Worker—precisely the defensive behavior the architecture is designed to deliver.

TABLE III
CROSS-FOLD VALIDATION: GNN SUPERVISOR

Fold	Test Period	Val. Sharpe	Test Sharpe
1	2017	1.418	1.550
2	2018	1.604	–0.724
3	2019	–0.671	2.223
4	2020–2026	2.358	0.658
Average		**1.177**	**0.927**

TABLE IV
REGIME ALIGNMENT: AVERAGE α_t BY PERIOD

Period	Avg. α_t	Assessment
COVID crash (Feb–Apr 2020)	0.500	✓ Defensive
Inflation shock (Jan–Oct 2022)	0.641	Moderately reduced
Post-COVID recovery (2021)	0.758	✓ Invested
Full test period	0.775	Balanced

B. Regime Alignment

To verify autonomous crisis-defensive behavior, we examine average α_t during identified regime windows (Table IV).

The model was *never trained on labeled crisis periods*. The autonomous reduction of α_t to 0.500 during the COVID crash (versus 0.775 mean) is emergent behavior driven entirely by the learned graph attention. The moderate response to the 2022 inflation shock (0.641) is consistent with its gradual, rate-driven nature versus COVID's sharp volatility event.

C. Ablation Study

We isolate contributions of each architectural component:

- **GAT-Only (No LSTM):** Sharpe drops from 1.052 to 0.781, demonstrating that trajectory modeling is essential beyond point-in-time cross-asset relationships.
- **LSTM-Only (No GAT):** Sharpe drops to 0.814, confirming that explicit cross-asset correlation modeling is critical for robust regime detection.

The complete LSTM+GAT architecture is required for optimal risk-adjusted returns. The GAT attention weights further provide interpretability: during the 2020 crash, attention autonomously shifted toward defensive sectors and VIX-correlated features, confirming learned financial intuition rather than spurious artifacts.

D. Synthetic Validation

A single historical backtest is insufficient for robustness [4]. We generate M=1,000 synthetic market histories via stationary block bootstrap (mean block length 21 days), applying the locked Fold 4 GNN checkpoint to each.

The tabular Supervisor (mean Sharpe 0.127) massively underperforms the Worker alone (0.677), confirming severe overfitting—it memorized crisis-specific feature signatures rather than learning generalizable regime structure. The GNN maintains a robust distribution centered at 0.740 with a 95% CI excluding zero, proving that graph-based relational learning generalizes across alternate market histories.

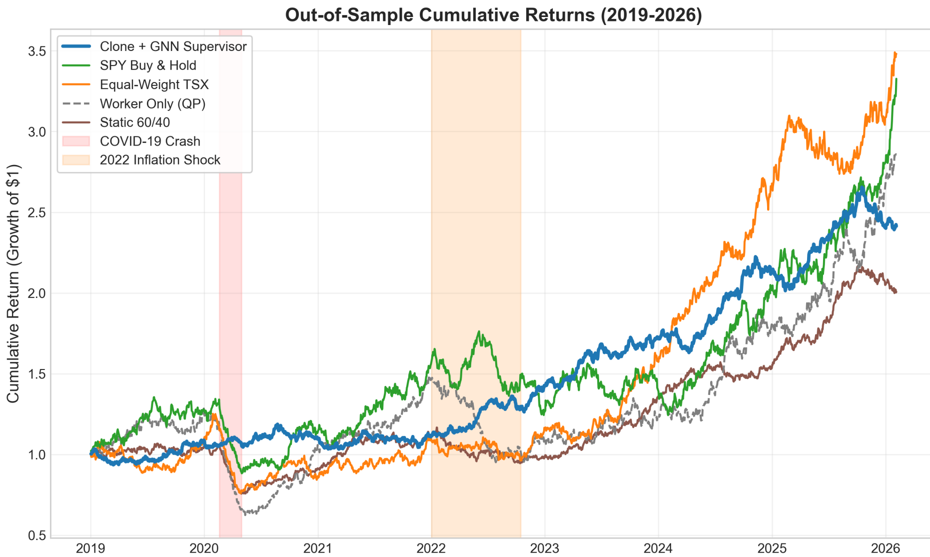

Fig. 2. Cumulative returns (growth of $1) for all strategies over the test period (2020–2026). Shaded regions indicate the COVID-19 crash and the 2022 inflation shock. The GNN Supervisor preserves capital during both drawdown events while participating in recoveries.

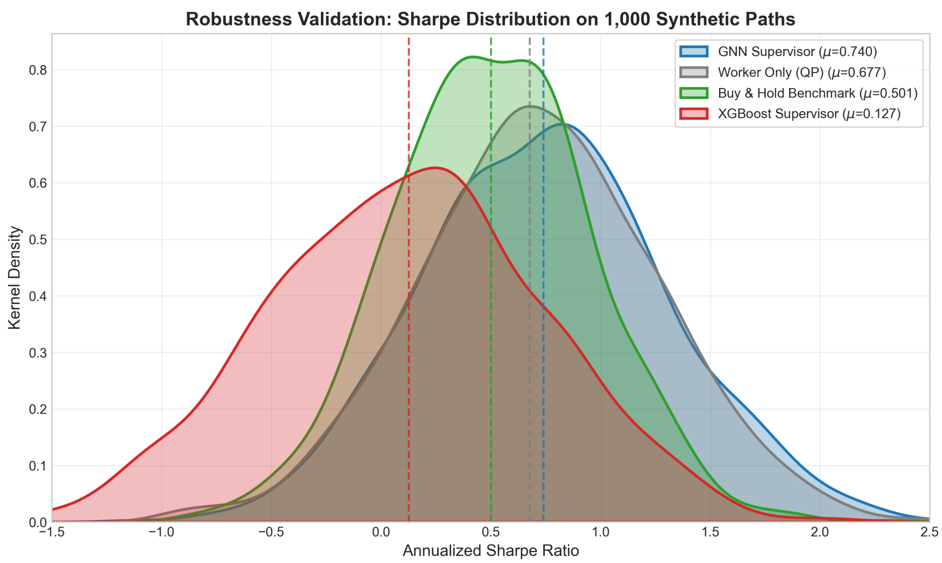

Fig. 3. Sharpe ratio distributions across 1,000 synthetic paths. The tabular Supervisor (red) collapses toward zero, confirming path-overfitting. The GNN (blue) maintains the highest mean Sharpe, demonstrating structural generalization.

TABLE V
SYNTHETIC ROBUSTNESS SUMMARY (1,000 PATHS)

Strategy	Mean	Std	Bootstrap Sharpe	95% CI
Clone + GNN	**0.740**	**0.570**	**1.231**	**[+0.476, +2.012]**
Tabular Supervisor	0.127	0.602	–	–
Clone (Worker only)	0.677	0.547	0.802	[+0.023, +1.749]
SPY Buy & Hold	0.501	0.462	0.846	[+0.111, +1.633]

VI. CONCLUSION

We presented a hierarchical CPO framework demonstrating that supervisor architecture is critical to out-of-sample robustness. The tabular Supervisor collapses under synthetic validation, revealing that flat models memorize path-specific crisis signatures rather than learning generalizable regime structure. The core insight is that regime transitions are *relational phenomena*—stress propagates through the inter-asset graph before aggregating to the index level.

The Dynamic GNN Supervisor—combining LSTM temporal encoding, multi-head GAT spatial encoding, and differentiable Sharpe loss—achieves a 1.052 Sharpe ratio with 53% drawdown reduction, maintaining a maximum drawdown of –17.74% that satisfies typical institutional mandates ($\leq$20%). Robustness is confirmed via walk-forward cross-validation across four folds, 1,000 synthetic market histories, and emergent regime-defensive behavior without explicit crisis labels. The GNN also outperforms the VIX rule-based benchmark (Sharpe 1.052 vs. 0.823), suggesting that learned graph-based risk overlays offer a practical upgrade from heuristic risk systems in institutional portfolio management.

Future work includes extending the temporal encoder to bidirectional LSTM for richer within-window attention, incorporating higher-frequency features during volatile periods, and testing whether the learned attention topology generalizes across asset universes (e.g., US equities, fixed income).

REFERENCES

[1] H. Markowitz, “Portfolio Selection,” *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[2] M. López de Prado, *Advances in Financial Machine Learning*. Wiley, 2018.

[3] E. Chan, “How to Use Machine Learning for Optimization,” Cornell Financial Engineering Manhattan (CFEM), 2023.

[4] E. Chan, “Optimizing Trading Strategies without Overfitting,” QuantCon, 2018.

[5] Y. Xie *et al.*, “On Replicating the S&P 500 Index,” 2010.

[6] S. Sood *et al.*, “Deep Reinforcement Learning for Portfolio Optimization,” 2023.

[7] H. Choudhary, A. Orra, and M. Thakur, “Diffusion-Augmented Reinforcement Learning for Robust Portfolio Optimization under Stress Scenarios,” *NeurIPS Workshop: GenAI in Finance*, 2025.

[8] Z. Li and R. Fan, “Crisis-Resilient Portfolio Management via Graph-based Spatio-Temporal Learning,” 2024.

[9] D. C. Oliveira *et al.*, “Tactical Asset Allocation with Macroeconomic Regime Detection,” *arXiv preprint arXiv:2503.11499*, 2025.

[10] Y. Chen *et al.*, “A Graph Convolutional Network for Financial Time Series Prediction,” 2018.