Metal Price Prediction on the London Metal Exchange: A Deep Learning Approach

ABSTRACT

Predicting metal prices accurately is crucial for stakeholders in the commodities market, especially those involved with the London Metal Exchange (LME), the world's largest market for industrial metals. This study aims to enhance the reliability of price forecasting for various metals traded on the LME by implementing a deep learning approach that integrates a broad spectrum of fundamental and technical data. We leverage the advanced capabilities of Temporal Fusion Transformers (TFT). This deep learning model employs gated inputs to assign varying degrees of importance to differentinput factors, thus optimizing the predictive accuracy. Our methodology encompasses a comprehensive dataset, including time series of currency pairs, interest rates, stock indices, Gross Domestic Product (GDP), inflation rates in countries significant to metal production and consumption, and sentiment analysis from relevant news sources. Applying the TFT model to predict the price of copper six months in advance demonstrated a promising result, achieving a mean absolute percentage error (MAPE) of approximately 2%. This level of precision in forecasting underscores the potential of combining deep learning models with diverse datasets to significantly improve the accuracy of metal price predictions on the LME. Our findings offer valuable insights for investors, policymakers, and industry practitioners seeking to navigate the volatile commodities market with better foresight and confidence.

I. DATA

For accurate time-series forecasts, choosing the right inputs is critical. Our method includes a range of crucial factors affecting the outcome, such as commodities prices indicating global trends, currency pairs showing forex dynamics, and stocks reflecting corporate health and investor mood. We also use macroeconomic indicators like GDP and inflation for a broader economic view, as well as calendar units and technical indicators from price movements to catch seasonal and temporal effects. These combined inputs offer insight into market drivers, improving our prediction accuracy.

Figure 1. Demonstration time series data for LME copper price prediction

The data preparation phase is a critical step that directly impacts the accuracy and effectiveness of the final analysis. One of the primary challenges encountered in this phase was the limitation imposed by the shortest time series available in our dataset. This constraint dictated the scope and granularity of the analysis we could perform, as it was essential to ensure consistency and comparability across the different datasets. To address this, we had to carefully assess the extent of available data, identify the shortest time series, and adjust our analysis framework accordingly.

Given the nature of the commodity being studied, the copper market, we choose a one-day granularity for our data. This decision was informed by the understanding that the copper market does not exhibit the same level of dynamism as financial markets, such as those for stocks or cryptocurrencies. Copper prices tend to be influenced by longer-term economic trends, industrial demand, and geopolitical events rather than the minute-to-minute fluctuations in more volatile markets. By opting for a daily granularity, we aimed to capture the essential movements and trends in the copper market without getting lost in the noise of inconsequential shortterm fluctuations. This approach allowed us to focus on the underlying factors driving the market and facilitated a more meaningful analysis of copper price trends over time.

Figure 2. Time window training approach

Figure 2 illustrates the approximate division of data into training, validation, and test datasets, providing an overview of how the data splits for different phases of model evaluation. During the training phase, the data is further segmented into time windows, with the initial segment serving as the model's input and the subsequent segment as the model's output. This methodological approach facilitates the training of models on sequential or time-series data by mapping specific inputs to their corresponding outputs, enhancing the model's predictive accuracy.

II. MODEL

TFT is an attention-based architecture that combines high-performance multi-horizon forecasting with interpretable insights into temporal dynamics. Designed to perform across a wide range of forecasting tasks, the Temporal Fusion Transformer excels in scenarios where data exhibits varying degrees of seasonality, trend, and irregularity. It can handle multivariate time series and predict multiple future time points in a single forward pass. Its versatility and robustness stem from combining different neural network architectures, each contributing its strengths to model complex temporal patterns effectively.

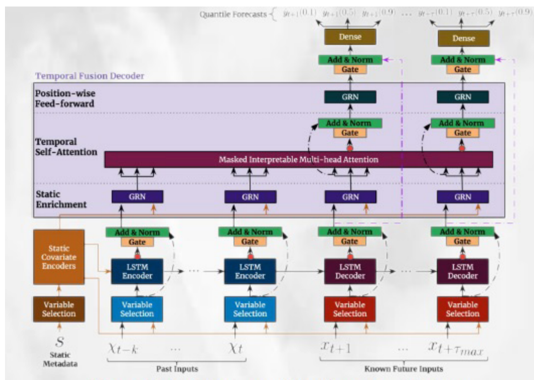

Figure 3. Architecture of the Temporal Fusion Transformer.

TFT leverages standard components to effectively generate feature representations for different input types - static, known, and observed - tailoring its architecture for optimal forecasting across a broad spectrum of challenges. Its design incorporates gating mechanisms for adaptable complexity, variable selection networks for dynamic input relevance, static covariate encoders to weave in unchanging features, and sophisticated temporal processing for understanding time-based relationships through a blend of sequence-to-sequence layers and a novel multi-head attention block. This setup adjusts to the demands of various datasets and scenarios and provides prediction intervals through quantile forecasts, outlining a range of possible outcomes for enhanced decision-making. The architecture's detailed components, illustrated and further elaborated in Figure 3, highlights TFT's capacity for high-precision forecasting in diverse settings.

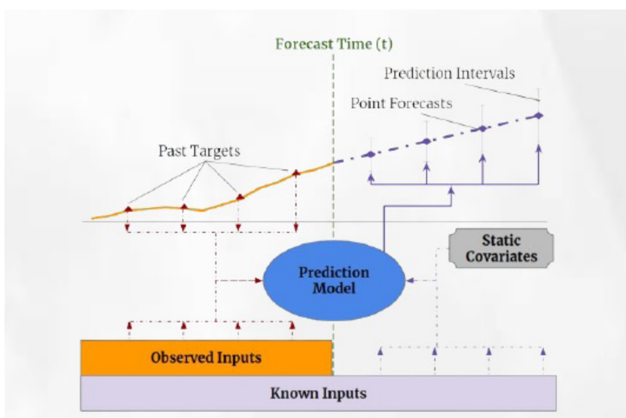

Figure 4. Multi-horizon forecasting with static covariates, past observed and a priori-known future time-dependent inputs

Multi-horizon forecasting applications frequently utilize diverse data sources, as illustrated in Figure 4, such as future events (for example, dates of future holidays), external time series data (like historical prices), and unchanging metadata (such as the seller location), even when the way these sources interrelate are not previously understood. The variety of these data sources, combined with a lack of insight into how they interact, significantly complicates the process of forecasting over multiple time horizons.

III. RESULTS

The figure 5 illustrates our model's predictive accuracy over a 6-month forecast horizon, employing a 1-year lag window to predict daily price movements up to 130 business days. This forecast is generated without direct exposure to the data within this forecast period, relying instead on historical data spanning the previous few years. This approach ensures that the predictions are unbiased, with the model's performance evaluated against past data to ascertain the precision of its forecasts.

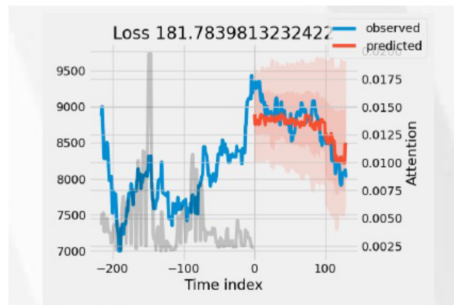

Figure 5. Half-year forecast with 1 year window.

The inclusion of 50% and 80% prediction intervals, represented by two distinct color schemes within the figure, showcases our model's capability not only to predict a median price line (highlighted in red) but also to provide a range within which prices are likely to fluctuate, enhancing our understanding of the model's confidence levels.

Time period	Baseline MAE	MAE	MAPE
6 months	490 USD	182 USD	2.10%
1 month	230 USD	74.5 USD	0.89%

Table 1. Absolute error comparison between two forecast intervals of 120 and 20 business days

Transformer technology, often regarded as a „black box" due to its complex and unclear nature, still has various ways to gain insights into its workings. One of the significant strides in this area is the ability to calculate variable importance within the model. This capability allows us to identify which parts of the input data the model considers most significant for making predictions, thereby shedding light on the decision-making process of transformers. Such approaches are not only helpful but also interpretable to some extent.

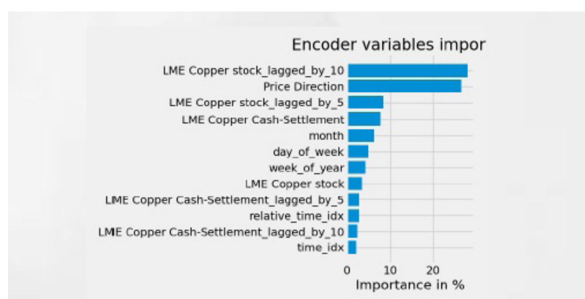

Figure 6. Variables importance for the encoder part of the model

CONCLUSION

The forecasting model for copper prices has been designed, utilizing the Temporal Fusion Transformer's capabilities to offer accurate and interpretable predictions tailored to the copper market's specific needs. By meticulously preparing data and choosing an appropriate granularity, the framework is adept at capturing fundamental market movements, further enhanced by the model's ability to provide detailed insights into variable significance and forecast with confidence intervals, ensuring robust and reliable predictions.

Member of CSG Aerospace