Knowledge & Data Engineering Systems

"ViFi-HMM: Fuzzy Regime-Switching State Inference for Uncertain Dynamic Systems"

Fatima Ghanduri, Christos Anagnostopoulos

School of Computer Science, University of Glasgow

University of Glasgow

VIA VERITAS VITA

About

Knowledge & Data Engineering Systems

School of Computer Science-University of Glasgow, UK

FATIMA GHANDURI
PhD Researcher - UNIVERSITY OF GLASGOW
Data & AI Technical Sales- Energy Economics– SLB (SCHLUMBERGER)

DR CHRISTOS ANAGNOSTOPOULOS
Associate Professor in Distributed Computing & Data Engineering
UNIVERSITY OF GLASGOW

Agenda

Knowledge & Data Engineering Systems

- Introduction Finance Problem Statement
- Problem Statement and Challenges
- Why current Solutions fall short
- Methodology
- Experiment & Results
- Conclusion & Next steps

This work addresses core challenges in quantitative finance: volatility forecasting, risk-aware allocation, and regime-triggered trading signals under uncertainty.

Financial Markets Live in Regimes

Asset prices alternate between **calm (low-volatility)** and **turbulent (high-volatility)** regimes. These transitions carry direct economic consequences.

there is no single model that can universally predict prices across all market conditions. A model that performs well in one regime will break in another.

WHEN TRANSITIONS ARE MISSED
- Delayed portfolio rebalancing
- Value-at-Risk (VaR) breaches
- Missed algorithmic trading signals
- Adverse exposure during crises (e.g., 2020 pandemic)

Key Insight
Regime boundaries are **fuzzy, not sharp** — assets can exhibit overlapping low- and high-volatility traits simultaneously.

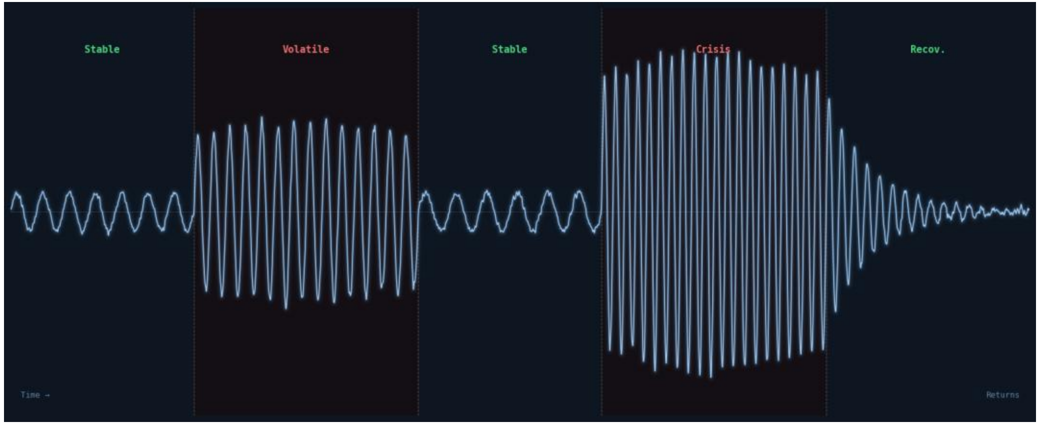

Knowledge & Data Engineering Systems

Problem Statement

Three Compounding Challenges

Given N financial time series → partition into K soft clusters, each linked to a regime-specific model

Non-Stationarity

Sequences shift statistical behavior over time. Volatility is not constant — parameters evolve with market conditions, macro events, and structural breaks.

Static models are fundamentally misspecified

Ambiguity in Boundaries

Regime transitions are not discrete events. Assets exhibit overlapping behavior — a stock can be simultaneously in a "transition" between low and high volatility.

Hard cluster labels misrepresent reality

Temporal Dependencies

State at time t depends on t−1. Volatility is autocorrelated (clustering). Transitions follow dynamics — today's turbulence predicts tomorrow's.

i.i.d. assumptions break completely

Core gap: No existing method simultaneously handles all three — temporal dynamics, soft membership, and joint probabilistic optimization.

Knowledge & Data Engineering Systems

Related Work

Why Existing Methods Fall Short

Knowledge & Data Engineering Systems

METHOD	TEMPORAL MODELING	FUZZY MEMBERSHIP	VARIATIONAL INF.	REGIME-AWARE	GAP
DTW / Euclidean	X	X	X	X	Ignores latent regime dynamics
Feature-based (ACF, Hurst)	X	X	X	X	Discards temporal ordering
Hard-HMM + k-means	✓	X	X	✓	Crisp labels; no uncertainty
FCM-VAR	X	✓	X	—	No latent state transitions
GARCH + FCM	Partial	✓	X	✓	Decoupled; no discrete states
Variational Clustering	X	✓	✓	X	Ignores temporal structure
Vifi-HMM (Ours)	✓	✓	✓	✓	Unified end-to-end framework

The key gap: No prior method jointly learns fuzzy sequence-level memberships and regime-switching HMMs within a single variational objective. ViFi-HMM is the first to do so.

Solution Design

What a Solution Must Deliver

Knowledge & Data Engineering Systems

X

Soft, Uncertainty-Aware Membership

Assets must not be forced into a single regime. A portfolio stock on the cusp of a volatility transition should hold partial membership in both the stable and turbulent clusters.

Fuzzy cluster memberships $U \in [0,1]$

Y

Temporal Latent State Modeling

Capture how volatility evolves within a series. Today's low-vol state influences tomorrow's — this autocorrelation is fundamental to financial risk modeling.

Per-cluster Hidden Markov Models

Z

End-to-End Joint Optimization

Clustering and dynamics must inform each other. Separate pipelines (cluster then model) lose the feedback loop that improves both regime detection and cluster quality.

Single variational objective $J(\Theta, U)$

With these requirements defined, we introduce **ViFi-HMM** *— the first framework to satisfy all three jointly.*

Proposed Framework

ViFi-HMM — Three Ideas, One Objective

One HMM per Cluster

Each of K clusters owns a full HMM $\theta_k = (\pi_k, A_k, B_k)$ capturing regime-specific latent state dynamics. Gaussian emissions model state-specific volatility (σ^2 per hidden state).

Fuzzy Membership Matrix U

$U \in [0,1]^{N\times K}$ with $\Sigma_k u_{ik} = 1$. Each series i belongs softly to all clusters. Fuzzifier $m > 1$ controls softness; $m\to 1$ recovers hard assignment.

Joint Variational Objective

Maximize $J(\Theta, U)$ which balances data fit, membership entropy, and inter-cluster separation via KL divergence — all optimized simultaneously.

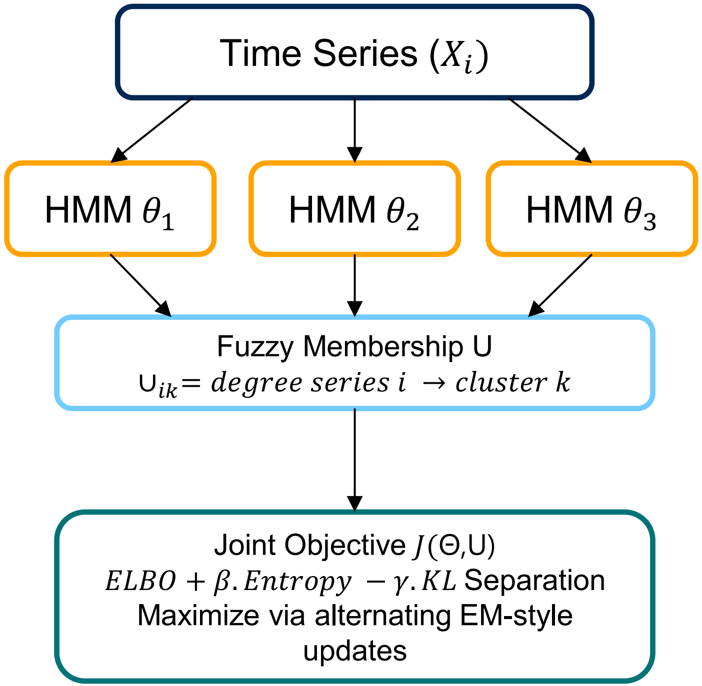

$$J(\Theta,U) = \Sigma\, u_{ik}^m \cdot ELBO_{ik} + \beta \cdot \Omega(U) - \lambda \cdot S\_KL(\Theta,U)$$

Fuzzy Variational Objective

Knowledge & Data Engineering Systems

EVIDENCE LOWER BOUND (ELBO)

For each (series i, cluster k), the exact HMM posterior is computed via forward–backward. This makes the ELBO tight — identical to the EM E-step.

Fuzzy Weighting — u_{ik}^m

The fuzzifier $m > 1$ controls membership softness. Each ELBO term is weighted by u_{ik}^m — higher membership → stronger pull toward that cluster's HMM.

Entropy Regularizer — $\Omega(U)$

Prevents degenerate memberships (all weight on one cluster). Weight $\beta \geq 0$ controls preference for high-entropy (more uncertain) partitions.

SOFT VS HARD ASSIGNMENT

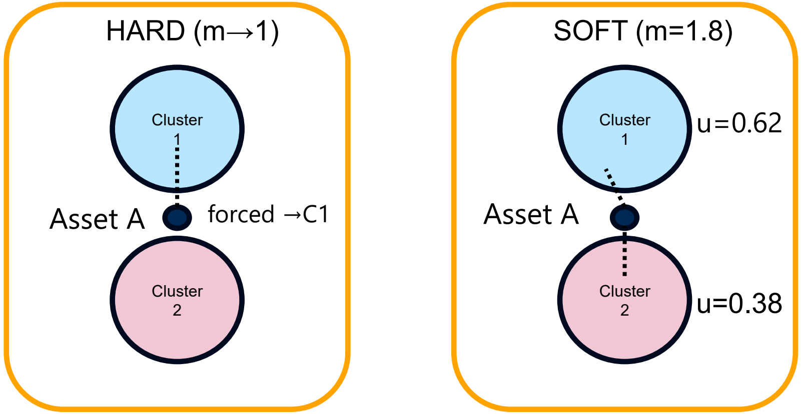

Key Departure from Mixture-of-HMMs

Standard EM uses likelihood-normalized responsibilities. ViFi-HMM uses entropy-regularized fuzzy variables controlled by (m, β) - enabling explicit uncertainty quantification over regime membership.

Methodology- Regularization

KL-Based Regime Separation

THE PROBLEM

Fuzzy ELBO alone does not force clusters to learn distinct dynamics. Without separation pressure, two HMMs may converge to nearly identical state posteriors ; offering no regime distinction.

Posterior State Profiles

For series i under cluster k, define $\gamma^{(s)}_{i,t,k} = P(s_{i,t}=s \mid X_i, \theta_k)$. This is the posterior probability of hidden state s at time t. Compute a membership-weighted average across clusters:

Separation Penalty S_KL

Penalize clusters that agree on state usage for the same series. Large S_KL = strong, distinct state dynamics per cluster.

EFFECT OF KL PENALTY ON CLUSTER SPECIALISATION

Without KL (λ=0) | **With KL (λ=0.7)**

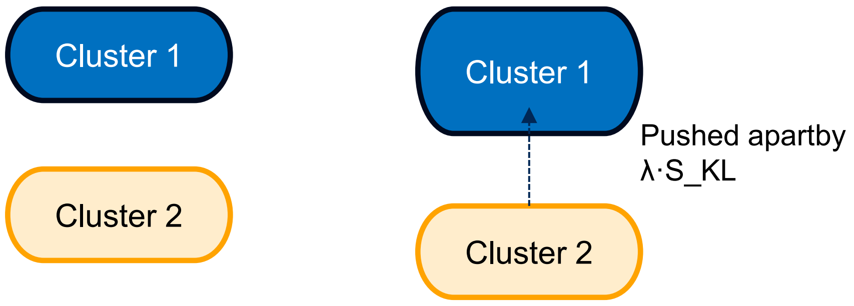

Similar posteriors but poor regime discrimination

Distinct dynamics with clear regime separation

Methodology- Algorithm

Optimization: EM-Style Alternating Scheme

Each outer iteration cycles through three coordinated steps

Step 1

POSTERIOR STEP (E-LIKE)

- Run **forward–back ward**for all (i, k) pairs
- Obtain $\gamma^{(s)}_{i,t,k}$ — state posteriors
- Obtain ξ — transition posteriors
- Evaluate ELBO_{ik} and S_KL(Θ,U)

Step 2

MEMBERSHIP UPDATE (VAR. E-STEP)

- Compute gradient $\partial J/\partial u_{ik}$ from ELBO, Ω, S_KL
- **Projected gradient ascent** (5–10 steps/iter)
- Project result onto simplex: $\Sigma_k\, u_{ik}=1,\ u_{ik}\geq 0$
- Step size $\eta \in \{10^{-2}, 10^{-3}\}$ from validation

Step 3

MODEL UPDATE (M-LIKE)

- **Fuzzy-weighted Baum–Welch** re-estimation
- π_k, A_k updated with weights u_{ik}^m
- Gaussian emissions from fuzzy sufficient stats
- Optional: small gradient step to reduce λS_KL

20–30
ITERATIONS TO CONVERGE

~13s
N=25, T=500
RUNTIME

O(NKT M²)
PER-ITERATION COMPLEXITY

Experimental Setup

Knowledge & Data Engineering Systems

Synthetic Dataset

100 series (50 per cluster), T=200 with known ground-truth HMM parameters

---

Cluster A: σ=[0.5, 2.0] (low/high vol)
Cluster B: σ=[1.0, 1.5] distinct transitions

Metrics: ARI, Purity, Parameter recovery error

Stock Index Returns

Daily log-returns of major indices (S&P 500, FTSE 100)

Period: Jan 2000 – Dec 2024
Setup: K=2 clusters, M=2 HMM states (stable vs turbulent)

Returns standardized; temporal holdout at 2011

FX Exchange Rates

15 currency pairs over 10-year window

Pairs: USD/EUR, USD/TRY, +13 others
Setup: K=3 clusters (low/med/high vol profiles)

Tests cross-sectional and temporal structure

Base Lines
Hard –HMM + K-means , DTW-HAC, GARCH+FCM, Variational Clustering, FCM-VAR

Evaluation Metrics

- Synthetic: ARI, Purity, Param.Error
- Real-world: Silhouette (KL), Partition,
- Financial: Sharpe, VaR, breach rate, Hedge P&L

Synthetic Data: Near-Perfect Regime Recovery

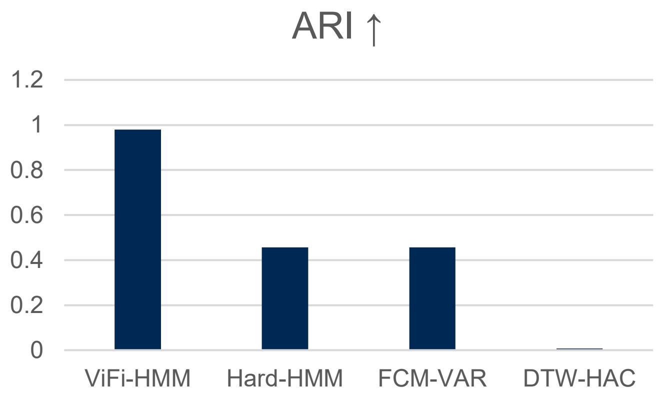

Method	ARI	Entropy	Purity
ViFi-HMM	**0.980**	**0.27**	**0.96**
Hard-HMM	0.457	–	0.89
FCM-VAR	0.457	4.16	0.83
DTW-HAC	0.007	–	0.71

Key Findings

• Soft memberships are **crisp within clusters**, ambiguous only near transition boundaries

• Parameter recovery error <0.1 (mean) — HMM parameters closely match ground truth

• Low entropy (0.27) = confident where confident, uncertain where appropriate

Why Such a Large Gap?

DTW fails entirely (ARI=0.007) because shape similarity is misleading — two series can have identical shape but vastly different volatility regimes. ViFi-HMM looks at latent states, not surface shape.

STATISTICAL RIGOR

All results reported as mean ± std over 5 independent runs. 95% CIs via paired bootstrap with 1000 resamples.

Results- Real World

Stock Indices & FX: Consistent Outperformance

Knowledge & Data Engineering Systems

METHOD	STOCK INDICES			FX PAIRS		
	SILHOUETTE ↑	PC ↑	PE ↓	SILHOUETTE ↑	PC ↑	PE ↓
ViFi-HMM ★	0.41*	0.88*	0.23*	0.37*	0.85*	0.29*
Variational	0.35±.04	0.80±.03	0.35±.04	0.30±.03	0.77±.04	0.38±.04
GARCH+FCM	0.31±.03	0.72±.04	0.44±.05	0.27±.03	0.71±.04	0.46±.04
Hard-HMM	0.28±.03	0.74±.04	—	0.22±.03	0.70±.04	—
FCM-VAR	0.25±.04	0.69±.05	0.51±.06	0.21±.03	0.67±.04	0.58±.05
DTW-HAC	0.09±.02	—	—	0.11±.02	—	—

★ Significantly better than ALL baselines (Paired Wilcoxon, Holm-corrected $p<0.05$, 5 independent runs)

Calibration
Posterior P(high-variance) correlates strongly with realised volatility: **$\rho \approx$ 0.7** for both S&P 500 and FTSE 100.

Regime Detection
High-volatility regime detection: **AUC ≈ 0.95** and **AP ≈ 0.80** on both equity indices.

No Overfitting
Temporal hold-out (train ≤ 2011, test > 2011) confirms robust generalisation across time periods.

Regime Detection Translates to Financial Value

Downstream portfolio and risk management evaluation — averaged across 5 market periods

1.24
SHARPE RATIO (RISK-PARITY)

4.1%
VAR BREACH RATE

+6.7%
HEDGE P&L IMPROVEMENT

0.72
ENTROPY SIGNAL HIT RATE

VS STATIC PORTFOLIO
Sharpe: 0.92 → **1.24** (+35%)
VaR breaches: 7.4% → **4.1%** (−45%)

Regime-aware rebalancing triggered by entropy-based signals from fuzzy memberships

VS GARCH+FCM BASELINE
Sharpe: 1.08 → **1.24** (+15%)
Hedge P&L: +3.2% → **+6.7%**
Hit rate: 0.58 → **0.72**

Earlier regime transition detection enables more timely hedging adjustments

Why This Works
Fuzzy memberships provide **early warning signals** — as an asset's membership weight shifts toward the high-vol cluster, the system alerts before the full regime transition occurs.

Most FX pairs show a dominant regime; a few show mixed memberships revealing transitional or hybrid behavior.

Ablation Study

Every Component Is Essential

PARTITION COEFFICIENT (PC ↑) — STOCK INDEX DATASET

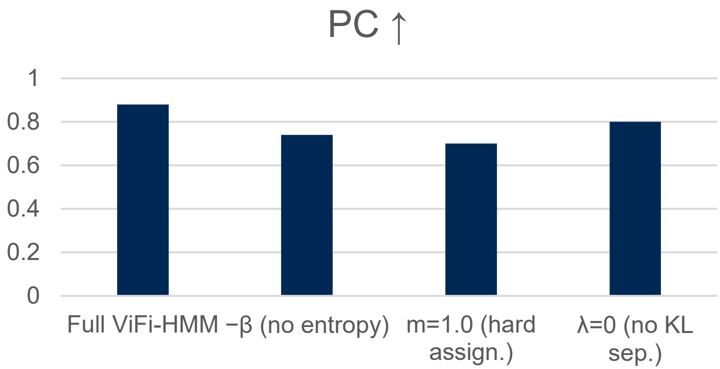

Variant	Silhouette ↑	PC ↑	PE ↓
Full ViFi-HMM	**0.41**	**0.88**	**0.23**
$-\beta$ (no entropy)	0.34	0.74	0.41
$m = 1.0$ (hard assign.)	0.29	0.70	0.39
$\lambda = 0$ (no KL)	0.36	0.80	0.30

−β: No Entropy Regularization

Memberships become ambiguous — Partition Entropy rises from 0.23 to 0.41. Without entropy pressure, memberships drift toward uniformity or collapse, losing meaningful regime structure.

m=1.0: Hard Assignments

Worst performance. Collapses ViFi-HMM to a hard mixture-of-HMMs. Crisp labels cannot represent the overlapping regime behavior inherent in financial markets.

λ=0: No KL Separation

Clusters fail to specialize. Without the divergence penalty, HMMs converge to similar state dynamics — reducing inter-cluster distinction and overall clustering quality.

Conclusion: All three components (entropy β, fuzzifier m, KL penalty λ) are complementary and each independently contributes to the model's robustness.

Knowledge & Data Engineering Systems

ViFi-HMM — A Unified Framework for Regime-Aware Finance

Soft regime modeling: Fuzzy memberships U allow assets to belong partially to multiple volatility clusters capturing the true ambiguity of financial regime transitions.

Temporal latent states: Per-cluster HMMs capture volatility autocorrelation and regime-switching dynamics directly from return sequences not from post-hoc features.

Principled joint optimization: Fuzzy-weighted ELBO + entropy regularization + KL separation are optimized simultaneously no decoupled phases, no post-hoc membership estimation.

State-of-the-art performance: ARI=0.98 on synthetic data; Silhouette=0.41, PC=0.88 on stocks significantly better than all 5 baselines (Wilcoxon, $p<0.05$).

Real financial value: +35% Sharpe improvement over static portfolio; VaR breach rate halved; earlier regime transitions detected via entropy-triggered signals.

0.98

ARI ON SYNTHETIC

0.95

AUC REGIME DETECTION

1.24

SHARPE (RISK-PARITY)

BROADLY APPLICABLE TO

Control systems · IoT monitoring
Medical time series · Energy grids

Next Steps for ViFi-HMM

Knowledge & Data Engineering Systems

Multivariate Extension

Extend to joint multivariate time series — modeling regime dynamics across correlated assets (e.g., equities + FX + rates). Enable cross-asset regime spillover detection.

Real-Time Deployment

Online ViFi-HMM variant for streaming market data sequential membership updates with constant-time complexity. Target: intraday regime detection for algorithmic trading.

Adaptive Capital Allocation

Regime-triggered portfolio rebalancing system. Fuzzy membership entropy as a continuous risk signal feeding directly into position-sizing and hedging algorithms.

Thank you!

F.ghanduri20@gmail.com
f.ghanduri.1@research.gla.ac.uk

SCAN ME

Add me on linkedin!