

Large Language Models for Time Series: A Survey

Xiyuan Zhang, Ranak Roy Chowdhury, Rajesh K. Gupta and Jingbo Shang

University of California, San Diego

{xiyuanzh, rrchowdh, rgupta, jshang}@ucsd.edu

Abstract

Large Language Models (LLMs) have seen significant use in domains such as natural language processing and computer vision. Going beyond text, image and graphics, LLMs present a significant potential for analysis of time series data, benefiting domains such as climate, IoT, healthcare, traffic, audio and finance. This survey paper provides an in-depth exploration and a detailed taxonomy of the various methodologies employed to harness the power of LLMs for time series analysis. We address the inherent challenge of bridging the gap between LLMs’ original text data training and the numerical nature of time series data, and explore strategies for transferring and distilling knowledge from LLMs to numerical time series analysis. We detail various methodologies, including (1) direct prompting of LLMs, (2) time series quantization, (3) alignment techniques, (4) utilization of the vision modality as a bridging mechanism, and (5) the combination of LLMs with tools. Additionally, this survey offers a comprehensive overview of the existing multimodal time series and text datasets and delves into the challenges and future opportunities of this emerging field. We maintain an up-to-date Github repository¹ which includes all the papers and datasets discussed in the survey.

1 Introduction

Time series analysis plays a critical role in a variety of fields, including climate modeling, traffic management, healthcare monitoring and finance analytics. Time series analysis comprises a wide range of tasks such as classification [Liu *et al.*, 2023d], forecasting [Gruver *et al.*, 2023], anomaly detection [Zhang *et al.*, 2023f], and imputation [Chen *et al.*, 2023]. Traditionally, these tasks have been tackled using classical signal processing techniques such as frequency analysis and decomposition-based approaches. More recently, deep learning approaches like Convolutional Neural Networks (CNNs), Long Short-Term Memory networks (LSTMs), and Transformers [Wen *et al.*, 2022] have revolutionized this field and




Figure 1: Large language models have recently been applied for various time series tasks in diverse application domains.

proved effective in extracting meaningful patterns from time series data, making them the primary approaches of time series analysis in various domains.

In recent years, Large Language Models (LLMs) have gained substantial attention particularly in the fields of Natural Language Processing (NLP) and Computer Vision (CV). Prominent models such as GPT-4 [OpenAI, 2023] have transformed the landscape of text processing by offering unprecedented accuracy in tasks such as text generation, translation, sentiment analysis, question answering and summarization. In the CV domain, LLMs have also facilitated advancements in image recognition, object detection, and generative tasks, leading to more intelligent and capable visual systems [Song *et al.*, 2023]. Inspired by these successes, researchers are now exploring the potential of LLMs in the realm of time series analysis, expecting further breakthroughs, as shown in Figure 1. While several surveys offer a broad perspective on large models for time series in general [Jin *et al.*, 2023b; Ma *et al.*, 2023], these do not specifically focus on LLMs or the key challenge of bridging modality gap, which stems from LLMs being originally trained on discrete textual data, in contrast to the continuous numerical nature of time series.

Our survey uniquely contributes to the existing literature by emphasizing how to bridge such modality gap and transfer knowledge from LLMs for time series analysis. Our survey also covers more diverse application domains, ranging from climate, Internet of Things (IoT), to healthcare, traffic management, and finance. Moreover, certain intrinsic properties of time series, like continuity, auto-regressiveness, and dependency on the sampling rate, are also shared by audio, speech, and music data. Therefore, we also present representative LLM-based works from these domains to explore how

¹<https://github.com/xiyuanzh/awesome-llm-time-series>

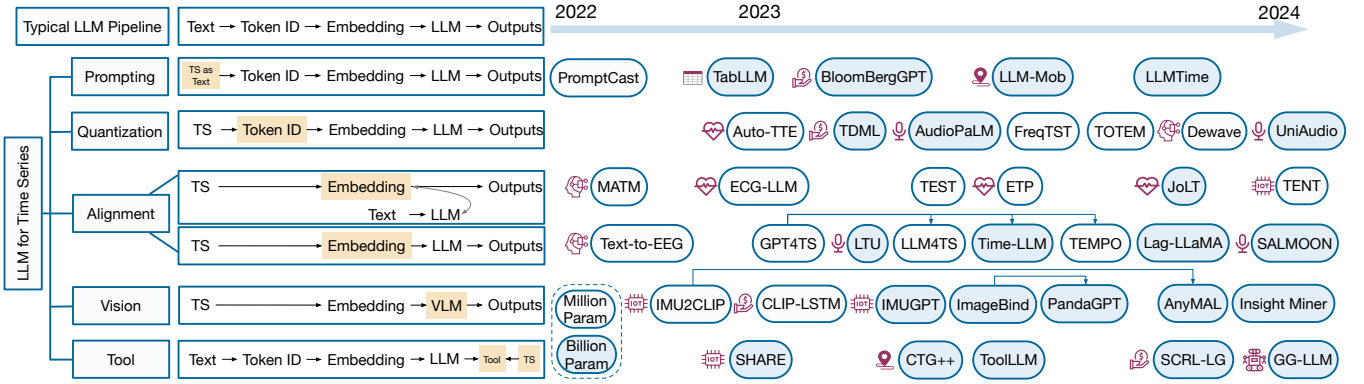


Figure 2: Left: Taxonomy of LLMs for time series analysis (prompting, quantization, alignment which is further categorized into two groups as detailed in Figure 4, vision as bridge, tool integration). For each category, key distinctions are drawn in comparison to the standard LLM pipeline shown at the top of the figure. Right: We present representative works for each category, sorted by their publication dates. The use of arrows indicates that later works build upon earlier studies. Dark(light)-colored boxes represent billion(million)-parameter models. Icons to the left of the text boxes represent the application domains of domain-specific models, with icons’ meanings illustrated in Figure 1.

we can use LLMs for other types of time series. We present a comprehensive taxonomy by categorizing these methodologies into five distinct groups, as shown in Figure 2. If we outline typical LLM-driven NLP pipelines in five stages - input text, tokenization, embedding, LLM, output - then each category of our taxonomy targets one specific stage in this pipeline. Specifically, (i) *Prompting* (input stage) treats time series data as raw text and directly prompts LLMs with time series; (ii) *Time Series Quantization* (tokenization stage) discretizes time series as special tokens for LLMs to process; (iii) *Alignment* (embedding stage) designs time series encoder to align time series embeddings with language space; (iv) *Vision as Bridge* (LLM stage) connects time series with Vision-Language Models (VLM) by employing visual representations as a bridge; (v) *Tool Integration* (output stage) adopts language models to output tools to benefit time series analysis. Beyond this taxonomy, our survey also compiles an extensive list of existing multimodal datasets that incorporate both time series and text. We conclude our paper by discussing future research directions in this emerging and promising field.

2 Background and Problem Formulation

Large language models are characterized by their vast number of parameters and extensive training data. They excel in understanding, generating, and interpreting human language and recently represent a significant advancement in artificial intelligence. The inception of LLMs can be traced back to models like GPT-2 [Radford *et al.*, 2019], BERT [Devlin *et al.*, 2018], BART [Lewis *et al.*, 2019], and T5 [Raffel *et al.*, 2020], which laid the foundational architecture. Over time, the evolution of these models has been marked by increasing complexity and capabilities, such as LLAMA-2 [Touvron *et al.*, 2023], PaLM [Chowdhery *et al.*, 2023], and GPT-4. More recently, researchers have developed multimodal large language models to integrate and interpret multiple forms of data, such as text, images, and time series, to achieve a more comprehensive understanding of information.

This survey focuses on how LLMs could benefit time series

analysis. We first define the mathematical formulation for the input and output, which may contain time series or (and) text depending on the downstream tasks, as well as the models.

Input: denoted as \mathbf{x} , composed of time series $\mathbf{x}_s \in \mathbb{R}^{T \times c}$ and optional text data \mathbf{x}_t represented as strings, where T, c represent the sequence length and the number of features.

Output: denoted as \mathbf{y} and may represent time series, text or numbers depending on the specific downstream task. For time series generation or forecasting task, \mathbf{y} represents generated time series \mathbf{y}_s or predicted k -step future time series $\mathbf{y}_s^{T+1:T+k}$. For text generation task, such as report generation, \mathbf{y} represents text data \mathbf{y}_t . For time series classification or regression task, \mathbf{y} represents numbers indicating the predicted classes or numerical values.

Model: We use f_θ parameterized by θ , g_ϕ parameterized by ϕ , and h_ψ parameterized by ψ to represent language, time series and vision models, where f_θ is typically initialized from pre-trained large language models. We optimize parameters θ, ϕ and ψ through loss function \mathcal{L} .

3 Taxonomy

In this section, we detail our taxonomy of applying LLMs for time series analysis, categorized by five groups. We summarize the representative works, mathematical formulation, advantages and limitations of each category in Table 2.

3.1 Prompting

Number-Agnostic Tokenization: The method treats numerical time series as raw textual data and directly prompts existing LLMs. For example, PromptCast [Xue and Salim, 2022] proposes prompt-based time series forecasting by converting numerical time series into text prompts and forecasting time series in a sentence-to-sentence manner. The input prompts are composed of context and questions following pre-defined templates. An illustrative prompt template for temperature forecasting, along with examples from other representative works, are showcased in Table 1. Similar prompting methods have been applied to forecast Place-of-Interest

Table 1: Examples of representative direct prompting methods.

Method	Example
PromptCast [Xue and Salim, 2022]	“From $\{t_1\}$ to $\{t_{\text{obs}}\}$, the average temperature of region $\{U_m\}$ was $\{x_t^m\}$ degree on each day. What is the temperature going to be on $\{t_{\text{obs}}\}$?”
Liu <i>et al.</i> [2023d]	“Classify the following accelerometer data in meters per second squared as either walking or running: 0.052,0.052,0.052,0.051,0.052,0.055,0.051,0.056,0.06,0.064”
TabLLM [Hegselmann <i>et al.</i> , 2023]	“The person is 42 years old and has a Master’s degree. She gained \$594. Does this person earn more than 50000 dollars? Yes or no? Answer:”
LLMTime [Gruver <i>et al.</i> , 2023]	“0.123, 1.23, 12.3, 123.0” \rightarrow “1 2 , 1 2 3 , 1 2 3 0 , 1 2 3 0 0”

(POI) customer flows (AuxMobLcast [Xue *et al.*, 2022]), energy load [Xue and Salim, 2023], and user’s next location (LLM-Mob [Wang *et al.*, 2023b]). Liu *et al.* [2023d] prompt PaLM-24B for health-related tasks such as activity recognition and daily stress estimate. TabLLM [Hegselmann *et al.*, 2023] prompts large language models with a serialization of the tabular data to a natural-language string for few-shot and zero-shot tabular data classification. Zhang *et al.* [2023f] prompt large language models to detect anomalous behaviors from mobility data. Xie *et al.* [2023a] extract historical price features such as open, close, high, and low prices to prompt ChatGPT in a zero-shot fashion.

Number-Specific Tokenization: More recently, LLM-Time [Gruver *et al.*, 2023] pointed out that Byte Pair Encoding (BPE) tokenization has the limitation of breaking a single number into tokens that don’t align with the digits, leading to inconsistent tokenization across different floating point numbers and complicating arithmetic operations [Spathis and Kawsar, 2023]. Therefore, following LLMs such as LLaMA and PaLM, they propose to insert spaces between digits to ensure distinct tokenization of each digit and use a comma (“,”) to separate each time step in a time series. They also scale time series to optimize token usage and keep fixed precision (e.g., two digits of precision) to efficiently manage context length. Meanwhile, BloomerGPT [Wu *et al.*, 2023] trains on financial data with text and numerical data and places each digit in its own chunk to better handle numbers. Using similar space-prefixed tokenization, Mirchandani *et al.* [2023] show that LLMs are general pattern machines capable of sequence transformation, completion and improvement.

3.2 Quantization

Quantization based method [Rabanser *et al.*, 2020] converts numerical data into discrete representations as input to LLMs. This approach can be further divided into two main categories based on the discretization technique employed.

Discrete Indices from VQ-VAE: The first type of quantization method transforms continuous time series into discrete indices as tokens. Among them one of the most popular methods is training a Vector Quantized-Variational AutoEncoder (VQ-VAE) [Van Den Oord *et al.*, 2017], which learns a codebook $\mathcal{C} = \{\mathbf{c}_i\}_{i=1}^K$ of K D -dimensional codewords $\mathbf{c}_i \in \mathbb{R}^D$ to capture the latent representations, as illustrated in Figure 3a. The method identifies the nearest neighbor

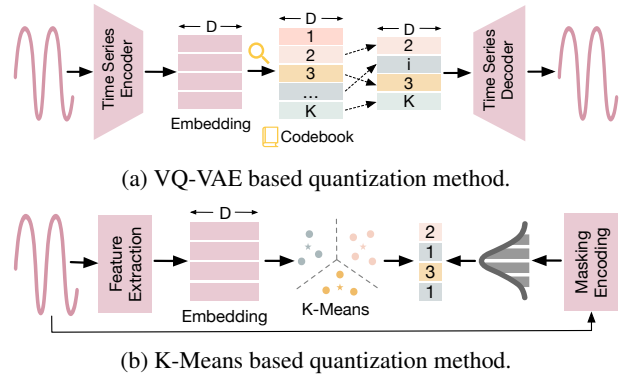


Figure 3: Two types of index-based quantization methods.

k_i of each step i of the encoded time series representation $g_\phi(\mathbf{x}_s) \in \mathbb{R}^{\frac{T}{S} \times D}$ in the codebook (S denotes the cumulative stride of VQ-VAE encoder), and uses the corresponding indices \mathbf{k} as the quantized input to language models:

$$\mathbf{q}_i = \mathbf{c}_{k_i}, k_i = \arg \min_j \|g_\phi(\mathbf{x}_s)_i - \mathbf{c}_j\|_2, \mathbf{k} = [k_i]_{i=1}^{\frac{T}{S}}. \quad (1)$$

Based on VQ-VAE, Auto-TTE [Chung *et al.*, 2023] quantizes ECGs into discrete formats and generates 12-lead ECG signals conditioned on text reports. DeWave [Duan *et al.*, 2023] adapts VQ-VAE to derive discrete codex encoding and aligns it with pre-trained BART for open-vocabulary EEG-to-text translation tasks. TOTEM [Anonymous, 2023c] also quantizes time series through VQ-VAE as input to Transformers for multiple downstream applications such as forecasting, classification, and translation. In the audio domain, UniAudio [Yang *et al.*, 2023] tokenizes different types of target audio using Residual Vector Quantization (RVQ) [Zeghidour *et al.*, 2021] (a hierarchy of multiple vector quantizers) and supports 11 audio generation tasks. ViOLA [Wang *et al.*, 2023a] unifies various crossmodal tasks involving speech and text by converting speech utterances to discrete tokens through RVQ. AudioGen [Kreuk *et al.*, 2022] learns discrete audio representations using vector quantization layers and generates audio samples conditioned on text inputs.

Discrete Indices from K-Means: Apart from employing VQ-VAE, researchers have also explored K-Means clustering for index-based tokenization, which uses the centroid in-

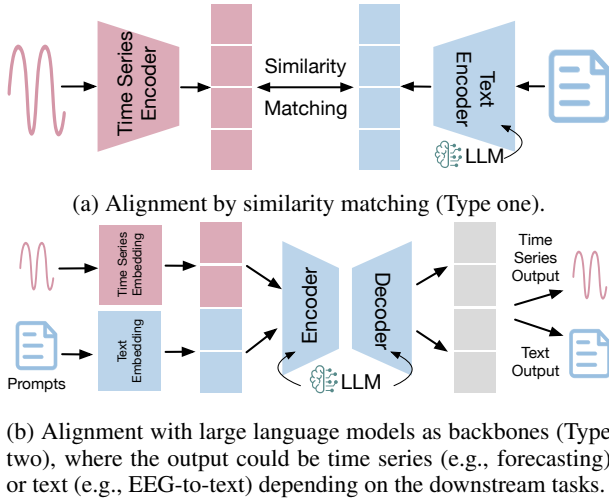


Figure 4: Two types of alignment based methods.

lices as discretized tokens [Hsu *et al.*, 2021], as shown in Figure 3b. Such methods are mostly applied in the audio domain. For example, SpeechGPT [Zhang *et al.*, 2023a] shows capability to perceive and generate multi-modal contents using K-Means based discrete unit extractor. AudioLM [Borsos *et al.*, 2023] discretizes codes produced by a neural audio codec using K-means clustering to achieve high-quality synthesis. It also combines discretized activations of language models pre-trained on audio using RVQ to capture long-term structure. Following the same quantization procedure, AudioPaLM [Rubenstein *et al.*, 2023] fuses PaLM-2 [Anil *et al.*, 2023] and AudioLM with a joint vocabulary that can represent speech and text with discrete tokens.

Discrete Indices from Other Techniques: Apart from the aforementioned time-domain quantization, FreqTST [Anonymous, 2023b] utilizes frequency spectrum as a common dictionary to discretize time series into frequency units with weights for downstream forecasting task.

Text Categories: The second type of quantization converts numerical data into pre-defined text categories, which is primarily adopted in financial domain. For example, TDML [Yu *et al.*, 2023] categorizes the weekly price fluctuations into 12 bins represented as “D $_i$ ” or “U $_i$ ”, where “D” indicates a decrease in price and “U” means an increase, and $i = 1, 2, 3, 4, 5, 5+$ represents the level of price change.

3.3 Alignment

The third type of works trains a separate encoder for time series, and aligns the encoded time series to the semantic space of language models. These works can be further categorized into two groups based on their specific alignment strategies, as illustrated in Figure 4.

Similarity Matching through Contrastive Loss: The first type of method aligns the time series embeddings with text embeddings through similarity matching, such as minimizing

the contrastive loss:

$$\mathcal{L} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\text{sim}(g_\phi(\mathbf{x}_{si}), f_\theta(\mathbf{x}_{ti})))^{\frac{1}{\gamma}}}{\sum_{k=1}^B \exp(\text{sim}(g_\phi(\mathbf{x}_{si}), f_\theta(\mathbf{x}_{tk})))^{\frac{1}{\gamma}}}, \quad (2)$$

where B, γ represent batch size and temperature parameter that controls distribution concentrations, and sim represents similarity score, typically computed as inner product:

$$\text{sim}(g_\phi(\mathbf{x}_{si}), f_\theta(\mathbf{x}_{ti})) = \langle g_\phi(\mathbf{x}_{si}), f_\theta(\mathbf{x}_{ti}) \rangle. \quad (3)$$

For instance, ETP [Liu *et al.*, 2023a; Li *et al.*, 2023a] integrates contrastive learning based pre-training to align electrocardiography (ECG) signals with textual reports. King *et al.* [2023] use similar contrastive framework to align 17 clinical measurements collected in Intensive Care Unit (ICU) to their corresponding clinical notes. TEST [Sun *et al.*, 2023] uses contrastive learning to generate instance-wise, feature-wise, and text-prototype-aligned time series embeddings to align with text embeddings. TENT [Zhou *et al.*, 2023b] aligns text embeddings with IoT sensor signals through a unified semantic feature space using contrastive learning. JoLT [Cai *et al.*, 2023] utilizes Querying Transformer (Q-Former) [Li *et al.*, 2023b] optimized with contrastive loss to align the time series and text representations.

Similarity Matching through Other Losses: Apart from contrastive loss, other loss functions are also employed to optimize similarity matching between time series embeddings and text embeddings. ECG-LLM [Qiu *et al.*, 2023] aligns the distribution between ECG and language embedding from ECG statements with an Optimal Transport based loss function to train an ECG report generation model. MTAM [Han *et al.*, 2022] uses various alignment techniques, such as Canonical Correlation Analysis and Wasserstein Distance, as loss functions to align electroencephalography (EEG) features with their corresponding language descriptions.

LLMs as Backbones: The second type of alignment method directly uses large language models as backbones following time series embedding layers. EEG-to-Text [Wang and Ji, 2022] feeds EEG embeddings to pre-trained BART for open vocabulary EEG-To-Text decoding and EEG-based sentiment classification. GPT4TS [Zhou *et al.*, 2023a] uses patching embeddings [Nie *et al.*, 2022] as input to frozen pre-trained GPT-2 where the positional embedding layers and self-attention blocks are retained during time series fine-tuning. The method provides a unified framework for seven time series tasks, including few-shot or zero-shot learning. Following GPT4TS, researchers further incorporated seasonal-trend decomposition (TEMPO [Cao *et al.*, 2023]), two-stage fine-tuning (LLM4TS [Chang *et al.*, 2023]), domain descriptions (UniTime [Liu *et al.*, 2023e]), graph attention mechanism (GATGPT [Chen *et al.*, 2023]), and spatial-temporal embedding module (ST-LLM [Liu *et al.*, 2024]). Time-LLM [Jin *et al.*, 2023a] reprograms time series data into text prototypes as input to LLaMA-7B. It also provides natural language prompts such as domain expert knowledge and task instructions to augment input context. Lag-Llama [Rasul *et al.*, 2023] builds univariate probabilistic time series forecasting model based on LLaMA architecture. In the audio, speech and music domains, researchers have also designed

Table 2: Summary of five major categories of aligning LLMs for time series analysis, including their respective subcategories, representative works, mathematical formulations, advantages and limitations. q and \mathbf{x}_v represent text-based quantization process and image data.

Method	Subcategory	Representative Works	Equations	Advantages	Limitations
Prompting	Number-Agnostic	PromptCast [Xue and Salim, 2022]	$\mathbf{y} = f_{\theta}(\mathbf{x}_s, \mathbf{x}_t)$	easy to implement; zero-shot capability	lose semantics;
	Number-Specific	LLMTime [Gruver <i>et al.</i> , 2023]			not efficient
Quantization	VQ-VAE	DeWave [Duan <i>et al.</i> , 2023]	$k_i = \arg \min_j \ g_{\phi}(\mathbf{x}_s)_i - \mathbf{c}_j\ _2$ $\mathbf{k} = [k_i]_{i=1}^T, \mathbf{y} = f_{\theta}(\mathbf{k}, \mathbf{x}_t)$	flexibility of index and time	may require
	K-Means	AudioLM [Borsos <i>et al.</i> , 2023]			two-stage
	Text Categories	TDML [Yu <i>et al.</i> , 2023]	$\mathbf{y} = f_{\theta}(q(\mathbf{x}_s), \mathbf{x}_t)$	series conversion	training
Alignment	Similarity Match	ETP [Liu <i>et al.</i> , 2023a]	$\mathbf{y} = g_{\phi}(\mathbf{x}_s)$	align semantics of different modalities;	complicated
		MATM [Han <i>et al.</i> , 2022]	$\mathcal{L} = \text{sim}(g_{\phi}(\mathbf{x}_s), f_{\theta}(\mathbf{x}_t))$		design and
	LLM Backbone	GPT4TS [Zhou <i>et al.</i> , 2023a]	$\mathbf{y} = f_{\theta}(g_{\phi}(\mathbf{x}_s), \mathbf{x}_t)$	end-to-end training	fine-tuning
Vision as Bridge	Paired Data	ImageBind [Girdhar <i>et al.</i> , 2023]	$\mathcal{L} = \text{sim}(g_{\phi}(\mathbf{x}_s), h_{\psi}(\mathbf{x}_v))$	additional visual knowledge	not hold
	TS Plots as Images	Wimmer and Rekabsaz [2023]			for all data
Tool	Code	CTG++ [Zhong <i>et al.</i> , 2023]	$z = f_{\theta}(\mathbf{x}_t)$	empower LLM with more abilities	optimization
	API	ToolLLM [Qin <i>et al.</i> , 2023]	$\mathbf{y} = z(\mathbf{x}_s)$		not end-to-end

dedicated encoders to embed speech (WavPrompt [Gao *et al.*, 2022], Speech LLaMA [Lakomkin *et al.*, 2023]), music (MU-LLaMA [Liu *et al.*, 2023c]), and general audio inputs (LTU [Gong *et al.*, 2023], SALMONN [Tang *et al.*, 2023]), and feed the embeddings to large language models.

3.4 Vision as Bridge

Time series data can be effectively interpreted or associated with visual representations, which align closer with textual data and have demonstrated successful integrations with large language models. Therefore, researchers have also leveraged vision modality as a bridge to connect time series with LLMs.

Paired Data: ImageBind [Girdhar *et al.*, 2023] uses image-paired data to bind six modalities (images, text, audio, depth, thermal, and Inertial Measurement Unit (IMU) time series) and learn a joint embedding space, enabling new emergent alignments and capabilities. PandaGPT [Su *et al.*, 2023] further combines the multimodal encoders from ImageBind and LLMs to enable visual and auditory instruction-following capabilities. IMU2CLIP [Moon *et al.*, 2022] aligns IMU time series with video and text, by projecting them into the joint representation space of Contrastive Language-Image Pre-training (CLIP) [Radford *et al.*, 2021]. AnyMAL [Moon *et al.*, 2023] builds upon IMU2CLIP by training a lightweight adapter to project the IMU embeddings into the text token embedding space of LLaMA-2-70B. It is also capable of transforming data from other modalities, such as images, videos, audio, into the same text embedding space.

Physics Relationships: IMUGPT [Leng *et al.*, 2023] generates IMU data from ChatGPT-augmented text descriptions. It first generates 3D human motion from text using pre-trained motion synthesis model T2M-GPT [Zhang *et al.*, 2023b]. Then it derives IMU data from 3D motion based on physics relationships of motion kinetics.

Time Series Plots as Images: CLIP-LSTM [Wimmer and Rekabsaz, 2023] transforms stock market data into sequences

of texts and images of price charts, and leverages pre-trained CLIP vision-language model to generate features for downstream forecasting. Insight Miner [Zhang *et al.*, 2023e] converts time series windows into images using lineplot, and feeds images into vision language model LLaVA [Liu *et al.*, 2023b] to generate time series trend descriptions.

3.5 Tool

This type of method does not directly use large language models to process time series. Instead, it applies large language models to generate indirect tools $z(\cdot)$, such as code and API calls, to benefit time series related tasks.

Code: CTG++ [Zhong *et al.*, 2023] applies GPT-4 to generate differentiable loss functions in a code format from text descriptions to guide the diffusion model to generate traffic trajectories. With this two-step translation, the LLM and diffusion model efficiently bridge the gap between user intent and traffic simulation.

API Call: ToolLLM [Qin *et al.*, 2023] introduces a general tool-use framework composed of data construction, model training, and evaluation. This framework includes API calls for time series tasks such as weather and stock forecasting.

Text Domain Knowledge: SHARE [Zhang *et al.*, 2023d] exploits the shared structures in human activity label names and proposes a sequence-to-sequence structure to generate label names as token sequences to preserve the shared label structures. It applies GPT-4 to augment semantics of label names. GG-LLM [Graule and Isler, 2023] leverages LLaMA-2 to encode world knowledge of common human behavioral patterns to predict human actions without further training. SCRL-LG [Ding *et al.*, 2023] leverages LLaMA-7B as stock feature selectors to extract meaningful representations from news headlines, which are subsequently employed in reinforcement learning for precise feature alignments.

Table 3: Summary of representative time series and text multimodal datasets.

Domain	Dataset	Size	Major Modalities	Task
Internet of Things	Ego4D ² [Grauman <i>et al.</i> , 2022]	3, 670h data, 3.85M narrations	text, IMU, video, audio, 3D	classification, forecasting
	DeepSQA ³ [Xing <i>et al.</i> , 2021]	25h data, 91K questions	text, imu	classification, question answering
Finance	PIXIU ⁴ [Xie <i>et al.</i> , 2023b]	136K instruction data	text, tables	5 NLP tasks, forecasting
	MoAT ⁵ [Anonymous, 2023a]	6 datasets, 2K timesteps in total	text, time series	forecasting
Healthcare	Zuco 2.0 ⁶ [Hollenstein <i>et al.</i> , 2019]	739 sentences	text, eye-tracking, EEG	classification, text generation
	PTB-XL ⁷ [Wagner <i>et al.</i> , 2020]	60h data, 71 unique statements	text, ECG	classification
	ECG-QA ⁸ [Oh <i>et al.</i> , 2023]	70 question templates	text, ECG	classification, question answering
Audio	OpenAQA-5M ⁹ [Gong <i>et al.</i> , 2023]	5.6M (audio, question, answer) tuples	text, audio	tagging, classification
Music	MusicCaps ¹⁰ [Agostinelli <i>et al.</i> , 2023]	5.5K music clips	text, music	captioning, generation
Speech	CommonVoice ¹¹ [Ardila <i>et al.</i> , 2019]	7, 335 speech hours in 60 languages	text, speech	ASR, translation

4 Comparison within the Taxonomy

We compare the five categories of our taxonomy and provide general guidelines for which category to choose based on considerations of data, model, efficiency and optimization.

Data: When no training data is available and the objective is to apply LLM for time series in an zero-shot fashion, it is preferable to use prompting-based methods. This is because direct prompting enables the utilization of pre-trained language models’ inherent capabilities without fine-tuning. However, representing numbers as strings can diminish the semantic value intrinsically tied to numerical data. Therefore, with adequate training data, quantization or alignment-based methods become more advantageous. As shown in Figure 2, these two categories are the most extensively studied ones in existing literature. Furthermore, if time series data can be interpreted or associated with visual representations, these representations can be incorporated to utilize the intrinsic knowledge embedded in the vision modality or pre-trained vision-language models.

Model: Prompting and tool integration methods tend to apply billion-parameter models as they often apply off-the-self LLMs without architectural modifications. By contrast, alignment and quantization methods vary from million to billion-parameter models, depending on the specific application requirements and available computational resources.

Efficiency: Prompting-based methods are not efficient for numerical data with high precision, as well as multivariate time series as it requires transforming each dimension into separate univariate time series, resulting in extremely long input. They are also less efficient for long-term predictions due to the computational demands of generating long sequences. These methods are more effective when dealing with simple numerical data that is richly interwoven with textual information, such as opening and closing stock prices in financial news articles. By contrast, quantization and alignment methods are more efficient to handle long sequences, as time series are typically down-sampled or segmented into patches before feeding into large language models.

Optimization: Depending on the specific discretization technique, quantization based method may require a two-stage training process (such as first training the VQ-VAE

model), which may result in sub-optimal performance compared with that achieved through end-to-end training in alignment methods. Using large language models as indirect tools empowers LLMs with more capabilities to manage numerical data, but also raises the level of complexity to optimize both LLMs and other components in an end-to-end fashion. Therefore, existing works of tool integration typically employ off-the-shelf LLMs without further fine-tuning.

5 Multimodal Datasets

Applying LLMs for time series benefits from the availability of multimodal time series and text data. In this section, we introduce representative multimodal time series and text datasets organized by their respective domains (Table 3).

Internet of Things (IoT): Human activity recognition is an important task in IoT domain, which identifies human activities given time series collected with IoT devices (such as IMU sensors). The corresponding text data are the labels or text descriptions of these activities. Ego4D [Grauman *et al.*, 2022] presents 3,670 hours of daily-life activity data across hundreds of scenarios, including household, outdoor, workplace, and leisure. The dataset is rich in modalities, including the IMU time series measurement, and dense temporally-aligned textual descriptions of the activities and object interactions, totaling 3.85 million sentences. Ego-Exo4D [Grauman *et al.*, 2023] further offers three kinds of paired natural language datasets including expert commentary, narrate-and-act descriptions provided by the participants, and atomic action descriptions similar as Ego4D. DeepSQA [Xing *et al.*, 2021] presents a generalized Sensory Question Answering (SQA)

²<https://ego4d-data.org/>

³<https://github.com/nesl/DeepSQA>

⁴<https://github.com/chancefocus/PIXIU>

⁵<https://openreview.net/pdf?id=uRXxnoqDHH>

⁶<https://osf.io/2urht/>

⁷<https://physionet.org/content/ptb-xl/1.0.3/>

⁸<https://github.com/Jwoo5/ecg-qa>

⁹<https://github.com/YuanGongND/ttu>

¹⁰<https://www.kaggle.com/datasets/googleai/musiccaps>

¹¹<https://commonvoice.mozilla.org/en/datasets>

framework to facilitate querying raw sensory data related to human activities using natural language.

Finance: PIXIU [Xie *et al.*, 2023b] presents multi-task and multi-modal instruction tuning data in the financial domain with 136K data samples. It contains both financial natural language understanding and prediction tasks, and covers 9 datasets of multiple modalities such as text and time series. MoAT [Anonymous, 2023a] constructs multimodal datasets with textual information paired with time series for each timestep, such as news articles extracted with relevant keywords, mostly covering finance related domains such as fuel, metal, stock and bitcoin.

Healthcare: Zuco 1.0 [Hollenstein *et al.*, 2018] and Zuco 2.0 [Hollenstein *et al.*, 2019] datasets contain simultaneous eye-tracking and EEG during natural reading and during annotation. PTB-XL [Wagner *et al.*, 2020] offers comprehensive metadata regarding ECG annotated by expert cardiologists, covering information such as ECG reports, diagnostic statements, diagnosis likelihoods, and signal-specific properties. Based on PTB-XL, ECG-QA [Oh *et al.*, 2023] introduces the first Question Answering (QA) dataset for ECG analysis, containing 70 question templates that cover a wide range of clinically relevant ECG topics.

Audio/Music/Speech: AudioSet [Gemmeke *et al.*, 2017] is a collection of 2 million 10-second audio clips excised from YouTube videos and labeled with the sounds that the clip contains from a set of 527 labels. OpenAQA-5M [Gong *et al.*, 2023] dataset consists of 1.9 million closed-ended and 3.7 million open-ended (audio, question, answer) tuples. MusicCaps [Agostinelli *et al.*, 2023] is a high-quality music caption dataset, including 5.5K music clips. MTG-Jamendo [Bogdanov *et al.*, 2019] is a dataset with 55,000 audio songs in various languages. Libri-Light [Kahn *et al.*, 2020] is an English dataset encompassing 60,000 hours of speech data. CommonVoice [Ardila *et al.*, 2019] is a multilingual speech dataset consisting of 7,335 validated hours in 60 languages.

6 Challenges and Future Directions

In this section, we introduce the challenges and promising future directions of applying LLMs for time series analysis.

6.1 Theoretical Understanding

Existing works empirically show the benefits of applying LLMs for time series analysis. For example, LIFT [Dinh *et al.*, 2022] empirically shows that language model fine-tuning can work for non-language tasks without changing the architecture or loss function; Gurnee and Tegmark [2023] empirically show that LLMs learn linear representations of space and time across multiple scales that are robust to prompting variations. Despite these empirical findings, there remains a gap in theoretical understanding of how models, primarily trained on textual data, can effectively interpret numerical time series. As a preliminary theoretical analysis, Yun *et al.* [2019] prove that Transformer models can universally approximate arbitrary continuous sequence-to-sequence functions on a compact domain. Additionally, GPT4TS [Zhou *et al.*, 2023a] theoretically shows that such generic capability of large language models can be related to Principal Component

Analysis (PCA), as minimizing the gradient with respect to the self-attention layer shares similarities with PCA. Further investigations on the generalizability of LLMs on numerical data is essential to establish solid understanding of the synergy between LLMs and time series analysis.

6.2 Multimodal and Multitask Analysis

Existing papers that apply LLMs for time series analysis mostly focus on single modality and single task at a time, such as forecasting, classification, text generation, and do not support simultaneous multimodal and multitask analysis. In computer vision and audio domains, models such as Unified-IO [Lu *et al.*, 2022] and UniAudio [Yang *et al.*, 2023] have unified multiple input modalities into a sequence of discrete vocabulary tokens to support multiple tasks within a single transformer-based architecture. More research into leveraging LLMs for multimodal and multitask analysis would lead to more powerful time series foundation models.

6.3 Efficient Algorithms

Time series, especially those that are multivariate or possess long history information may increase the computational complexity for existing large language models. Patching [Nie *et al.*, 2022] has been a widely adopted strategy to improve performance as well as reduce complexity, but large patches may obscure the semantic information of time series and negatively impact the performance. Therefore, developing more efficient algorithms is crucial for facilitating large-scale time series analysis and enhancing interactions with end users.

6.4 Combining Domain Knowledge

Combining existing statistical domain knowledge with LLMs may further boost the model’s capability for time series analysis. For example, TEMPO [Cao *et al.*, 2023] applies time series seasonal-trend decomposition and treats decomposed components as different semantic inductive biases as input to the pre-trained transformer. FreqTST [Anonymous, 2023b] leverages insights from the frequency domain by tokenizing single time series into frequency units with weights for downstream forecasting. Further incorporating domain knowledge, such as wavelet decomposition, auto-correlation analysis, and empirical mode decomposition may augment LLMs’ capabilities in analyzing time series data.

6.5 Customization and Privacy

Existing works on large language models and time series analysis typically train a global model for all end users. Training customized models for different users based on the global model may bring further benefits and flexibility. Another important consideration is privacy, especially as many time series data are collected in private settings for clinical purposes or smart home applications. As an initial attempt, FedAlign [Zhang *et al.*, 2023c] leverages federated learning frameworks and uses the expressive natural language class names as a common ground to align the latent spaces across different clients. Advancing research into model customization and user privacy preservation would broaden the scope and utility of LLM-empowered time series analysis.

7 Conclusion

We present the first survey that systematically analyzes the categorization of transferring knowledge from large language models for numerical time series analysis: direct prompting, time series quantization, alignment, the use of the vision modality to connect text and time series, and the integration of large language models with other analytical tools. For each category, we introduce their mathematical formulation, representative works, and compare their advantages and limitations. We also introduce representative multimodal text and time series datasets in various domains such as healthcare, IoT, finance, and audio. Concluding the paper, we outline the challenges and emerging directions for potential future research of LLM-empowered time series analysis.

References

- Andrea Agostinelli, Timo I Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, et al. Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*, 2023.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- Anonymous. MoAT: Multi-modal augmented time series forecasting. In *Submitted to The Twelfth International Conference on Learning Representations*, 2023. under review.
- Anonymous. Modeling time series as text sequence a frequency-vectorization transformer for time series forecasting. In *Submitted to The Twelfth International Conference on Learning Representations*, 2023. under review.
- Anonymous. Time series modeling at scale: A universal representation across tasks and domains. In *Submitted to The Twelfth International Conference on Learning Representations*, 2023. under review.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*, 2019.
- Dmitry Bogdanov, Minz Won, Philip Tovstogan, Alastair Porter, and Xavier Serra. The mtg-jamendo dataset for automatic music tagging. *ICML*, 2019.
- Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, et al. Audioldm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- Yifu Cai, Mononito Goswami, Arjun Choudhry, Arvind Srinivasan, and Artur Dubrawski. Jolt: Jointly learned representations of language and time-series. In *Deep Generative Models for Health Workshop NeurIPS 2023*, 2023.
- Defu Cao, Furong Jia, Serkan O Arik, Tomas Pfister, Yixiang Zheng, Wen Ye, and Yan Liu. Tempo: Prompt-based generative pre-trained transformer for time series forecasting. *arXiv preprint arXiv:2310.04948*, 2023.
- Ching Chang, Wen-Chih Peng, and Tien-Fu Chen. Llm4ts: Two-stage fine-tuning for time-series forecasting with pre-trained llms. *arXiv preprint arXiv:2308.08469*, 2023.
- Yakun Chen, Xianzhi Wang, and Guandong Xu. Gatgpt: A pre-trained large language model with graph attention network for spatiotemporal imputation. *arXiv preprint arXiv:2311.14332*, 2023.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- Hyunseung Chung, Jiho Kim, Joon-myung Kwon, Ki-Hyun Jeon, Min Sung Lee, and Edward Choi. Text-to-ecg: 12-lead electrocardiogram synthesis conditioned on clinical text reports. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Yujie Ding, Shuai Jia, Tianyi Ma, Bingcheng Mao, Xiuze Zhou, Liuliu Li, and Dongming Han. Integrating stock features and global information via large language models for enhanced stock return prediction. *arXiv preprint arXiv:2310.05627*, 2023.
- Tuan Dinh, Yuchen Zeng, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Pappaliopoulos, and Kangwook Lee. Lift: Language-interfaced fine-tuning for non-language machine learning tasks. *Advances in Neural Information Processing Systems*, 35:11763–11784, 2022.
- Yiqun Duan, Charles Zhou, Zhen Wang, Yu-Kai Wang, and Chin-teng Lin. Dewave: Discrete encoding of eeg waves for eeg to text translation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Heting Gao, Junrui Ni, Kaizhi Qian, Yang Zhang, Shiyu Chang, and Mark Hasegawa-Johnson. Wavprompt: Towards few-shot spoken language understanding with frozen language models. *arXiv preprint arXiv:2203.15863*, 2022.
- Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan

- Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2023.
- Yuan Gong, Hongyin Luo, Alexander H Liu, Leonid Karlinsky, and James Glass. Listen, think, and understand. *arXiv preprint arXiv:2305.10790*, 2023.
- Moritz A Graule and Volkan Isler. Gg-llm: Geometrically grounding large language models for zero-shot human activity forecasting in human-aware task planning. *arXiv preprint arXiv:2310.20034*, 2023.
- Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. *arXiv preprint arXiv:2311.18259*, 2023.
- Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew Gordon Wilson. Large language models are zero-shot time series forecasters. *arXiv preprint arXiv:2310.07820*, 2023.
- Wes Gurnee and Max Tegmark. Language models represent space and time. *arXiv preprint arXiv:2310.02207*, 2023.
- William Han, Jieli Qiu, Jiacheng Zhu, Mengdi Xu, Douglas Weber, Bo Li, and Ding Zhao. An empirical exploration of cross-domain alignment between language and electroencephalogram. *arXiv preprint arXiv:2208.06348*, 2022.
- Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. Tabllm: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*, pages 5549–5581. PMLR, 2023.
- Nora Hollenstein, Jonathan Rotsztein, Marius Troendle, Andreas Pedroni, Ce Zhang, and Nicolas Langer. Zuco, a simultaneous eeg and eye-tracking resource for natural sentence reading. *Scientific data*, 5(1):1–13, 2018.
- Nora Hollenstein, Marius Troendle, Ce Zhang, and Nicolas Langer. Zuco 2.0: A dataset of physiological recordings during natural reading and annotation. *arXiv preprint arXiv:1912.00903*, 2019.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuanfang Li, Shirui Pan, et al. Time-llm: Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728*, 2023.
- Ming Jin, Qingsong Wen, Yuxuan Liang, Chaoli Zhang, Siqiao Xue, Xue Wang, James Zhang, Yi Wang, Haifeng Chen, Xiaoli Li, et al. Large models for time series and spatio-temporal data: A survey and outlook. *arXiv preprint arXiv:2310.10196*, 2023.
- Jacob Kahn, Morgane Rivièrè, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazaré, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, et al. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673. IEEE, 2020.
- Ryan King, Tianbao Yang, and Bobak J Mortazavi. Multimodal pretraining of medical time series and notes. In *Machine Learning for Health (ML4H)*, pages 244–255. PMLR, 2023.
- Felix Kreuk, Gabriel Synnaeve, Adam Polyak, Uriel Singer, Alexandre Défossez, Jade Copet, Devi Parikh, Yaniv Taigman, and Yossi Adi. Audiogen: Textually guided audio generation. *arXiv preprint arXiv:2209.15352*, 2022.
- Egor Lakomkin, Chunyang Wu, Yassir Fathullah, Ozlem Kalinli, Michael L Seltzer, and Christian Fuegen. End-to-end speech recognition contextualization with large language models. *arXiv preprint arXiv:2309.10917*, 2023.
- Zikang Leng, Hyeokhyen Kwon, and Thomas Plötz. Generating virtual on-body accelerometer data from virtual textual descriptions for human activity recognition. *arXiv preprint arXiv:2305.03187*, 2023.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- Jun Li, Che Liu, Sibó Cheng, Rossella Arcucci, and Shenda Hong. Frozen language model helps ecg zero-shot learning. *arXiv preprint arXiv:2303.12311*, 2023.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- Che Liu, Zhongwei Wan, Sibó Cheng, Mi Zhang, and Rossella Arcucci. Etp: Learning transferable ecg representations via ecg-text pre-training. *arXiv preprint arXiv:2309.07145*, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- Shansong Liu, Atin Sakkeer Hussain, Chenshuo Sun, and Ying Shan. Music understanding llama: Advancing text-to-music generation with question answering and captioning. *arXiv preprint arXiv:2308.11276*, 2023.

- Xin Liu, Daniel McDuff, Geza Kovacs, Isaac Galatzer-Levy, Jacob Sunshine, Jiening Zhan, Ming-Zher Poh, Shun Liao, Paolo Di Achille, and Shwetak Patel. Large language models are few-shot health learners. *arXiv preprint arXiv:2305.15525*, 2023.
- Xu Liu, Junfeng Hu, Yuan Li, Shizhe Diao, Yuxuan Liang, Bryan Hooi, and Roger Zimmermann. Unitime: A language-empowered unified model for cross-domain time series forecasting. *arXiv preprint arXiv:2310.09751*, 2023.
- Chenxi Liu, Sun Yang, Qianxiong Xu, Zhishuai Li, Cheng Long, Ziyue Li, and Rui Zhao. Spatial-temporal large language model for traffic prediction. *arXiv preprint arXiv:2401.10134*, 2024.
- Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Motlaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *arXiv preprint arXiv:2206.08916*, 2022.
- Qianli Ma, Zhen Liu, Zhenjing Zheng, Ziyang Huang, Siying Zhu, Zhongzhong Yu, and James T Kwok. A survey on time-series pre-trained models. *arXiv preprint arXiv:2305.10716*, 2023.
- Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. Large language models as general pattern machines. *arXiv preprint arXiv:2307.04721*, 2023.
- Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Alireza Dirafzoon, Aparajita Saraf, Amy Bearman, and Babak Damavandi. Imu2clip: Multimodal contrastive learning for imu motion sensors from egocentric videos and text. *arXiv preprint arXiv:2210.14395*, 2022.
- Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Tushar Nagarajan, Matt Smith, Shashank Jain, Chun-Fu Yeh, Prakash Murugesan, Peyman Heidari, Yue Liu, et al. Any-mal: An efficient and scalable any-modality augmented language model. *arXiv preprint arXiv:2309.16058*, 2023.
- Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.
- Jungwoo Oh, Seongsu Bae, Gyubok Lee, Joon-myung Kwon, and Edward Choi. Ecg-qa: A comprehensive question answering dataset combined with electrocardiogram. *arXiv preprint arXiv:2306.15681*, 2023.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023.
- Jielin Qiu, William Han, Jiacheng Zhu, Mengdi Xu, Michael Rosenberg, Emerson Liu, Douglas Weber, and Ding Zhao. Transfer knowledge from natural language to electrocardiography: Can we detect cardiovascular disease through language models? *arXiv preprint arXiv:2301.09017*, 2023.
- Stephan Rabanser, Tim Januschowski, Valentin Flunkert, David Salinas, and Jan Gasthaus. The effectiveness of discretization in forecasting: An empirical study on neural time series models. *arXiv preprint arXiv:2005.10111*, 2020.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Arian Khorasani, George Adamopoulos, Rishika Bhagwatkar, Marin Bilos, Hena Ghonia, Nadhir Vincent Hassen, Anderson Schneider, et al. Lag-llama: Towards foundation models for time series forecasting. *arXiv preprint arXiv:2310.08278*, 2023.
- Paul K Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, et al. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*, 2023.
- Shezheng Song, Xiaopeng Li, and Shasha Li. How to bridge the gap between modalities: A comprehensive survey on multimodal large language model. *arXiv preprint arXiv:2311.07594*, 2023.
- Dimitris Spathis and Fahim Kawsar. The first step is the hardest: Pitfalls of representing and tokenizing temporal data for large language models. *arXiv preprint arXiv:2309.06236*, 2023.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.
- Chenxi Sun, Yaliang Li, Hongyan Li, and Shenda Hong. Test: Text prototype aligned embedding to activate llm’s ability for time series. *arXiv preprint arXiv:2308.08241*, 2023.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. Salmonn: Towards generic hearing abilities for large language models. *arXiv preprint arXiv:2310.13289*, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al.

- Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Patrick Wagner, Nils Strodthoff, Ralf-Dieter Boussejot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):154, 2020.
- Zhenhailong Wang and Heng Ji. Open vocabulary electroencephalography-to-text decoding and zero-shot sentiment classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5350–5358, 2022.
- Tianrui Wang, Long Zhou, Ziqiang Zhang, Yu Wu, Shujie Liu, Yashesh Gaur, Zhuo Chen, Jinyu Li, and Furu Wei. Viola: Unified codec language models for speech recognition, synthesis, and translation. *arXiv preprint arXiv:2305.16107*, 2023.
- Xinglei Wang, Meng Fang, Zichao Zeng, and Tao Cheng. Where would i go next? large language models as human mobility predictors. *arXiv preprint arXiv:2308.15197*, 2023.
- Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey. *arXiv preprint arXiv:2202.07125*, 2022.
- Christopher Wimmer and Navid Rekasaz. Leveraging vision-language models for granular market change prediction. *arXiv preprint arXiv:2301.10166*, 2023.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabrovolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564*, 2023.
- Qianqian Xie, Weiguang Han, Yanzhao Lai, Min Peng, and Jimin Huang. The wall street neophyte: A zero-shot analysis of chatgpt over multimodal stock movement prediction challenges. *arXiv preprint arXiv:2304.05351*, 2023.
- Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. Pixiu: A large language model, instruction data and evaluation benchmark for finance. *arXiv preprint arXiv:2306.05443*, 2023.
- Tianwei Xing, Luis Garcia, Federico Cerutti, Lance Kaplan, Alun Preece, and Mani Srivastava. Deepsqa: Understanding sensor data via question answering. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, pages 106–118, 2021.
- Hao Xue and Flora D Salim. Promptcast: A new prompt-based learning paradigm for time series forecasting. 2022.
- Hao Xue and Flora D Salim. Utilizing language models for energy load forecasting. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 224–227, 2023.
- Hao Xue, Bhanu Prakash Voutharoja, and Flora D Salim. Leveraging language foundation models for human mobility forecasting. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pages 1–9, 2022.
- Dongchao Yang, Jinchuan Tian, Xu Tan, Rongjie Huang, Songxiang Liu, Xuankai Chang, Jiatong Shi, Sheng Zhao, Jiang Bian, Xixin Wu, et al. Uniaudio: An audio foundation model toward universal audio generation. *arXiv preprint arXiv:2310.00704*, 2023.
- Xinli Yu, Zheng Chen, Yuan Ling, Shujing Dong, Zongyi Liu, and Yanbin Lu. Temporal data meets llm-explainable financial time series forecasting. *arXiv preprint arXiv:2306.11025*, 2023.
- Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank J Reddi, and Sanjiv Kumar. Are transformers universal approximators of sequence-to-sequence functions? *arXiv preprint arXiv:1912.10077*, 2019.
- Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:495–507, 2021.
- Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities. *arXiv preprint arXiv:2305.11000*, 2023.
- Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. *arXiv preprint arXiv:2301.06052*, 2023.
- Jiayun Zhang, Xiyuan Zhang, Xinyang Zhang, Dezhi Hong, Rajesh K. Gupta, and Jingbo Shang. Navigating Alignment for Non-identical Client Class Sets: A Label Name-Anchored Federated Learning Framework. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, aug 2023.
- Xiyuan Zhang, Ranak Roy Chowdhury, Jiayun Zhang, Dezhi Hong, Rajesh K. Gupta, and Jingbo Shang. Unleashing the power of shared label structures for human activity recognition. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 3340–3350, 2023.
- Yunkai Zhang, Yawen Zhang, Ming Zheng, Kezhen Chen, Chongyang Gao, Ruian Ge, Siyuan Teng, Amine Jelloul, Jinqiang Rao, Xiaoyuan Guo, et al. Insight miner: A time series analysis dataset for cross-domain alignment with natural language. In *NeurIPS 2023 AI for Science Workshop*, 2023.
- Zheng Zhang, Hossein Amiri, Zhenke Liu, Andreas Züfle, and Liang Zhao. Large language models for spatial trajectory patterns mining. *arXiv preprint arXiv:2310.04942*, 2023.

Ziyuan Zhong, Davis Rempe, Yuxiao Chen, Boris Ivanovic, Yulong Cao, Danfei Xu, Marco Pavone, and Baishakhi Ray. Language-guided traffic simulation via scene-level diffusion. *arXiv preprint arXiv:2306.06344*, 2023.

Tian Zhou, Peisong Niu, Xue Wang, Liang Sun, and Rong Jin. One fits all: Power general time series analysis by pretrained lm. *arXiv preprint arXiv:2302.11939*, 2023.

Yunjiao Zhou, Jianfei Yang, Han Zou, and Lihua Xie. Tent: Connect language models with iot sensors for zero-shot activity recognition. *arXiv preprint arXiv:2311.08245*, 2023.