

AI Trading Bots Review: What Works in 2025 CHEATSHEET

Profit From AI Revolution

Wealth From AI

Based on: AI Trading Bots Review: What Works in 2025

Overview

- **Evaluation Framework for AI Trading Bots**
 - Use Case Definition:** Specifies the trading strategy (e.g., high-frequency arbitrage, trend following) to guide model choice and data requirements.
 - Model Architecture:** Transformer-based LLMs for signal generation, recurrent networks for time-series forecasting.
 - Dataset Requirements:** Includes market symbols, order-book depth, macroeconomic indicators, and alternative data (news sentiment, satellite imagery).
 - Data Pipeline:** Ingests raw ticks, normalizes features, creates tokenized sequences for model input.
 - Benchmarking:** Compares AI model performance against baseline statistical models (e.g., ARIMA).
 - Key Metrics:** Sharpe ratio, max drawdown, turnover, latency (market update to signal), throughput (symbols processed/sec).
 - Fine-tuning:** Optimizes parameters (learning rate, batch size) on a validation walk-forward set to prevent look-ahead bias.
- **Key Technical Components and Integrations**
 - Model Hosting:** Hugging Face Inference Endpoints, custom PyTorch serving (TorchServe).
 - Experimentation Frameworks:** LangChain for chaining data retrieval, prompt construction, and model calls, especially for LLMs.
 - Data Streaming:** Apache Kafka, AWS Kinesis for real-time market tick data.
 - Feature Stores:** Feast for managing embeddings (dense vectors from raw prices/volumes).
 - Inference Service:** REST or gRPC API for model output.
 - Execution Subsystem:** Lightweight SDK (Python/Go) to subscribe to inference API.
- **Deployment and Performance Considerations**
 - Containerization:** Docker for packaging applications.
 - Orchestration:** Kubernetes for managing containerized applications.
 - Autoscaling:** Based on queue depth to maintain latency thresholds (e.g., <10 ms for ultra-low-latency).
 - Monitoring:** Tracks model drift by comparing prediction distributions to training data.
 - Retraining Trigger:** Statistical significance tests on model drift.
 - Performance Example:** Fine-tuned DistilBERT with price ticks and FinBERT sentiment achieved an annualized Sharpe ratio of 1.

Want more? Visit <https://wealthfromai.com>

Get the full article, plus exclusive guides and resources.