

A Unified Multimodal State Representation Framework for Deep Reinforcement Learning in Financial Trading

Gbemileke Micah ^a, Neda Ahmadi ^{a,*}, Mehrbakhsh Nilashi ^b

^a School of Computer Science and Mathematics, Kingston University, London, United Kingdom

^b Centre for Business Informatics and Industrial Management (CBIIM), UCSI Graduate Business School, UCSI University, Kuala Lumpur, Malaysia

* Corresponding author email address: n.ahmadi@kingston.ac.uk

Abstract

Financial markets are inherently non-stationary, noisy, and driven by complex interactions between price dynamics, news sentiment, and macroeconomic shifts. Traditional reinforcement learning (RL) trading systems predominantly rely on raw price data or technical indicators, often overlooking rich exogenous information. This paper presents a unified multimodal state representation framework designed for deep reinforcement learning trading agents under realistic transaction cost constraints. The proposed pipeline systematically integrates high-frequency market data from Binance (BTCUSD) and OANDA (EURUSD) with offline-engineered features into a standardized 30-minute bar-end state matrix. Specifically, the framework combines price-derived indicators, latent price embeddings and directional probabilities extracted via a 1D Convolutional Neural Network (CNN), event-based media attention and sentiment features from the Global Database of Events, Language, and Tone (GDELT), and resampled macroeconomic indicators (such as the VIX and central bank interest rates). To prevent look-ahead bias and model distortion, all non-price variables are rigorously lagged, winsorized, and scaled using robust normalization fitted exclusively on historical training partitions. By decoupling feature engineering from online RL optimization, this modular architecture ensures absolute reproducibility and provides a robust, low-dimensional state foundation for policy optimization algorithms like Proximal Policy Optimization (PPO).

Keywords: Multimodal Representation Learning, Deep Reinforcement Learning, Feature Engineering, Financial Time Series, Convolutional Neural Networks (CNN), GDELT Event Data, Algorithmic Trading

1. Introduction

Financial markets are complex, non-stationary systems shaped by the convergence of asset price dynamics, macroeconomic policies, and rapid information flow. Historically, automated trading strategies relied heavily on rule-based technical indicators or linear time-series models. However, these traditional methods struggle to adapt to sudden regime shifts and structural volatility. In response, deep reinforcement learning (DRL) has emerged as a powerful paradigm for algorithmic trading. By framing trading as a sequential decision-making process under uncertainty, DRL agents can learn adaptive trading policies through direct interaction with market environments rather than relying on explicit rules [1].

Despite the growing popularity of reinforcement learning in quantitative finance, a critical bottleneck remains in how market states are represented. The vast majority of existing DRL trading literature relies exclusively on raw OHLCV (Open, High, Low, Close, Volume) price data or basic technical indicators. While price series capture immediate micro-structural activity, they omit crucial exogenous drivers - such as geopolitical events, global market sentiment, and monetary policy shifts - that profoundly impact asset valuations and market volatility. Consequently, single-source price models frequently fail to generalize across changing market conditions. To address these limitations, recent studies have begun exploring multimodal architectures that fuse heterogeneous information sources, including news sentiment and macroeconomic indicators, into the learning process [2]. Yet, integrating diverse data types introduces significant technical challenges: