



PDF Download
3711896.3736567.pdf
31 March 2026
Total Citations: 14
Total Downloads: 3584

Latest updates: <https://dl.acm.org/doi/10.1145/3711896.3736567>

TUTORIAL

Multi-modal Time Series Analysis: A Tutorial and Survey

YUSHAN JIANG, University of Connecticut, Storrs, CT, United States

KANGHUI NING, University of Connecticut, Storrs, CT, United States

ZIJIE PAN, University of Connecticut, Storrs, CT, United States

XUYANG SHEN, University of Connecticut, Storrs, CT, United States

JINGCHAO NI, University of Houston, Houston, TX, United States

WENCHAO YU, NEC Laboratories America, Inc., Princeton, NJ, United States

[View all](#)

Open Access Support provided by:

University of Connecticut

Morgan Stanley, USA

NEC Laboratories America, Inc.

University of Houston

Published: 03 August 2025

Citation in BibTeX format

KDD '25: The 31st ACM SIGKDD
Conference on Knowledge Discovery and
Data Mining
August 3 - 7, 2025
Toronto ON, Canada

Conference Sponsors:
SIGMOD
SIGKDD

Multi-modal Time Series Analysis: A Tutorial and Survey

Yushan Jiang*
University of Connecticut
Storrs, CT, USA
yushan.jiang@uconn.edu

Xuyang Shen
University of Connecticut
Storrs, CT, USA
xuyang.shen@uconn.edu

Anderson Schneider
Morgan Stanley
New York, NY, USA
anderson.schneider@morganstanley.com

Kanghui Ning*
University of Connecticut
Storrs, CT, USA
kanghui.ning@uconn.edu

Jingchao Ni
University of Houston
Houston, TX, USA
jni7@uh.edu

Haifeng Chen[†]
NEC Laboratories America
Princeton, NJ, USA
haifeng@nec-labs.com

Dongjin Song[†]
University of Connecticut
Storrs, CT, USA
dongjin.song@uconn.edu

Zijie Pan*
University of Connecticut
Storrs, CT, USA
zijie.pan@uconn.edu

Wenchao Yu
NEC Laboratories America
Princeton, NJ, USA
wyu@nec-labs.com

Yuriy Nevmyvaka[†]
Morgan Stanley
New York, NY, USA
yuriy.nevmyvaka@morganstanley.com

Abstract

Multi-modal time series analysis has recently emerged as a prominent research area, driven by the increasing availability of diverse data modalities, such as text, images, and structured tabular data from real-world sources. However, effective analysis of multi-modal time series is hindered by data heterogeneity, modality gap, misalignment, and inherent noise. Recent advancements in multi-modal time series methods have exploited the multi-modal context via cross-modal interactions based on deep learning methods, significantly enhancing various downstream tasks. In this tutorial and survey, we present a systematic and up-to-date overview of multi-modal time series datasets and methods. We first state the existing challenges of multi-modal time series analysis and our motivations, with a brief introduction of preliminaries. Then, we summarize the general pipeline and categorize existing methods through a unified cross-modal interaction framework encompassing fusion, alignment, and transference at different levels (*i.e.*, input, intermediate, output), where key concepts and ideas are highlighted. We also discuss the real-world applications of multi-modal analysis for both standard and spatial time series, tailored to general and specific domains. Finally, we discuss future research directions to

help practitioners explore and exploit multi-modal time series. The up-to-date resources are provided in the GitHub repository¹.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Information systems** → **Data mining**.

Keywords

Multi-modal Time Series, Deep Learning, Data Mining

ACM Reference Format:

Yushan Jiang, Kanghui Ning, Zijie Pan, Xuyang Shen, Jingchao Ni, Wenchao Yu, Anderson Schneider, Haifeng Chen, Yuriy Nevmyvaka, and Dongjin Song. 2025. Multi-modal Time Series Analysis: A Tutorial and Survey. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '25), August 3–7, 2025, Toronto, ON, Canada*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3711896.3736567>

1 Introduction

Time series analysis is a fundamental task in machine learning and data science, driven by the proliferation of sequential data exhibiting rich temporal dynamics across diverse real-world systems. With the advent of deep learning, various methods have been proposed to effectively model complex temporal relationships within time series [10, 60, 63, 65, 79, 81, 89], facilitating downstream tasks in diverse domains, including healthcare [21, 33, 95], finance [66, 93], transportation [24, 30, 94] and environmental sciences [8, 9].

In practice, time series are often associated with external contexts beyond their temporal dynamics [6, 78]. Such contexts are multi-modal, encompassing a variety of representations, such as texts [42, 75], images [19, 72], tables [6], and graphs [69], which carry rich semantic information for time series analysis. As such, incorporating the multi-modal contexts allows models to have a

* equal contribution

[†] corresponding authors: Haifeng Chen, Yuriy Nevmyvaka, and Dongjin Song

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1454-2/2025/08
<https://doi.org/10.1145/3711896.3736567>

¹<https://github.com/UConn-DSIS/Multi-modal-Time-Series-Analysis>

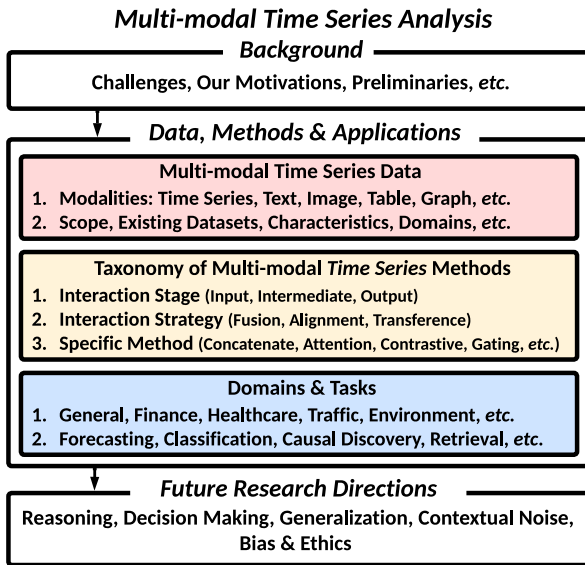


Figure 1: The framework of our tutorial and survey.

comprehensive view of underlying systems, capture subtle dependencies, and explain complex temporal behaviors more accurately.

Effective analysis of multi-modal time series, however, is hindered by several key challenges in terms of *data heterogeneity*, *modality gap* and *contextual relevance*. First, different modalities exhibit distinct statistical properties, structures, and dimensionalities, leading to discrepancies in feature distributions and semantic meanings. For instance, while time series data is sequentially ordered with temporal dependencies, textual and image data contains rich contextual semantics and correlations. Aligning these heterogeneous data into a unified representation space is non-trivial. Second, the textual, tabular, or visual contexts may appear at different timesteps or granularities. Such temporal misalignment may impede meaningful cross-modal interactions. Third, real-world data is inevitably noisy with irrelevant information that may mislead correlation learning, resulting in suboptimal performance. For example, in finance, news articles related to stock market prediction often contain much redundant or speculative narratives that does not reflect actual market conditions. Therefore, the focus of multi-modal time series analysis is to effectively capture complementary and relevant information from multi-modal context and leverage it for predictive or analytical tasks.

More recently, an increasing number of multi-modal methods have shown promise in exploiting contextual information from diverse data sources, which boosts performance in wide tasks ranging from forecasting [42, 53], classification [43, 46], anomaly detection [83] to retrieval [3] and causal discovery [70, 98]. Despite the promising results of multi-modal time series methods, they are tailored for their own tasks with domain-specific applications. The existing literature lacks a comprehensive and systematic review that provides a unified perspective on the underlying principles and pipelines for multi-modal time series learning. In this survey, we provide a systematic and up-to-date overview of existing methods for multi-modal time series analysis. As shown in Figure 1, we discuss the challenges, motivations, and preliminaries of multi-modal time series. Then we introduce the general pipeline for multi-modal

time series analysis and propose three types of interactions for cross-modal modeling between time series and other modalities – *fusion*, *alignment*, and *transference* – at the input, intermediate and output level, respectively. We also discuss the applications of multi-modal time series across multiple domains. Furthermore, we provide Table 2 to comprehensively summarize representative methods, encapsulating the modalities, fine-grained cross-modal interactions, real-world domains and tasks. Finally, we highlight potential future research opportunities to further advance time series analysis with multi-modal data. In summary, the major contributions of our survey are:

- We systematically catalog over 40 multi-modal time series methods with the corresponding open-source datasets.
- We uniquely categorize the existing methods into a unified cross-modal interaction framework, highlighting fusion, alignment, and transference at the input/intermediate/output levels.
- We discuss real-world applications of multi-modal time series and identify promising future directions, encouraging researchers and practitioners to explore and exploit multi-modal time series.

2 Background and Our Scope

2.1 Multi-modal Machine Learning

Recent advancements in multi-modal machine learning have significantly enhanced models’ ability to process and integrate data from diverse modalities, such as language, acoustic, vision, and tabular data [26, 68, 92]. With the development of deep learning architectures and sophisticated interaction designs, models are able to learn, infer, and reason by integrating multiple communicative modalities. Current research in multi-modal machine learning spans multiple key areas, including (1) representing multi-modal data to encode joint and individual characteristics, (2) identifying interconnections between modality elements, (3) transferring knowledge across modalities, and (4) theoretically and empirically analyzing the underlying learning process in a quantitative manner. We refer the audiences to the recent surveys [2, 51] for a more detailed overview of general multi-modal machine learning research. Building upon these advancements, we investigate multi-modal time series analysis with a focus on modeling temporal dependencies and leveraging the data interactions across heterogeneous modalities for predictive and analytical tasks.

2.2 Multi-modal Time Series Analysis

Multi-modal time series analysis aims to model time series data in combination with other complementary modalities. By leveraging cross-modal interactions, this approach yields deeper insights and more robust solutions for a wide range of predictive and analytical tasks across diverse real-world contexts.

This survey aims to provide a unique and systematic perspective on *effectively leveraging cross-modal interactions from relevant real-world contexts to advance multi-modal time series analysis*, addressing both foundational principles and practical solutions. Our assessment is threefold: (1) reviewing multi-modal time series data (Section 3), (2) analyzing cross-modal interactions between time series and other modalities (Section 4), and (3) revealing the impact of multi-modal time series analysis in applications across diverse domains (Section 5).

To resolve ambiguities, we define the scope of our survey by clarifying the types of time series considered and the criteria for multi-modal time series methods. First, we mainly consider standard time series and spatial time series. For the latter, spatial structures (often represented as graphs) are inherently paired with temporal data rather than treated as a separate modality. Second, we focus on methods that leverage **multi-modal inputs** from real-world contexts to provide complementary information, but for generation and retrieval tasks, the focus is more on transforming the input modality to another output modality. We acknowledge recent research on representing time series as a **single modality** (e.g., time series as images [16, 48, 64, 101, 102], time series as tabular data [28]) for downstream tasks. However, as these approaches are less relevant to our scope, we refer readers to their respective works.

Besides, we would like to highlight the difference between our survey and recent related survey and position papers. Ni *et al.* [59] focuses on imaging-based transformations of time series and subsequent visual modeling techniques, where the discussion on multi-modal models is limited to those involving vision modalities. Kong *et al.* [40] concentrates on the use of multi-modal large language models (LLMs) for enhancing reasoning capabilities (e.g., causal reasoning, QA, planning, *etc.*) with multi-modal context. In contrast, our survey provides a broader and structured framework by delivering a systematic and unified perspective of multi-modal time series analysis, not limited to a specific modality or task type.

3 Multi-modal Time Series Data

3.1 Modalities in Multi-modal Time Series Data

Multi-modal time series data often originate from diverse sources, each exhibiting unique characteristics that influence how they are processed and analyzed. Besides **Time Series**, *i.e.*, continuous or discrete measurements recorded over time, such as sensor readings, financial metrics, or physiological signals, their modalities often include: **1) Tabular**: Time-indexed records that are inherently organized in a tabular format, such as event logs, transaction records, or demographic information. **2) Text**: Time-stamped or domain-specific textual information — like clinical notes, financial reports, news articles, or social media posts — that provides contextual or interpretative insights. **3) Image**: Visual data acquired as images over time, such as photographs, medical images (e.g., X-rays, MRI), satellite imagery, or visual representations generated from time series data. **4) Graph**: Relational data representing interactions or structural dependencies among entities that evolve. They are typically modeled as networks or graphs, where the connections may change dynamically. Although **audio** is widely studied as an independent modality in multi-modal research, we consider it a special form of time series in this survey and briefly discuss representative works within this scope.

3.2 Common Datasets and Benchmarks

Healthcare: In this domain, physiological signals (e.g., ECG, EEG) are extensively analyzed alongside textual data such as clinical notes, patient demographics, and tabular data including vital signs and laboratory results. Common datasets include MIMIC-III [36], a comprehensive dataset containing electronic health records (EHRs) of ICU patients with physiological measurements, clinical notes,

and diagnostic information, widely used for tasks like patient monitoring, mortality prediction, and clinical decision support. MIMIC-IV [35] is an enhanced version of MIMIC-III, providing a larger and more modular dataset with expanded patient coverage and new data sources such as electronic medication records. Other notable healthcare datasets include ICBHI [67], which contains respiratory sound recordings paired with clinical annotations for respiratory disease classification; Coswara [4], which provides respiratory audio samples and rich metadata for COVID-19 detection tasks; KAUH [22], which comprises audio records and corresponding annotations for healthcare analytics; PTB-XL [73], a large-scale ECG dataset annotated with diagnostic labels for cardiac monitoring and diagnosis; ZuCo [15, 27], which consists of simultaneous EEG and textual data from reading comprehension tasks, being useful for cognitive neuroscience studies; and Image-EEG [23], which pairs EEG signals with images of objects on a natural background, aiding studies in visual neuroscience and computer vision.

Finance: Datasets that combine time series data with financial news and reports are vital for financial modeling and analysis. Notable examples include ACL18 [86], CIKM18 [80], and DOW30 [12]. These datasets focus on high-trade-volume stocks from the U.S. stock markets, providing historical stock price data, such as opening, high, low, and closing prices; alongside related textual information, including tweets or financial news. Another large-scale dataset, FNSPID [18], contains stock prices and time-aligned financial news records for over 4,000 companies from 1999 to 2023. FinMultiTime [85] extends the scope to four modalities: stock price time series, structured financial tables, financial news text, and K-line technical charts (as images). It covers 5,105 stocks from 2009 to 2025 across the U.S. and China.

Multi-domain: Datasets featuring general-purpose numerical time series combined with textual data are suitable for broad analytical applications. Examples include Time-MMD [55], which encompasses nine primary data domains: Agriculture, Climate, Economy, Energy, Environment, Health, Security, Social Good, and Traffic, while ensuring fine-grained alignment between time series and textual data; TimeCAP [43] compiles seven real-world time series datasets across three domains: weather, finance, and healthcare. To generate textual descriptions for each time series, a large language model (LLM) agent is employed, leveraging contextual information and domain-specific knowledge. NewsForecast [75] integrates task-specific time series data with verified public news reports across various domains, including finance, energy, traffic, and cryptocurrency; TTC [38] is a meticulously curated, time-aligned dataset designed for multimodal forecasting. It consists of paired time series and text data synchronized to timestamps, spanning two distinct domains: climate science and healthcare; CiK [78] is a dataset comprising 71 forecasting tasks across seven real-world domains, including climatology, economics, energy, mechanics, public safety, transportation, and retail. The TSQA [39] dataset consists of 200k question-answer pairs derived from time series data across 12 domains, supporting forecasting, imputation, anomaly detection, classification, and open-ended reasoning. MTBench [7] provides a benchmark for forecasting, trend analysis, and question answering in finance and weather domains. MoTime [100] collects eight multimodal time series datasets across diverse domains, featuring time series, text, and image modalities for multimodal forecasting.

Table 1: Representative open-source multi-modal time series datasets and across domains.

Domain	Dataset (Superscripts include the URLs to the datasets)	Modalities
Healthcare	MIMIC-III [36] ^[1] , MIMIC-IV [35] ^[2] ICBHI [67] ^[3] , Coswara [4] ^[4] , KAUH [22] ^[5] , PTB-XL [73] ^[6] , ZuCo [15, 27] ^[7] Image-EEG [23] ^[8]	TS, Text, Table TS, Text TS, Image
Finance	FNSPID [18] ^[9] , AlphaFin [47] ^[10] , ACL18 [86] ^[11] , CIKM18 [80] ^[12] , DOW30 [12] ^[13] FinMultiTime [85] ^[14]	TS, Text TS, Image, Text, Table
Retail	VISUELLE [72] ^[15]	TS, Image, Text
IoT	LEMMA-RCA [41] ^[16]	TS, Text
Speech	LRS3 [1] ^[17] , VoxCeleb2 [14] ^[18]	TS (Audio), Image
Traffic	NYC-taxi, NYC-bike [50] ^[19]	ST, Text
Environment	Terra [11] ^[20]	ST, Text
Multi-domain	Time-MMD [55] ^[21] , TimeCAP [43] ^[22] , NewsForecast [75] ^[23] , TTC [38] ^[24] , CiK [78] ^[25] , TSQA [39] ^[26] , MTBench [7] ^[27] MoTime [100] ^[28]	TS, Text TS, Text, Image

Other domains: Beyond the previously discussed major sectors, multi-modal time series analysis extends to various other domains. In *Retail*, datasets such as VISUELLE [72] integrate numerical sales data with product images and textual descriptions, facilitating thorough analyses of consumer behavior and inventory management. The *Internet of Things (IoT)* domain benefits from datasets such as LEMMA-RCA [41], which combine time series sensor data with textual metadata, enabling enhanced monitoring and more robust and secure methodologies that ensure the high performance of modern systems. In the *Speech* domain, datasets like LRS3 [1] and VoxCeleb2 [14] integrate audio recordings with corresponding visual data, supporting advancements in speech recognition and speaker identification technologies. In the *Traffic* domain, datasets like NYC-Taxi, NYC-Bike [50] contain spatial-temporal (ST) data alongside associated textual metadata. These integrations allow LLMs to effectively capture and utilize spatial-temporal contextual signals. In the *Environment* domain, Terra [11] collect 45 years of global geographic spatial-temporal data with textual descriptions.

4 Cross-modal Interactions with Time Series

We define three fundamental types of interactions between time series and other modalities, including *fusion*, *alignment*, and *transference*, which occur at different stages within a framework – *input*, *intermediate* (i.e., representations or intermediate outputs), and *output*. Representative examples are provided in Figure 2, and domain-specific adaptations are discussed in Section 5. The detailed taxonomy is provided in Table 2.

4.1 Fusion

Fusion refers to the process of integrating heterogeneous modalities in a way that captures complementary information across diverse sources to improve time series modeling. To fuse multi-modal inputs, a common practice is to directly integrate time series, tabular data and texts into a unified textual prompt, then use it to query LLMs for downstream tasks. This is typically facilitated by instruction fine-tuning for task-oriented analysis [20, 25, 39, 50, 58, 75, 91]. Some works also leverage the zero-shot reasoning and inference capability of pretrained LLMs (e.g., GPT-4 and its variants) [74, 82, 90]. Recent research efforts like TaTS [49] attempt to integrate paired text embedding as an additional variable of time series for temporal modeling, yielding competitive task performance.

Most existing methods perform cross-modal fusion at the intermediate stage, such as adding and concatenating multi-modal representations, where each individual modal encoder first maps the raw data into a shared latent space. Addition combines time series and other modalities by summing up encoded representations, effectively blending shared information while preserving their interconnections in the latent space [6, 31, 53, 87, 96, 98]. On the other hand, concatenation stacks multi-modal representations along the same dimension, retaining modality-specific characteristics and allowing models to capture joint relationships between the modalities [17, 37, 37, 38]. To effectively leverage cross-modal information, existing methods often incorporate alignment designs after fusing representations [5, 12, 13, 19, 32, 34, 42, 43, 45, 56, 57, 62, 71, 72, 76, 87, 97], which will be detailed in Section 4.2.

When fusion is performed at the output level, different modalities contribute separately to the final output, allowing each modality to retain its unique predictive signal [32, 42, 43, 55]. Time-MMD [55] provides a paradigm that fuses predictions from both state-of-the-art forecasters and a pretrained language model with a projection layer, in an end-to-end manner. MOAT [42] introduces a two-stage framework for multi-modal time series forecasting. In the first stage, the model is optimized to generate forecasts from decomposed time series and text embeddings. In the second stage, an offline synthesis via MLP is applied to dynamically fuse different components, yielding the final forecast based on their relative contributions. Beyond fusing outputs from a single model, TimeCAP [43] enhances performance by combining predictions from both a multi-modal predictor and a pretrained LLM, which synergizes the gradient-based method and LLM agents reasoning on real-world contexts. Output fusion gains advantage of design flexibility and robustness, but it may not fully utilize the complementary relationship between modalities without additional countermeasures. In general, cross-modal fusion relies on well-aligned multi-modal data for effective exploitation of the contextual information. However, ideally-aligned data may not be given in real-world scenarios. As such, existing methods also leverage alignment mechanisms to mitigate the challenge.

4.2 Alignment

Alignment ensures that the relationships between different modalities are preserved and semantically coherent when integrated into a unified learning framework. At the input level, we primarily refer

Table 2: Taxonomy of representative multi-modal time series methods. *Modality* refers to the different data modalities involved in each method. *TS* represents standard time series, *ST* denotes spatial time series. The *Method* column lists the techniques used for each interaction, separated by semicolons, where each interaction may include one or more techniques, separated by commas. Superscripts in the *Code* column include the URLs to Github repositories.

Method	Modality	Domain	Task	Cross-Modal Interaction				Method	Large Model	Year	Code
				Stage	Fusion	Align.	Trans.				
Time-MMD [55]	TS, Text	General	Forecasting	Output	✓	✗	✗	Addition	Multiple	2024	Yes ^[1]
Wang et al. [75]	TS, Text	General	Forecasting	Input Intermediate	✓ ✓	✗ ✓	✗ ✗	Prompt Prompt; LLM Reasoning	LLaMa2 GPT-4 Turbo	2024	Yes ^[2]
GPT4MTS [31]	TS, Text	General	Forecasting	Intermediate	✓	✓	✗	Addition; Self-attention	GPT-2	2024	No
TimeCMA [53]	TS, Text	General	Forecasting	Input Intermediate	✗ ✓	✗ ✓	✓ ✗	Meta-description Addition; Cross-attention	GPT-2	2025	Yes ^[3]
MOAT [42]	TS, Text	General	Forecasting	Intermediate Output	✓ ✓	✓ ✗	✗ ✗	Concat.; Self-attention Offline Synthesis (MLP)	S-Bert	2024	No
TimeCAP [43]	TS, Text	General	Classification	Input Intermediate Output	✗ ✓ ✓	✗ ✓ ✗	✓ ✗ ✗	LLM Generation Concat.; Self-attention, Retrieval Addition	Bert, GPT-4	2024	No
TimeXL [32]	TS, Text	General	Classification Forecasting	Intermediate Output	✓ ✓	✓ ✗	✗ ✗	Concat., Prompt; LLM Reasoning Addition	Bert, S-Bert GPT-4o	2025	No
Hybrid-MMF [38]	TS, Text	General	Forecasting	Intermediate	✓	✗	✗	Concat.	GPT-4o	2024	Yes ^[4]
Time-LLM [34]	TS, Text	General	Forecasting	Input Intermediate	✗ ✓	✗ ✓	✓ ✗	Meta-description Concat.; Self-attention	LLaMA, GPT-2	2024	Yes ^[5]
Time-VLM [99]	TS, Text, Image	General	Forecasting	Input Intermediate	✗ ✓	✗ ✓	✓ ✗	Feat. Imaging, Meta-description Addition; Gating, Cross-attention	ViLT, CLIP BLIP-2	2025	No
Unitime [57]	TS, Text	General	Forecasting	Input Intermediate	✗ ✓	✗ ✓	✓ ✗	Meta-description Concat.; Self-attention	GPT-2	2024	Yes ^[6]
TESSA [52]	TS, Text	General	Annotation	Intermediate	✓	✓	✓	Prompt; RL; LLM Generation	GPT-4o	2024	No
InstrucTime [13]	TS, Text	General	Classification	Intermediate	✓	✓	✗	Concat.; Self-attention	GPT-2	2025	Yes ^[7]
MATMCD [70]	TS, Text, Graph	General	Causal Discovery	Intermediate	✓	✓	✓	Prompt; LLM Reasoning; Supervision	Multiple	2025	No
STG-LLM [56]	ST, Text	General	Forecasting	Intermediate	✓	✓	✗	Concat.; Self-attention	GPT-2	2024	No
TableTime [74]	TS, Text	General	Classification	Input	✓	✗	✓	Prompt; Reformulate	Multiple	2024	Yes ^[8]
ContextFormer [6]	TS, Table	General	Forecasting	Intermediate	✓	✓	✗	Addition; Cross-attention	No	2025	No
Time-MQA [39]	TS, Text	General	Multiple	Input	✓	✗	✗	Prompt	Multiple	2025	Yes ^[9]
MAN-SF [69]	TS, Text, Graph	Finance	Classification	Intermediate	✓	✓	✗	Bilinear; Graph Convolution	USE	2020	No
Bamford et al. [3]	TS, Text TS, Image	Finance	Retrieval	Intermediate Output	✗ ✗	✓ ✗	✗ ✓	Supervision	S-Bert	2024	No
Chen et al. [12]	TS, Text, Graph	Finance	Classification	Input Intermediate	✗ ✓	✗ ✓	✓ ✗	LLM Generation Concat.; Graph Convolution	ChatGPT	2023	No
Xie et al. [82]	TS, Text	Finance	Classification	Input	✓	✗	✗	Prompt	ChatGPT	2023	No
Yu et al. [90]	TS, Text	Finance	Forecasting	Input	✓	✗	✗	Prompt	GPT-4, Open LLaMA	2023	No
MedTsLLM [5]	TS, Text, Table	Healthcare	Multiple	Intermediate	✓	✓	✗	Concat.; Self-attention	Llama2	2024	Yes ^[10]
RespLLM [96]	TS (Audio), Text	Healthcare	Classification	Intermediate	✓	✓	✗	Addition, Self-attention	OpenBioLLM-8B	2024	Yes ^[11]
METS [46]	TS, Text	Healthcare	Classification	Output	✗	✓	✗	Contrastive	ClinicalBERT	2023	No
Wang et al. [77]	TS, Text	Healthcare	Classification	Intermediate	✗	✗	✓	Supervision	Bart, Bert, RoBERTa	2021	No
EEG2TEXT [54]	TS, Text	Healthcare	Generation	Output	✗	✗	✓	Self-supervision, Supervision	Bart	2024	No
MEDHMP [76]	TS, Text	Healthcare	Classification	Intermediate	✓	✓	✗	Concat.; Self-attention, Contrastive	ClinicalT5	2023	Yes ^[12]
Deznabi et al. [17]	TS, Text	Healthcare	Classification	Intermediate	✓	✗	✗	Concat.	Bio+Clinical Bert	2021	Yes ^[13]
Niu et al. [62]	TS, Text	Healthcare	Classification	Intermediate	✓	✓	✗	Concat.; Cross-attention	BioBERT	2023	No
Yang et al. [87]	TS, Text	Healthcare	Classification	Intermediate	✓	✓	✗	Concat., Addition; Gating	ClinicalBERT	2021	Yes ^[14]
Liu et al. [58]	TS, Text	Healthcare	Classification Regression	Input	✓	✗	✗	Prompt	PaLM	2023	Yes ^[15]
xTP-LLM [25]	ST, Text	Traffic	Forecasting	Input	✓	✗	✓	Prompt; Meta-description	Llama2-7B-chat	2024	Yes ^[16]
UrbanGPT [50]	ST, Text	Traffic	Forecasting	Input	✓	✗	✓	Prompt; Meta-description	Vicuna-7B	2024	Yes ^[17]
CityGPT [20]	ST, Text	Mobility	Multiple	Input	✓	✗	✗	Prompt	Multiple	2025	Yes ^[18]
MULAN [98]	TS, Text, Graph	IoT	Causal Discovery	Intermediate	✓	✓	✓	Addition; Contrastive; Supervision	No	2024	No
MIA [83]	TS, Image	IoT	Anomaly Detection	Intermediate	✓	✓	✗	Addition; Cross-attention, Gating	No	2023	No
Ekambaram et al. [19]	TS, Image, Text	Retail	Forecasting	Intermediate	✓	✓	✗	Concat.; Self & Cross-attention	No	2020	Yes ^[19]
Skenderi et al. [72]	TS, Image, Text	Retail	Forecasting	Intermediate	✓	✓	✗	Concat.; Cross-attention	No	2024	Yes ^[20]
VIMTS [97]	ST, Image	Environment	Imputation	Intermediate	✓	✓	✗	Concat.; Supervision	No	2022	No
LITE [45]	ST, Text, Image	Environment	Forecasting	Intermediate	✓	✓	✗	Concat.; Self-attention	LLaMA-2-7B	2024	Yes ^[21]
AV-HuBERT [71]	TS (Audio), Image	Speech	Classification	Intermediate	✓	✓	✗	Concat.; Self-attention	HuBert	2022	Yes ^[22]
SpeechGPT [91]	TS(Audio), Text	Speech	Generation	intermediate	✓	✓	✗	Concat.; Self-attention	LLaMA-13B	2023	Yes ^[23]
LA-GCN [84]	ST, Text	Vision	Classification	Intermediate	✗	✓	✗	Supervision	Bert	2023	Yes ^[24]

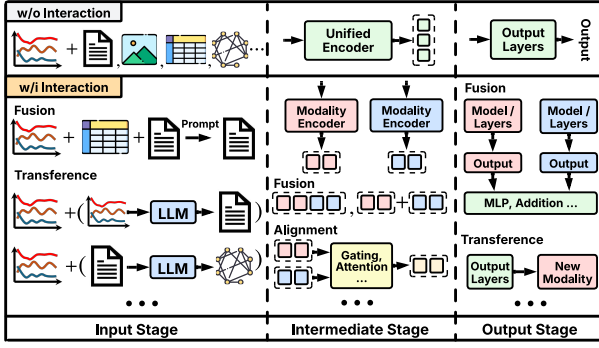


Figure 2: Categorization of cross-modal interaction methods and representative examples.

alignment to data preprocessing techniques that aim at mitigating temporal misalignment caused by missing values, irregular sampling intervals, and differing granularities across modalities. This process is crucial for ensuring that data from multiple sources are properly synchronized before fusion, where domain knowledge is usually needed to handle such inconsistencies [11, 55, 75]. In addition, none of the existing methods we reviewed explicitly perform output alignment. However, the aforementioned output fusion can be easily adapted to alignment through the incorporation of a gating or attention mechanism that we will introduce shortly.

Alignment at the intermediate stage plays a crucial role in multi-modal interactions. We first introduce the alignment of multi-modal representations, spanning techniques from model design to learning objective. The common components include self-attention [5, 13, 31, 34, 42, 43, 45, 56, 57, 71], cross-attention [6, 19, 53, 62, 72, 83, 87, 99] and gating mechanisms [83, 99]. Self-attention is often used to fuse multi-modal representations. It enables a joint and undirected alignment across all modalities by dynamically attending to important features. Given multi-modal embeddings $E_{mm} \in \mathbb{R}^{n \times d}$, where n is the number of modality tokens and d is the embedding dimension, self-attention is computed as:

$$\text{Attention}(E_{mm}) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V$$

where the queries Q , keys K , and values V are linear projections of E_{mm} : $Q = E_{mm}W_Q$, $K = E_{mm}W_K$, $V = E_{mm}W_V$ with learnable weights of dimensionality d_k : $W_Q, W_K, W_V \in \mathbb{R}^{d \times d_k}$.

In cross-attention, time series serves as the query modality to get contextualized by other modalities, providing a directed alignment that ensure auxiliary modalities contribute relevant contextual information while preserving the temporal structure of time series. Given a query embedding $E_{ts} \in \mathbb{R}^{n \times d}$ and that of auxiliary modalities $E_c \in \mathbb{R}^{n \times d}$ as keys and values:

$$\text{CrossAttention}(E_{ts}, E_c) = \text{softmax} \left(\frac{Q_{ts}K_c^T}{\sqrt{d_k}} \right) V_c$$

where the query, key and value are denoted as $Q_{ts} = E_{ts}W_Q$, $K_c = E_cW_K$, $V_c = E_cW_V$. Note that existing methods adopt multi-head attentions, which is omitted here for simplicity.

Similarly, the gating mechanism is a parametric filtering operation that explicitly regulates the influence of time series and other

modalities on the fused embeddings in E :

$$G = \sigma(W_g[E_{ts}; E_c] + b_g), \quad E = G \odot E_{ts} + (1 - G) \odot E_c$$

where $\sigma(\cdot)$ denotes the sigmoid function, the learnable weight and bias are denoted as $W_g \in \mathbb{R}^{2d \times d}$ and $b_g \in \mathbb{R}^d$, respectively.

When a graph modality is available from external contexts, the underlying topological insights can be leveraged for graph-based alignment [12, 69]. Unlike the above methods that rely solely on feature interactions, it explicitly aligns multi-modal representations with relational structures through graph convolution, enabling context-aware feature propagation across modalities.

Representation alignments can also be achieved by learning objectives [3, 46, 84, 97, 98]. For example, MULAN [98] extracts modality-invariant and modality-specific representations from multi-modal time series. It employs contrastive learning to enhance cross-modal alignment by maximizing the similarity between invariant representations across modalities while minimizing the similarity between invariant and specific representations of the same modality. Moreover, Bamford *et al.* [3] align cross-modal representations by using the mean of uni-modal cosine similarities as the target similarity and optimizing cross-modal similarity via cross-entropy loss, which effectively connects both modalities in a shared latent space for time series retrieval tasks. In general, this branch of methods is effective as it directly integrates the alignment objective into the optimization process, ensuring that meaningful representations are explicitly learned.

Lastly, we introduce the intermediate alignment of component outputs within a framework, extending beyond representation alignment within a model. The most recent studies explore the synergy between time series models and LLM agents, leveraging the strong reasoning capabilities of pretrained LLMs to provide contextual understanding and calibration in real-world scenarios [32, 43, 52, 70, 75]. We briefly discuss a few representative examples for demonstration. TimeCAP [43] utilizes the embedding space of a trained multi-modal encoder to retrieve in-context examples with the highest cosine similarity. These retrieved examples with ground truth labels are then fed, along with the query text, into an LLM to provide contextual guidance and improve outcome prediction. TimeXL [32] incorporates a multi-modal prototype-based encoder to generate explainable case-based rationales for both time series and texts, integrating three LLM agents, where prediction, reflection, and refinement LLMs collaborate to iteratively enhance prediction accuracy, identify textual inconsistencies or noise, and calibrate textual contexts, yielding more accurate predictions and explanations. NewsForecast [75] also employs reflection in language agents to iteratively select relevant news from a large database, enhancing alignment of textual information for text-based forecasting. Similarly, MATMCD [70] ensures alignment between statistical causal discovery on time series and LLM reasoning on textual context by leveraging iterative self-reflective tool-calling to structure textual context, which is then used to explain and refine causal constraints.

In a nutshell, alignment aims to calibrate real-world contexts and effectively capture relevant multi-modal elements for a semantically coherent time series modeling. It enhances task performance, robustness and explanation, ensuring that models leverage meaningful contextual information for improved decision-making.

4.3 Transference

Transference refers to the process of mapping between different modalities. It allows one modality to be inferred, translated, or synthesized from another. This concept plays a crucial role across different stages of multi-modal time series analysis. The input-level transference typically serves for modality augmentation. It helps introduce contextual priors, enrich training samples, and provide alternative representations. This is particularly useful in scenarios of data scarcity and imbalance. In existing literature, a common practice is to use meta information to describe the narrative of real-world contexts (e.g., domain, data statistics and granularity, variable descriptions, other co-variates, etc.) [20, 25, 34, 50, 53, 57, 74, 99] or leverage pretrained LLMs to generate fine-grained textual contexts [43] or graphs [12] for real-world time series, serving as an augmented modality. In addition to texts, time series can also be transformed into high-dimensional images via feature imaging, such as stacking the original data with frequency and periodicity features [99]. Alternatively, time series can be represented in tabular form, transforming time series analysis into a table understanding task [74]. Note that the aforementioned uni-modal methods for transforming time series into other single modalities can also be integrated into multi-modal time series frameworks [16, 28, 48, 64, 101, 102]. The exploitation of input-level transference is two-fold. First, the embedding of generated modality can serve as semantic anchors that guides time series modeling via representation alignment, improving downstream supervised tasks [34, 53, 57, 99]. Second, it provides additional contextual guidance for pretrained LLMs through input fusion and prompting [20, 25, 50, 74].

At the intermediate [12, 52, 70, 77, 98] and output [3, 54] levels, transference are more task-oriented. The output-level transference typically refers to the end-to-end generation of new modalities, such as text-based and image-based time series retrieval, where users provide textual descriptions or sketched trends to query relevant time series data [3]. This also includes EEG-to-text conversion, enabling direct transformation from physiological signals to human-readable narratives [54].

The output of intermediate transference typically serves as an initial solution to be refined for modality generation tasks [52, 70, 98] or a medium to be inferred for predictive tasks [77], facilitating downstream reasoning and further alignment within the multi-modal framework. MATMCD [70] generates an initial causal graph from time series, achieving modality transference in the intermediate level. Subsequently, it incorporates textual modality to refine the causal graph, ensuring improved alignment and interpretability. Moreover, Wang *et al.* [77] adopt a two-stage mechanism for sentiment classification based on EEG data, where the model first converts EEG signals into reading texts and then employs a pretrained LLM for classification, achieving impressive zero-shot results.

5 Multi-modal Time Series Applications

5.1 Standard Time Series

Healthcare. Recent studies in healthcare highlight the multi-modal analysis of diverse medical data sources, such as EHRs (Electronic Health Records, containing lab values and clinical reports, etc.), audio, EEG (Electroencephalogram), ECG (Electrocardiogram), and other wearable and medical sensor recordings, for better disease

diagnosis and patient monitoring. For multi-modal analysis on EHR data, a common modeling strategy involves the interaction between lab values and clinical reports, including the concatenation [17] and attention mechanisms [62, 87] on modality embeddings. Moreover, existing methods explore different modeling techniques to better exploit the clinical notes, via domain-specific text encoders (e.g., ClinicalBERT [29, 87] and BioBERT [44, 62]) and different processing strategies. For example, text embeddings can be modeled separately based on patient groups [88] or through a decaying mechanism based on time intervals [37] before interacting with time series embeddings, which leads to improved mortality prediction.

In addition to EHRs, multi-modal modeling methods have been tailored for audio [96], ECG [46], and EEG [77]. Zhang *et al.* [96] focus on a respiratory health classification task by integrating both audio and textual descriptions. Li *et al.* [46] propose a multi-modal contrastive learning framework, constructing positive and negative samples by pairing patients' report texts with corresponding ECG signals for self-supervised pretraining. The classification task is then performed by computing the cosine similarity between different text representations and the target ECG representation. Wang *et al.* [77] propose a two-stage method for zero-shot EEG-based sentiment classification. First, a pretrained BART model is used for EEG-to-text decoding, followed by a trained text sentiment classifier that converts the generated text into sentiment categories. Similarly, Liu *et al.* [58] fuses physiological and behavioral time-series sensor data with real-world contextual information to effectively harness LLMs for wellness assessment. By fine-tuning the models with few-shot question-answer pairs that include contextual details, they improve performance on various healthcare tasks (e.g., cardiac signal analysis, physical activity recognition *etc.*).

Finance. Recently, multi-modal time series analysis has received increasing attention in financial applications. Yu *et al.* [90] and Xie *et al.* [82] focus on stock prediction by integrating stock price movements, company profiles, and news directly into structured LLM prompts, enabling models to perform reasoning over multiple modalities. Yu *et al.* [90] applies GPT-4 and Open LLaMA to forecast NASDAQ-100 stock returns through instruction-based prompting and fine-tuning, demonstrating that structured LLM-driven inference can outperform traditional econometric models. Meanwhile, Xie *et al.* [82] conducts a zero-shot analysis of ChatGPT's capabilities for multi-modal stock movement prediction, incorporating CoT prompting to assess the impact of social media sentiment on stock trends. Chen *et al.* [12] and Sawhney *et al.* [69] also incorporate graph structures for stock movement prediction. For instance, Chen *et al.* [12] uses ChatGPT to infer dynamic stock relationship graphs from news, which reflects market conditions and enhances prediction accuracy. Beyond the predictive tasks, Bamford *et al.* [3] proposes a multi-modal retrieval framework, where the model aligns both modalities in a shared latent space through contrastive learning. This framework enables flexible financial time series search via text or sketches, with faster and more accurate retrieval than SQL-based methods.

Others. Multi-modal time series analysis also exists in retail, IoT, computer vision and audio. In the retail sector, Ekambaram *et al.* [19] utilizes product images and textual descriptions, including attributes like color, pattern, and sleeve style, while incorporating temporal and exogenous features for new product sales forecasting.

More recently, Skenderi *et al.* [72] integrates additional modality data, including product images and text descriptions, along with Google Trends data for sales forecasting. In IoT applications, MIA [83] enhances power transformer fault diagnosis by integrating multi-modal data, including dissolved gas analysis (DGA) and infrared images, to improve accuracy and efficiency. MULAN [98] converts log sequences into time-series data using a log-tailored language model and employs contrastive learning to leverage multi-modal data, facilitating root-cause discovery for system failures. In computer vision, LA-GCN [84] utilizes textual embeddings of joint names and action labels to generate faithful structural priors, enhancing skeleton-based action modeling and improving recognition tasks. In speech applications, AV-HuBERT [71] employs a self-supervised representation learning framework to leverage correlated audio and visual information [91], while SpeechGPT [71, 91] integrates audio and text to enhance generation performance.

5.2 Spatial Time Series

Transportation and Mobility. Several recent studies on traffic prediction highlight the importance of multi-modal contexts to enhance forecasting accuracy. Guo *et al.* [25] transforms California traffic data into structured LLM prompts. The method uses LLaMA models and instruction fine-tuning to improve spatial-temporal learning. Meanwhile, Li *et al.* [50] employs a spatial-temporal dependency encoder to align numerical New York City traffic data with LLMs, incorporating weather, geographic context, and historical flow patterns to refine predictions. Similarly, Feng *et al.* [20] proposes CityGPT, enhancing LLMs' spatial cognition for urban reasoning, mobility prediction, and navigation by integrating urban mobility data, road networks, and human behavior through instruction tuning. These studies demonstrate that LLM-based multi-modal fusion not only enhances traffic forecasting but also improves model interpretability and adaptability across diverse urban scenarios.

Environment. Integrating multi-modal information benefits environmental studies, particularly by addressing the prevalent challenge of missing values. VIMTS [97] utilizes a structured variational approximation technique to impute missing high-dimensional modalities (stream image data) by transforming them into low-dimensional features derived from simpler, related modalities (meteorological time series records), ensuring cross-modal correlations and interpretability of the imputation process. Additionally, LITE [45] addresses the imputation of missing features through a sparse Mixture of Experts framework. It integrates and encodes various environmental variables through a unified encoder. Directed by domain-specific instructions, a language model is utilized to merge these multi-modal representations, thereby improving the accuracy of environmental spatial-temporal predictions.

6 Future Research Directions

Reasoning with Multi-modal Time Series. Enhancing reasoning with multi-modal time series is pivotal for the development of intelligent systems. Future research should focus on creating a unified framework that can seamlessly integrate temporal reasoning with contextual understanding, enabling models to handle multiple time series tasks with interpretability. One potential path is to incorporate external knowledge bases and real-world context,

such as developing a retrieval-augmented generation (RAG) [61] system, to enhance the reasoning process and allow models to make informed inferences beyond the immediate data. It is also promising to synergize time series model and language agents to provide more faithful and reliable reasoning on real-world contexts [32, 70].

Decision Making. Multi-modal time series analysis presents a promising future direction to enhance decision-making processes, which is crucial in high-stakes applications. By leveraging predictive signals and explanations from multi-modal contexts, future research can develop more adaptive, interpretable, and reliable decision-support systems, to facilitate the downstream optimization tasks such as resource allocation and risk management.

Domain Generalization. One key challenge in multi-modal time series analysis is domain generalization, which enables a model trained on one or more source domains to effectively generalize to unseen target domains, ensuring robustness against distribution shifts. In multi-modal time series, distribution shifts can be multifaceted, stemming not only from time series, but also from other modalities. Therefore, it is crucial to develop specialized domain generalization methods for effective multi-modal time series analysis, including strategies to identify and preserve domain-invariant components across modalities while capturing modality-specific variations for rapid adaptation. Additionally, disentangling the effects of each modality is essential to better understand their individual contributions and mitigate cross-modal interference.

Robustness to Missing and Noisy Modalities. Multi-modal time series analysis often frequently encounters messy real-world contexts with incomplete or noisy data. Existing methods employ an iterative context-refinement algorithm [32] that filters out less relevant information, thereby enhancing the predictive insights derived from multi-modal time series. Nonetheless, effectively dealing with missing and noisy modalities still demands further exploration. In particular, developing strategies for modality-specific imputation, noise reduction, and relevance quantification will be crucial to improving the real-world applicability of existing methods.

Ethical Considerations and Bias Mitigation. In light of potential biases in multi-modal time series datasets, future research should integrate fairness-aware techniques (*e.g.*, fairness constraints, counterfactual analysis, adversarial debiasing). These methods should be paired with robust assessments to detect and mitigate bias, ensuring fair and responsible outcomes.

7 Conclusion

In this survey, we provide a comprehensive overview of existing multi-modal time series methods. We first discuss the multi-modal time series used in existing methods. Then, we propose a taxonomy based on cross-modal interactions between time series and other modalities. The existing methods are categorized and summarized accordingly. We also discuss the real-world applications and highlight future research directions in this promising area.

Acknowledgments

This research was supported in part by the National Science Foundation (NSF) CAREER IIS-2338878, as well as by generous research gifts from NEC Labs America Inc. and Morgan Stanley.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. 2018. LRS3-TED: a large-scale dataset for visual speech recognition. arXiv:1809.00496 [cs.CV] <https://arxiv.org/abs/1809.00496>
- [2] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multi-modal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* 41, 2 (2018), 423–443.
- [3] Tom Bamford, Andrea Coletta, Elizabeth Fons, Sriram Gopalakrishnan, Svitlana Vytrenko, Tucker Balch, and Manuela Veloso. 2023. Multi-Modal Financial Time-Series Retrieval Through Latent Space Projections. In *Proceedings of the Fourth ACM International Conference on AI in Finance* (Brooklyn, NY, USA) (ICAF '23). Association for Computing Machinery, New York, NY, USA, 498–506. doi:10.1145/3604237.3626901
- [4] Dhananjay Bhattacharya, Nayan K Sharma, Debottam Dutta, Srikanth R Chetupalli, Prashant Mote, Sriram Ganapathy, Jyothi Bhat, Shreyas Ramoji, Pravin Ghosh, Aswin Subramanian, et al. 2023. Coswara: a respiratory sounds and symptoms dataset for remote screening of SARS-CoV-2 infection. *Scientific Data* 10, 1 (2023), 397.
- [5] Nimesha Chan, Felix Parker, William Bennett, Tianyi Wu, Mung Yao Jia, James Fackler, and Kimia Ghobadi. 2024. Medtsllm: Leveraging llms for multimodal medical time series analysis. *arXiv preprint arXiv:2408.07773* (2024).
- [6] Sameep Chattopadhyay, Pulkit Paliwal, Sai Shankar Narasimhan, Shubhankar Agarwal, and Sandeep P. Chinchali. 2025. Context Matters: Leveraging Contextual Features for Time Series Forecasting. arXiv:2410.12672 [cs.LG] <https://arxiv.org/abs/2410.12672>
- [7] Jialin Chen, Aosong Feng, Ziyu Zhao, Juan Garza, Gaukhar Nurbek, Cheng Qin, Ali Maatouk, Leandros Tassiulas, Yifeng Gao, and Rex Ying. 2025. MTBench: A Multimodal Time Series Benchmark for Temporal Reasoning and Question Answering. *arXiv preprint arXiv:2503.16858* (2025).
- [8] Shengyu Chen, Yiqun Xie, Xiang Li, Xu Liang, and Xiaowei Jia. 2023. Physics-guided meta-learning method in baseflow prediction over large regions. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*. SIAM, 217–225.
- [9] Shengyu Chen, Jacob A Zwart, and Xiaowei Jia. 2022. Physics-guided graph meta learning for predicting water temperature and streamflow in stream networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2752–2761.
- [10] Si-An Chen, Chun-Liang Li, Sercan O Arik, Nathanael Christian Yoder, and Tomas Pfister. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. *Transactions on Machine Learning Research* (2023).
- [11] Wei Chen, Xixuan Hao, Yuankai Wu, and Yuxuan Liang. 2024. Terra: A Multimodal Spatio-Temporal Dataset Spanning the Earth. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 66329–66356. https://proceedings.neurips.cc/paper_files/paper/2024/file/7a6a7fbd1ee0c9684b3f919f79d129ef-Paper-Datasets_and_Benchmarks_Track.pdf
- [12] Zihan Chen, Lei Nico Zheng, Cheng Lu, Jialu Yuan, and Di Zhu. 2023. ChatGPT Informed Graph Neural Network for Stock Movement Prediction. Available at SSRN 4464002 (2023).
- [13] Mingyue Cheng, Yiheng Chen, Qi Liu, Zhiding Liu, Yucong Luo, and Enhong Chen. 2025. InstrucTime: Advancing Time Series Classification with Multimodal Language Modeling. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining* (Hannover, Germany) (WSDM '25). Association for Computing Machinery, New York, NY, USA, 792–800. doi:10.1145/3701551.3703499
- [14] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. 2018. VoxCeleb2: Deep Speaker Recognition. In *Interspeech 2018* (Interspeech Proceedings). ISCA. doi:10.21437/Interspeech.2018-1929
- [15] Helena Cousijn, Patricia Feeney, Daan Lowenberg, Elisa Presani, and Natasha Simons. 2019. A data citation roadmap for scholarly data repositories. *Scientific Data* 6, 1 (2019), 28.
- [16] Mayank Daswani, Mathias MJ Bellaiche, Marc Wilson, Desislav Ivanov, Mikhail Papkov, Eva Schneider, Jing Tang, Kay Lamerigts, Gabriela Botea, Michael A Sanchez, et al. 2024. Plots Unlock Time-Series Understanding in Multimodal Models. *arXiv preprint arXiv:2410.02637* (2024).
- [17] Iman Deznabi, Mohit Iyyer, and Madalina Fiterau. 2021. Predicting in-hospital mortality by combining clinical notes with time-series data. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 4026–4031. doi:10.18653/v1/2021.findings-acl.352
- [18] Zihan Dong, Xinyu Fan, and Zhiyuan Peng. 2024. FNSPID: A Comprehensive Financial News Dataset in Time Series. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Barcelona, Spain) (KDD '24). Association for Computing Machinery, New York, NY, USA, 4918–4927. doi:10.1145/3637528.3671629
- [19] Vijay Ekambaram, Kushagra Manglik, Sumanta Mukherjee, Surya Shravan Kumar Sajja, Satyam Dwivedi, and Vikas Raykar. 2020. Attention based multimodal new product sales time-series forecasting. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 3110–3118.
- [20] Jie Feng, Yuwei Du, Tianhui Liu, Siqi Guo, Yuming Lin, and Yong Li. 2024. Citygpt: Empowering urban spatial cognition of large language models. *arXiv preprint arXiv:2406.13948* (2024).
- [21] Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Alicia Curth, Stefan Bauer, Niki Kilbertus, Isaac S Kohane, and Mihaela van der Schaar. 2024. Causal machine learning for predicting treatment outcomes. *Nature Medicine* 30, 4 (2024), 958–968.
- [22] Mohammad Fraiwan, Luay Fraiwan, Basheer Khassawneh, and Ali Ibnian. 2021. A dataset of lung sounds recorded from the chest wall using an electronic stethoscope. *Data in Brief* 35 (2021), 106913. doi:10.1016/j.dib.2021.106913
- [23] Alessandro T. Gifford, Kshitij Dwivedi, Gemma Roig, and Radoslaw M. Cichy. 2022. A large and rich EEG dataset for modeling human visual object recognition. *NeuroImage* 264 (2022), 119754. doi:10.1016/j.neuroimage.2022.119754
- [24] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 922–929.
- [25] Xusen Guo, Qiming Zhang, Junyue Jiang, Mingxing Peng, Meixin Zhu, and Hao Frank Yang. 2024. Towards explainable traffic flow prediction with large language models. *Communications in Transportation Research* 4 (2024), 100150.
- [26] Paul Hager, Martin J Menten, and Daniel Rueckert. 2023. Best of both worlds: Multimodal contrastive learning with tabular and imaging data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 23924–23935.
- [27] Nora Hollenstein, Marius Troendle, Ce Zhang, and Nicolas Langer. 2020. ZuCo 2.0: A Dataset of Physiological Recordings During Natural Reading and Annotation. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, Marseille, France, 138–146. <https://aclanthology.org/2020.lrec-1.18/>
- [28] Shi Bin Hoo, Samuel Müller, David Salinas, and Frank Hutter. 2025. The tabular foundation model TabPFN outperforms specialized time series forecasting models based on simple features. *arXiv preprint arXiv:2501.02945* (2025).
- [29] Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2020. ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission. arXiv:1904.05342 [cs.CL] <https://arxiv.org/abs/1904.05342>
- [30] Jiahao Ji, Jingyuan Wang, Chao Huang, Junjie Wu, Boren Xu, Zhenhe Wu, Junbo Zhang, and Yu Zheng. 2023. Spatio-temporal self-supervised learning for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 4356–4364.
- [31] Furong Jia, Kevin Wang, Yixiang Zheng, Defu Cao, and Yan Liu. 2024. GPT4MTS: Prompt-based Large Language Model for Multimodal Time-series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 21 (Mar. 2024), 23343–23351. doi:10.1609/aaai.v38i21.30383
- [32] Yushan Jiang, Wenchao Yu, Geon Lee, Dongjin Song, Kijung Shin, Wei Cheng, Yanchi Liu, and Haifeng Chen. 2025. Explainable Multi-modal Time Series Prediction with LLM-in-the-Loop. arXiv:2503.01013 [cs.LG] <https://arxiv.org/abs/2503.01013>
- [33] Bo Jin, Haoyu Yang, Leilei Sun, Chuanren Liu, Yue Qu, and Jianing Tong. 2018. A treatment engine by predicting next-period prescriptions. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1608–1616.
- [34] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y. Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen. 2024. Time-LLM: Time Series Forecasting by Reprogramming Large Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=Unb5CVPtat>
- [35] Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. 2021. MIMIC-IV (version 1.0). doi:10.13026/s6n6-xd98
- [36] Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data* 3 (2016), 160035.
- [37] Swaraj Khadanga, Karan Aggarwal, Shafiq Joty, and Jaideep Srivastava. 2019. Using Clinical Notes with Time Series Data for ICU Management. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 6432–6437. doi:10.18653/v1/D19-1678
- [38] Kai Kim, Howard Tsai, Rajat Sen, Abhimanyu Das, Zihao Zhou, Abhishek Tanpure, Mathew Luo, and Rose Yu. 2024. Multi-Modal Forecaster: Jointly Predicting Time Series and Textual Data. arXiv:2411.06735 [cs.AI] <https://arxiv.org/abs/2411.06735>

- org/abs/2411.06735
- [39] Yaxuan Kong, Yiyuan Yang, Yuntae Hwang, Wenjie Du, Stefan Zohren, Zhangyang Wang, Ming Jin, and Qingsong Wen. 2025. Time-MQA: Time Series Multi-Task Question Answering with Context Enhancement. arXiv:2503.01875 [cs.CL]. <https://arxiv.org/abs/2503.01875>
- [40] Yaxuan Kong, Yiyuan Yang, Shiyu Wang, Chenghao Liu, Yuxuan Liang, Ming Jin, Stefan Zohren, Dan Pei, Yan Liu, and Qingsong Wen. 2025. Position: Empowering Time Series Reasoning with Multimodal LLMs. *arXiv preprint arXiv:2502.01477* (2025).
- [41] Dongjie Wang, Chengyuan Deng, Reon Matsuoka, Lecheng Zheng, Zhengzhang Chen, and Haifeng Chen. 2024. LEMMA-RCA: A Large Multi-modal Multi-domain Dataset for Root Cause Analysis. arXiv:2406.05375 [cs.AI]
- [42] Geon Lee, Wenchao Yu, Wei Cheng, and Haifeng Chen. 2024. MoAT: Multi-Modal Augmented Time Series Forecasting. <https://openreview.net/forum?id=uRXnoqDHH>
- [43] Geon Lee, Wenchao Yu, Kijung Shin, Wei Cheng, and Haifeng Chen. 2025. TimeCAP: Learning to Contextualize, Augment, and Predict Time Series Events with Large Language Model Agents. In *AAAI*.
- [44] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan So, and Jaewoo Kang. 2019. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics (Oxford, England)* 36 (09 2019). doi:10.1093/bioinformatics/btz682
- [45] Haoran Li, Junqi Liu, Zexian Wang, Shiyuan Luo, Xiaowei Jia, and Huaxiu Yao. 2024. LITE: Modeling Environmental Ecosystems with Multimodal Large Language Models. *arXiv preprint arXiv:2404.01165* (2024).
- [46] Jun Li, Che Liu, Sibao Cheng, Rossella Arcucci, and Shenda Hong. 2023. Frozen Language Model Helps ECG Zero-Shot Learning. In *Medical Imaging with Deep Learning*.
- [47] Xiang Li, Zhenyu Li, Chen Shi, Yong Xu, Qing Du, Mingkui Tan, and Jun Huang. 2024. AlphaFin: Benchmarking Financial Analysis with Retrieval-Augmented Stock-Chain Framework. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. 773–783.
- [48] Zekun Li, Shiyang Li, and Xifeng Yan. 2023. Time series as images: Vision transformer for irregularly sampled time series. *Advances in Neural Information Processing Systems* 36 (2023), 49187–49204.
- [49] Zihao Li, Xiao Lin, Zhining Liu, Jiaru Zou, Ziwei Wu, Lecheng Zheng, Dongqi Fu, Yada Zhu, Hendrik Hamann, Hanghang Tong, and Jingrui He. 2025. Language in the Flow of Time: Time-Series-Paired Texts Weaved into a Unified Temporal Narrative. arXiv:2502.08942 [cs.LG]. <https://arxiv.org/abs/2502.08942>
- [50] Zhonghang Li, Lianghao Xia, Jiabin Tang, Yong Xu, Lei Shi, Long Xia, Dawei Yin, and Chao Huang. 2024. UrbanGPT: Spatio-Temporal Large Language Models. *ArXiv abs/2403.00813* (2024). <https://api.semanticscholar.org/CorpusID:268230972>
- [51] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2024. Foundations & trends in multimodal machine learning: Principles, challenges, and open questions. *Comput. Surveys* 56, 10 (2024), 1–42.
- [52] Minhua Lin, Zhengzhang Chen, Yanchi Liu, Xujiang Zhao, Zongyu Wu, Junxiang Wang, Xiang Zhang, Suhang Wang, and Haifeng Chen. 2024. Decoding Time Series with LLMs: A Multi-Agent Framework for Cross-Domain Annotation. arXiv:2410.17462 [cs.AI]. <https://arxiv.org/abs/2410.17462>
- [53] Chenxi Liu, Qianxiong Xu, Hao Miao, Sun Yang, Lingzheng Zhang, Cheng Long, Ziyue Li, and Rui Zhao. 2025. TimeCMA: Towards LLM-Empowered Multivariate Time Series Forecasting via Cross-Modality Alignment. In *AAAI*.
- [54] Hanwen Liu, Daniel Hajjaligol, Benny Antony, Aiguo Han, and Xuan Wang. 2024. EEG2Text: Open Vocabulary EEG-to-Text Translation with Multi-View Transformer. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE Computer Society, Los Alamitos, CA, USA, 1824–1833. doi:10.1109/BigData62323.2024.10825980
- [55] Haoxin Liu, Shangqing Xu, Zhiyuan Zhao, Lingkai Kong, Harshavardhan Kamarthi, Aditya B. Sasanur, Megha Sharma, Jiaming Cui, Qingsong Wen, Chao Zhang, and B. Aditya Prakash. 2024. Time-MMD: Multi-Domain Multimodal Dataset for Time Series Analysis. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=fuD0h4R1IL>
- [56] Lei Liu, Shuo Yu, Runze Wang, Shenxun Ma, and Yanming Shen. 2024. How can large language models understand spatial-temporal data? *arXiv preprint arXiv:2401.14192* (2024).
- [57] Xu Liu, Junfeng Hu, Yuan Li, Shizhe Diao, Yuxuan Liang, Bryan Hooi, and Roger Zimmermann. 2024. UniTime: A Language-Empowered Unified Model for Cross-Domain Time Series Forecasting. In *Proceedings of the ACM Web Conference 2024 (Singapore, Singapore) (WWW '24)*. Association for Computing Machinery, New York, NY, USA, 4095–4106. doi:10.1145/3589334.3645434
- [58] Xin Liu, Daniel McDuff, Geza Kovacs, Isaac Galatzer-Levy, Jacob Sunshine, Jiening Zhan, Ming-Zher Poh, Shun Liao, Paolo Di Achille, and Shwetak Patel. 2023. Large Language Models are Few-Shot Health Learners. *arXiv preprint arXiv:2305.15525* (2023). <https://arxiv.org/abs/2305.15525>
- [59] Jingchao Ni, Ziming Zhao, ChengAo Shen, Hanghang Tong, Dongjin Song, Wei Cheng, Dongsheng Luo, and Haifeng Chen. 2025. Harnessing Vision Models for Time Series Analysis: A Survey. *arXiv preprint arXiv:2502.08869* (2025).
- [60] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.
- [61] Kanghui Ning, Zijie Pan, Yu Liu, Yushan Jiang, James Y. Zhang, Kashif Rasul, Anderson Schneider, Lintao Ma, Yuriy Nevmyvaka, and Dongjin Song. 2025. TS-RAG: Retrieval-Augmented Generation based Time Series Foundation Models are Stronger Zero-Shot Forecaster. arXiv:2503.07649 [cs.LG]. <https://arxiv.org/abs/2503.07649>
- [62] K. Niu, K. Zhang, X. Peng, Y. Pan, and N. Xiao. 2023. Deep Multi-Modal Intermediate Fusion of Clinical Record and Time Series Data in Mortality Prediction. *Frontiers in Molecular Biosciences* 10 (2023), 1136071. doi:10.3389/fmolb.2023.1136071
- [63] Zijie Pan, Yushan Jiang, Sahil Garg, Anderson Schneider, Yuriy Nevmyvaka, and Dongjin Song. 2024. S²IP-LLM: Semantic Space Informed Prompt Learning with LLM for Time Series Forecasting. In *Forty-first International Conference on Machine Learning*.
- [64] Vinay Prithyani, Mohsin Mohammed, Richa Gadgil, Ricardo Buitrago, Vinija Jain, and Aman Chadha. 2024. On the Feasibility of Vision-Language Models for Time-Series Classification. *arXiv preprint arXiv:2412.17304* (2024).
- [65] Yao Qin, Dongjin Song, Haifeng Cheng, Wei Cheng, Guofei Jiang, and Garrison W Cottrell. 2017. A dual-stage attention-based recurrent neural network for time series prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 2627–2633.
- [66] Hadi Rezaei, Hamidreza Faaljou, and Gholamreza Mansourfar. 2021. Stock price prediction using deep learning and frequency decomposition. *Expert Systems with Applications* 169 (2021), 114332.
- [67] Bruno M Rocha, Dimitris Filos, Luis Mendes, Gorkem Serbes, Sezer Ulukaya, Yasemin P Kahya, Nikša Jakovljevic, Tatjana L Turukalo, Ioannis M Vogiatzis, Eleni Perantoni, Evangelos Kaimakamis, Pantelis Natsiavas, Ana Oliveira, Cristina Jácome, Alda Marques, Nicos Maglaveras, Rui Pedro Paiva, Ioanna Chouvarda, and Paulo de Carvalho. 2019. An open access database for the evaluation of respiratory sound classification algorithms. *Physiological Measurement* 40, 3 (2019), 035001.
- [68] Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. 2023. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10219–10228.
- [69] Ramit Sawhney, Shivam Agarwal, Arnav Wadhwa, and Rajiv Ratn Shah. 2020. Deep Attentive Learning for Stock Movement Prediction From Social Media Text and Company Correlations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 8415–8426. doi:10.18653/v1/2020.emnlp-main.676
- [70] ChengAo Shen, Zhengzhang Chen, Dongsheng Luo, Dongkuan Xu, Haifeng Chen, and Jingchao Ni. 2024. Exploring Multi-Modal Integration with Tool-Augmented LLM Agents for Precise Causal Discovery. arXiv:2412.13667 [cs.LG]. <https://arxiv.org/abs/2412.13667>
- [71] Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed. 2022. Learning Audio-Visual Speech Representation by Masked Multimodal Cluster Prediction. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Z1Qm11uOM>
- [72] Geri Skenderi, Christian Joppi, Matteo Denitto, and Marco Cristani. 2024. Well googled is half done: Multimodal forecasting of new fashion product sales with image-based google trends. *Journal of Forecasting* 43, 6 (2024), 1982–1997.
- [73] Patrick Wagner, Nils Strodthoff, Ralf-Dieter Boussejot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific Data* 7, 1 (2020), 154.
- [74] Jiahao Wang, Mingyue Cheng, Qingyang Mao, Yitong Zhou, Feiyang Xu, and Xin Li. 2025. TableTime: Reformulating Time Series Classification as Training-Free Table Understanding with Large Language Models. arXiv:2411.15737 [cs.AI]. <https://arxiv.org/abs/2411.15737>
- [75] Xinlei Wang, Maiké Feng, Jing Qiu, JINJIN GU, and Junhua Zhao. 2024. From News to Forecast: Integrating Event Analysis in LLM-Based Time Series Forecasting with Reflection. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 58118–58153. https://proceedings.neurips.cc/paper_files/paper/2024/file/6aef8bffb372096ee73d98da30119f89-Paper-Conference.pdf
- [76] Xiaochen Wang, Junyu Luo, Jiaqi Wang, Ziyi Yin, Suhan Cui, Yuan Zhong, Yaqing Wang, and Fenglong Ma. 2023. Hierarchical Pretraining on Multimodal Electronic Health Records. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 2839–2852. doi:10.18653/v1/2023.emnlp-main.171

- [77] Zhenhailong Wang and Heng Ji. 2021. Open Vocabulary Electroencephalography-To-Text Decoding and Zero-shot Sentiment Classification. In *AAAI Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:244909027>
- [78] Andrew Robert Williams, Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Jithendara Subramanian, Roland Riachi, James Requeima, Alexandre Lacoste, Irina Rish, Nicolas Chapados, et al. 2024. Context is key: A benchmark for forecasting with essential textual information. *arXiv preprint arXiv:2410.18959* (2024).
- [79] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *The Eleventh International Conference on Learning Representations*.
- [80] Huizhe Wu, Wei Zhang, Weiwei Shen, and Jun Wang. 2018. Hybrid Deep Sequential Modeling for Social Text-Driven Stock Prediction. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (Torino, Italy) (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 1627–1630. doi:10.1145/3269206.3269290
- [81] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 753–763.
- [82] Qianqian Xie, Weiguang Han, Yanzhao Lai, Min Peng, and Jimin Huang. 2023. The wall street neophyte: A zero-shot analysis of chatgpt over multimodal stock movement prediction challenges. *arXiv preprint arXiv:2304.05351* (2023).
- [83] Zhikai Xing and Yigang He. 2023. Multi-modal information analysis for fault diagnosis with time-series data from power transformer. *International Journal of Electrical Power & Energy Systems* 144 (2023), 108567. doi:10.1016/j.ijepes.2022.108567
- [84] Haojun Xu, Yan Gao, Zheng Hui, Jie Li, and Xinbo Gao. 2023. Language Knowledge-Assisted Representation Learning for Skeleton-Based Action Recognition. arXiv:2305.12398 [cs.CV] <https://arxiv.org/abs/2305.12398>
- [85] Wenyan Xu, Dawei Xiang, Yue Liu, Xiyu Wang, Yanxiang Ma, Liang Zhang, Chang Xu, and Jiaheng Zhang. 2025. FinMultiTime: A Four-Modal Bilingual Dataset for Financial Time-Series Analysis. arXiv:2506.05019 [cs.CE] <https://arxiv.org/abs/2506.05019>
- [86] Yumo Xu and Shay B. Cohen. 2018. Stock Movement Prediction from Tweets and Historical Prices. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, Melbourne, Australia, 1970–1979. doi:10.18653/v1/P18-1183
- [87] Bo Yang and Lijun Wu. 2021. How to Leverage the Multimodal EHR Data for Better Medical Prediction?. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 4029–4038. doi:10.18653/v1/2021.emnlp-main.329
- [88] Haiyang Yang, Li Kuang, and FengQiang Xia. 2021. Multimodal Temporal-Clinical Note Network for Mortality Prediction. *Journal of Biomedical Semantics* 12, 1 (2021), 1–14. doi:10.1186/s13326-021-00235-3
- [89] Kun Yi, Qi Zhang, Wei Fan, Shoujin Wang, Pengyang Wang, Hui He, Ning An, Defu Lian, Longbing Cao, and Zhendong Niu. 2024. Frequency-domain MLPs are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems* 36 (2024).
- [90] Xinli Yu, Zheng Chen, and Yanbin Lu. 2023. Harnessing LLMs for Temporal Data - A Study on Explainable Financial Time Series Forecasting. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Mingxuan Wang and Imed Zitouni (Eds.). Association for Computational Linguistics, Singapore, 739–753. doi:10.18653/v1/2023.emnlp-industry.69
- [91] Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023. SpeechGPT: Empowering Large Language Models with Intrinsic Cross-Modal Conversational Abilities. arXiv:2305.11000 [cs.CL] <https://arxiv.org/abs/2305.11000>
- [92] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. 2024. Vision-language models for vision tasks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024).
- [93] Liheng Zhang, Charu Aggarwal, and Guo-Jun Qi. 2017. Stock price prediction via discovering multi-frequency trading patterns. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 2141–2149.
- [94] Xiyue Zhang, Chao Huang, Yong Xu, Lianghao Xia, Peng Dai, Liefeng Bo, Junbo Zhang, and Yu Zheng. 2021. Traffic flow forecasting with spatial-temporal graph diffusion network. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35, 15008–15015.
- [95] Xiang Zhang, Lina Yao, Manqing Dong, Zhe Liu, Yu Zhang, and Yong Li. 2020. Adversarial representation learning for robust patient-independent epileptic seizure detection. *IEEE journal of biomedical and health informatics* 24, 10 (2020), 2852–2859.
- [96] Yuwei Zhang, Tong Xia, Aaqib Saeed, and Cecilia Mascolo. 2024. RespLLM: Unifying Audio and Text with Multimodal LLMs for Generalized Respiratory Health Prediction. arXiv:2410.05361 [cs.LG] <https://arxiv.org/abs/2410.05361>
- [97] Xiaohu Zhao, Kebin Jia, Benjamin Letcher, Jennifer Fair, Yiqun Xie, and Xiaowei Jia. 2022. VIMTS: Variational-based Imputation for Multi-modal Time Series. In *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 349–358.
- [98] Lecheng Zheng, Zhengzhang Chen, Jingrui He, and Haifeng Chen. 2024. MULLAN: multi-modal causal structure learning and root cause analysis for microservice systems. In *Proceedings of the ACM Web Conference 2024*, 4107–4116.
- [99] Siru Zhong, Weilin Ruan, Ming Jin, Huan Li, Qingsong Wen, and Yuxuan Liang. 2025. Time-VLM: Exploring Multimodal Vision-Language Models for Augmented Time Series Forecasting. arXiv:2502.04395 [cs.CV] <https://arxiv.org/abs/2502.04395>
- [100] Xin Zhou, Weiqing Wang, Francisco J Baldán, Wray Buntine, and Christoph Bergmeir. 2025. MoTime: A Dataset Suite for Multimodal Time Series Forecasting. *arXiv preprint arXiv:2505.15072* (2025).
- [101] Zihao Zhou and Rose Yu. 2024. Can LLMs Understand Time Series Anomalies? *arXiv preprint arXiv:2410.05440* (2024).
- [102] Jiaxin Zhuang, Leon Yan, Zhenwei Zhang, Ruiqi Wang, Jiawei Zhang, and Yuantao Gu. 2024. See it, Think it, Sorted: Large Multimodal Models are Few-shot Time Series Anomaly Analyzers. *arXiv preprint arXiv:2411.02465* (2024).