Axyon AI

WHITE PAPER

Learning-to-Rank for Quantitative Investing

Table of Contents

About Us 3

What is Learning-to-Rank? 4

Types of LTR Algorithms 5

Pointwise LTR 5

Pairwise LTR 5

Listwise LTR 6

Ranking Financial Assets with ML 7

Case study 8

About Us

We are at the forefront of AI-powered investment management, leveraging the latest AI/ML advancements.

Through close partnerships with CINECA, major European research centres, and world-renowned universities specialising in AI and computer science, we gain access to state-of-the-art HPC infrastructure and AI talent. This enables us to drive pioneering R&D initiatives that seamlessly translate into practical, high-impact solutions.

Founded in 2016, Axyon AI is an Italian fintech company committed to revolutionising the investment management industry through the power of advanced Predictive AI.

Our mission is to empower investors with scalable research-led AI technologies that deliver timely signals and predictive insights.

Our Solutions

The power of advanced Predictive AI solutions for Alpha generation:

AI-BASED SIGNALS

AI-based assets performance signals that identify under and outperformers within a given investment universe & time horizon.

AI-POWERED MODEL STRATEGIES

Off-the-shelf and bespoke AI-powered model strategies and indices on several asset classes and investment universes.

Introduction

The diversity problem is a significant challenge in ensemble learning, a machine learning technique that combines multiple models to improve overall performance. This diversity helps to average out errors and improve the ensemble's ability to generalise, or perform well on unseen data.

At Axyon AI, the issue is amplified by the need for accurate and resilient predictions in financial asset management. This challenge becomes critical when ensembles fail to leverage distinct patterns in the data, resulting in weak prediction quality and increased risk for end-users. Moreover, ensuring diversity in asset ranking environments poses significant challenges especially in ensemble learning frameworks.

About the problem

Chances are you have used an AI-based Learning-to-Rank (LTR) algorithm in the last hour without even knowing about it.

Google Search and most of today's search engines are equipped with LTR technologies, providing you with accurate results among hundreds of billions of web pages.[1]
In this whitepaper, we introduce the concept of Learning-to-Rank and see how it can be applied to financial assets, with the end goal of building quantitative investment strategies.

What is Learning-to-Rank?

Learning-to-Rank (LTR) is a family of methods that apply supervised learning algorithms to solve ranking problems given a set of items and their relevancy to a certain query. As mentioned, LTR algorithms are commonly applied to search engines, where users provide queries through a search bar and the system returns websites ranked by their estimated relevancy.

However, applications of LTR methods span way beyond web search. In computational biology, these algorithms are used to rank candidate 3D structures of proteins or in recommender systems to suggest the most relevant articles/ movies/ products to a user based on their past preferences.

The first step for training Learning-to-Rank models is preparing a set of queries with an ordered list of results for each query (and, optionally, a relevancy score for each result). This labelled set of data is called the "ground truth" dataset, and we can use it to train our algorithms and evaluate our predictions against it.

Once trained, a good ML model will output accurate rank predictions not only for the ground truth data but also for new, previously unseen instances.

1. https://www.google.com/search/howsearchworks/how-search-works/organizing-information/

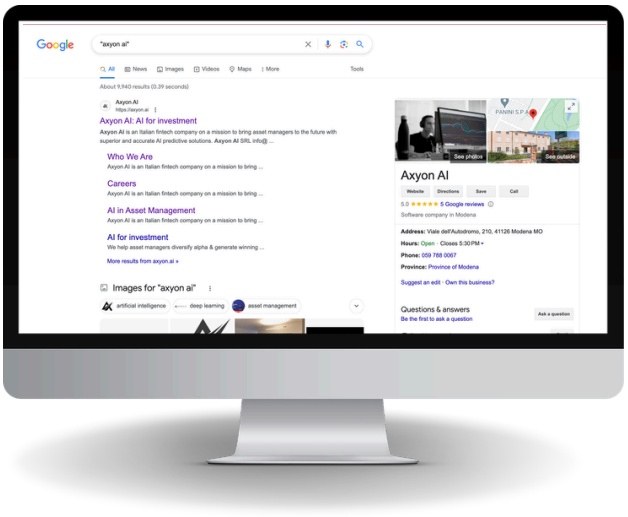

Figure 1:
An example of Learning-to-Rank applied to Web Search. When a user enters the query "axyon ai" on Google Search, the system returns a list of websites ranked by relevance.

Types of LTR Algorithms

LTR techniques can be broadly classified into three following categories: pointwise, pairwise, and listwise ranking algorithms, based on the set of modelling choices and underlying assumptions. A summarized comparison of advantages and disadvantages of these methods is shown in Table 1.

Pointwise LTR

The simplest modelling approach to rank items is looking at them one at a time and producing an estimated relevancy score based on the given query, then sorting all items according to these estimates. We can train a regressor to predict the true score, or quantize the ground truth relevancy and use a classifier instead.

A remarkable limitation of this simpleapproach is that the prediction for one item is produced independently of other items: each item is considered a "point" for the ranking, independent of other "points" (hence the name).

Pairwise LTR

Algorithms in this class look at two items at a time and decide which is more relevant with respect to a given query. This can be framed as either a binary classification problem or as a regression problem by predicting the relevancy difference (if available). Compared to Pointwise LTR, this approach is more computationally costly, as requires evaluating all possible pairs of items.

Listwise LTR

In Listwise LTR, algorithms are provided with the entire list of items and output the optimal ordering with respect to the given query. This means that the model is not trained to solve a surrogate task (e.g. regression of single-item relevance score, or binary classification of item pairs), but rather optimizes the whole ranking. For this reason, specialized ML models are required.

	Pointwise LTR	Pairwise LTR	Listwise LTR
PROS	• Can use any supervised classifier/ regressor • No post processing required • O(n) computational scaling • Can be used with mini-batches	• Accounts for interactions between itemsCan be used with mini-batches	• Accounts for interactions between items • O(n) computational scaling • Simple to evaluate
CONS	• Does not account for interactions between items	• Requires data preprocessing and postprocessing (handling item pairs) • High computational loadO(n^2) scaling	• Requires especialized models • Cannot be used with mini-batches, requires the full item list

Table 1
Advantages and disadvantages of LTR techniques.

Ranking Financial Assets with ML

Besides web search, computational biology or recommendation systems, LTR can also be applied to investment management. Specifically, we can consider a pool of N financial assets, such as the 50 constituents of the Europe Large/Mid Cap index and ask ourselves how they will perform in terms of return over a certain time horizon, let's say four weeks.

Some of these assets might exhibit a positive return, others might end up in the negative territory, and their market cap-weighted average will yield the monthly return of the Europe Large/Mid Cap index.

If we were able to forecast the return-based ranking of such assets, we could build an optimised investment strategy (e.g. buy the winners) and outperform the index return. This is where LTR enters the game.

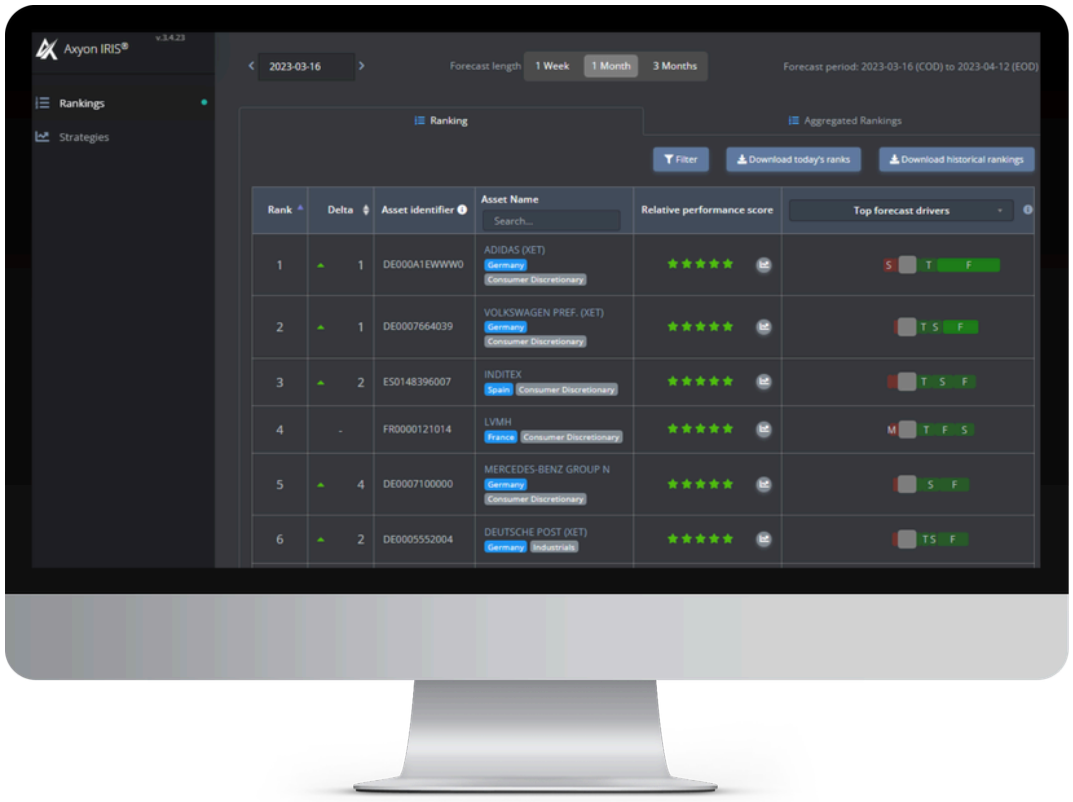

Figure 2
An example of Learning-to-Rank applied to Quantitative Investing. Here, our query can be formulated as "which of the Europe Large/Mid Cap constituents will outperform the index in the next month?", and Axyon IRIS® returns a list of stocks ranked by forecasted performance.

Case Study

Now that we know about LTR algorithms, we should consider whether to approach this problem in a pointwise, pairwise, or listwise fashion, based on our domain knowledge, the advantages and disadvantages of each method, and the available data. In Figures 3, 4, and 5, we compare the three approaches graphically, using three stocks as a running example: Volkswagen AG (VOW), Inditex (ITX), and Adidas (ADS).

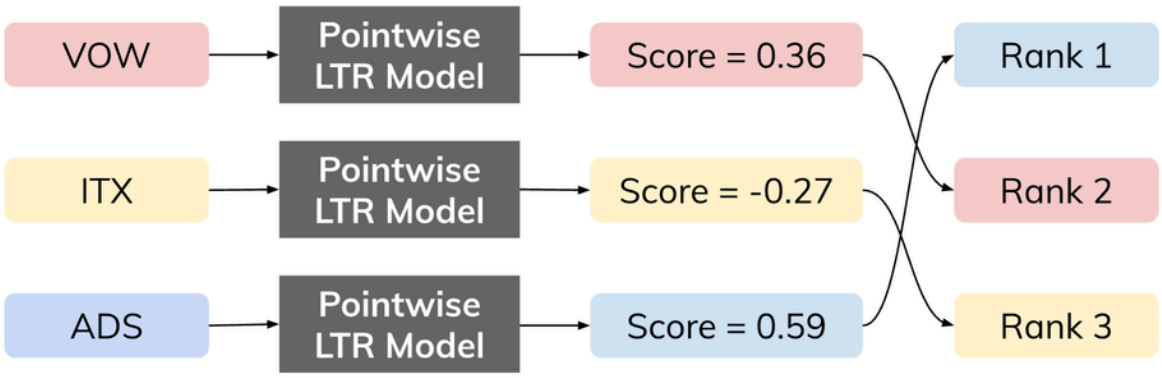

Figure 3
In pointwise LTR, our ML model looks at one "point" at a time to produce a ranking score.

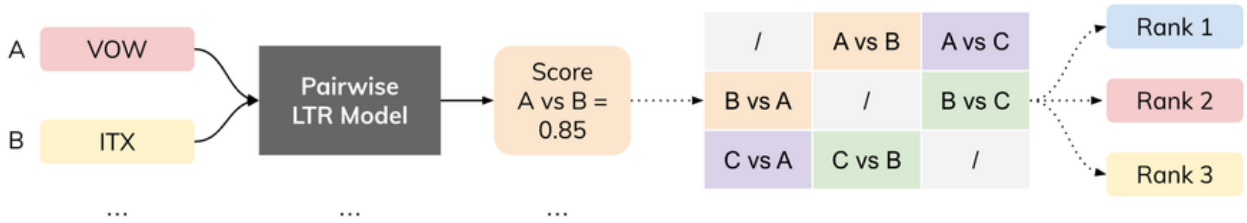

Figure 4
In pairwise LTR, our ML model looks at a "pair" at a time and produces a prediction on it. From the cross-sectional matrix of all predictions for all pairs, we can obtain a predicted ranking.

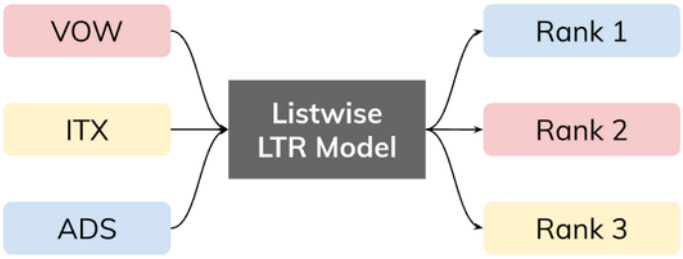

Figure 5
In listwise LTR, our ML model looks at the entire "list" to produce a ranking prediction.

At Axyon AI, we have experimented with multiple LTR approaches, and since February 2021 we have been continuously operating a pairwise LTR model on the European Large/Mid Cap index, powered by an ensemble of Siamese Neural Networks[2] trained on hundreds of input variables to predict the monthly return-based ranking of the target stocks. One way to use such a model is building a very simple Long-Only Equally Weighted strategy that invests in a selection of the top 10 stocks, as predicted by the model, and updates its weights daily.

So, how has such a basic strategy stood the test of time?

If we look at Figure 6, comparing the Axyon AI European Large/Mid Cap Top 10 strategy to the European Large/Mid Cap index, we can see that the former has clearly outperformed the latter in terms of annual return (CAGR of 23% vs 12%), with comparable annual volatility (around 19%).

Therefore, is pairwise LTR always the best solution?

As we have seen, different methods have different pros and cons, and ultimately it is our job as AI experts to collaborate with domain experts; in this case, Asset Managers and Quant Analyst select the best technique for the task at hand, depending on the business problem being addressed.

2. Taigman, Yaniv, et al. "Deepface: Closing the gap to human-level performance in face verification." Proceedings of the IEEE conference on computer vision and pattern recognition. 2014.

Figure 6
A simple Long-Only Top 10 Equally Weighted strategy on European Large/Mid Cap constituents (solid light blue line) compared with the European Large/Mid Cap index (dashed green line) from Feb 2021 to Apr 2023.