

Fish vs. AI

Results and demonstration from 2024/25:



Figure 1: Implementation of a YOLO object detection model on a real fish to map its positioning to buy/sell orders of £30 total.




Figure 2: Team Portfolio Performance over a trading period. Three competing agents: Piranhas - Quantitative regression model using historical performance and correlation variables to the 2016 US Election-related stock movements; Autobots - NLP based model evaluating sentiment analysis of Bluesky social media data; Benji - Control agent acting as baseline for a random walk.

(Project brief sent prior to results found above):

For help contact:


email: dev.thakkar@warwick.ac.uk

discord: devt_

Fish vs. AI

Machine Learning Trading

(Group Project)



High-Level Explanation

In **Fish vs. AI**, you'll dive into the world of **algorithmic trading** and **machine learning** with a twist: you'll be competing against a fish making random trades! Inspired by the viral video from Michael Reeves, this course lets you develop trading algorithms that will face off against trades chosen by pure chance. You'll get **hands-on experience building** and **testing** your own **trading models using real market data**, learning how to **apply machine learning techniques** to **predict trends** and **execute trades**.

Python will be the main tool here, alongside **libraries like Scikit-learn and PyTorch**, which make it possible to build machine-learning models capable of navigating complex data and market fluctuations. With a foundation in ML, you'll see how AI can enhance decision-making and adapt to market dynamics, allowing you to create smarter, data-driven trading strategies that try to outsmart the fish.

Success Criteria

For the course to be a success, each team will aim to:

1. **Develop a working trading algorithm** that **processes market data** and **makes trade decisions**.
2. Implement **at least one machine learning technique**, like regression or decision trees, to inform trade choices.
3. Successfully **backtest** their **algorithm** on historical data to check its performance.
4. Compete in the **Fish vs. AI challenge**, where your model's performance is compared to trades chosen by the fish.
5. **Present your algorithm and results**, discussing how well it performed and how you'd improve it - this entails making a small youtube video going over your project as a team and what you achieved (maker portfolio video).

For help contact:

email: dev.thakkar@warwick.ac.uk

discord: devt_

Resources Required

Here's what we'll be working with:

Tech Stack:

- **Python:** with open-source libraries, specifically Scikit-learn, PyTorch, and Pandas for data processing.
- **Market Data Access:** Use free resources, such as Yahoo Finance API or Kaggle datasets, to access historical stock market data.
- **Backtesting in Python:** Basic backtesting frameworks can be set up directly in Python with Pandas to simulate how strategies would have performed historically without using paid platforms.

Proposed Architecture

Each team's model will follow a clear and accessible structure:

1. **Data Ingestion:** Pull in and clean real or historical market data from free sources, setting it up for use in your algorithm.
2. **Feature Engineering:** Create useful features, such as moving averages, price trends, or volatility metrics, that help the model make informed trade decisions.
3. **Model Training:** Use Scikit-learn or PyTorch to build and train an ML model that adapts to trends in the data.
4. **Backtesting:** Simulate your model's performance on historical data using custom-built backtesting functions in Python, with results that help you evaluate performance and make improvements.
5. **Competition Deployment:** Final models will run in a simulated environment, where each team's algorithm will be benchmarked against random trades made by the fish.

Alternatives Considered

1. **Other Algorithm Types:** Rule-based strategies (like moving averages) were considered for simplicity but didn't offer the same flexibility or excitement as machine learning.
2. **Higher-End Trading Software:** Platforms like MetaTrader offer a professional setup, but we kept it accessible with open-source or free options.
3. **Stretch Goals:**
 - For advanced teams, there's room to try out reinforcement learning.
 - If budget permits, strategies could even be tested on live data.

For help contact:

email: dev.thakkar@warwick.ac.uk

discord: devt_

Training Resources

To make sure everyone has the tools they need, we've curated a few beginner-friendly resources:

1. **Books and Docs:**
 - *"Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow"* by Géron A.
2. **Videos and Websites:**
 - *YouTube Series on Algorithmic Trading* to cover key skills and concepts.
3. **Google Collab:**
 - Google Collab document to help you get started

Conclusion

With these resources, anyone can jump into *Fish vs. AI*, learn some machine learning, and give the fish a run for its money. Dive in and see if your algorithm has what it takes!

Resources

- [I Gave My Goldfish \\$50,000 to Trade Stocks](#) by Michael Reeves
- [Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow](#) by Géron A.
- [YouTube Series on Algorithmic Trading](#)
-  **Goldfish vs AI - Getting Started.ipynb**

For help contact:

email: dev.thakkar@warwick.ac.uk

discord: devt_