

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TS-REASONER: ALIGNING TIME SERIES FOUNDATION MODELS WITH LLM REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Time series reasoning is crucial to decision-making in diverse domains, including finance, energy usage, traffic, weather, and scientific discovery. While existing time series foundation models (TSFMs) can capture low-level dynamic patterns and provide accurate forecasting, further analysis usually requires additional background knowledge and sophisticated reasoning, which are lacking in most TSFMs but can be achieved through large language models (LLMs). On the other hand, without expensive post-training, LLMs often struggle with the numerical understanding of time series data. Although it is intuitive to integrate the two types of models, developing effective training recipes that align the two modalities for reasoning tasks is still an open challenge. To this end, we propose TS-REASONER that aligns the latent representations of TSFMs with the textual inputs of LLMs for downstream understanding/reasoning tasks. Specifically, we propose a simple yet effective method to curate diverse, synthetic pairs of time series and textual captions for alignment training. We then develop a two-stage training recipe that applies instruction finetuning after the alignment pretraining. Unlike existing works that train an LLM to take time series as inputs, we leverage a pretrained TSFM and freeze it during training. Extensive experiments on several benchmarks demonstrate that TS-REASONER not only outperforms a wide range of prevailing LLMs, Vision Language Models (VLMs), and Time Series LLMs, but also achieves this with remarkable data efficiency, e.g., using less than half the training data.

1 INTRODUCTION

Time series analysis has long been fundamental to various real-world applications in finance, energy, weather, traffic, and other domains (Prakarsha & Sharma, 2022; Xu et al., 2023; Nie et al., 2024). Its ability to model dynamics and predict future states based on historical data makes it an indispensable tool for informed decision-making and strategic planning. While numerical attributes form the bedrock of time series analysis, human decision-making is often complemented by rich prior knowledge and qualitative contextual information, including news articles, social media trends, and expert assessments. This gap prevents analytical models from achieving a deeper, more contextualized understanding of the events and dynamics driving the numerical data. By enabling machines to understand both contextual information and numerical time series patterns, we can empower them as automated systems that assist humans in gaining deeper insights into complex phenomena.

Recent advances in Time Series Foundation Models (TSFMs) have significantly enhanced the understanding of time series data through large-scale pretraining. These models are capable of generalizing across a wide variety of time series tasks and domains. Although TSFMs (Goswami et al., 2024; Das et al., 2024) demonstrate strong modeling capabilities, most are pre-trained exclusively on unimodal numerical time series and cannot therefore comprehend or integrate textual information. On the other hand, large Language Models (LLMs) and Vision Language Models (VLMs) can take texts and images as input context, and have demonstrated remarkable reasoning and problem-solving abilities across various tasks (Wei et al., 2022; Yao et al., 2023; Hao et al., 2023; Yu et al., 2024), sparking interest in transferring their capabilities to time series analysis. Some studies (Gruver et al., 2023; Liu et al., 2024c; Jia et al., 2024) transform numerical time series into string form and perform time series forecasting on LLMs by prompting them with the strings. However, despite their strong reasoning abilities, LLMs struggle to capture temporal dependencies due to their inherent lack of temporal understanding (Fons et al., 2024; Merrill et al.,

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107
2024) and limited ability to interpret numerical values. These limitations hinder their understanding of time series data. As shown in the figure 1, TSFM and LLM have complementary strengths; the former specializes in temporal understanding, while the latter excels at text understanding.

To combine the complementary strengths of TSFMs and LLMs while overcoming their respective limitations, we propose TS-REASONER, a Time Series Large Language Model (TSLLM) designed to enhance time series reasoning by aligning a TSFM with an LLM. Specifically, we first employ the TSFM to extract rich temporal representations from numerical time series data. To effectively incorporate this temporal information into the LLM, TS-REASONER introduces a TS-to-Text adapter, which projects the TSFM-extracted temporal features into the LLM’s input embedding space. This enables seamless integration of the TSFM’s temporal understanding with the LLM’s powerful linguistic and reasoning capabilities. Our training framework consists of two stages: pretraining and fine-tuning. In the pretraining stage, we finetune TS-REASONER to produce textual captions of input time series and achieve a fundamental alignment. To this end, we propose a simple yet effective prompting strategy to curate high-quality captions for diverse time series data using advanced LLMs/VLMs. In the fine-tuning stage, we further enhance the model’s reasoning abilities through instruction tuning, ensuring robust performance in downstream tasks.

Our work makes unique contributions to a recent line of research combining TSFMs and LLMs. First, our formulation sets up the connection between LLMs and TSFMs, facilitating time series reasoning through the integration of rich contextual information and LLM reasoning. Second, we address a critical data bottleneck by a simple yet effective time series captioning method, which diversifies the training data for aligning LLMs and TSFMs. Finally, we offer new empirical insights into the strengths and limitations of existing approaches.

We evaluate the understanding and reasoning capabilities of our approach on two standard benchmarks: TimeSeriesExam (Cai et al., 2024a) and MTBench (Chen et al., 2025). TS-REASONER significantly outperforms a wide range of baseline models, including LLMs, VLMs, and the TSLLMs, as shown in Figure 2. Finally, comprehensive analyses, including extensive ablation studies, validate the effectiveness of our key designs and establish the superiority of TS-REASONER in generalization performance, scalability, and training data efficiency.

2 RELATED WORK

LLMs for Time Series. LLMs have recently garnered significant interest in time series analysis. Traditional time series forecasting relies on statistical models (RB, 1990) or data-driven neural networks (Liu et al., 2021; Lim et al., 2021; Wu et al., 2021; Zhou et al., 2022; Li et al., 2023b; 2024b) for tasks like weather and stock prediction. Recent efforts explore LLMs for this task, with some designing prompts to elicit forecasting abilities (Cao et al., 2023; Chuang* et al., 2024). Others focus on enabling LLMs to understand time series data by converting it into textual sequences or aligning its embeddings with language model embeddings via prompting or semantic information (Jin et al., 2023; Sun et al., 2023; Pan et al., 2024). In addition, multimodal vision-based LLMs are being



Figure 1: Time series forecasting vs. reasoning. The time series reasoning task requires both contextual reasoning (e.g., news) by LLMs and numerical understanding by TSFM.

projects the TSFM-extracted temporal features into the LLM’s input embedding space. This enables seamless integration of the TSFM’s temporal understanding with the LLM’s powerful linguistic and reasoning capabilities. Our training framework consists of two stages: pretraining and fine-tuning. In the pretraining stage, we finetune TS-REASONER to produce textual captions of input time series and achieve a fundamental alignment. To this end, we propose a simple yet effective prompting strategy to curate high-quality captions for diverse time series data using advanced LLMs/VLMs. In the fine-tuning stage, we further enhance the model’s reasoning abilities through instruction tuning, ensuring robust performance in downstream tasks.




Figure 2: Results on time series understanding and reasoning benchmarks. TS-REASONER demonstrates a consistent advantage over the prevailing LLMs, VLMs, and TSLLMs.

108 investigated for time series prediction (Chen et al., 2024c; Zhong et al., 2025). Though LLMs exhibit
 109 non-trivial performance on some forecasting tasks, Merrill et al. (Merrill et al., 2024) indicate that
 110 LLMs struggle to reason time series. To tackle this challenge, several works (Chow et al., 2024;
 111 Zhang et al., 2025a; Xie et al., 2024) propose to enable LLMs to understand the time series with
 112 context. TS-REASONER lies in this direction, distinguishing itself by employing a pre-trained Time
 113 Series Foundation Model to ground the LLM’s reasoning in robust temporal features.

114 **Modality Alignment.** Modality alignment methods are widely studied in the multimodal domain (Li
 115 et al., 2022; Lai et al., 2024; Li et al., 2023a; Liu et al., 2024b). Inspired by the success of multimodal
 116 alignment, recent works treat time series as another modality and align it to the LLM (Xie et al., 2024;
 117 Zhang et al., 2025a). Though they achieve a certain degree of time series understanding, they focus
 118 on narrow domains (e.g., electricity) and tasks (e.g., time series understanding), and train time series
 119 encoders from scratch. In contrast, we adapt the successful training paradigm in VLMs, identify
 120 and address the key challenges (e.g., Integration of characteristics of time series into LLMs, and the
 121 shortage of time series-text pairs) faced in applying this paradigm to the unique modality of time
 122 series, exploring pre-trained time series foundation models to exploit rich time series knowledge.

123 **Time Series Foundation Models.** Recent advancements in pre-training methods are significantly
 124 contributing to the development of foundation models for time series analysis. Early efforts, such as
 125 TST (Zerveas et al., 2021) and PatchTST (Nie et al., 2022), applied BERT-like masked pretraining
 126 techniques, focusing on point-level and patch-level masking, respectively. A separate line of work,
 127 exemplified by models like TimesFM (Das et al., 2024), Timer (Liu et al., 2024e), TTMS (Ekambaram
 128 et al., 2024), Chronos (Ansari et al., 2024), and Time-MoE (Shi et al., 2024), Moirai (Liu et al.,
 129 2024d), TimesBERT (Zhang et al., 2025b), and Sundial (Liu et al., 2025) demonstrates the advantages
 130 of large-scale pre-training for improving forecasting performance. Exploring diverse pre-training
 131 objectives, MOMENT (Goswami et al., 2024) leverages a T5 encoder to achieve strong downstream
 132 multi-task capabilities. ChronoSteer (Wang et al., 2025a) also explores the alignment between TSFMs
 133 and LLMs, yet it leverages the LLM’s revisions to enhance TSMF’s forecasting capability.

134 3 TS-REASONER FOR TEMPORAL REASONING

135 As illustrated in Figure 3, TS-REASONER is composed of (1) a pretrained TSFM that encodes
 136 normalized, non-overlapping patches of input time series into compact embeddings; (2) a pretrained
 137 LLM, and (3) a TS-to-Text adapter that projects the TSFM’s output embedding to the input space of
 138 the LLM. The LLM concatenates the sequence of projected time series features with the sequence
 139 of embeddings for input text tokens, with the former demarcated by special tokens “⟨ts⟩⟨ts⟩”. The
 140 training of TS-REASONER consists of two stages: (1) a pretraining stage to align time series features
 141 from the TSFM with the LLM, using time series caption data synthesized by an advanced LLM/VLM,
 142 and (2) an instruction tuning stage to enhance complex reasoning capabilities on downstream tasks.

143 3.1 MODEL ARCHITECTURE

144 Given a natural language context \mathcal{X} and a corresponding set of time series $\mathcal{S} = \{\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_K\}$, we
 145 first project both into a shared embedding space. Specifically, for each time series $\mathcal{T}_i \in \mathbb{R}^{L_i}$, where
 146 L_i is the length of the series, we first apply instance normalization to standardize its distribution
 147 to zero mean and unit variance. This preprocessing step ensures that the model is robust to shifts
 148 and scales in the input data. Subsequently, we partition the normalized time series into a sequence
 149 of non-overlapping patches, each of a fixed length P . This patching strategy yields a sequence of
 150 $N_i = \lfloor L_i/P \rfloor$ patches, transforming the time series into a tensor $\mathcal{T}_i^P \in \mathbb{R}^{N_i \times P}$. These patches are
 151 then encoded using the TSFM, which acts as our time series feature extractor. The TSFM processes
 152 the sequence of patches and produces a sequence of embedding vectors:

$$153 \mathcal{Z}_i^T = \text{TSFM}(\mathcal{T}_i^P) \in \mathbb{R}^{N_i \times d_{ts}}, \quad (1)$$

154 where d_{ts} denotes the dimension of the time series embeddings. Concurrently, the natural language
 155 context \mathcal{X} is tokenized and fed into the pre-trained LLM’s embedding layer. This process converts
 156 the textual input into a sequence of contextualized token embeddings:

$$157 \mathcal{Z}^L = \text{LLM}_{\text{embed}}(\mathcal{X}) \in \mathbb{R}^{M \times d_{\text{text}}}, \quad (2)$$

158 where M is the number of tokens in the instruction, and d_{text} is the dimensionality of the LLM’s
 159 hidden states. To align the dimension and semantics of embeddings between LLM and TSFM, we




Figure 3: Overview of TS-REASONER architecture and training pipeline. To perform reasoning, a time series is first encoded by a pretrained Time Series Foundation Model (TSFM). Its output features are then projected into the LLM’s input embedding space by a trainable TS-to-Text Adapter and subsequently processed by the LLM. The model is trained in two stages: (1) a pretraining stage that aligns the TSFM outputs with the LLM inputs using both template-based (code-synthesized) and LLM-generated captions, as described in §3.2, and (2) an instruction-tuning stage to improve complex reasoning capabilities.

use a multilayer perceptron (MLP) as a TS-to-Text Adapter to transform the time series embedding into the text embedding space:

$$\mathcal{H}_i^T = \text{MLP}(Z_i^T) \in \mathbb{R}^{N_i \times d_{\text{text}}}, \quad (3)$$

To form a unified input sequence for the LLM that accommodates multiple time series, we structure the natural language instruction \mathcal{X} to include K indicators, $\{K \cdot \langle \text{ts} \rangle \langle \text{ts} \rangle\}$. The i -th placeholder $\langle \text{ts} \rangle \langle \text{ts} \rangle$ marks the insertion point for the corresponding i -th time series \mathcal{T}_i .

Let $\{\mathcal{H}_i^T \in \mathbb{R}^{N_i \times d_{\text{text}}}\}_{i=1}^K$ be the set of projected time series embeddings. The final input sequence H is constructed by sequentially inserting the embedding to each $\langle \text{ts} \rangle \langle \text{ts} \rangle$ with its corresponding time series embedding sequence \mathcal{H}_i^T . This substitution process results in a composite sequence where language and time series representations are interleaved. The total length of this fused sequence is $M + \sum_{i=1}^K N_i$. The final tensor fed to the LLM’s transformer layers is therefore: $\mathcal{H} \in \mathbb{R}^{(M + \sum_{i=1}^K N_i) \times d_{\text{text}}}$. This strategy enables the LLM to process multiple, arbitrarily placed time series within a single, coherent context and capture complex inter-series and text-series dependencies. After the combination, the input embedding H is fed to LLM to produce the final prediction \mathcal{Y} .

3.2 TRAINING RECIPE

Our training process consists of two sequential stages: the first stage aligns time series data with the LLM to establish a foundational understanding of temporal-textual relationships, while the second stage refines the LLM’s reasoning capabilities to interpret and analyze these aligned representations. Throughout both stages, we keep the parameters of the TSFM frozen to preserve its pretrained temporal knowledge, while allowing the LLM’s parameters to remain trainable, ensuring adaptive learning without compromising the integrity of the encoded time series features.




Figure 4: Workflow for our attribute-aware caption synthesis, designed to curate training data for alignment in stage 1. It enriches basic instructions with key attributes and generates diverse paraphrases, yielding the high-fidelity captions to train TS-REASONER effectively.

Stage 1: Pre-training for Language-Timeseries Alignment. In this stage, our primary objective is to align temporal data with textual information. We initially leverage synthesized data from (Xie et al., 2024), which provides predefined templates to describe time series attributes. However, while this template-based data offers accurate numerical information, its focus on specific time series patterns limits diversity, and the caption structure is monotonous. This lack of diversity can lead to overfitting to the templates, encouraging the model to learn shallow patterns and resulting in poor generalization ability (Dong et al., 2025; Choi et al., 2024). To alleviate this problem, we draw inspiration from captioning techniques in multimodal LLMs (Chen et al., 2024a). We synthesize comprehensive captions using advanced LLMs (e.g., GPT-4.1) to enrich our alignment data. Specifically, we collect time series from two sources: (Merrill et al., 2024), which includes contextual information, and synthetic data from Chronos (Ansari et al., 2024), which provides pure numerical time series.

Attribute-aware Captioning. Caption generation has been extensively investigated in visual domains (Cheng et al., 2023; 2025; Chen et al., 2024b), playing a crucial role in multimodal alignment. However, time series captioning remains largely underexplored, presenting a significant impediment to achieving comprehensive alignment. To address this gap, we introduce a straightforward approach for generating scalable time series captions, as shown in Figure 4.

Given a time series \mathcal{T} with a temporal context \mathcal{C} , we begin by defining a fundamental captioning instruction, denoted as $\mathcal{I}_{\text{base}}$. To facilitate enhanced comprehension by LLMs, we transform the time series into an image plot via Python code, $I_{TS} = \Phi(\mathcal{T})$. As evidenced in Table 1 (Section 4.1), presenting the time series as an image to advanced LLMs (e.g., GPT-4.1) demonstrates a substantial advantage in understanding compared to providing it as a raw numerical series.

To enrich the generated captions, we first identify a set of G pertinent attributes of the time series, denoted as $\{a_1, a_2, \dots, a_G\}$ (e.g., trend, frequency, periodicity, noise, local variations). These attributes are then incorporated into the basic instruction, yielding an augmented instruction $\mathcal{I}' = \mathcal{I}_{\text{base}} \cup \{a_1, a_2, \dots, a_G\}$. To further promote caption diversity, we leverage the LLM to paraphrase \mathcal{I}' into R distinct instructions, forming a candidate set of prompts $\mathcal{P} = \{\mathcal{I}'_1, \mathcal{I}'_2, \dots, \mathcal{I}'_R\}$. For each time series \mathcal{T} , a single prompt \mathcal{I}'' is uniformly sampled from this set. The final caption is then generated conditioned on the sampled prompt and the time series visualization:

$$\text{Caption} = \text{LLM}(\mathbf{c} | \mathcal{I}'', I_{TS}), \quad (4)$$

where $\mathcal{I}'' \sim \mathcal{U}(\mathcal{P})$. The prompts are shown in the Figure 10 in Appendix E. We randomly sample 10K time series from each of two distinct sources: the Chronos synthetic dataset (Ansari et al., 2024), which contains purely numerical time series, and a dataset of text-attributed time series from Merrill et al. (Merrill et al., 2024), which provides contextual backgrounds. The construction of data offers two benefits: (1) Pure time series data enables the model to build a foundational understanding of temporal patterns by focusing solely on the intrinsic characteristics of the data. (2) Context-augmented time series enhances domain-specific comprehension by linking numerical trends to real-world scenarios, thereby improving the model’s ability to generalize across diverse applications.

Stage 2: Instruction Finetuning for Time Series Reasoning. To elevate the model’s capabilities from foundational understanding to complex reasoning, we employ an instruction fine-tuning stage

270 based on the instruction tuning dataset (Xie et al., 2024), which encompasses a wide range of Q&As
 271 and instruction-following tasks. This training facilitates TS-REASONER with two critical abilities:
 272 the fidelity to adhere to complex instructions and structured response formats, and the capacity for
 273 nuanced, context-driven reasoning on time series-specific queries.

276 4 EXPERIMENTS

277 **Datasets** To assess the capabilities of TS-REASONER, we conduct comparative experiments against
 278 various baselines on benchmarks tailored for time series reasoning. Our evaluation incorporates
 279 TimeSeriesExam (Cai et al., 2024a), a comprehensive multiple-choice question answering dataset.
 280 TimeSeriesExam is specifically engineered to systematically evaluate a model’s time series under-
 281 standing and reasoning abilities across several key aspects: Pattern Recognition (PR), which addresses
 282 identifying trends, cycles, and stationarity; Noise Understanding (NU), focused on recognizing noise
 283 types such as white noise and random walks; Anomaly Detection (AD), for detecting unusual patterns;
 284 Similarity Analysis (SA), which involves comparing the shape and distribution of two time series;
 285 and Causality Analysis (CA), assessing the recognition of Granger Causality between time series.
 286 Furthermore, we evaluate on MTBench (Chen et al., 2025), a large-scale benchmark for evaluating
 287 time series reasoning in the real-world financial and weather domains, featuring questions that span
 288 both short-term (7-day) and long-term (14-day) temporal horizons.

289 **Baselines and Evaluation Metrics** We compare our method against three types of baselines:
 290 closed-source LLMs / VLMs, open-source LLMs / VLMs, and TSLLMs. Specifically, for closed-
 291 source models, we include GPT-4o, GPT-4.1 (gpt, 2024; Achiam et al., 2023), Claude-Sonnet-
 292 3.7 (The), and DeepSeek-Chat (Liu et al., 2024a). For open source LLM, we evaluated LLama-3.1-8B-
 293 Instruct (Grattafiori et al., 2024), Qwen-2.5-7B-Instruct (Yang et al., 2024), GLM-4-9B-Chat (GLM
 294 et al., 2024), InternLM3-8B-Instruct (Cai et al., 2024b), and Ministrail-8B-Instruct (Jiang et al., 2024).
 295 Time series are transformed into textual sequences of numbers for LLMs. For open-source VLM
 296 models, we compare Qwen2.5-VL-7B (Bai et al., 2025), Phi-4-Multimodal-Instruct (Abouelenin
 297 et al., 2025), Llama3-LLaVA-Next-8B (Li et al., 2024a), InternVL3-8B (Zhu et al., 2025), and
 298 MinICPM-V-2.6 (Yao et al., 2024). Time series are transformed into plots via code for VLMs. For
 299 TSLLMs models, we compare with ChatTime-7B (Wang et al., 2025b), ChatTS-14B (Xie et al.,
 300 2024), and we use the official training data and code to fine-tune a 7B model for a fair comparison.
 301 As all benchmarks are multiple-choice Q&As, we use accuracy as the evaluation metric. See more
 302 implementation details in Appendix A.

304 4.1 MAIN RESULTS


305 Table 1 presents the performance of all models on the two benchmarks. The best results are bolded, and
 306 the second-best results are underlined. Based on the results, we have the following key observations:

307 **(i) TS-REASONER achieves superior overall performance on all benchmarks among models of
 308 the same size.** Specifically, TS-REASONER demonstrates superior performance, surpassing the best-
 309 performing LLM by 7.60% overall, the best VLM by 5.25% overall, and the TSLLM by 3.54% overall
 310 on the TimeSeriesExam benchmark. Compared to the backbone model, TS-REASONER improves
 311 on our backbone LLM performance by a substantial 16.29%. TS-REASONER also excels the best
 312 baseline on MTBench by around 2%. In addition, TS-REASONER performs even competently with
 313 ChatTS-14B, which has a larger base model. The notable improvement demonstrates the effectiveness
 314 of our model in various time series reasoning scenarios by introducing the temporal information of
 315 TSFM for the LLM.

316 **(ii) TS-REASONER delivers consistent gains in most time series understanding and reasoning
 317 subtasks.** In particular, it surpasses the second best baseline with absolute improvements of 2.16%
 318 in *Pattern Recognition*, 8.05% in *Noise Understanding*, 2.33% in *Anomaly Detection*, and 2.65%
 319 in *Similarity Analysis*, while also yielding around 2% improvements on both financial and weather
 320 reasoning tasks. These gains stem from two key factors: (1) aligning time series with text during
 321 training substantially strengthens TS-REASONER’ understanding of time series patterns; and (2)
 322 this improved understanding further enhances its ability to reason over time series in context when
 323 combined with textual information.

324
325
326
327 Table 1: Performance of LLMs, VLMs, TSLLMs, and proprietary models on time series understanding
328 and reasoning benchmarks. Our baselines also include ChatTS-14B, which uses a larger base model.
329

Model	TimeSeriesExam (Cai et al., 2024a)						MTBench (Chen et al., 2025)			
	PR	NU	AD	SA	CA	OA	Finance (long)	Finance (short)	Weather (long)	Weather (short)
<i>Proprietary models</i>										
DeepSeek-Chat	65.23	55.17	52.71	63.71	42.86	59.89	89.15	90.02	59.75	58.76
Claude-Sonnet-3.7	62.26	55.17	48.06	72.57	50.79	59.63	84.11	88.56	51.24	47.91
GPT-4o	59.03	55.17	53.49	62.83	31.75	55.96	84.30	82.69	48.07	48.22
GPT-4o (vision)	67.12	62.07	62.79	64.60	26.98	62.12	84.11	80.65	46.43	48.53
GPT-4.1 (vision)	69.81	68.97	68.22	75.22	41.27	67.89	93.41	91.45	56.04	55.35
<i>Open-source Large Language Models</i>										
Llama-3.1-8B-Instruct	37.73	37.93	30.23	36.28	28.57	35.52	63.37	35.52	40.25	40.00
Qwen2.5-7B-Instruct	47.17	47.13	41.86	53.10	41.27	46.66	87.98	89.41	57.14	58.44
GLM-4-9B-chat	41.78	39.08	37.21	47.79	38.09	41.28	71.31	77.19	50.27	50.85
InternLM3-8B-Instruct	43.93	51.72	26.35	52.21	34.92	42.33	71.70	71.08	45.05	46.67
Ministral-8B-Instruct	43.13	37.93	39.53	44.25	36.51	41.55	46.32	50.71	39.15	40.93
<i>Open-source Vision Language Models</i>										
Qwen2.5-VL-7B-Instruct	25.34	32.18	19.38	42.48	12.70	26.61	81.98	86.35	52.06	46.82
Phi-4-Multimodal-Instruct	36.39	34.48	30.23	38.94	14.28	33.68	70.35	74.54	48.35	49.77
Llama3-LLaVA-Next-8B	31.27	35.63	29.46	30.09	38.09	31.85	52.14	51.50	47.53	47.29
InternVL3-8B	50.13	52.87	43.41	54.87	38.09	49.01	88.95	90.00	53.71	54.88
MinicPM-V-2.6	29.11	39.08	27.13	51.33	31.75	33.42	81.78	83.09	48.63	45.12
<i>Time Series Large Language Models</i>										
ChatTime-7B	42.85	49.42	35.65	44.24	34.92	41.94	25.97	28.10	47.80	42.79
ChatTS-7B	50.13	50.57	50.38	61.95	34.92	50.72	87.60	88.01	58.92	58.75
ChatTS-14B*	59.30	54.02	51.16	62.83	41.27	56.36	89.22	91.22	59.61	59.22
TS-REASONER-7B (ours)	52.29	60.92	52.71	64.60	41.27	54.26	90.12	91.85	60.44	61.24
Δ Over Best 7B	+2.16	+8.05	+2.33	+2.65	+0.00	+3.54	+1.17	+1.85	+1.52	+2.49



346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
Figure 5: Data scaling and efficiency of TS-REASONER. The top (bottom) row illustrates how the
378 performance of TS-REASONER varies when increasing the training data for alignment (instruction
379 tuning). The columns correspond to sub-tasks in TimeSeriesExam. ChatTS-7B (Xie et al., 2024) is
380 included for reference, denoted by the gray triangle.
381

4.2 ANALYSIS OF DATA SCALING AND EFFICIENCY

382 Figure 5 presents our data scaling analysis on the TimeSeriesExam benchmark. TS-REASONER
383 demonstrates remarkable data efficiency compared to the ChatTS-7B baseline. For the alignment
384 stage, TS-REASONER achieves superior overall accuracy using just 60K samples, less than half
385 the data required by the baseline. This efficiency is even more stark in the instruction tuning stage,
386 where 10K samples suffice to outperform ChatTS-7B. This significant reduction in data dependency
387 stems from our pre-trained TSFM and effective alignment, which equips the LLM with a robust
388 temporal foundation. Consequently, TS-REASONER develops advanced reasoning capabilities with
389 a substantially smaller amount of data, marking a key advantage for practical deployments where
390 labeled data is scarce.

4.3 CHOICES OF CAPTIONING MODEL FOR ALIGNMENT

391 The quality of the generated captions is a critical factor in the efficacy of our time-series-language
392 alignment. To validate this, we conducted an experiment where we trained TS-REASONER using
393 three distinct sets of captioning data, each generated by a model with varying capabilities:




Figure 6: Performance of TS-REASONER and their associated LLM backbones (Qwen2.5 series). The top row and bottom row report the performance on TimeSeriesExam and MTBench, respectively.

the state-of-the-art GPT-4.1, and two VLMs, InternVL3-8B and Qwen2.5-VL-7B-Instruct. As illustrated in Figure 7, the results demonstrate that the performance of TS-REASONER is directly correlated with the fidelity of the captioning model. A distinct performance hierarchy emerges across both benchmarks: the model trained on GPT-4.1 captions consistently outperforms the one trained on InternVL3-8B captions, which in turn surpasses the one trained on Qwen2.5-VL-7B-Instruct captions. The higher performance gain from GPT-4.1 is attributed to its advanced capability in time series understanding. It is not surprising that the captions generated by InternVL3-8B achieve higher performance than Qwen2.5-VL-7B-Instruct, as its better time series understanding capability is shown in Table 1.

4.4 CHOICES OF TSFM AND LLM IN TS-REASONER

Different choices of TSFMs. To investigate the performance of TS-REASONER with different TSFMs, we replaced TimesFM (200M) with MOMENT-1-base (200M), a TSFM of the same size, and re-evaluated its performance on the TimeSeriesExam benchmark. The results presented in Table 2 reveal a substantial performance degradation when using MOMENT, with overall accuracy falling from 54.26% to 45.74%. The result is expected because TimesFM outperforms Moment on various time series forecasting benchmarks (Shi et al., 2024; Mulayim et al., 2024). This suggests that TimesFM provides better time series representations, enabling TS-REASONER to better understand and reason for the time series.

Different choices of LLMs. To investigate the scalability and robustness of our approach with different LLM backbones, we evaluate TS-REASONER against across three distinct sizes of the




Figure 7: Comparison of multimodal LLMs used to generate time series captions for training TS-REASONER. **Left:** performance on TimeSeriesExam. **Right:** Performance on MTBench.

Table 2: Comparison of TS-REASONER using different TSFMs on the TimeSeriesExam benchmark.

Model	PR	NU	AD	SA	CA	OA
MOMENT	46.90	47.13	41.86	54.87	28.57	45.74
TimesFM	52.29	60.92	52.71	64.60	41.27	54.26

432 Table 3: Ablation study results of different components in TS-REASONER.
433

434 Model	435 TimeSeriesExam (Cai et al., 2024a)						436 MTBench (Chen et al., 2025)			
	437 PR	438 NU	439 AD	440 SA	441 CA	442 OA	443 Finance (long)	444 Finance (short)	445 Weather (long)	446 Weather (short)
437 TS-REASONER-7B (ours)	438 52.29	439 60.92	440 52.71	441 64.60	442 41.27	443 54.26	444 90.12	445 91.85	446 60.44	447 61.24
<i>Ablation on Training Data</i>										
- LLM-caption	51.21	56.32	52.71	56.54	36.51	51.25	88.67	89.40	58.24	59.69
- Attributes	52.02	57.47	48.83	62.83	39.68	52.69	89.71	89.20	57.28	59.07
<i>Ablation on Training Stages</i>										
- Stage 1	47.98	54.02	37.98	57.52	30.16	46.92	80.24	83.71	52.88	55.34
- Stage 2	33.42	28.73	13.95	25.67	1.59	25.81	88.07	86.76	56.86	58.60
<i>Ablation on Model Architecture</i>										
- TSFM	51.48	52.87	51.16	63.71	38.09	51.76	89.43	89.70	58.65	60.62

444 Qwen-2.5-Instruct backbone: 0.5B, 3B, and 7B. The results, shown in Figure 6, confirm that TS-
445 REASONER is both highly effective and robustly scalable. We observe a clear positive scaling law
446 for both TS-REASONER and a baseline. More importantly, TS-REASONER maintains a consistent
447 and significant lead across all models, with Overall Accuracy improvements of +3.15% (29.49% vs.
448 26.34%), +9.70% (47.71% vs. 38.01%), and +7.61% (54.26% vs. 46.65%) for the 0.5B, 3B, and 7B
449 models, respectively. This demonstrates that our approach performs robustly across different LLM
450 backbones for complex time series reasoning.

451 4.5 ABLATION STUDIES

452 To further demonstrate the effectiveness of TS-REASONER, we conduct ablation studies to analyze
453 the impact of individual components. Table 3 summarizes our component-wise ablations from both
454 training and model architecture perspectives:

455 (1) **Attribute-aware captioning is critical for robust language-timeseries alignment.** (1) Removing
456 our attribute-aware captioning data entirely degrades overall accuracy by 3.01% on TimeSeriesExam
457 and 2% on MTBench. (2) Removing the attributes from the captioning instructions still results in a
458 performance drop of 1.57% and 2.09%, respectively. These results confirm that fine-grained details
459 are vital for learning nuanced temporal patterns. The quality of these captions is confirmed through a
460 quantitative analysis in Appendix C. We provide a qualitative case study in Appendix D to illustrate
461 how attribute-rich captions provide crucial details for model comprehension.

462 (2) **Absence of any training stage significantly harms the performance.** When removing stage 1
463 and training only with instruction tuning data, the performance on both benchmarks drops to a large
464 extent due to the weak time series understanding ability. Lack of stage 2 leads to a performance drop
465 by 28.45% on the TimeSeriesExam benchmark and 3.34% on MTBench. The significant performance
466 gap is attributed to the weak ability to understand time series instructions. We observe that removing
467 stage 1 (alignment) leads to a larger drop on MTBench, while removing stage 2 (instruction tuning)
468 causes a larger drop on TimeSeriesExam. This difference comes from the task characteristics of the
469 benchmarks: Removing Stage 1 hurts MTBench more because its tasks require reasoning across both
470 time series and textual news, a skill entirely dependent on the cross-modal alignment learned in Stage
471 1. In contrast, removing Stage 2 impacts TimeSeriesExam more severely because it directly tests the
472 model’s ability to follow specific analytical commands, which is precisely the skill taught in Stage 2.

473 (3) **Pretrained TSFM is crucial for effective time series feature extraction.** We remove the
474 pretrained TSFM and repurpose the TS-to-Text adapter to directly project time series patches into the
475 LLM’s embedding space. As shown in Table 3, this modification leads to a performance decrease
476 of 2.50% on the TimeSeriesExam benchmark and 1.31% on MTBench. This result underscores the
477 importance of the TSFM as a powerful temporal feature extractor.

480 5 CONCLUSION

481 We introduce TS-REASONER, a framework that advances the ability of LLMs to understand and
482 reason about time series via bridging with the TSFM. To mitigate the intrinsic semantic gap, we further
483 developed an attribute-aware captioning method that enriches time-series alignment data, fostering
484 a more robust alignment. Extensive experiments demonstrate that TS-REASONER substantially
485 outperforms a wide range of baselines on time series understanding and reasoning benchmarks.

486 ETHICS AND REPRODUCIBILITY STATEMENTS
487488 (1) **Ethics:** Our work aims to improve the time series understanding and reasoning ability, and the
489 experiments conducted in this paper adopt open-source data only for research purposes. It is far from
490 exceeding the understanding of humanity, which does not anticipate any ethical concerns with this
491 work.492 (3) **Reproducibility:** Sections 3 and 4 describe our methods and experiments. Further experiment
493 details and results are available in Appendix C. Finally, we include our code repository in the
494 supplemental materials.
495496 REFERENCES
497498 The claude 3 model family: Opus, sonnet, haiku. URL <https://api.semanticscholar.org/CorpusID:268232499>.
499500 Openai gpt-4o., 2024. URL <https://openai.com/index/hello-gpt-4o/>.
501502 Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin
503 Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, et al. Phi-4-mini technical
504 report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint*
505 *arXiv:2503.01743*, 2025.506 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
507 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
508 *arXiv preprint arXiv:2303.08774*, 2023.
509510 Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen,
511 Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al.
512 Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.513 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
514 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
515 2025.516 Yves Bestgen. Measuring lexical diversity in texts: The twofold length problem. *Language Learning*,
517 74(3):638–671, 2024.519 Yifu Cai, Arjun Choudhry, Mononito Goswami, and Artur Dubrawski. Timeseriesexam: A time
520 series understanding exam. *arXiv preprint arXiv:2410.14752*, 2024a.
521522 Zheng Cai, Maosong Cao, Haojong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen,
523 Zhi Chen, Pei Chu, et al. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*, 2024b.524 Defu Cao, Furong Jia, Sercan O Arik, Tomas Pfister, Yixiang Zheng, Wen Ye, and Yan Liu. Tempo:
525 Prompt-based generative pre-trained transformer for time series forecasting. *arXiv preprint*
526 *arXiv:2310.04948*, 2023.527 Jialin Chen, Aosong Feng, Ziyu Zhao, Juan Garza, Gaukhar Nurbek, Cheng Qin, Ali Maatouk,
528 Leandros Tassiulas, Yifeng Gao, and Rex Ying. Mtbench: A multimodal time series benchmark
529 for temporal reasoning and question answering. *arXiv preprint arXiv:2503.16858*, 2025.
530531 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahu Lin.
532 Sharegpt4v: Improving large multi-modal models with better captions. In *European Conference*
533 *on Computer Vision*, pp. 370–387. Springer, 2024a.534 Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan,
535 Zhenyu Tang, Li Yuan, et al. Sharegpt4video: Improving video understanding and generation with
536 better captions. *Advances in Neural Information Processing Systems*, 37:19472–19495, 2024b.
537538 Mouxiang Chen, Lefei Shen, Zhuo Li, Xiaoyun Joy Wang, Jianling Sun, and Chenghao Liu. Visions:
539 Visual masked autoencoders are free-lunch zero-shot time series forecasters. *arXiv preprint*
540 *arXiv:2408.17253*, 2024c.

540 Kanzhi Cheng, Wenpo Song, Zheng Ma, Wenhao Zhu, Zixuan Zhu, and Jianbing Zhang. Beyond
 541 generic: Enhancing image captioning with real-world knowledge using vision-language pre-
 542 training model. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp.
 543 5038–5047, 2023.

544 Kanzhi Cheng, Wenpo Song, Jiaxin Fan, Zheng Ma, Qiushi Sun, Fangzhi Xu, Chenyang Yan, Nuo
 545 Chen, Jianbing Zhang, and Jiajun Chen. Caparena: Benchmarking and analyzing detailed image
 546 captioning in the llm era. *arXiv preprint arXiv:2503.12329*, 2025.

548 Juhwan Choi, Junehyoung Kwon, JungMin Yun, Seunguk Yu, and YoungBin Kim. Voldoger:
 549 Llm-assisted datasets for domain generalization in vision-language tasks. *arXiv preprint*
 550 *arXiv:2407.19795*, 2024.

551 Winnie Chow, Lauren Gardiner, Haraldur T Hallgrímsson, Maxwell A Xu, and Shirley You Ren.
 552 Towards time series reasoning with llms. *arXiv preprint arXiv:2409.11376*, 2024.

554 Yu-Neng Chuang*, Songchen Li*, Jiayi Yuan*, Guanchu Wang*, Kwei-Herng Lai*, Leisheng Yu,
 555 Sirui Ding, Chia-Yuan Chang, Qiaoyu Tan, Daochen Zha, and Xia Hu. Understanding different
 556 design choices in training large time series models. *arXiv preprint arXiv:2406.14045*, 2024.

557 Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for
 558 time-series forecasting. In *Forty-first International Conference on Machine Learning*, 2024.

559

560 Hao Dong, Moru Liu, Kaiyang Zhou, Eleni Chatzi, Juho Kannala, Cyrill Stachniss, and Olga Fink.
 561 Advances in multimodal adaptation and generalization: From traditional approaches to foundation
 562 models. *arXiv preprint arXiv:2501.18592*, 2025.

563 Vijay Ekambaram, Arindam Jati, Pankaj Dayama, Sumanta Mukherjee, Nam Nguyen, Wesley M
 564 Gifford, Chandra Reddy, and Jayant Kalagnanam. Tiny time mixers (ttms): Fast pre-trained models
 565 for enhanced zero/few-shot forecasting of multivariate time series. *Advances in Neural Information
 566 Processing Systems*, 37:74147–74181, 2024.

567

568 Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso, and
 569 Svitlana Vyetrenko. Evaluating large language models on time series feature understanding: A
 570 comprehensive taxonomy and benchmark. *arXiv preprint arXiv:2404.16563*, 2024.

571 Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas,
 572 Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to
 573 glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024.

574

575 Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski.
 576 Moment: A family of open time-series foundation models. In *International Conference on
 577 Machine Learning*, 2024.

578 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 579 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
 580 models. *arXiv preprint arXiv:2407.21783*, 2024.

581

582 Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew G Wilson. Large language models are zero-shot
 583 time series forecasters. *Advances in Neural Information Processing Systems*, 36:19622–19635,
 584 2023.

585 Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu.
 586 Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*,
 587 2023.

588 Furong Jia, Kevin Wang, Yixiang Zheng, Defu Cao, and Yan Liu. Gpt4mts: Prompt-based large
 589 language model for multimodal time-series forecasting. In *Proceedings of the AAAI Conference on
 590 Artificial Intelligence*, volume 38, pp. 23343–23351, 2024.

591

592 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris
 593 Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al.
 Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

594 Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, et al. Time-llm: Time series forecasting by reprogramming
 595 large language models. *arXiv preprint arXiv:2310.01728*, 2023.

596

597 Zhengfeng Lai, Haotian Zhang, Bowen Zhang, Wentao Wu, Haoping Bai, Aleksei Timofeev, Xianzhi
 598 Du, Zhe Gan, Jiulong Shan, Chen-Nee Chuah, et al. Veclip: Improving clip training via visual-
 599 enriched captions. In *European Conference on Computer Vision*, pp. 111–127. Springer, 2024.

600

601 Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan
 602 Zhang, Ziwei Liu, and Chunyuan Li. Llava-next: Stronger llms supercharge multi-
 603 modal capabilities in the wild, May 2024a. URL <https://llava-v1.github.io/blog/2024-05-10-llava-next-stronger-llms/>.

604

605 Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-
 606 training for unified vision-language understanding and generation. In *International conference on
 607 machine learning*, pp. 12888–12900. PMLR, 2022.

608

609 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 610 pre-training with frozen image encoders and large language models. In *International conference
 611 on machine learning*, pp. 19730–19742. PMLR, 2023a.

612

613 Yuxin Li, Wenchao Chen, Bo Chen, Dongsheng Wang, Long Tian, and Mingyuan Zhou. Prototype-
 614 oriented unsupervised anomaly detection for multivariate time series. In *International Conference
 615 on Machine Learning*, pp. 19407–19424. PMLR, 2023b.

616

617 Yuxin Li, Wenchao Chen, Xinyue Hu, Bo Chen, Mingyuan Zhou, et al. Transformer-modulated
 618 diffusion models for probabilistic multivariate time series forecasting. In *The Twelfth International
 619 Conference on Learning Representations*, 2024b.

620

621 Bryan Lim, Sercan Ö Arik, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for
 622 interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):
 1748–1764, 2021.

623

624 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
 625 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint
 626 arXiv:2412.19437*, 2024a.

627

628 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in
 629 neural information processing systems*, 36, 2024b.

630

631 Haixin Liu, Zhiyuan Zhao, Jindong Wang, Harshavardhan Kamarthi, and B Aditya Prakash. Lst-
 632 prompt: Large language models as zero-shot time series forecasters by long-short-term prompting.
arXiv preprint arXiv:2402.16132, 2024c.

633

634 Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X Liu, and Schahram Dust-
 635 dar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and
 forecasting. In *International Conference on Learning Representations*, 2021.

636

637 Xu Liu, Juncheng Liu, Gerald Woo, Taha Aksu, Yuxuan Liang, Roger Zimmermann, Chenghao
 638 Liu, Silvio Savarese, Caiming Xiong, and Doyen Sahoo. Moirai-moe: Empowering time series
 639 foundation models with sparse mixture of experts. *arXiv preprint arXiv:2410.10469*, 2024d.

640

641 Yong Liu, Haoran Zhang, Chenyu Li, Xiangdong Huang, Jianmin Wang, and Mingsheng Long. Timer:
 642 Generative pre-trained transformers are large time series models. *arXiv preprint arXiv:2402.02368*,
 2024e.

643

644 Yong Liu, Guo Qin, Zhiyuan Shi, Zhi Chen, Caiyin Yang, Xiangdong Huang, Jianmin Wang, and
 645 Mingsheng Long. Sundial: A family of highly capable time series foundation models. *arXiv
 646 preprint arXiv:2502.00816*, 2025.

647 Mike A Merrill, Mingtian Tan, Vinayak Gupta, Tom Hartvigsen, and Tim Althoff. Language models
 648 still struggle to zero-shot reason about time series. *arXiv preprint arXiv:2404.11757*, 2024.

648 Ozan Baris Mulayim, Pengrui Quan, Liying Han, Xiaomin Ouyang, Dezhi Hong, Mario Bergés, and
 649 Mani Srivastava. Are time series foundation models ready to revolutionize predictive building
 650 analytics? In *Proceedings of the 11th ACM International Conference on Systems for Energy-*
 651 *Efficient Buildings, Cities, and Transportation*, pp. 169–173, 2024.

652 Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64
 653 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

654 Yuqi Nie, Yaxuan Kong, Xiaowen Dong, John M Mulvey, H Vincent Poor, Qingsong Wen, and Stefan
 655 Zohren. A survey of large language models for financial applications: Progress, prospects and
 656 challenges. *arXiv preprint arXiv:2406.11903*, 2024.

657 Zijie Pan, Yushan Jiang, Sahil Garg, Anderson Schneider, Yuriy Nevmyvaka, and Dongjin Song. s^2 IP-
 658 LLM: Semantic space informed prompt learning with LLM for time series forecasting. In Ruslan
 659 Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and
 660 Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*,
 661 volume 235 of *Proceedings of Machine Learning Research*, pp. 39135–39153. PMLR, 21–27 Jul
 662 2024. URL <https://proceedings.mlr.press/v235/pan24c.html>.

663 Kandukuri Ratna Prakarsha and Gaurav Sharma. Time series signal forecasting using artificial neural
 664 networks: An application on ecg signal. *Biomedical Signal Processing and Control*, 76:103705,
 665 2022.

666 CLEVELAND RB. Stl: A seasonal-trend decomposition procedure based on loess. *J Off Stat*, 6:
 667 3–73, 1990.

668 Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, and Ming Jin. Time-
 669 moe: Billion-scale time series foundation models with mixture of experts. *arXiv preprint*
 670 *arXiv:2409.16040*, 2024.

671 Chenxi Sun, Hongyan Li, Yaliang Li, and Shenda Hong. Test: Text prototype aligned embedding to
 672 activate llm’s ability for time series. *arXiv preprint arXiv:2308.08241*, 2023.

673 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 674 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 675 *systems*, 30, 2017.

676 Chengsen Wang, Qi Qi, Zhongwen Rao, Lujia Pan, Jingyu Wang, and Jianxin Liao. Chronosteer:
 677 Bridging large language model and time series foundation model via synthetic data. *arXiv preprint*
 678 *arXiv:2505.10083*, 2025a.

679 Chengsen Wang, Qi Qi, Jingyu Wang, Haifeng Sun, Zirui Zhuang, Jinming Wu, Lei Zhang, and
 680 Jianxin Liao. Chattime: A unified multimodal time series foundation model bridging numerical
 681 and textual data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp.
 682 12694–12702, 2025b.

683 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 684 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 685 *neural information processing systems*, 35:24824–24837, 2022.

686 Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers
 687 with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing*
 688 *Systems*, 34:22419–22430, 2021.

689 Zhe Xie, Zeyan Li, Xiao He, Longlong Xu, Xidao Wen, Tieying Zhang, Jianjun Chen, Rui Shi, and
 690 Dan Pei. Chatts: Aligning time series with llms via synthetic data for enhanced understanding and
 691 reasoning. *arXiv preprint arXiv:2412.03104*, 2024.

692 Jingwen Xu, Fei Lyu, and Pong C Yuen. Density-aware temporal attentive step-wise diffusion model
 693 for medical time series imputation. In *Proceedings of the 32nd ACM International Conference on*
 694 *Information and Knowledge Management*, pp. 2836–2845, 2023.

702 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 703 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2.5 technical report. *arXiv preprint*
 704 *arXiv:2412.15115*, 2024.

705 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan.
 706 Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural*
 707 *information processing systems*, 36:11809–11822, 2023.

708 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li,
 709 Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint*
 710 *arXiv:2408.01800*, 2024.

711 Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. Flow of reasoning: Training
 712 llms for divergent problem solving with minimal examples. *arXiv preprint arXiv:2406.05673*,
 713 2024.

714 George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff.
 715 A transformer-based framework for multivariate time series representation learning. In *Proceedings*
 716 *of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 2114–2124,
 717 2021.

718 Haochuan Zhang, Chunhua Yang, Jie Han, Liyang Qin, and Xiaoli Wang. Tempogpt: Enhancing
 719 temporal reasoning via quantizing embedding. *arXiv preprint arXiv:2501.07335*, 2025a.

720 Haoran Zhang, Yong Liu, Yunzhong Qiu, Haixuan Liu, Zhongyi Pei, Jianmin Wang, and Mingsheng
 721 Long. Timesbert: A bert-style foundation model for time series understanding. *arXiv preprint*
 722 *arXiv:2502.21245*, 2025b.

723 Siru Zhong, Weilin Ruan, Ming Jin, Huan Li, Qingsong Wen, and Yuxuan Liang. Time-vlm:
 724 Exploring multimodal vision-language models for augmented time series forecasting. *arXiv*
 725 *preprint arXiv:2502.04395*, 2025.

726 Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Fre-
 727 quency enhanced decomposed transformer for long-term series forecasting. *arXiv preprint*
 728 *arXiv:2201.12740*, 2022.

729 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
 730 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
 731 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

732 Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. Texygen:
 733 A benchmarking platform for text generation models. In *The 41st international ACM SIGIR*
 734 *conference on research & development in information retrieval*, pp. 1097–1100, 2018.

742 A IMPLEMENTATION DETAILS

743 TS-REASONER uses the Qwen-2.5-7B-
 744 Instruct as the LLM backbone across all the
 745 experiments with an embedding dimension of
 746 5120, and uses the TimesFM-1.0-200M (Das
 747 et al., 2024) as our backbone TSFM with
 748 an embedding dimension of 1080. All the
 749 parameters of the backbone are finetuned
 750 during training. The detailed derivation of
 751 these time series embeddings from TimesFM
 752 can be found in the Appendix B. All training
 753 and inference procedures for TS-REASONER
 754 were conducted locally on 8 × L40s GPUs.
 755 Comprehensive training parameters are
 756 further detailed in Table 4.

757 Table 4: Training details of TS-REASONER.

	Stage-1	Stage-2
Patch Size	32	32
Dataset	Captions	Instructions
#Samples	120K	30K
TSFM	TimesFM-1.0-200M	
LLM Backbone	Qwen2.5-7B-Instruct	
Trainable Params.	7.3B	7.3B
Batch Size	64	32
Learning Rate:	1×10^{-5}	2×10^{-5}
Epoch	1	2

756 **B TIMESFM FOR TIME SERIES EMBEDDING**
757

758 Given a time series $\mathcal{T} \in \mathbb{R}^L$, where L is the length of the time series. We first normalize it to have a
759 mean of zero and a variance of one. We then segment \mathcal{T} into consecutive, non-overlapping patches
760 of fixed length P , resulting in a total of $N = \lfloor L/P \rfloor$ patches. This yields a patched time series
761 $\mathcal{T}_p \in \mathbb{R}^{N \times P}$.

762 Following the approach of (Das et al., 2024), j -th patch \mathcal{T}_p^j is passed through a residual block to
763 project it into the model dimension. This block is implemented as a two-layer MLP with a skip
764 connection, processing each patch independently. The input token for the j -th patch is computed as:
765

$$\mathcal{E}_p^j = \text{InputResidualBlock}(\mathcal{T}_p^j) + \text{PE}_j, \quad (5)$$

766 where PE_j is the position encoding for the j -th patch, as defined in the original transformer (Vaswani
767 et al., 2017). These encoded patch representations are then fed into an M -layer stacked Transformer
768 to produce the final sequence of time series features:
769

$$\mathcal{Z}_T = \text{StackedTransformer}([\mathcal{E}_p^{(0)}, \mathcal{E}_p^{(1)}, \dots, \mathcal{E}_p^{(N)})]), \quad (6)$$

770 where $\mathcal{Z}_T \in \mathbb{R}^{N \times d_{\text{ts}}}$ and d_{ts} denotes the embedding dimension for each time series patch. Refer to
771 more details of TimesFM in (Das et al., 2024).
772

773 **C CAPTION ANALYSIS**
774

775 A critical limitation of synthetic datasets
776 is the risk of models learning spurious correlations from similar templates. To mitigate this, our attribute-aware generation
777 process is designed to produce captions that are both lexically diverse. To quantitatively validate the richness of our approach,
778 we compare it against the template-based method. We evaluate both lexical diversity
779 using the Measure of Textual Lexical Diversity (MTLD) (Bestgen, 2024) and Self-BLEU-4 (Zhu
780 et al., 2018) on a random sample of 1K captions from each dataset. The results presented in Table 5
781 show that our attribute-aware captions achieve an MTLD score of 133.30, a nearly 3 times increase
782 over the template-based score of 42.95. Furthermore, the Self-BLEU-4 score is almost halved from
783 0.82 to 0.45. This substantial improvement in lexical diversity confirms that our method generates a
784 significantly more expressive and diverse set of captions, crucial for training robust and generalizable
785 models.
786

787 To ensure comprehensive data coverage, we curated time series with context from a wide range of
788 domains. The distribution of these domains is visualized in Figure 8.
789

790 **D QUALITATIVE ANALYSIS: A CASE STUDY**
791

792 To qualitatively evaluate the distinct advantages of our approach, we conduct a case study comparing
793 three methods: (1) our proposed attribute-aware captioning, which leverages visual time series plots
794 and explicit attribute guidance; (2) a basic captioning baseline that operates on visual plots but lacks
795 attribute guidance; and (3) LLM prompted with the raw textual (numerical) time series data. Our
796 analysis, illustrated in Figure 9, yields two key insights.
797

798 (i) **Attribute-Aware Captions Provide Semantically Richer Descriptions.** A primary limitation
799 of basic captioning is its tendency to produce superficial, chronological narrations of the data. As
800 shown in Figure 9, the captioner describes the series' movements (e.g., "the value increases, then
801 decreases sharply") but fails to extract deeper, underlying characteristics. While factually correct, this
802 description omits properties crucial for a comprehensive understanding. In contrast, our attribute-
803 aware captioning enriches this chronological account with critical semantic attributes. It not only
804 captures the temporal dynamics but also identifies and articulates the series' overall trend, periodicity,
805 and noise level. This multifaceted analysis provides a more holistic understanding of the time series,
806 which is essential for TS-REASONER to conduct reasoning on downstream tasks.
807




Figure 8: Domain distribution of LLM-generated time series with context.

(ii) **Visual Representation is Crucial for Capturing Global Temporal Patterns.** When comparing our visually-grounded method to an LLM processing raw numerical data, a significant gap emerges in the ability to identify global patterns. The text-based LLM, while capable of discerning local features like high-frequency oscillations or noise within a limited window, consistently fails to recognize the overarching periodicity of the entire series. We hypothesize that this failure stems from the inherent inefficiency of representing long numerical sequences as text. The exceeding length of input may distract the LLM, preventing it from observing the complete pattern. Conversely, a time series plot serves as a highly compressed, holistic representation. It enables the model to perceive the entire sequence as a single input, making global structures like periodicity visually salient and readily identifiable.

E PROMPTS FOR ATTRIBUTE-AWARE CAPTIONING

Below, we detail all the prompts we used. Figure 10 shows our captioning prompts.


F THE USE OF LARGE LANGUAGE MODELS

In this work, LLMs were only used to review and polish small sections of writing

LIMITATION

Our current experimental setup utilizes LLMs with up to 7B parameters, primarily due to existing resource constraints. While we anticipate that our findings will extend to larger models, future work should aim to validate these results with more substantial LLMs. Additionally, exploring alternative architectural designs, such as different encoder architectures, and investigating more effective training strategies for time series text alignment would be valuable avenues for future research. Future research could also explore incorporating RL approaches to elicit reflection behaviors.

864
865
866
867
868
869
870
871
872
873
874
875
876
877



Input: Create a detailed description of the time series in one paragraph.

Output: The time series plot shows a fluctuating trend with distinct phases of increase and decrease over time. Initially, the series begins near zero and exhibits a sharp upward movement, peaking around the 10-20 time mark, where it reaches values greater than 1. This is followed by a volatile but generally declining trend that persists through to approximately the 60th time unit, dipping below -1. From this low point, the series remains subdued and oscillates around a low mean until about time step 100. After that, a gradual upward recovery begins, with the series rising back toward positive values and ending slightly above zero near time step 130. The fluctuations suggest a non-stationary process with alternating periods of high volatility and relative stability.

Input with visualized time series: Create a detailed description of the time series in one paragraph, including its trend, frequency properties, periodicity, noise, local variations, and other relevant characteristics.

Output: This time series displays a **prominent, low-frequency pattern without a clear long-term linear trend**. The series appears **non-stationary**, as its mean value shifts significantly over time, completing roughly one full oscillation over the observed interval of approximately 125 time units. The primary cycle involves a rise to a peak around time unit 20, followed by a decline to a trough near time unit 85, and a subsequent recovery. Superimposed on this dominant wave are higher-frequency, smaller-amplitude fluctuations, which introduce **considerable noise** and give the line a jagged, irregular texture. The volatility seems relatively constant, although there are **local variations in the magnitude and sharpness of these smaller oscillations**. Strong positive autocorrelation is evident, as consecutive data points are closely related, resulting in a continuous, flowing path rather than random jumps.

Input with textual time series: Here is a time series: [0.13, 0.05, -0.12, -0.24, ...]. Create a detailed description of the time series in one paragraph, including its trend, frequency properties, periodicity, noise, local variations, and other relevant characteristics.

Output: The time series exhibits a strong **quasi-periodic behavior** without a discernible long-term trend, characterized by a distinct cycle of a sharp rise to a peak, a steep fall into a trough, and a subsequent gradual recovery. This dominant low-frequency pattern is superimposed with considerable high-frequency noise and local variations, resulting in fluctuations that **make each cycle similar** but not identical. Overall, the data represents a noisy cyclical signal, where **a clear underlying periodicity** is modulated by random, short-term movements.

✗

907
908
909
910
911
912
913
914
915
916
917

Figure 9: Case study comparing time series captioning methods. Basic visual captioning (row 1) gives a chronological description. Our attribute-aware visual captioning (row 2) provides a richer description including key attributes. Textual time series captioning (row 3) identifies some attributes but fails to capture the periodicity pattern.

918
919
920
921
922
923
924
925
926

927 Instructions for prompting LLMs to generate time series captions.
928

1. Write a paragraph that analyzes the time series, covering its local behaviors, noise levels, periodic structures, overall trend, frequency content, and any other characteristics you consider important.
2. Create a detailed description of the time series in one paragraph, including its trend, frequency properties, periodicity, noise, local variations, and other relevant characteristics.
3. Provide a paragraph summarizing the time series characteristics such as noise, periodic patterns, long-term trends, frequency behavior, local anomalies, and any other significant features.
4. Compose a detailed caption describing the frequency characteristics, noise, trends, local variations, periodic structures, and any other meaningful patterns you observe in the time series.
5. Craft a one-paragraph summary of the time series, noting local fluctuations, periodic behavior, frequency features, trend, noise content, and any other insights you find important.
6. Generate a descriptive paragraph detailing the time series' key attributes, including frequency structure, noise patterns, trend direction, local features, periodic elements, and other notable aspects.
7. Give a thorough one-paragraph explanation of the time series, addressing periodicity, noise, frequency components, trend, local variations, and other relevant characteristics.
8. Write a narrative paragraph explaining the time series, focusing on noise, frequency characteristics, periodicity, localized structures, the overall trend, and other important features you identify.
9. Summarize the time series in a paragraph, describing its fluctuations, recurring patterns, noise levels, frequency-domain features, trend direction, and any additional traits you find significant.,
10. Develop a paragraph that captures the key features of the time series, such as frequency traits, trend, noise, periodic components, local behaviors, and other characteristics worth noting.
11. Provide a one-paragraph caption analyzing the time series data in terms of noise, trend, periodicity, local features, frequency-related behavior, and any additional characteristics of interest.
12. Create a rich paragraph description of the time series, including its trend, local anomalies, periodic activity, noise artifacts, spectral content, and other important descriptive elements.
13. Write a descriptive paragraph for the time series, highlighting frequency properties, trend behavior, periodic patterns, local structures, noise, and other characteristics you consider relevant.
14. Generate a compact yet thorough paragraph explaining the time series in terms of periodicity, trend movement, noise level, frequency details, local dynamics, and any other key aspects.
15. Construct a one-paragraph analysis of the time series by examining its local variations, noise, trend, periodic elements, frequency spectrum, and other notable features you deem important.
16. Write a summary paragraph that discusses the time series' periodic features, trend behavior, local patterns, noise levels, frequency domain signals, and other characteristics worth mentioning.,
17. Create a detailed one-paragraph commentary on the time series that outlines its noise characteristics, periodicity, frequency content, trends, localized behaviors, and other useful insights.
18. Prepare a paragraph-long description of the time series covering its trend, noise, frequency-related traits, local fluctuations, periodic structures, and any additional attributes of note.,
19. Offer a one-paragraph interpretation of the time series, highlighting its frequency features, periodic nature, local patterns, noise, trend line, and any other important characteristics you observe.
20. Compose a detailed summary in one paragraph focusing on the time series' periodic behavior, frequency spectrum, localized fluctuations, overall trend, noise, and other relevant descriptive elements.

962
963 Figure 10: The list of instructions for attributes-aware time series captioning.
964
965
966
967
968
969
970
971