Education

ELTE

Multimodal Forecasting of Stock Prices using GPT-2 Embeddings and Dynamic Graph Networks.

Luyanda Mjiyakho

Supervisor: Adrián Csiszárik

PRESENTATION

RECAP OF THE ARCHITECTURE AND GOAL OF THE PROJECT

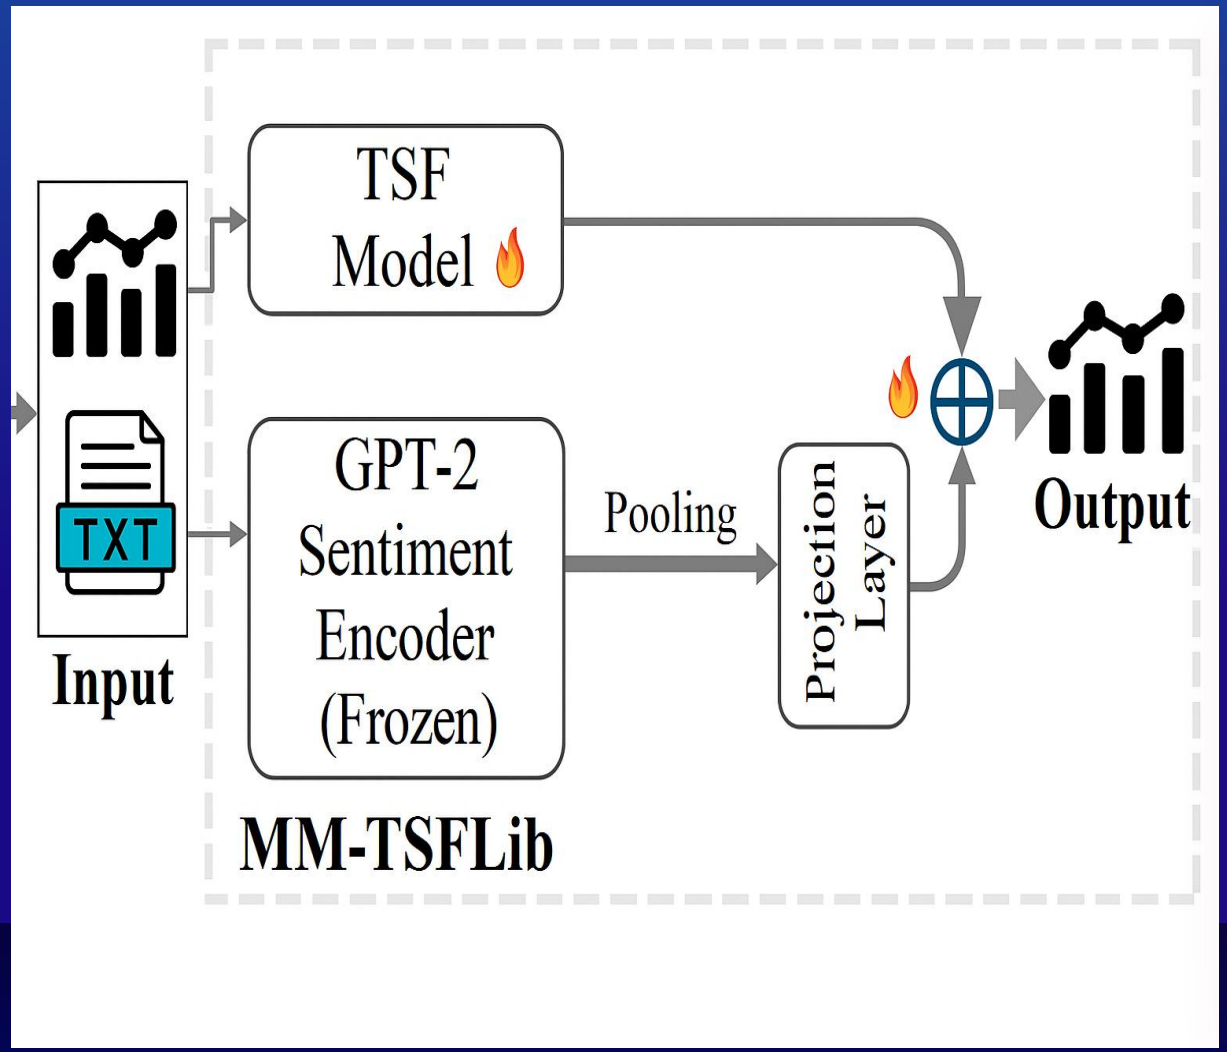

RESULTS FROM LAST SEMESTER.

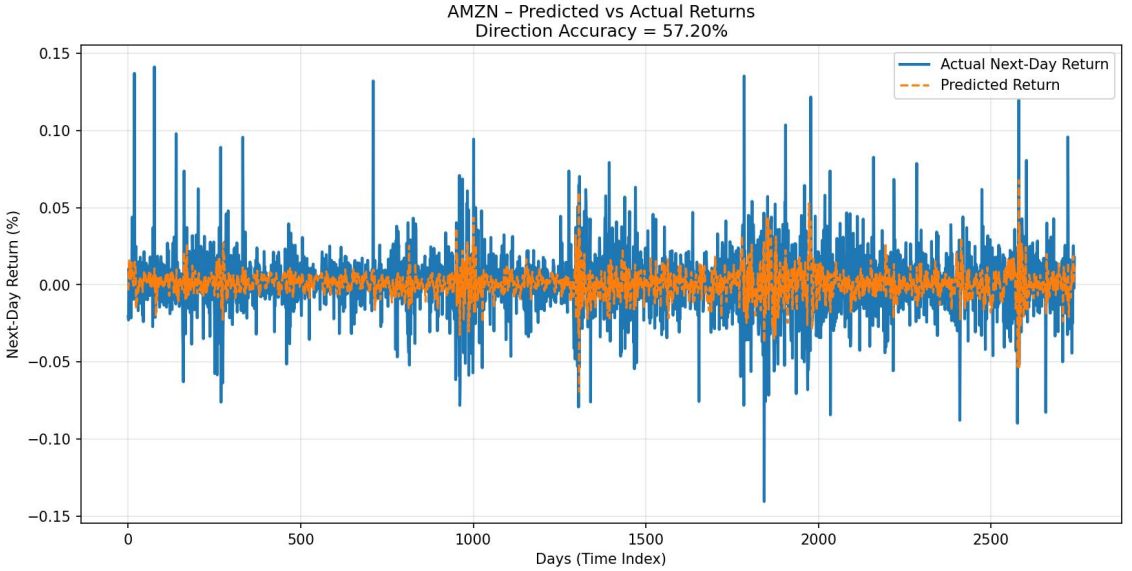

1.Tested whether the MM-TSFlib multimodal framework can forecast single-stock price movements by combining news sentiment and historical market data.
2.Extracted daily news sentiment using a frozen GPT-2 model, converted it into mean-pooled sentiment vectors, and aligned it with S&P 500 data from 2015–2025.
3.Achieved a test RMSE of 0.0148 (1.48%) and 57.20% directional accuracy on target stocks such as AMZN.

Shifting from Asset Isolation to Market Interdependence

01

Stock prices are connected

when major news impacts one company, the effects often spread across related businesses and industries.

02

Project expansion

expanded the project from analyzing individual stocks to studying the entire Nasdaq market using data from January 2020 to May 2026

03

framework learns how companies influence one another over time.

04

The goal is to uncover hidden market relationships while improving next-day stock return predictions.

The new DEEP LEARNING ARCHITECTURE.

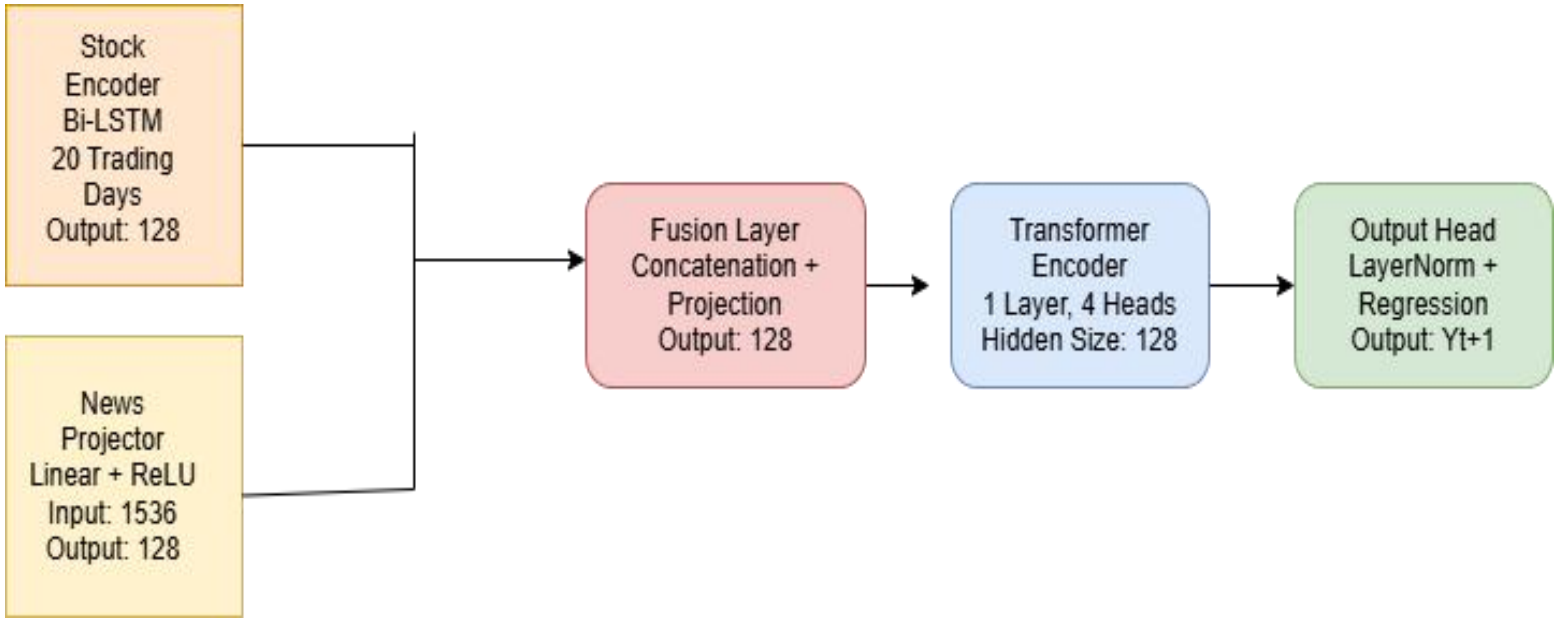

- **Feature Extraction: A Bi-LSTM processes the previous 20 trading days of stock data, while news articles are converted into text embeddings to capture market sentiment.**
- **Multimodal Fusion: Stock data, news sentiment, and categorical information are combined into a shared representation, allowing the model to learn from multiple sources simultaneously.used representation = activation(stock features + news features + company features)**
- **Relational Learning: A 4-head Transformer Encoder analyzes interactions between different stocks, helping the model capture market-wide relationships and dependencies.**

Enforcing Mathematical Rigor & Temporal Sanity

- **Stage 1 – Feature Extraction: The model learns patterns from the previous 20 trading days while also capturing information from financial news articles.**
- **Stage 2 – Multimodal Fusion: Market data, news sentiment, and company-specific information are combined into a single representation of each stock.**
- **Stage 3 – Relational Learning: A 4-head Transformer analyzes how stocks influence one another, helping the model capture market-wide relationships and dependencies.**

This Semester's Financial Network Baseline Results

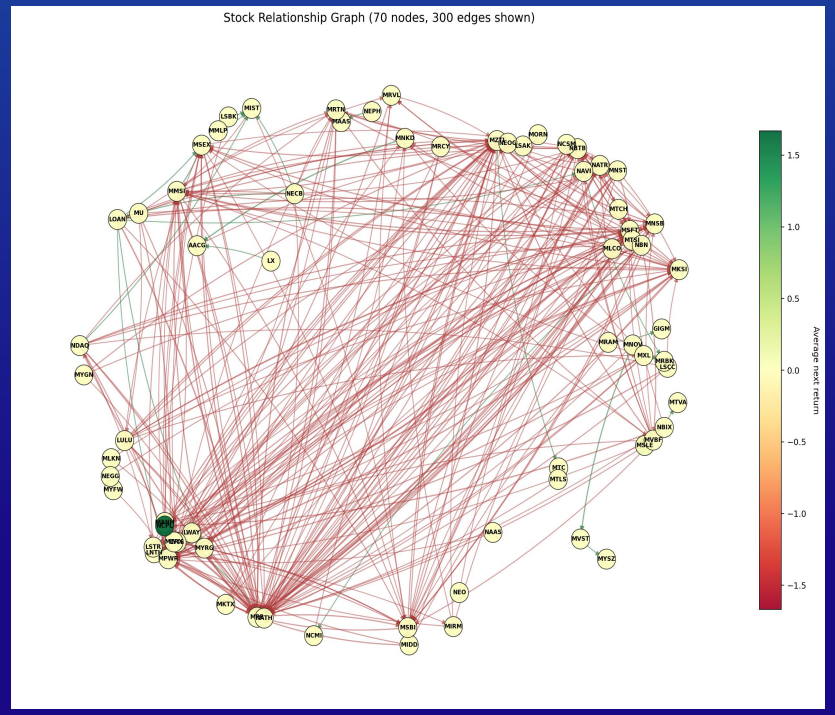

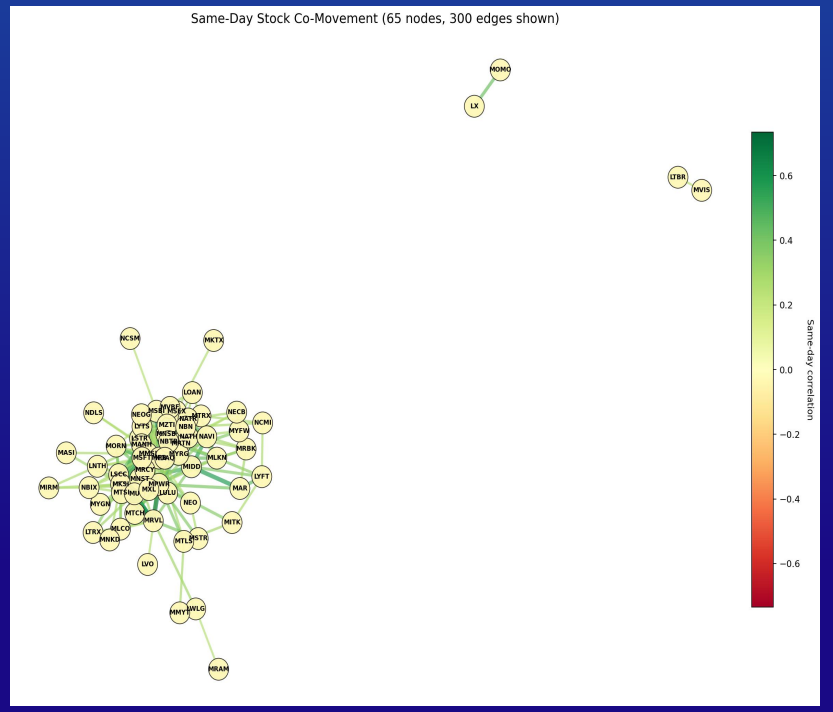

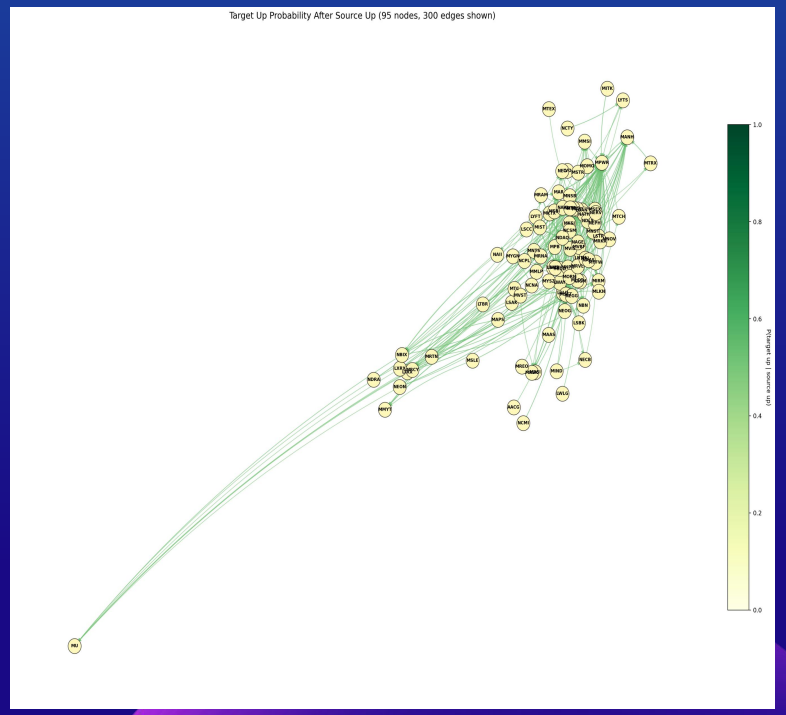

Minimum overlapping trading history > 60 days.

Minimum overlapping trading history > 60 days.

Minimum overlapping trading history > 60 days.

Shifting from Asset Isolation to Market Interdependence

01 Minimum overlapping trading history > 60 days.

02 Same-Day Co-movement (Synchronous Pearson Correlation) > 0.25.

03 Lead-Lag Trajectories (Asynchronous Predictive Probability)> 0.05.

04 Node constraint: Outward connections capped at the top 10 strongest relationships per stock.

Future work

Understand market relationships

- Transition the data pipeline from loading static, historical .csv records to a live streaming API engine for continuous inference.

Move to real-time forecasting

- Extract and visualize the internal dynamic attention weight matrices directly from the trained 4-head Transformer Encoder layers. This will allow us to mathematically cross-examine neural-learned weights against our classical empirical baseline graphs.

Detect unusual market behavior

- Program an automated structural anomaly detection engine. The system will flag automated mathematical triggers if established cross-ticker connection weights suddenly break down or if brand-new, unmapped dependencies emerge in real-time.

Thank you For Listening and watching

AI WAS USED TO ASSIST WITH THE CODING PART AND PROJECT LAYOUT ACCORDING TO HOW I WANTED IT TO BE AS WELL AS SUMMARIZING THE REPORTS AND SOME IMPORTANT POINTS TO ENSURE THAT THE MAIN POINTS THAT I WANTED RELAYED REGARDING THE PROJECT ARE HIGHLIGHTED.