

Binary Trigger Signals for Deep Reinforcement Learning in Equity Trading

Tuesday, 31 March 2026 12:00 (2 hours)

This study introduces a novel binary trigger-based state representation for deep reinforcement learning (DRL) in stock trading. Unlike conventional approaches using continuous technical indicators (MACD, RSI, CCI, ADX), we encode market state via binary signals: MVX (moving-average crossover) and BOLLX (Bollinger band breakout). We also propose trigger-date filtering, which trains only on dates when triggers fire, reducing training data by 50-70%.

Evaluating 27 configurations (three algorithms: A2C, PPO, SAC across nine indicator variants) on Dow Jones 30 daily data (Jan-Nov 2025), we discover a strong algorithm-indicator dependency: A2C with MVX yields +30.85% improvement, PPO with BOLLX achieves +16.09%, while SAC remains robust to both. The best configuration (A2C with filtered MVX) achieves 31.90% cumulative return, a Sharpe ratio of 1.41, and outperforms the DJIA baseline by 154%.

A systematic review of papers (2015-2025) suggests both contributions are novel: no prior work employs binary trigger signals or trigger-date filtering in DRL trading. Results partially validate RL over traditional strategies (37% of models beat DJIA) while showing trigger-date filtering benefits A2C but hurts PPO/SAC. Limitations include the 11-month test period and absence of LSTM temporal modeling, suggesting future work on recurrent architectures and multi-market validation.

This study introduces a novel binary trigger-based state representation for deep reinforcement learning (DRL) in stock trading. Unlike conventional approaches using continuous technical indicators (MACD, RSI, CCI, ADX), we encode market state via binary signals: MVX (moving-average crossover) and BOLLX (Bollinger band breakout). We also propose trigger-date filtering, which trains only on dates when triggers fire, reducing training data by 50-70%.

Evaluating 27 configurations (three algorithms: A2C, PPO, SAC across nine indicator variants) on Dow Jones 30 daily data (Jan-Nov 2025), we discover a strong algorithm-indicator dependency: A2C with MVX yields +30.85% improvement, PPO with BOLLX achieves +16.09%, while SAC remains robust to both. The best configuration (A2C with filtered MVX) achieves 31.90% cumulative return, a Sharpe ratio of 1.41, and outperforms the DJIA baseline by 154%.

A systematic review of papers (2015-2025) suggests both contributions are novel: no prior work employs binary trigger signals or trigger-date filtering in DRL trading. Results partially validate RL over traditional strategies (37% of models beat DJIA) while showing trigger-date filtering benefits A2C but hurts PPO/SAC. Limitations include the 11-month test period and absence of LSTM temporal modeling, suggesting future work on recurrent architectures and multi-market validation.

Student

Yes

Primary author: MONTOYA BAYARDO, Juan Manuel

Co-authors: Prof. BORGELT, Christian (Uni Salzburg); Dr HIRLÄNDER, Simon (Uni Salzburg)

Presenters: MONTOYA BAYARDO, Juan Manuel; Dr HIRLÄNDER, Simon (Uni Salzburg)

Session Classification: Poster session