# Temporal Fusion Transformers For High-Accuracy Forecasting Of Global Raw Material Trade

**Matteo Bianchi**
**Phd Student at Department of Data Science and Economics, Bocconi University, Milan, Italy**

**Dr. Lina Philip**
**Center for Trade Analytics, Qatar Computing Research Institute, Doha, Qatar**

**Published Date: 19 December 2024 // Page no.:- 26-32**

## ABSTRACT

Accurate forecasting of international raw material trade flows is critical for effective policy-making, strategic supply chain management, and mitigating risks in an increasingly volatile global economy characterized by polycrises and supply chain disruptions. Traditional forecasting methods, while valuable, often struggle to capture the complex temporal dynamics and interplay of diverse influencing factors. This article explores the application of the Temporal Fusion Transformer (TFT), a state-of-the-art deep learning model, for achieving high-accuracy predictions of international raw material trade flows. We outline a conceptual framework for utilizing the TFT, highlighting its ability to leverage multiple time series inputs, incorporate static and dynamic exogenous variables, and provide interpretable insights into the drivers of trade. By comparing its potential performance against established models like ARIMA, Prophet, LSTM, and Graph Neural Networks (GNNs), we demonstrate the theoretical advantages of the TFT for this challenging forecasting task. The discussion emphasizes the implications of improved forecasting accuracy for enhancing resilience in global value chains and navigating turbulent times. While acknowledging data requirements and model complexity, this article posits that the Temporal Fusion Transformer represents a significant advancement in the toolkit for predicting international raw material trade, offering both enhanced accuracy and crucial interpretability.

**Keywords:** International Trade Forecasting, Raw Materials, Temporal Fusion Transformer, Time Series Analysis, Deep Learning, Supply Chain Resilience, Global Value Chains, Interpretable AI.

## INTRODUCTION

International trade in raw materials forms the backbone of global manufacturing and economic activity. The reliable flow of these essential resources is vital for industrial production, infrastructure development, and overall economic stability. However, forecasting international trade flows, particularly for raw materials, is a complex endeavor influenced by a myriad of factors including geopolitical events, economic policies, technological shifts, supply and demand dynamics, and environmental considerations [10, 11, 13, 14, 15]. The current era, often described as a "polycrisis," characterized by interconnected global risks and uncertainties, further exacerbates the challenge of accurate prediction [10, 11, 12, 14]. Disruptions to global value chains have highlighted the vulnerability of economies dependent on stable raw material supplies [9, 15, 16].

Traditional economic models, such as the gravity model, have been foundational in explaining bilateral trade flows based on factors like economic size and geographical distance [4, 23, 24, 25, 26]. While providing valuable structural insights, these models are often static or less adept at capturing high-frequency temporal dynamics and the impact of sudden shocks. Time series forecasting methods like Autoregressive Integrated Moving Average (ARIMA) [2, 27, 30, 31, 32] and Prophet [3, 28, 29, 33, 34] have been widely applied to trade data, offering robust approaches for capturing historical patterns and seasonality. More recently, deep learning models, including Long Short-Term Memory (LSTM) networks [5, 52] and Graph Neural Networks (GNNs) [35, 36, 37, 38, 39, 40, 53], have shown promise in modeling the complex, non-linear relationships and network structures inherent in global trade [5, 35, 36, 37, 38, 39, 40, 41, 42].

Despite these advancements, accurately predicting international raw material trade flows with high precision remains a significant challenge. The need for models that can effectively integrate diverse data sources (temporal, static, and exogenous), capture long-range dependencies, and provide interpretable insights into the drivers of forecasts is paramount for navigating the current turbulent global economic climate [10, 11, 12, 14, 15].

The Temporal Fusion Transformer (TFT) [46] is a novel deep learning architecture specifically designed for multi-horizon time series forecasting. It combines the strengths of recurrent neural networks and attention mechanisms to handle complex temporal patterns, incorporate various types of input data, and offer interpretability through attention weights. Recent studies have demonstrated the

TFT's superior performance and interpretability in various forecasting domains, including energy consumption [44, 45], tourism demand [8, 47], and economic systems [48].

This article proposes that the Temporal Fusion Transformer is particularly well-suited for the task of high-accuracy prediction of international raw material trade flows. By leveraging its unique architectural features, the TFT can potentially overcome some of the limitations of existing methods, providing more accurate, robust, and interpretable forecasts essential for enhancing resilience in global raw material supply chains [19, 20, 21, 22]. We outline a conceptual framework for applying the TFT to this problem and discuss the potential benefits and implications.

## 2. Methods

The proposed methodology for achieving high-accuracy prediction of international raw material trade flows utilizes the Temporal Fusion Transformer (TFT) model [46]. This section outlines the conceptual approach, data requirements, model architecture considerations, and evaluation strategy.

### 2.1. Data Requirements

Effective application of the TFT requires a comprehensive dataset encompassing various types of information relevant to international raw material trade. This includes:

- Historical Trade Flow Data: Time series data on the volume or value of specific raw material exports and imports between pairs of countries over a significant period [17, 18]. This forms the primary target variable for forecasting.
- Temporal Exogenous Variables: Time-varying factors that influence trade flows, such as:
    - o Economic indicators (GDP, industrial production indices, inflation rates) of trading partners.
    - o Exchange rates.
    - o Commodity prices.
    - o Shipping costs and logistics indices.
    - o Policy changes (tariffs, trade agreements, export restrictions) [17].
    - o Global events (pandemics, conflicts, natural disasters) [10, 11, 12].
- Static Exogenous Variables: Time-invariant or slowly changing factors relevant to trade, such as:
    - o Geographical distance between countries (a key component of gravity models) [4, 23, 24, 25, 26].
    - o Shared borders, common language, colonial ties (from gravity models) [25, 26].
    - o Institutional factors and trade agreements.
    - o Raw material production and consumption capacities of countries.

Data collection would involve compiling these variables from reputable international trade databases, economic statistics agencies, and relevant policy sources (e.g., OECD inventory of export restrictions [17]). The data should be structured as multiple time series, with each series representing a specific trade flow (e.g., exports of iron ore from Country A to Country B) and associated temporal and static features.

### 2.2. Temporal Fusion Transformer (TFT) Architecture

The TFT architecture [46] is specifically designed for multi-horizon time series forecasting with interpretability. Its key components include:

- Gating Mechanisms: These allow the model to selectively process relevant information and discard irrelevant inputs, enhancing robustness to noisy data.
- Variable Selection Networks: These learn the importance of different input features at each time step, contributing to interpretability.
- LSTM Encoders: These process the historical time series data, capturing temporal dependencies and patterns [52].
- Transformer Interpretable Multi-Head Attention: This mechanism allows the model to attend to relevant past time steps and identify important correlations across different time series, providing insights into which historical periods or variables are driving the forecast [57].
- Decoder: This component generates forecasts for multiple future time steps simultaneously (multi-horizon forecasting).
- Quantile Outputs: The TFT can be trained to predict different quantiles of the forecast distribution, providing uncertainty estimates alongside point forecasts.

For predicting raw material trade flows, the historical trade data and temporal exogenous variables would be fed into the LSTM encoders. Static exogenous variables would be incorporated through the gating and variable selection mechanisms. The multi-head attention would allow the model to learn complex interactions between different raw materials, trading partners, and influencing factors over time.

### 2.3. Model Training and Evaluation

The dataset would be split into training, validation, and testing sets. The TFT model would be trained to minimize a loss function (e.g., quantile loss for quantile forecasts or Mean Squared Error for point forecasts) using an optimization algorithm. Hyperparameter tuning would be performed using the validation set.

Model evaluation would be conducted on the unseen test set using standard time series forecasting metrics, including:

- Mean Absolute Error (MAE)
- Root Mean Squared Error (RMSE)
- Mean Absolute Percentage Error (MAPE)
- Weighted Pinball Loss (for quantile forecasts)

The performance of the TFT would be compared against established baseline models commonly used in trade forecasting, such as:

- ARIMA [2, 27, 30, 31, 32]
- Prophet [3, 28, 29, 33, 34]
- LSTM networks [5, 52]
- Graph Neural Networks (GNNs) [35, 36, 37, 38, 39, 40, 41, 42, 53]
- Gravity models (potentially as a benchmark or integrated as features) [4, 23, 24, 25, 26]

Interpretability features of the TFT, such as variable importance scores and attention weights, would also be analyzed to understand which factors the model deems most influential in predicting specific trade flows.

### 2.4. Implementation Considerations

Implementing the TFT requires access to deep learning frameworks (e.g., TensorFlow, PyTorch) and potentially specialized libraries for time series modeling. The computational resources needed for training can be substantial, especially for a large number of time series and features. Careful data preprocessing, including handling missing values and scaling, is essential. The dynamic nature of global events necessitates a strategy for incorporating new information and potentially retraining or fine-tuning the model periodically [10, 11, 12]. Prompt-based learning paradigms [59, 60], while more common in NLP, could potentially inspire future adaptations for incorporating qualitative information or expert knowledge into the forecasting process.

## 3. RESULTS

Applying the Temporal Fusion Transformer (TFT) to the task of forecasting international raw material trade flows is expected to yield significant improvements in prediction accuracy compared to traditional and other deep learning methods. Based on the architectural advantages of the TFT and its demonstrated performance in other complex time series domains [8, 44, 45, 47, 48], the following results are anticipated:

### 3.1. Superior Forecasting Accuracy

The primary expected result is that the TFT model will achieve higher accuracy on the chosen evaluation metrics (MAE, RMSE, MAPE, Weighted Pinball Loss) when predicting raw material trade flows across multiple future horizons compared to benchmark models such as ARIMA [2, 27], Prophet [3, 28], LSTM [5, 52], and GNNs [35, 36, 53]. This superior performance is attributed to the TFT's ability to effectively model complex, non-linear temporal dynamics, capture long-range dependencies through its attention mechanisms, and leverage diverse static and dynamic exogenous variables simultaneously [46]. While LSTMs and GNNs can capture some of these aspects [5, 35, 36], the TFT's integrated architecture for handling multiple input types and its interpretable attention mechanism provide an advantage in discerning relevant patterns in noisy and volatile trade data.

### 3.2. Robustness to Volatility and Shocks

International raw material trade is particularly susceptible to sudden shocks arising from geopolitical events, policy changes, or supply chain disruptions [10, 11, 13, 14, 15, 17]. The TFT's gating mechanisms and variable selection networks are hypothesized to contribute to its robustness by allowing it to selectively focus on the most relevant information and potentially down-weight noisy or less informative inputs during periods of high volatility. While no forecasting model can perfectly predict unprecedented events, the TFT's architecture is better equipped to adapt to changing patterns influenced by dynamic factors compared to models that rely solely on historical time series decomposition or simpler linear relationships.

### 3.3. Enhanced Interpretability

Beyond point forecasts, a key result from using the TFT is the generation of interpretable insights. The interpretable multi-head attention mechanism [46, 57] allows for visualizing which historical time steps and input variables are most influential in generating a particular forecast. For example, the attention weights could reveal that the predicted volume of copper exports from Country X to Country Y is heavily influenced by recent industrial production data in Country Y, combined with a historical trade agreement signed several years ago, and is currently paying significant attention to global energy prices. Variable importance scores derived from the variable selection networks would quantify the overall relevance of each static and dynamic feature across all forecasts. This interpretability is a significant advantage over many other black-box deep learning models [5, 35], providing decision-makers with actionable insights into the drivers of predicted trade flows [8, 47, 48].

### 3.4. Multi-Horizon Forecasting Capability

The TFT's design for multi-horizon forecasting allows for generating predictions for multiple future time steps simultaneously [46]. This is crucial for strategic planning in raw material trade, which often requires forecasts spanning several months or even years. The model is expected to maintain relatively high accuracy even for longer forecast horizons compared to models that forecast one step at a time or rely on iterative prediction, where errors can accumulate.

3.5. Quantifiable Uncertainty Estimates

The ability of the TFT to output quantile forecasts provides essential information about the uncertainty associated with the predictions. This allows stakeholders to assess the range of possible outcomes and implement risk management strategies accordingly. For raw material trade, where price and volume volatility can be high, understanding the potential variability in future trade flows is as important as the point estimate itself.

In summary, the expected results from applying the Temporal Fusion Transformer to international raw material trade forecasting include demonstrably higher accuracy, increased robustness to market fluctuations, valuable interpretable insights into forecast drivers, reliable multi-horizon predictions, and quantifiable uncertainty estimates, collectively representing a significant advancement in the field.

## 4. DISCUSSION AND CONCLUSION

The accurate forecasting of international raw material trade flows is not merely an academic exercise but a strategic imperative for governments, industries, and organizations operating within the intricate web of global value chains [9, 15, 16, 19, 20, 21, 22]. In an era defined by interconnected crises and heightened uncertainty [10, 11, 12, 14], the ability to anticipate shifts in the supply and demand of critical raw materials is paramount for ensuring economic stability and resilience [9, 15, 16]. This study posits that the Temporal Fusion Transformer (TFT) offers a powerful new approach to meet this challenge, promising enhanced accuracy and much-needed interpretability.

The expected results, highlighting the TFT's potential for superior accuracy compared to established methods like ARIMA [2, 27], Prophet [3, 28], LSTM [5, 52], and GNNs [35, 36, 53], underscore its capability to model the complex, non-linear dynamics and the influence of diverse factors that govern raw material trade. Unlike simpler time series models, the TFT's architecture is designed to handle the interplay between historical patterns, time-varying exogenous influences (economic indicators, policies, global events), and static characteristics (geography, trade agreements) [46]. This integrated approach is crucial for capturing the multifaceted nature of international trade, which is not solely driven by past trends but also by evolving external conditions [4, 23, 24, 25, 26].

Furthermore, the interpretability offered by the TFT's attention mechanisms provides a critical advantage over many black-box deep learning models [5, 35]. Understanding why a particular forecast is made—which factors are most influential at a given time—empowers decision-makers to not only react to predictions but also to understand the underlying drivers. This insight is invaluable for formulating effective trade policies, optimizing logistics and inventory, identifying potential bottlenecks, and developing strategies to mitigate supply chain risks [9, 15, 16, 17]. For instance, if the model highlights increasing export restrictions [17] in a key producing country and rising industrial demand in a major consuming region as primary drivers for a predicted price surge, stakeholders can proactively seek alternative suppliers or adjust production plans.

While the potential benefits are substantial, the successful implementation of the TFT for raw material trade forecasting requires addressing certain considerations. The model's effectiveness is heavily reliant on the availability and quality of comprehensive, granular data across all relevant temporal and static variables. Compiling and maintaining such a dataset for a wide range of raw materials and trading partners is a significant undertaking. The computational resources required for training large TFT models can also be considerable.

Future research should focus on several areas. Empirically validating the TFT's performance on diverse raw material trade datasets is essential. Exploring methods for incorporating real-time or near real-time data streams into the forecasting process would enhance responsiveness to sudden market shifts. Further investigation into the interpretability features of the TFT, perhaps combined with causal inference techniques, could deepen our understanding of the complex causal mechanisms driving global trade flows [8, 58]. Adapting prompt-based learning paradigms [59, 60] might offer novel ways to integrate qualitative information, such as expert assessments of geopolitical risks, into the quantitative forecasting framework.

In conclusion, the Temporal Fusion Transformer holds significant promise for advancing the field of international raw material trade forecasting. Its ability to deliver high-accuracy, multi-horizon predictions while providing interpretable insights into the drivers of trade makes it a valuable tool for navigating the complexities and uncertainties of the global economy. By embracing such sophisticated deep learning architectures, stakeholders can enhance the resilience of their supply chains, make more informed strategic decisions, and better prepare for the challenges and opportunities in the dynamic world of raw material trade. The era of polycrisis demands advanced analytical tools, and the TFT appears well-equipped to be a cornerstone of the next generation of trade forecasting systems.

## REFERENCES

Aamir S, Mughal SF, Kayani AJ, Yousuf MZ, Rastgar OA, Syed AA (2024). Impact of generative AI in revolutionizing education. In: 2024 8th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT), 1–6.

Saboia, J.: Autoregressive integrated moving average (ARIMA) models for birth forecasting. J. Am. Stat. Assoc. 72(358), 264–270 (1977).

Sean, J.: Taylor and Benjamin Letham. Forecasting at scale. Am. Stat. 72(1), 37–45 (2018).

Tiits, M., Kalvet, T., Ounoughi, C., Yahia, S.B.: Relatedness and product complexity meet gravity models of international trade. J. Open Innov.: Technol. Mark. Complex. 10(2), 100288 (2024).

Shen, M.-L., Lee, C.-F., Liu, H.-H., Chang, P.-Y., Yang, C.-H.: Effective multinational trade forecasting using LSTM recurrent neural network. Expert Syst. Appl. 182, 115199 (2021).

Rincon-Yanez, D., Ounoughi, C., Sellami, B., Kalvet, T., Tiits, M., Senatore, S., Yahia, S.B.: Accurate prediction of international trade flows: leveraging knowledge graphs and their embeddings. J. King Saud Univ. Comput. Inf. Sci. 35(10), 101789 (2023).

Binrong, W., Wang, L.: Two-stage decomposition and temporal fusion transformers for interpretable wind speed forecasting. Energy 288, 129728 (2024).

Binrong, W., Wang, L., Zeng, Y.-R.: Interpretable tourism demand forecasting with temporal fusion transformers amid covid-19. Appl. Intell. 53(11), 14493–14514 (2023).

Jicha, F., Hassink, R., Chu, H.: Advancing regional economic resilience theorizing by integrating global value chain resilience. In: The Changing Economic Geography of Companies and Regions in Times of Risk, Uncertainty, and Crisis, pp. 32–46. Routledge (2024).

Lawrence, M., Homer-Dixon, T., Janzwood, S., Rockstörm, J., Renn, O., Donges, J.F.: Global polycrisis: the causal mechanisms of crisis entanglement. Global Sustain. 7, 6 (2024).

Davies, M., Hobson, C.: An embarrassment of changes: international relations and the covid-19 pandemic. Austral. J. Int. Aff. 77(2), 150–168 (2023).

United Nations Development Programme Regional Bureau for Asia and the Pacific. Polycrisis and long-term thinking. Technical report, United Nations Development Programme (2022).

World Economic Forum: The global risks report 2023. Technical report, World Economic Forum, Geneva (2023).

Neise, T., Verfürth, P., Franz, M.: The Changing Economic Geography of Companies and Regions in Times of Risk, Uncertainty, and Crisis. Routledge (2024).

Baldwin, R., Freeman, R.: Risks and global supply chains: what we know and what we need to know. Annu. Rev. Econ. 14, 153–180 (2022).

Xing, Y., Wang, R., Dollar, D.: Global value chain development report 2023: Resilient and sustainable GVCS in turbulent times. Technical report, World Trade Organization (2023).

OECD. OECD inventory of export restrictions on industrial raw materials 2024: monitoring the use of export restrictions amid market and policy tensions. Technical report, OECD Publishing, Paris, (2024).

Georgitzikis, K.: Trade codes of non-food, non-fuel raw materials and their products. (JRC136041) (2023).

Schimpf, S., Sturm, F., Correa, V., Bodo, B., Keane, C.: The world of raw materials 2050: scoping future dynamics in raw materials through scenarios. Energy Procedia, 125:6–13 (2017). European Geosciences Union General Assembly 2017, EGU Division Energy, Resources & Environment (ERE).

Nassar, N.T., Brainard, J., Gulley, A., Manley, R., Matos, G., Lederer, G., Bird, L.R., Pineault, D., Alonso, E., Gambogi, J., Fortier, S.M.: Evaluating the mineral commodity supply risk of the U.S. manufacturing sector. Sci. Adv. 6(8), eaay8647 (2020).

Carrara, S., Bobba, S., Blagoeva, D., Alves Dias, P., Cavalli, A., Georgitzikis, K., Grohol, M., Itul, A., Kuzov, T., Latunussa, C., Lyons, L., Malano, G., Maury, T., Prior Arce, A., Somers, J., Telsnig, T., Veeh, C., Wittmer, D., Black, C., Pennington, D., Christou, M.: Supply chain analysis and material demand forecast in strategic technologies and sectors in the eu - a foresight study. Policy assessment, Anticipation and foresight, Risk assessment KJ-NA-31-437-EN-N (online), KJ-NA-31-437-EN-C (print), Luxembourg (Luxembourg) (2023).

Tiits, M., Karo, E., Kalvet, T.: Small countries facing the technological revolution: fostering synergies between economic complexity and foresight research. Compet. Rev. 35(1), 53–75 (2025).

Capoani, L.: Review of the gravity model: origins and critical analysis of its theoretical development. SN Bus. Econ. 3(5), 95 (2023).

Sharma, P., Rohatgi, S., Jasuja, D.: Scientific mapping of gravity model of international trade literature: a bibliometric analysis. J. Scientometr. Res. 11(3), 447–57 (2023).

Baldwin, R., Taglioni, D.: Gravity for dummies and dummies for gravity equations. NBER Working Paper 12516 (2006).

Yotov, Y.V., Piermartini, R., Larch, M. et al.: An advanced guide to trade policy analysis: the structural gravity model. WTO iLibrary (2016).

Box, G., Pierce, D.A.: Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. J. Am. Stat. Assoc. 65(332), 1509–1526 (1970).

Sean, J.: Taylor and Benjamin Letham. Forecasting at scale.

PeerJ 5, e3190v2 (2017).

Yenidoğan, I., Çayir, A., Kozan, O., Dağ, T., Arslan, C.: Bitcoin forecasting using ARIMA and prophet. In: 2018 3rd International Conference on Computer Science and Engineering (UBMK). IEEE, pp. 621–624 (2018).

Habib Ahmed Elsayir: Forecasting ARIMA model for foreign trade statistics. Stand. Sci. Res. Essays 2(12), 636–648 (2014).

Sharma, P.: Extrapolative modeling of international trade using ARIMA model: a performance comparison between India and China during covid-19. SAMVAD: SIBM Pune Res. J. XXI, 71–91 (2020).

Aliul Islam, Md., Sumy, M., Uddin, M.A., Hossain, M.S.: Fitting ARIMA model and forecasting for the tea production, and internal consumption of tea (per year) and export of tea. Int. J. Mater. Math. Sci. 2(1), 8–15 (2020).

Güleryüz, D., Özden, E.: The prediction of brent crude oil trend using LSTM and facebook prophet. Avrupa Bilim ve Teknoloji Dergisi 20, 1–9 (2020).

Aytaç, E.: Forecasting turkey’s hazelnut export quantities with facebook’s prophet algorithm and box-cox transformation. ADCAIJ: Adv. Distrib. Comput. Artif. Intell. J. 10(1), 33 (2021).

Sellami, B., Ounoughi, C., Kalvet, T., Tiits, M., Rincon-Yanez, D.: Harnessing graph neural networks to predict international trade flows. Big Data Cogn. Comput. 8(6), 65 (2024).

Minakawa, N., Izumi, K., Sakaji, H.: Bilateral trade flow prediction by gravity-informed graph auto-encoder. In: 2022 IEEE International Conference on Big Data (Big Data). IEEE, pp. 2327–2332 (2022).

Monken, A., Haberkorn, F., Gopinath, M., Freeman, L., Batarseh, F.: Graph neural networks for modeling causality in international trade 34, 04 (2021).

Ahmed, F., Cui, Y., Fu, Y., Chen, E.: A graph neural network approach for product relationship prediction. In: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, vol. 85383, pp. V03AT03A036. American Society of Mechanical Engineers (2021).

Panford-Quainoo, K.,Bose, A.J., Defferrard, M.: Bilateral trade modelling with graph neural networks. In: ICLR Workshop on Practical ML for Developing Countries (2020).

Verstyuk, S., Douglas, M.R.: Machine learning the gravity equation for international trade. Available at SSRN 4053795 (2022).

Yang, C.-H., Lee, C.-F., Chang, P.-Y.: Export-and import-based economic models for predicting global trade using deep learning. Expert Syst. Appl. 218, 119590 (2023).

Liu, H., Dong, Z., Jiang, R., Deng, J., Deng, J., Chen, Q., Song, X.:. Spatio-temporal adaptive embedding makes vanilla transformer sota for traffic forecasting. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, pp. 4125–4129 (2023).

Li, F., Wan, Z., Koch, T., Zan, G., Li, M., Zheng, Z., Liang, B.: Improving the accuracy of multi-step prediction of building energy consumption based on EEMD-PSO-Informer and long-time series. Comput. Electr. Eng. 110, 108845 (2023).

Wang, Z., Chen, Z., Yang, Y., Liu, C., Li, X., Jinran, W.: A hybrid autoformer framework for electricity demand forecasting. Energy Rep. 9, 3800–3812 (2023).

Lim, B., Arık, S.Ö., Loeff, N., Pfister, T.: Temporal fusion transformers for interpretable multi-horizon time series forecasting. Int. J. Forecast. 37(4), 1748–1764 (2021).

Binrong, W., Wang, L., Tao, R., Zeng, Y.-R.: Interpretable tourism volume forecasting with multivariate time series under the impact of covid-19. Neural Comput. Appl. 35(7), 5437–5463 (2023).

Han, Y., Tian, Y., Liangliang, Yu., Gao, Y.: Economic system forecasting based on temporal fusion transformers: multi-dimensional evaluation and cross-model comparative analysis. Neurocomputing 552, 126500 (2023).

Breiman, L.: Random forests. Mach. Learn. 45, 5–32 (2001).

Hochreiter, S.: Long Short-Term Memory. Neural Computation MIT-Press (1997).

Scarselli, F., Marco Gori, A., Tsoi, C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. IEEE Trans. Neural Netw. 20(1), 61–80 (2008).

Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 11106–11115 (2021).

Haixu, W., Jiehui, X., Wang, J., Long, M.: Autoformer: decomposition transformers with auto-correlation for long-term series forecasting. Adv. Neural Inf. Process. Syst. 34, 22419–22430 (2021).

Li, S., Jin, X., Xuan, Y., Zhou, X., Chen, W., Wang, Y.-X., Yan, X.: Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. Advances in neural information processing systems, 32 (2019).

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems, 30 (2017).

Kattel, R., Randma-Liiv, T., Kalvet, T.: Small states, innovation and administrative capacity. In: Bekkers, V., Edelenbos, J., Steijn, B. (eds.) Innovation in the Public Sector: Linking Capacity and Leadership, pp. 61–81. Palgrave Macmillan UK, London (2011).

Tiits, M., Kalvet, T.: Nordic Small Countries in the Global High-tech Value Chains: The Case of Telecommunications Systems Production in Estonia. Technical report, Working Papers in Technology Governance and Economic Dynamics (2012).

Lai, Z., Tangjie, W., Fei, X., Ling, Q.: Bert4st: fine-tuning pre-trained large language model for wind power forecasting. Energy Convers. Manag. 307, 118331 (2024).

Xue, H., Salim, F.D.: Promptcast: a new prompt-based learning paradigm for time series forecasting. IEEE Trans. Knowl. Data Eng. 36, 6851 (2023).