Abstract

Vice Chancellor’s Awards and Academic Sessions 2026
University of Ruhuna

AI-Assisted Autonomous Trading System for Cryptocurrency Tokens

Artificial Intelligence (AI)-assisted Cryptocurrency Trading System is an end-to-end framework designed to support informed, emotion-free trading decisions in volatile cryptocurrency markets. The system integrates on-chain analytics, off-chain sentiment, and quantitative risk management into an automated decision-making pipeline. A multi-model Natural Language Processing (NLP) architecture, combining FinBERT, DeBERTa-based Named Entity Recognition (NER), and Valence Aware Dictionary and sEntiment Reasoner (VADER), continuously extracts token mentions and contextual sentiment from Twitter. These outputs are aggregated into token-level scores with confidence measures and evidence counts, enabling inspection of underlying tweets that drive each signal.

In parallel, a Trader Selection Framework (TSF) statistically identifies reliable trader personas by analysing profitability, risk-adjusted returns, drawdowns, and behavioural consistency from on-chain data. Wallets that satisfy minimum data and stability criteria are retained as candidate signal sources. A specialised Wallet Tracker module then monitors a configurable subset of influential “whale” wallets in real time, detecting coordinated buying patterns that may signal emerging market moves or accumulation phases before large price changes are visible in Open, High, Low, Close, and Volume (OHLCV) data..

When sentiment scores, whale activity, and market-level indicators such as Exponential Moving Average (EMA), Relative Strength Index (RSI), Average True Range (ATR), and volume spikes align within predefined rules, an orchestrator initiates risk-managed autonomous trades with adaptive position sizing and stop-loss/take-profit controls. Trades are logged with the active signals, thresholds, and wallet events that triggered them, supporting downstream analysis and audit trails..

Developed using a modular, service-oriented architecture, the system exposes APIs for data ingestion, signal retrieval, and order execution and can integrate with external trading APIs on Solana. Docker-based deployment provides isolated, reproducible runtime environments and simplifies migration between development, paper-trading, and production. The project shows how combining AI models, on-chain behaviour analysis, and blockchain-native market data can support systematic, auditable autonomous trading for cryptocurrency tokens..

Keywords: *Artificial Intelligence, Autonomous Trading, Blockchain, Cryptocurrency, Natural Language Processing,.*

AI-Assisted Autonomous Trading System for Cryptocurrency Tokens

1. Introduction and research problem/issue

The rapid expansion of cryptocurrency markets, particularly on high-throughput chains such as Solana, has created an environment where prices can move dramatically within minutes in response to whale activity, social media narratives, and on-chain governance or DeFi innovations [1]-[4]. Human traders and conventional rule-based bots struggle to process this heterogeneous information at scale, leading to delayed, emotionally biased, or poorly risk-managed decisions. Existing trading systems frequently treat on-chain analytics, social sentiment, and risk controls as separate modules, rarely aligning whale behaviour, market microstructure, and community sentiment in a unified decision framework [2]-[5]. They also offer limited transparency regarding why a trade was executed, undermining user trust and making systematic improvement difficult.

This project addresses the research problem of how to design an AI-assisted autonomous trading system that fuses multi-source data to generate explainable, risk-aware trading actions on Solana in real time. The system needs to consume noisy, high-frequency data streams, standardise them into stable features, and combine them in a way that reduces false positives while still reacting early to market-moving events. The work focuses on portfolio-level risk constraints and token-level trade filters rather than pure profit maximisation, so that trading logic can be analysed and adapted across different operating conditions.

The specific objectives are to:

- (i) Model social and community sentiment using domain-aware NLP and token-level aggregation.
- (ii) Identify high-quality trader personas from on-chain data using risk-adjusted performance metrics.
- (iii) Detect coordinated whale accumulation events from wallet transfer patterns.
- (iv) Integrate sentiment, whale activity, and technical indicators into a single decision engine.
- (v) Implement the system as a modular, production-ready architecture suitable for real-world deployment, with clear interfaces for extension and monitoring.

2. Materials and Methods

The system is implemented as a modular AI-assisted trading architecture centred on the Solana blockchain. Off-chain sentiment is captured through a multi-model NLP pipeline combining FinBERT for finance-specific sentiment classification, DeBERTa-based Named Entity Recognition for extracting token symbols and project names, and VADER for polarity analysis on short messages [5], [6]. Data are collected from Twitter through keyword and address-based filters, de-duplicated, and normalised before being aggregated into token-level sentiment scores with associated confidence values, evidence counts, and positive/negative tweet ratios.

A Retrieval-Augmented Generation layer is used as a validation component that can be queried with a token address or symbol and a time window. It retrieves recent social content and sentiment summaries from a curated store and checks whether the computed sentiment is consistent with the underlying messages, which helps reduce spurious correlations caused by short-lived noise or off-topic mentions.

On-chain analytics are driven by a Trader Selection Framework that ranks wallets using return, Sharpe ratio, volatility, drawdown, and behavioural consistency. The framework applies minimum trade-count and look-back window filters before computing scores. A dedicated Wallet Tracker, implemented to monitor preselected whale wallets, normalises token transfer events and detects coordinated buying patterns within configurable time windows and wallet thresholds. Market-level features are derived from OHLCV data and indicators including EMA, RSI, ATR, and volume-spike detection, which are calculated over rolling windows and aligned to the same time grid as the sentiment and wallet signals [1]-[4].

All features are fed into a central orchestration layer that applies rule-based and model-driven strategies, calculates position sizing, and enforces risk limits. Token signals are received through Server-Sent Events (SSE), while OHLCV data flows through WebSockets to provide real-time communication with downstream services. Docker-based deployment ensures isolated, reproducible environments and enables integration with external trading APIs, while a shared data store persists signals and executed orders for later analysis.

3. Results and Discussion

The implemented system delivers an end-to-end AI-assisted autonomous trading framework that fuses sentiment analysis, on-chain behaviour, and quantitative indicators for Solana-based tokens. The modular design decouples data acquisition, analytics, coordination detection, decision-making, and execution, which simplifies maintenance and allows components to evolve without destabilising the system. The Wallet Tracker service continuously monitors influential wallets, stores transfer events, and detects coordinated accumulation when multiple tracked wallets buy the same token (*Fig. 1*) within a configurable window. This behaviour-focused signal complements technical analysis by emphasising the actions of historically profitable traders rather than raw price movements alone.

The Trader Selection Framework (TSF) refines signal quality by scoring wallets on profitability (*Fig. 2*), risk, and behavioural stability. By filtering out erratic or low-quality traders, the system reduces the likelihood of copying noise or manipulative behaviour. Combined with coordinated trade detection, this enables early identification of tokens undergoing accumulation by credible whales (*Fig. 3*), which is valuable in illiquid or newly launched markets where order books are thin and direct price signals are unstable [3], [4]. The framework also records which wallets are currently treated as signal sources, making it easier to review and adjust selection rules when market conditions change or new actors emerge.

Trader details are supplied by a wallet extraction pipeline (*Fig. 4*) that continuously scrapes candidate tokens from “DexScreener”. For each update cycle, the Token Extractor compares the current snapshot with the previous one and partitions tokens into three queues: newly listed tokens, rank-moved tokens, and removed or unranked tokens. These groups are assigned priorities for downstream processing, with newly added tokens treated as highest priority. The Trader Details Extractor consumes these queues, queries “DexCheck” AI for each token, and extracts the top traders associated with recent activity. The resulting trader addresses and statistics are stored in a database and later consumed by TSF, ensuring that the wallet universe tracks fast-changing markets while keeping computational cost bounded.

On the sentiment side, the multi-model NLP pipeline generates token-level outputs [5], [6] (*Fig. 5*) that include aggregate sentiment scores, confidence levels, evidence counts, and positive and negative tweet distributions. The system also returns representative sample texts and polarity classifications, enabling inspection of why a token’s sentiment score was assigned. This structured output improves interpretability, supports post-trade diagnostics, and strengthens auditability. Domain-specific models such as FinBERT help reduce misclassification when handling financial jargon, sarcasm, or rapidly evolving crypto terminology. In practice, the sentiment module acts as an additional filter that can suppress trades when social activity is low or highly polarised, even if price-based indicators appear favourable.

The orchestrator's integration of sentiment signals, whale activity, and indicators like EMA, RSI, ATR, and volume spikes leads to a more nuanced decision surface than any single signal in isolation. Trades are only executed when multiple conditions align (*Fig. 6*) within predefined risk thresholds, and are wrapped in safeguards such as adaptive position sizing, stop-loss and take-profit levels, token exclusion lists, and configurable coordination parameters. These controls are crucial in a market that is prone to pump-and-dump schemes, rug pulls (*Fig. 7*), and extreme intraday volatility. The orchestrator also logs each trade together with the active signals, thresholds, and wallet events that led to execution, allowing users to reconstruct decision paths and to compare realised outcomes with the conditions that were expected at entry.

The end application exposes these decisions through a mission-control dashboard (*Fig. 8*). Tokens selected by the on-chain analysis pipeline arrive in a trade queue via a WebSocket server. For each candidate token, the sentiment module performs a real-time check; only tokens that pass this test move into the active position panel and receive entry, stop-loss, and take-profit levels. The top bar displays cumulative profit and loss (P&L) across all trades, while the queue shows per-token sentiment percentiles and realised P&L, providing an operational view of how upstream signals translate into live positions.

Overall, the observed behaviour of the integrated framework indicates that combining wallet-level coordination, sentiment structure, and technical indicators yields trade recommendations that are more selective than pure momentum or copy-trading rules. In periods of conflicting signals, the system tends to stay out of the market, which directly addresses the objective of prioritising risk control over constant position-taking. These properties suggest that the architecture can serve as a basis for live trading and as a research platform for evaluating alternative fusion rules, ranking schemes, and execution policies for cryptocurrency tokens.

4. Conclusions

This project presents an AI-assisted autonomous trading system that unifies social sentiment, on-chain trader behaviour, and quantitative market indicators to support risk-aware trading of Solana-based tokens. By coupling a production-ready Wallet Tracker for coordinated whale activity, a Trader Selection Framework for identifying credible personas, and a multi-model NLP and sentiment aggregation stack, the system moves beyond conventional rule-based bots toward a data-driven decision framework in which each signal can be examined and revised.

The modular architecture, together with Dockerised deployment and well-defined APIs, facilitates integration with external trading platforms and downstream analytics tools, and allows individual services to be updated without redesigning the entire system. While the current implementation focuses on Solana and emphasises architectural robustness and explainability, the underlying design principles are portable to other blockchains and asset classes that expose sufficient on-chain and market data. Future work will include formal backtesting across multiple market regimes, evaluation of latency and slippage effects in live markets, incorporation of reinforcement learning or adaptive policy optimisation on top of the existing signal layer, and expansion of risk analytics to cover cross-chain exposures and links to macroeconomic variables [7].

References (Selected)

[1] N. Jagannath *et al*., "An On-Chain Analysis-Based Approach to Predict Ethereum Prices," *IEEE Access*, vol. 9, pp. 167972–167989, 2021, doi: 10.1109/ACCESS.2021.3135620.

[2] G. Kim, D.-H. Shin, J. G. Choi, and S. Lim, "A Deep Learning-Based Cryptocurrency Price Prediction Model That Uses On-Chain Data," *IEEE Access*, vol. 10, pp. 56232–56248, 2022, doi: 10.1109/ACCESS.2022.3177888.

[3] A. Muminov, O. Sattarov, and J. Cho, "Forecasting Bitcoin Volatility Through On-Chain and Whale-Alert Tweet Analysis Using the Q-Learning Algorithm," *IEEE Access*, vol. 11, 2023, doi: 10.1109/ACCESS.2023.3282170.

[4] D. Herremans and K. W. Low, "Forecasting Bitcoin Volatility Spikes from Whale Transactions and CryptoQuant Data Using Synthesizer Transformer Models," *IEEE Access*, vol. 13, pp. 117788–117807, 2025, doi: 10.1109/ACCESS.2025.3459022.

[5] K. Jahanbin and M. A. Z. Chahooki, "Aspect-Based Sentiment Analysis of Twitter Influencers to Predict the Trend of Cryptocurrencies Based on Hybrid Deep Transfer Learning Models," *IEEE Access*, vol. 11, pp. 121656–121670, 2023, doi: 10.1109/ACCESS.2023.3327060.

[6] W. Jung and S. Jang, "Enhancing Financial Sentiment Analysis Ability of Language Model via Targeted Numerical Change-Related Masking," *IEEE Access*, vol. 12, pp. 7170–7198, 2024, doi: 10.1109/ACCESS.2024.3364515.

[7] T. Kabbani and E. Duman, "Deep Reinforcement Learning Approach for Trading Automation in the Stock Market," *IEEE Access*, vol. 10, pp. 93564–93574, 2022, doi: 10.1109/ACCESS.2022.3203114

Annex

Tables and Figures


```
: keepalive

: keepalive

event: data
data: {
  "tokenAddress": "CpqEUjzrco1eu4Te9JQk7iNtL9EQCGFRLrjYBN6zpump",
  "uniqueWalletCount": 4,
  "walletAddresses": ["T5474UjB6yBUUYLzBzKDWS1ZvTsxnLdn5dCJ0KlV9cZQ", "ZWaY82nNpRY2gDRpwsguTGWt0xEn7TdrHjMN0LjeVdWR",
"fxdrxlvu1AP3X9H1b0pL2xJRTiL2V8U51QdYdLGk9kzr", "hV7VKNW4gsSl3E2cacy4GUfzoNRvGv7lMrQIO86mYTnV"],
  "windowStart": "2025-11-14T06:51:11.967Z",
  "windowEnd": "2025-11-14T06:51:16.967Z",
  "triggeredAt": "2025-11-14T06:51:14.678Z"
}

: keepalive

: keepalive

: keepalive

: keepalive

: keepalive

event: data
data: {
  "tokenAddress": "2R5awbjoGYhzdXU5gErRtAQK3hyKmRfnJBeJeVvMpump",
  "uniqueWalletCount": 4,
  "walletAddresses": ["sw3W6IWbYCQpjj3s8IEMnAR25TjeFUghLZJwT05eQ14B", "KBeRG8htCcinGgrV9fLdyzOQpxrFhTQtkyWM5gdPhive",
"1u6Mi9fZRhkv03OOHbemMpNc4wGi7kPM1bELSoWTYOlz", "uiJoFO6HLkq3NYu9BsAkMD4F4wtRjuLkvZpXCa7S7ZZj"],
  "windowStart": "2025-11-14T06:51:21.969Z",
  "windowEnd": "2025-11-14T06:51:26.969Z",
  "triggeredAt": "2025-11-14T06:51:24.171Z"
}

: keepalive

: keepalive

: keepalive
```


Fig. 1. Real-time stream of transfer events for tracked wallets delivered by the Wallet Tracker over Server-Sent Events (SSE), highlighting a coordinated accumulation event on a single token.

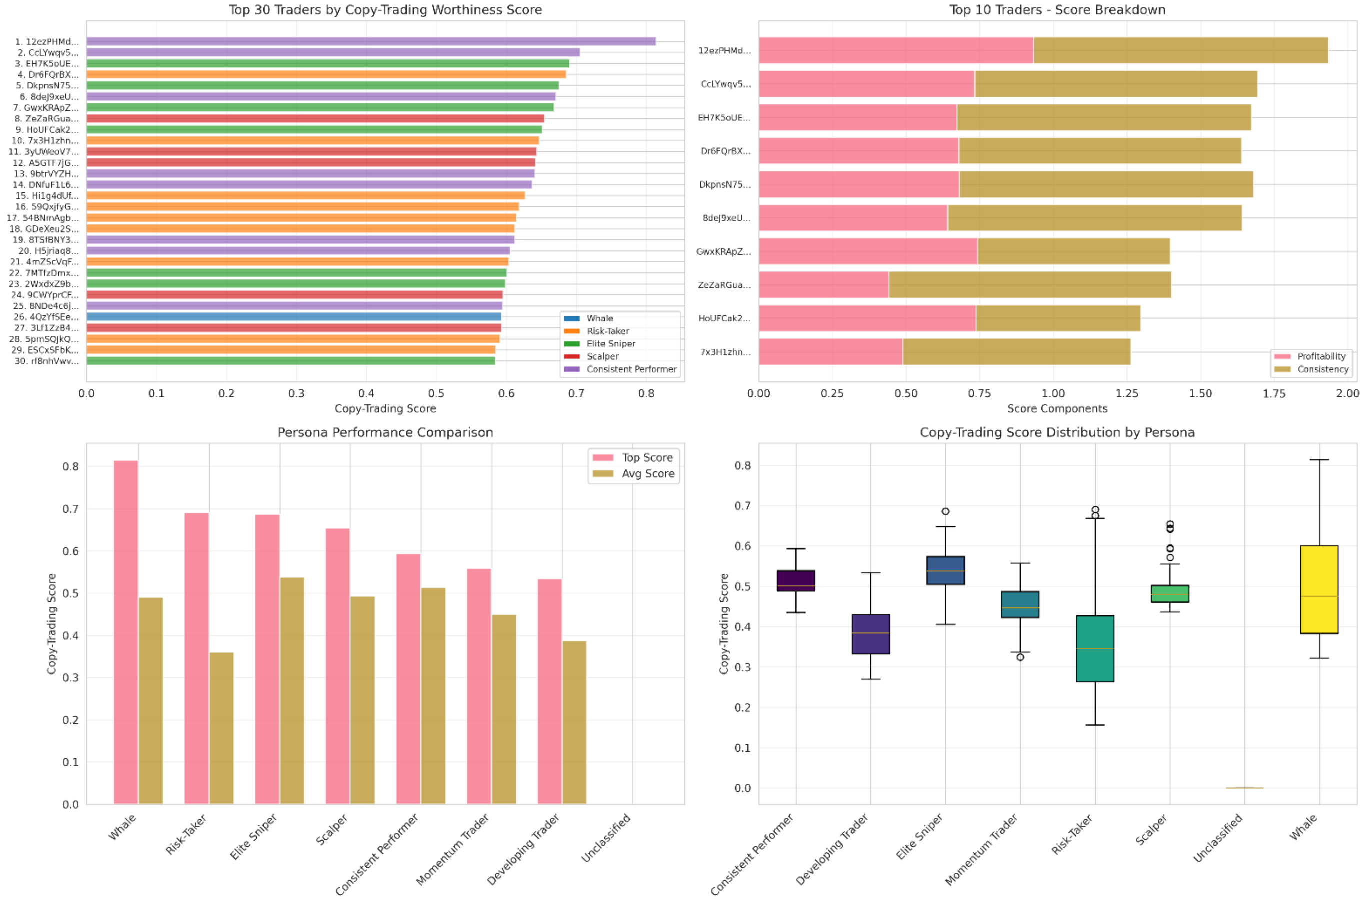

Fig. 2. Output of the Trader Selection Framework (TSF) ranking wallets by quality score, profitability, risk-adjusted return, win rate, and trade count, used to auto-classify "whale" trading personas

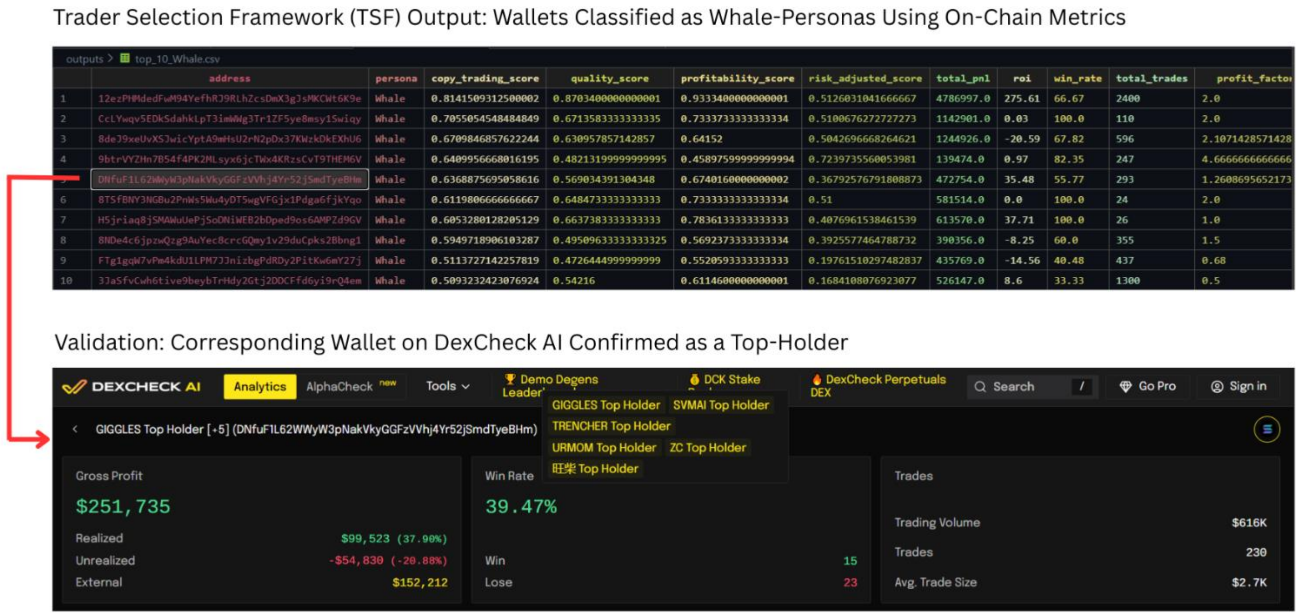

	address	persona	copy_trading_score	quality_score	profitability_score	risk_adjusted_score	total_pnl	roi	win_rate	total_trades	profit_factor
1	12ezPHMdedFwM94YefhRJ9RLhZcsDmX3gJsMKCWt6K9e	Whale	0.8141509312500002	0.8703400000000001	0.9333400000000001	0.5126031041666667	4786997.0	275.61	66.67	2400	2.0
2	CcLYwqv5EDkSdahkLpT3imWWg3Tr1ZF5ye8msy1Swiqy	Whale	0.7055054548484849	0.6713583333333335	0.7333733333333334	0.5100676272727273	1142901.0	0.03	100.0	110	2.0
3	8deJ9xeUvXSJwicYptA9mHsU2rN2pDx37KWzkDkEXhU6	Whale	0.6709846857622244	0.630957857142857	0.64152	0.5042696668264621	1244926.0	-20.59	67.82	596	2.1071428571428
4	9btrVYZHn7B54f4PK2MLsyx6jcTWx4KRzsCvT9THEM6V	Whale	0.6409956668016195	0.4821319999999995	0.4589759999999994	0.7239735560053981	139474.0	0.97	82.35	247	4.6666666666666
5	DNfuF1L62WWyW3pNakVkyGGFzVVhj4Yr52jSmdTyeBHm	Whale	0.6368875695058616	0.569034391304348	0.6740160000000002	0.36792576791808873	472754.0	35.48	55.77	293	1.260869565217З
6	8TSfBNY3NGBu2PnWs5Wu4yDT5wgVFGjx1Pdga6fjkYqo	Whale	0.6119806666666667	0.6484733333333333	0.7333333333333334	0.51	581514.0	0.0	100.0	24	2.0
7	H5jriaq8jSMAWuUePjSoDNiWEB2bDped9os6AMPZd9GV	Whale	0.6053280128205129	0.6637383333333333	0.7836133333333333	0.4076961538461539	613570.0	37.71	100.0	26	1.0
8	8NDe4c6jpzwQzg9AuYec8crcGQmy1v29duCpks2Bbng1	Whale	0.5949718906103287	0.4950963333333325	0.5692373333333334	0.3925577464788732	390356.0	-8.25	60.0	355	1.5
9	FTg1gqW7vPm4kdU1LPM7JJnizbgPdRDy2PitKw6mY27j	Whale	0.5113727142257819	0.4726444999999999	0.5520593333333333	0.19761510297482837	435769.0	-14.56	40.48	437	0.68
10	3JaSfvCwh6tive9beybTrHdy2Gtj2DDCFfd6yi9rQ4em	Whale	0.5093232423076924	0.54216	0.6114600000000001	0.1684108076923077	526147.0	8.6	33.33	1300	0.5

Fig. 3. Cross-check of a TSF-classified whale wallet against DexCheck AI, confirming that the same address appears as a top holder and thereby validating the credibility of TSF-generated whale classifications.

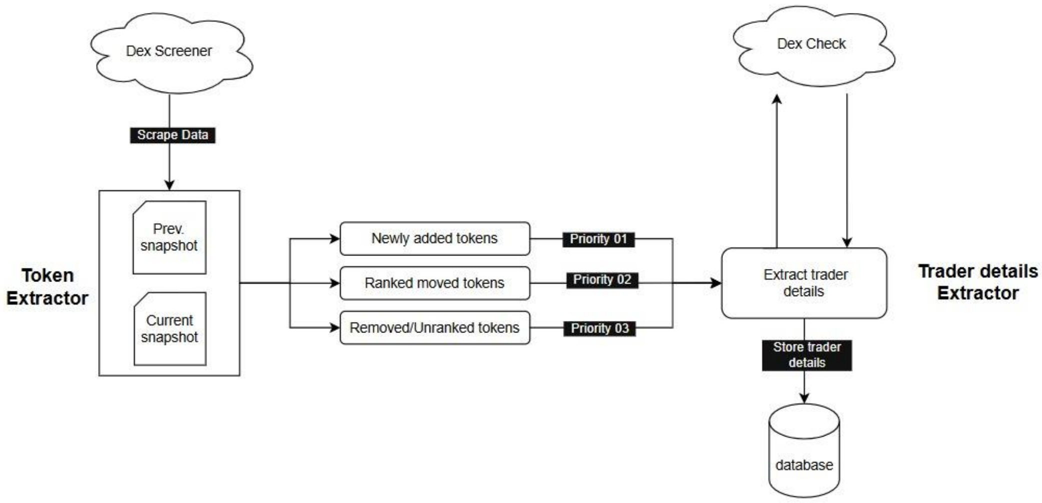

Fig. 4. Wallet extraction and trader discovery pipeline. Dex Screener snapshots are diffed to identify newly listed, rank-moved, and removed tokens, which are queued by priority and forwarded to DexCheck for trader detail extraction.

```
{
  "coin": "BTC",
  "found": true,
  "score": 0.0357,
  "why": "twitter_sent=+0.71",
  "mode": "static-weights",
  "evidence": 3,
  "confidence": 0.5,
  "sources": [
    "twitter_sentiment"
  ],
  "twitter_details": {
    "pos": 40,
    "neg": 5,
    "total": 45,
    "pos_pct": 88.89,
    "neg_pct": 11.11,
    "shrunk_sentiment": 71.42857142857142        sentiment after rag
  },
  "raw": {
    "news_sent": 0.12,        News sentiment value
    "general_sent": 0.18,        General sentiment value
    "focus_sent": 0.21,        Focus group sentiment value
    "flow": 0.11,        Flow Analysis value
    "mentions": 67,        Total tweets including neutral
    "twitter_pos": 40,        positive tweets count
    "twitter_neg": 5,        negative  tweets count
    "twitter_total": 45,        total tweets count without neutral
    "twitter_pos_pct": 88.89,        positive twiter precentage
    "twitter_neg_pct": 11.11,        negative twiter precentage
    "twitter_sent": 0.7142857142857142,        sentiment after rag
    "sources": [
      "twitter_sentiment"
    ],
    "score": 0.0357,        pipeline score value
    "evidence": 3,        direct pipeline coin mentions
    "confidence": 0.5,
    "score_breakdown": {
      "news_sent": 0.12,        News sentiment Z value
      "general_sent": 0.18,        General sentiment z value
      "focus_sent": 0.21,        Focus group sentiment Z value
      "flow_z": 0,
      "mentions_z": 0,
      "twitter_sent": 0.7142857142857142,
      "_mode": "static-weights"        Weight mode selction
    }
  },
  "sample_texts": [        Sample tweets
    {
      "text": "Good morning, fam \n\nTwo key levels for $BTC this week: we need to flip $107.2k;
also, be aware of the CME gap left behind.",
      "sentiment": "positive",
      "color": "green"
    },
    ....
    ....
  ]
}
```


Fig. 5. Example token-level sentiment record produced by the multi-model NLP pipeline, showing BTC's aggregate score, confidence, evidence count, Twitter sentiment breakdown, and sampled source tweets.

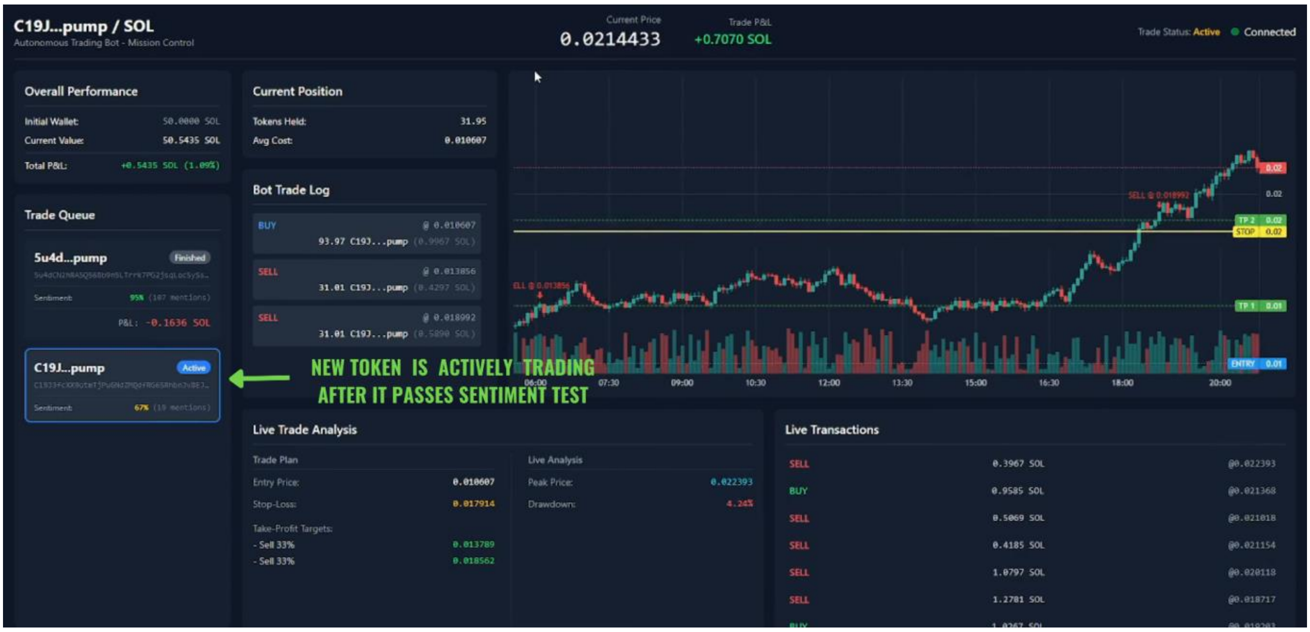

Fig. 6. Mission-control dashboard during an active trade: a new token promoted from the on-chain queue after passing the sentiment test receives entry, stop-loss, and take-profit levels once on-chain, sentiment, and technical signals align

Fig. 7. Example of a memecoin rug pull, illustrating the characteristic pattern of rapid liquidity removal and price collapse observed in market and wallet-level data.

Fig. 8. Integrated AI-assisted trading framework linking wallet extraction, TSF scoring, sentiment analysis, orchestration logic, and execution engine into a single modular pipeline.