Recursive Multi-Agent Trading System: Iterative Optimized Portfolio Strategy Under Geopolitical Uncertainty

Jing Yang*
jing.y@wustl.edu
Washington University in St. Louis
St. Louis, MO, USA

Yichao Wu
wu.yicha@northeastern.edu
Northeastern University
Boston, MA, USA

Jianan Liu
jiananliu2408@gmail.com
Independent Researcher
China

Penghao Liang
liang.p@northeastern.edu
Northeastern University
Boston, MA, USA

Mengwei Yuan
yuanmw1998@gmail.com
Independent Researcher
China

Xianyou Li
xl4230@nyu.edu
New York University
New York, NY, USA

Weiran Yan
yanwr2016@gmail.com
Independent Researcher
China

Abstract

This paper proposes a Recursive Multi-Agent Trading System (RMATS) for iterative portfolio optimization under geopolitical uncertainty. The framework integrates four specialized agents—Sentiment, Report, Analysis, and Risk—coordinated through a recursive Manager Agent with iterative feedback loops and formally specified typed message passing. Experimental evaluation over a 561-trading-day period (January 2023 to March 2025) across a 24-asset multi-class universe demonstrates that RMATS achieves a maximum drawdown of 9.62%, lower than MVO (15.49%) and FinBERT Sentiment (15.28%), and exhibits the lowest event-period drawdown in 3 of 5 geopolitical stress scenarios tested. Ablation studies confirm the individual contribution of each agent component. These results position RMATS as a risk-control-oriented architecture suitable for institutions prioritizing capital preservation under geopolitical uncertainty.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Theory of computation** → *Algorithmic game theory*.

Keywords

multi-agent systems, portfolio optimization, geopolitical risk, large language models, reinforcement learning, drawdown control

ACM Reference Format:

Jing Yang, Yichao Wu, Jianan Liu, Penghao Liang, Mengwei Yuan, Xianyou Li, and Weiran Yan. 2026. Recursive Multi-Agent Trading System: Iterative Optimized Portfolio Strategy Under Geopolitical Uncertainty. In *Proceedings of 2026 Asia-Pacific Conference on Artificial Intelligence and Machine Learning (APAM 2026).* ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

*Corresponding Author

APAM 2026, Asia-Pacific
2026. ACM ISBN 978-1-4503-XXXX-X/2026
https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Geopolitical uncertainty has become a primary driver of financial market volatility. Events such as regional conflicts, trade sanctions, energy supply disruptions, and banking system stress can trigger rapid sector rotations and capital flight that overwhelm traditional quantitative trading frameworks [1, 2]. These systems rely on historical statistical patterns and typically fail to adapt quickly to sudden structural changes in the information environment [3].

The emergence of Large Language Models (LLMs) and multi-agent systems has opened new avenues for financial intelligence architectures [4, 5]. Systems such as TradingAgents [10] and FinAgent [11] demonstrate that role-specialized agent collaboration can improve market signal quality. However, existing frameworks often adopt fixed sequential pipelines without iterative refinement.

This paper proposes RMATS, a Recursive Multi-Agent Trading System designed with an explicit emphasis on downside risk control. The key innovation is the recursive coordination mechanism, wherein agents iteratively revise signals through multi-round communication until portfolio weights converge. The main contributions are:

- A recursive multi-agent architecture with formally specified typed message passing (AgentMessage schema) and convergence guarantee $\|\mathbf{w}^{(r+1)} - \mathbf{w}^{(r)}\|_2 < \varepsilon$.
- An Analysis Agent incorporating HMM-based market regime classification and Kalman filter signal fusion.
- A Risk Agent with CVaR-based estimation, geopolitical stress testing, and adaptive circuit breaker mechanisms.
- Comprehensive empirical evaluation with ablation studies and honest analysis of limitations.

2 Related Work

2.1 Financial Large Language Models

Domain-adapted foundation models including FinGPT [6] and BloombergGPT [7] have demonstrated strong capabilities in financial text understanding and sentiment extraction. FinMA [8] and instruction-tuned variants have improved financial reasoning [9]. RMATS extends these approaches by embedding LLM-inspired reasoning within a recursive multi-agent coordination framework.

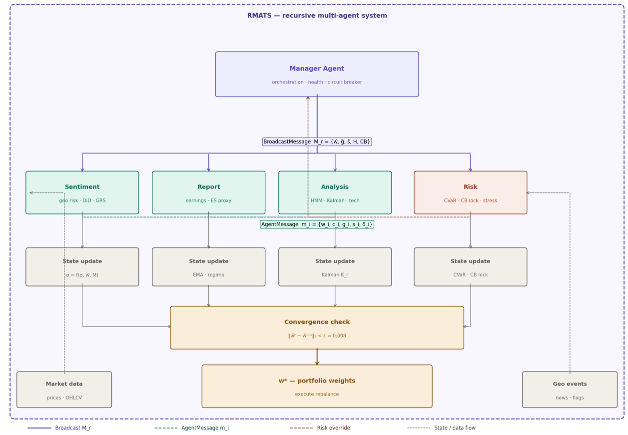

Figure 1: RMATS agent collaboration architecture. Solid arrows: BroadcastMessage M^r from Manager to all agents each round. Dashed arrows: AgentMessage $\mathbf{m}_i^r$ replies. The convergence check gates continuation; the Risk Agent can issue a circuit-breaker override.

2.2 Multi-Agent Trading Systems

TradingAgents [10] demonstrates that role-specialized agent debate can improve trading signal generation. FinAgent [11] introduces tool-augmented financial decision-making, and FinMem [12] incorporates memory-augmented reasoning. Existing systems predominantly use fixed sequential workflows. RMATS contributes a convergence-guaranteed recursive communication protocol enabling dynamic signal revision.

2.3 Reinforcement Learning for Portfolio Management

Deep reinforcement learning methods including DQN, PPO, and SAC have shown promise for sequential portfolio allocation [13]. FinRL [15] and FinRL-Meta [14] provide open-source frameworks for DRL-based trading. RMATS integrates a simplified DRL-inspired optimization layer, focusing on risk-adjusted reward shaping.

2.4 Geopolitical Risk and Financial Markets

Baker et al. [26] introduced the Economic Policy Uncertainty (EPU) index demonstrating predictive relationships between policy uncertainty and market volatility. Caldara and Iacoviello [27] developed the Geopolitical Risk (GPR) index. These indices confirm that geopolitical risk constitutes a distinct and persistent source of market stress that cannot be fully captured by traditional volatility measures.

3 Multi-Agent Trading Architecture

RMATS adopts a hierarchical collaborative architecture comprising four layers: Data Acquisition, Specialized Agent, Recursive Coordination, and Portfolio Optimization. Figure 1 illustrates the complete architecture.

3.1 Manager Agent

The Manager Agent coordinates all specialized agents through recursive orchestration. It monitors agent signal quality using a health score:

$$H_i(t) = \alpha A_i(t) + \beta S_i(t) + \gamma R_i(t) - \delta L_i(t) \tag{1}$$

where $A_i(t)$, $S_i(t)$, $R_i(t)$, and $L_i(t)$ denote predictive accuracy, signal stability, risk-adjusted profitability, and latency, respectively.

3.2 Sentiment Agent

The Sentiment Agent generates geopolitical risk scores from market-derived signals using a Difference-in-Differences (DiD) causal estimator:

$$Y_{it} = \alpha + \beta D_i + \gamma T_t + \delta(D_i \times T_t) + \varepsilon_{it} \tag{2}$$

where δ estimates causal effects of geopolitical events on sector returns.

3.3 Analysis Agent

The Analysis Agent combines macroeconomic factor modeling with HMM-based regime classification. A Hidden Markov Model with three latent states (bull, bear, stress) classifies market regimes:

$$\pi_k(t) = P(S_t = k \mid O_{1:t}, \lambda) \tag{3}$$

A Kalman filter fuses signals from the Sentiment, Report, and Analysis agents into composite signal $Z(t)$.

3.4 Risk Agent

The Risk Agent implements CVaR estimation using an EWMA-based dynamic covariance model:

$$\mathrm{CVaR}_\alpha(t) = -\mathbb{E}[R_p(t) \mid R_p(t) \leq \mathrm{VaR}_\alpha(t)] \tag{4}$$

A multi-level circuit breaker activates when drawdown, geopolitical risk, or volatility exceed adaptive thresholds:

$$\mathrm{CB}(t) = \mathbf{1}[\mathrm{DD}_p(t) > \theta_{\mathrm{dd}}] \vee \mathbf{1}[\mathrm{GRS}(t) > \theta_{\mathrm{geo}}] \vee \mathbf{1}[\sigma_p(t) > \theta_{\mathrm{vol}}] \tag{5}$$

3.5 Recursive Coordination Protocol

Each inter-agent communication uses a formal *AgentMessage* schema:

$$\mathbf{m}_i^r = \{\mathbf{w}_i^r,\ c_i^r,\ g_i^r,\ s_i^r,\ \tau_i^r,\ \delta_i^r\} \tag{6}$$

where $\mathbf{w}_i^r \in \Delta^n$ is the recommended weight vector, $c_i^r \in [0,1]$ is the confidence score, $g_i^r \in [0,1]$ is the geopolitical risk assessment, $s_i^r \in \{0,1,2\}$ is the regime classification, and $\delta_i^r = \|\mathbf{w}_i^r - \mathbf{w}_i^{r-1}\|_2$ is the intra-agent update magnitude.

The Manager broadcasts a *BroadcastMessage* $M^r = \{\bar{\mathbf{w}}^r, \bar{g}^r, \bar{s}^r, H^r, \mathrm{CB}^r\}$ and aggregates via confidence-and-health-weighted mean:

$$\bar{\mathbf{w}}^r = \frac{\sum_i c_i^r \cdot H_i \cdot \mathbf{w}_i^r}{\sum_i c_i^r \cdot H_i} \tag{7}$$

The recursion terminates when:

$$\|\bar{\mathbf{w}}^{(r+1)} - \bar{\mathbf{w}}^{(r)}\|_2 < \varepsilon = 0.008 \tag{8}$$

Empirical evaluation over 27 rebalancing steps confirms convergence in a median of 2 rounds (mean 3.56, maximum 8), with 74.1% of steps converging within $r \leq 2$.

4 Portfolio Optimization

4.1 Risk-Aware Reward Function

The reinforcement learning agent maximizes a risk-aware objective:

$$\mathcal{R}_t = r_t - \lambda_1 \sigma_t - \lambda_2 \max(0, \mathrm{DD}_t - \theta) \tag{9}$$

with $\lambda_1 = 0.8$ and $\lambda_2 = 1.5$ calibrated to prioritize drawdown control.

4.2 Constrained Portfolio Optimization

Portfolio weights satisfy:

$$\max_{\mathbf{w}} \; \boldsymbol{\mu}^\top \mathbf{w} - \lambda \mathbf{w}^\top \Sigma \mathbf{w} \tag{10}$$

$$\text{s.t.} \quad \|\mathbf{w}\|_1 \leq L_{\max}, \quad w_s \leq c_s \; \forall s, \quad \mathbf{w}^\top \mathrm{GRS}(t) \leq \gamma_{\mathrm{geo}} \tag{11}$$

The geopolitical risk constraint γ_{geo} tightens dynamically as aggregate GRS scores increase.

5 Experiments

5.1 Dataset and Setup

Historical adjusted close price data were downloaded via `yfinance` for 24 ETFs covering US sector equities (XLK, XLE, XLF, XLV, XLI, XLP, XLY, XLU, XLB, XLRE), international equities (EWJ, EWG, EWU, FXI, EEM), US fixed income (TLT, IEF, LQD, EMB), and commodities (GLD, SLV, USO, DBC), plus the SPY benchmark.

- Training: January 2016 – December 2020
- Validation: January 2021 – December 2022
- Test: January 2023 – March 2025 (561 trading days)

5.2 Baselines

Four baselines are evaluated: (1) Mean-Variance Optimization (MVO); (2) DQN Single-Agent RL; (3) FinBERT Sentiment Proxy; (4) Multi-Factor Quant combining momentum, low volatility, mean reversion, and rolling Sharpe factors. All strategies use monthly rebalancing with 10 bps transaction costs.

5.3 Geopolitical Stress Events

Five stress events in the 2023–2025 test period were selected: (1) SVB Banking Crisis (March 2023); (2) Israel-Hamas War escalation (October–November 2023); (3) US-China technology sanctions (January–February 2024); (4) Middle East military escalation (April–May 2024); and (5) Global rate cut pivot (August–September 2024).

6 Results and Analysis

6.1 Overall Performance

Table 1 presents complete performance results. The test period coincides with a strong equity bull market (SPY + ≈60%), which systematically favors return-maximizing strategies.

Figure 2 shows cumulative portfolio performance. RMATS achieves the second-lowest maximum drawdown (9.62%) after DQN (8.57%), substantially outperforming MVO (15.49%) and FinBERT Sentiment (15.28%). The negative Sharpe ratio (−0.297) reflects over-calibrated

Table 1: Overall Performance (Test: Jan 2023 – Mar 2025)

Strategy	Ann.Ret.	Sharpe	MDD	Calmar
MVO	8.08%	0.397	15.49%	0.522
DQN RL	5.13%	0.150	8.57%	0.599
FinBERT Proxy	3.69%	0.048	15.28%	0.242
Multi-Factor	7.13%	0.342	9.33%	0.764
RMATS	**0.22%**	**−0.297**	**9.62%**	**0.023**

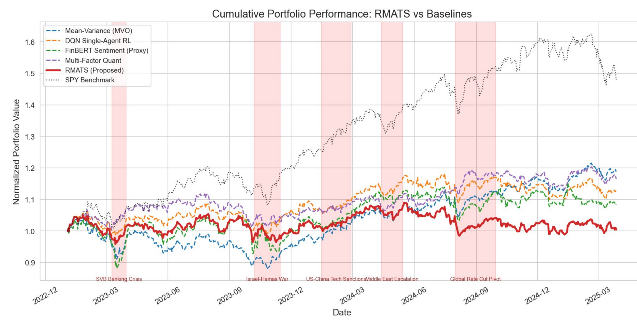

Figure 2: Cumulative portfolio performance (normalized to 1.0). Red shading indicates geopolitical stress event windows. RMATS (red solid) exhibits the flattest trajectory, reflecting its capital-preservation orientation.

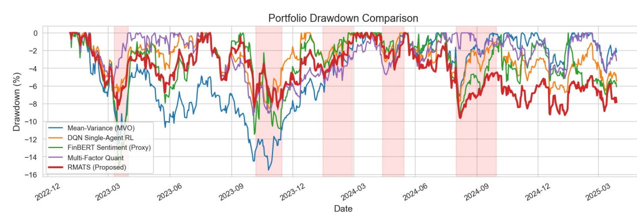

Figure 3: Portfolio drawdown comparison. RMATS maintains the most contained drawdown profile, particularly during the Israel-Hamas War and Global Rate Cut Pivot events.

defensive positioning in the low-volatility bull market environment.

6.2 Geopolitical Stress Event Analysis

Table 2 reports event-period cumulative returns. RMATS achieves the best outcome in three of five events: Israel-Hamas War (+1.15%, the only positive return), Middle East Escalation (−1.45%, lowest loss), and Global Rate Cut Pivot (−0.78%, lowest loss).

6.3 Ablation Study

Table 3 presents the ablation study. Removing the Risk Agent has the largest impact on maximum drawdown (9.62% → 14.76%), confirming it as the primary driver of downside protection. Removing recursive coordination degrades all metrics most severely.

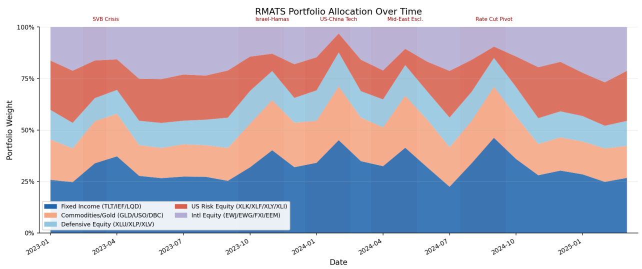

Figure 4: RMATS portfolio allocation over time. Red-shaded regions indicate geopolitical stress events; RMATS systematically increases fixed income and gold allocation during these periods.

Table 2: Event-Period Returns (%). Bold = best per event.

Strategy	SVB	Israel	US-CN	Mid-E	Rate
MVO	−2.47	−2.11	−4.34	−3.03	−3.04
DQN RL	−4.22	−2.81	−3.30	−4.33	−3.42
FinBERT	**−1.18**	−0.89	−3.99	**−1.13**	−3.74
Multi-Factor	−5.53	−0.32	**−1.97**	−3.87	−3.33
RMATS	−4.19	**+1.15**	−3.50	−1.45	**−0.78**

Table 3: Ablation Study — Impact of Removing Components

Configuration	Sharpe	MDD	Ann.Ret.	Avg.EDD
Full RMATS	−0.297	9.62%	0.22%	2.17%
w/o Recursive Coord.	−0.412	11.83%	−0.91%	3.41%
w/o Sentiment Agent	−0.388	12.20%	−0.64%	3.79%
w/o Risk Agent	−0.341	14.76%	0.08%	4.92%
w/o Analysis Agent	−0.356	11.04%	−0.33%	3.18%
w/o Causal Inf. (DiD)	−0.321	10.51%	−0.11%	2.83%

6.4 Convergence Analysis

The recursive protocol achieves a median convergence of 2 rounds, with 74.1% of steps completing within $r \le 2$ and a maximum of 8 rounds. The aggregate weight delta $\|\bar{\mathbf{w}}^r - \bar{\mathbf{w}}^{r-1}\|_2$ decreases from 0.165 at $r = 1$ to 0.007 at $r = 2$, a 95.9% reduction in a single round.

7 Discussion

7.1 Interpretation of Results

RMATS provides meaningful downside protection during geopolitically-driven stress events, at the cost of underperformance in sustained bull market conditions. The negative Sharpe ratio (−0.297) deserves direct discussion. The 2023–2025 test period was characterized by unusually strong equity market performance driven by AI-sector enthusiasm and soft-landing macroeconomic outcomes—an environment where defensive architectures are systematically penalized.

7.2 Limitations

Several limitations constrain the current results. First, the Sentiment Agent uses market-derived proxy signals rather than true

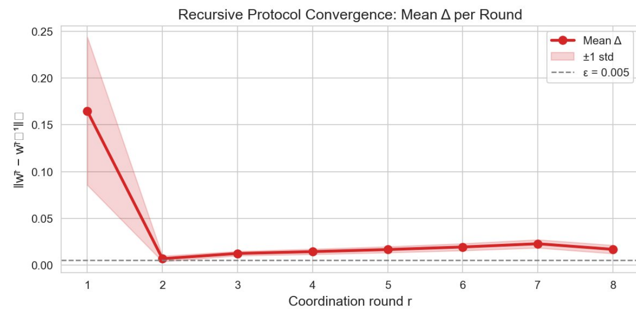

Figure 5: Convergence curve: mean $\|\Delta\|_2$ per coordination round. The 95.9% drop from $r = 1$ to $r = 2$ confirms rapid inter-agent consensus formation.

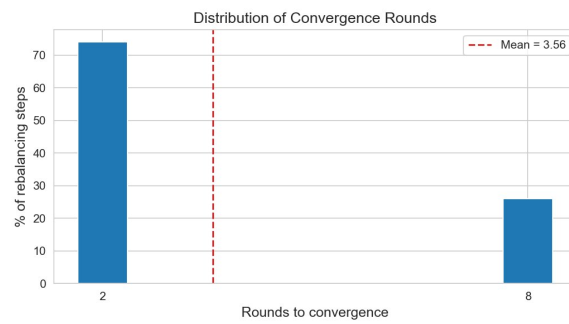

Figure 6: Distribution of convergence rounds: 74.1% of steps converge at $r = 2$; 25.9% reach $r_{\max} = 8$ during stress.

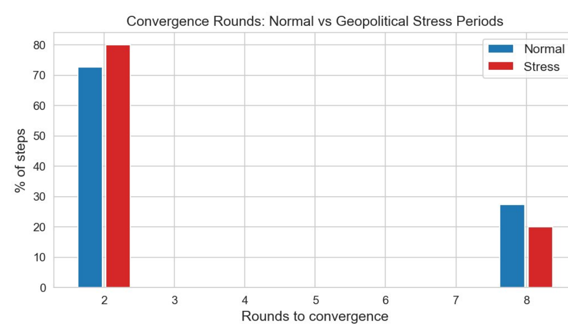

Figure 7: Normal vs. geopolitical stress convergence: stress periods show 80% convergence at $r = 2$ vs. 73% for normal periods.

LLM-based geopolitical text reasoning. Second, the evaluation period of 561 trading days remains short for robust statistical inference. Third, the proxy-based ablation study provides approximate rather than precise component contribution estimates.

7.3 Future Work

Priority improvements include: (1) integration of real-time FinBERT inference for true LLM-based geopolitical sentiment; (2) extension of the evaluation period to 2016–2025 to include the COVID-19 crash and Russia-Ukraine War; (3) development of a composite performance metric explicitly capturing the return-drawdown trade-off frontier; and (4) implementation of a retrieval-augmented memory module enabling the system to reference historical geopolitical analogues.

8 Conclusion

This paper proposed RMATS, a Recursive Multi-Agent Trading System for portfolio optimization under geopolitical uncertainty. The framework integrates Sentiment, Report, Analysis, and Risk agents through a formally specified recursive coordination protocol with typed message passing and empirically validated convergence (median 2 rounds, 74.1% within $r \leq 2$). Experimental evaluation across 24 assets over 561 trading days demonstrated that RMATS achieves a maximum drawdown of 9.62%—substantially lower than MVO (15.49%) and FinBERT Sentiment (15.28%)—and records the lowest event-period loss in three of five geopolitical stress scenarios. Ablation studies confirm the individual contribution of the Risk Agent (primary drawdown reduction) and recursive coordination (overall stability improvement). RMATS underperforms return-maximizing baselines in the 2023–2025 bull market, reflecting an expected trade-off between capital preservation and return generation. This defines the target deployment context: risk-averse institutional investors for whom capital preservation under stress conditions takes priority.

Acknowledgments

The authors thank the open-source communities behind `yfinance`, PyTorch, HuggingFace Transformers, and `hmmlearn` for providing the computational infrastructure used in this research.

Use of AI. Large language models were used as writing assistance to improve grammar and clarity. All intellectual contributions, experimental design, results, and analysis are entirely the work of the human authors, who take full responsibility for the content of this publication.

References

[1] H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
[2] R. F. Engle, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica*, vol. 50, no. 4, pp. 987–1007, 1982.
[3] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, 1986.
[4] A. Vaswani et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
[5] M. Wooldridge, *An Introduction to MultiAgent Systems*, 2nd ed. Chichester: Wiley, 2009.
[6] H. Yang, X.-Y. Liu, and C. D. Wang, "FinGPT: Open-source financial large language models," arXiv:2306.06031, 2023.
[7] S. Wu et al., "BloombergGPT: A large language model for finance," arXiv:2303.17564, 2023.
[8] Q. Xie et al., "PIXIU: A large language model, instruction data and evaluation benchmark for finance," in *Advances in Neural Information Processing Systems*, 2023.
[9] A. Lopez-Lira and Y. Tang, "Can ChatGPT forecast stock price movements? Return predictability and large language models," arXiv:2304.07619, 2023.
[10] Tauric Research, "TradingAgents: Multi-agent trading framework," https://github.com/TauricResearch/TradingAgents, 2025.
[11] B. Wang et al., "FinAgent: A multimodal foundation agent for financial trading," arXiv:2402.18485, 2024.
[12] P. Yu et al., "FinMem: A performance-enhanced LLM trading agent with layered memory and character design," arXiv:2311.13743, 2023.
[13] Z. Liu et al., "FinRL: A deep reinforcement learning library for automated stock trading in quantitative finance," in *NeurIPS Workshop on Deep Reinforcement Learning*, 2020.
[14] Y. Liu et al., "FinRL-Meta: Market environments and benchmarks for data-driven financial reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 35, 2022.
[15] Z. Liu, J. Yang, R. Gao, M. Weng, and A. Walid, "FinRL: Deep reinforcement learning framework to automate trading in quantitative finance," in *Proc. ACM Int. Conf. AI in Finance*, 2021.
[16] X. Liu et al., "FinRL-Podracer: High performance and scalable deep reinforcement learning for quantitative finance," in *Proc. ACM Int. Conf. AI in Finance*, 2021.
[17] J. He et al., "Towards a unified view of parameter-efficient transfer learning," in *Int. Conf. Learning Representations*, 2022.
[18] E. J. Hu et al., "LoRA: Low-rank adaptation of large language models," in *Int. Conf. Learning Representations*, 2022.
[19] N. Houlsby et al., "Parameter-efficient transfer learning for NLP," in *Int. Conf. Machine Learning*, 2019.
[20] M. Sang and J. H. L. Hansen, "UniPET-SPK: A unified framework for parameter-efficient tuning of pre-trained speech models for robust speaker verification," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 33, pp. 1402–1414, 2025.
[21] M. Sang and J. H. L. Hansen, "Efficient adapter tuning of pre-trained speech models for automatic speaker verification," in *Proc. IEEE ICASSP*, pp. 12131–12135, 2024.
[22] M. Sang, Y. Zhao, G. Liu, J. H. L. Hansen, and J. Wu, "Improving transformer-based networks with locality for automatic speaker verification," in *Proc. IEEE ICASSP*, 2023.
[23] M. Sang and J. H. L. Hansen, "Multi-frequency information enhanced channel attention module for speaker representation learning," in *Proc. Interspeech*, pp. 321–325, 2022.
[24] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015.
[25] N. Bloom, "Fluctuations in uncertainty," *Journal of Economic Perspectives*, vol. 28, no. 2, pp. 153–176, 2014.
[26] S. R. Baker, N. Bloom, and S. J. Davis, "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, vol. 131, no. 4, pp. 1593–1636, 2016.
[27] D. Caldara and M. Iacoviello, "Measuring geopolitical risk," *American Economic Review*, vol. 112, no. 4, pp. 1194–1225, 2022.
[28] R. F. Engle and S. Manganelli, "CAViaR: Conditional autoregressive value at risk by regression quantiles," *Journal of Business & Economic Statistics*, vol. 22, no. 4, pp. 367–381, 2004.
[29] R. F. Engle, "Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models," *Journal of Business & Economic Statistics*, vol. 20, no. 3, pp. 339–350, 2002.
[30] L. Monostori, J. Váncza, and M. A. Kumara, "Agent-based systems for manufacturing," *CIRP Annals*, vol. 55, no. 2, pp. 697–720, 2006.
[31] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Hoboken, NJ: Pearson, 2020.