

CaTS-Bench: Can Language Models Describe Time Series?

Luca Zhou^{*†‡}, Pratham Yashwante^{*‡}, Marshall Fisher[‡], Alessio Sampieri[§],
Zihao Zhou[‡], Fabio Galasso[†], Rose Yu[‡]

[†]Sapienza University of Rome, [‡]UC San Diego, [§]ItalAI Labs

Correspondence: luca.zhou@uniroma1.it

Abstract

Time series captioning, the task of describing time series in natural language, requires numeric and temporal reasoning, trend interpretation, and contextual understanding. Existing benchmarks, however, often rely on fully synthetic or generic captions, and typically neglect metadata and visual representations. We introduce **CaTS-Bench**, a comprehensive benchmark for **Context-aware Time Series** reasoning across 11 diverse domains, centered on a gold-standard evaluation set of 1746 human-written captions that measure how effectively models translate numeric trends into immediately interpretable narratives. To address the scarcity of human-annotated data, we also propose a scalable pipeline for generating high-fidelity synthetic captions, the quality of which we validate. We evaluate leading Vision-Language Models on our benchmark, revealing that even proprietary models struggle to capture numeric nuances in temporal descriptions, while finetuning open-source models on synthetic data yields substantial performance gains. Finally, we release a diagnostic suite of 910 multiple-choice questions and use tailored numeric metrics to gauge time-series-specific reasoning capabilities, establishing CaTS-Bench as a reliable foundation for grounded, multimodal text generation in numeric domains.

1 Introduction

Effective interpretation of time series data underpins decision-making in domains such as finance, healthcare, and climate analysis. Although raw numeric sequences are abundant, turning them into concise, human-readable summaries is labor-intensive and demands domain expertise and statistical literacy. In practice, analysts, policymakers, journalists, and practitioners rarely inspect raw sequences directly, instead relying on textual descriptions to quickly assess trends, anomalies, and

changes before deciding whether deeper analysis is needed. Automating this process through *time series captioning* (TSC) not only accelerates insight discovery but also democratizes access to complex temporal analytics, which enables non-experts to ask natural-language questions and receive explanations without writing code or inspecting raw charts.

We further discuss the motivation and practical utility of TSC in Appendix A, including: (i) its distinction from deterministic statistical analysis, (ii) its role in generating human-interpretable, context-aware narratives, and (iii) its use as a diagnostic task for evaluating free-form language generation.

Recent advances in Large Language Models (LLMs) and Vision-Language Models (VLMs) have transformed grounded language generation, yet these models show critical weaknesses in precise reasoning and factual grounding on numeric time series (Tang et al., 2025; Cao and Wang, 2024). Progress in TSC is bottlenecked by the lack of high-quality, linguistically diverse datasets and benchmarks: existing resources are largely synthetic, template-based, or omit contextual metadata (e.g., units, domain-specific labels) that humans rely on, leading models to produce generic captions that miss the underlying narrative.

To address this, we introduce **CaTS-Bench**¹, a context-rich multimodal benchmark for *context-aware* time series captioning. An overview of the benchmark is shown in Figure 1. Here, “context-aware” denotes synthesizing information from numeric series, visual plots, and contextual metadata. Rather than assuming that access to multiple modalities implies effective multimodal reasoning, we explicitly probe whether VLMs properly use these signals. Unlike previous efforts that rely solely on simplistic or synthetic data, our work centers on a new **gold-standard evaluation set** of 1746 human-

^{*}Equal contribution.

¹The dataset is available at <https://huggingface.co/datasets/mhfisher/CaTSBench>

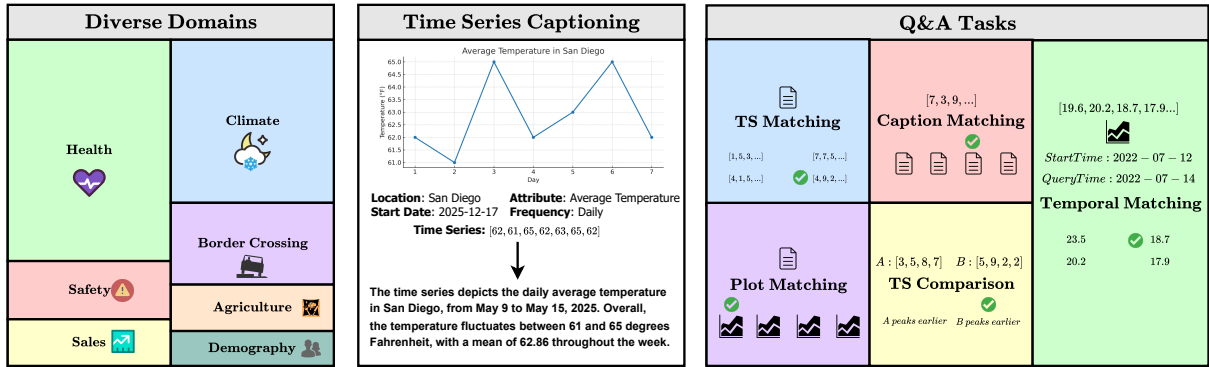


Figure 1: **Overview of CaTS-Bench.** It features 11 domains organized into 7 macro domains, provides human-written and synthetic data, and formulates six challenging tasks, with time series captioning as the primary one.

rewritten (HR) captions. These captions provide a rigorous baseline to measure linguistic richness, factual and numeric precision, and human-like reasoning in temporal descriptions.

Recognizing the prohibitive cost of large-scale human annotation, another key contribution of this work is a **scalable data augmentation pipeline** that uses an oracle LLM, grounded in factual metadata and validated via human verification studies, to generate high-fidelity synthetic training captions. Our experiments show that zero-shot VLMs struggle with the complexity of HR descriptions, whereas finetuning on our synthetic data yields substantial gains, particularly in semantic and numeric metrics on the HR set, positioning the pipeline as an effective solution to data scarcity for TSC.

Beyond captioning, CaTS-Bench provides a diagnostic suite of 910 multiple-choice questions (MCQs) targeting numeric precision and multimodal alignment, together with tailored evaluation metrics that prioritize numeric fidelity. Our results reveal that current VLMs often fail to exploit visual cues for TSC and struggle even on simple diagnostic MCQ tasks, such as correctly mapping values to time stamps and comparing basic statistical properties, highlighting persistent gaps in grounded multimodal reasoning. Overall, our contributions are:

1. **CaTS-Bench:** A multimodal, 11-domain benchmark featuring a 1746-sample human-written evaluation set for TSC and 910 MCQs designed to rigorously measure linguistically rich and factually grounded free-form text generation and temporal reasoning.
2. **Scalable Augmentation Pipeline:** A framework for generating high-fidelity synthetic data. We empirically demonstrate that finetuning on this synthetic data improves model performance

on our human-written test set.

3. **Diagnostic Analysis & Metrics:** A suite of tailored numeric fidelity metrics that reveal systematic failures in visual grounding, temporal alignment, and multimodal reasoning.

2 Related Work

LLMs for Time Series Analysis. LLMs are increasingly repurposed for temporal reasoning, with early efforts focused on forecasting and anomaly detection (Zhang et al., 2024; Liu et al., 2024a). Methodological approaches include specialized prompt engineering (Liu et al., 2024a; Chatzigeorgakidis et al., 2024), modality alignment (Liu et al., 2024b; Sun et al., 2023; Liu et al., 2024c; Pan et al., 2024; Cai et al., 2024), numeric discretization (Ansari et al., 2024; Jin et al., 2024), and parameter-efficient finetuning (Zhou et al., 2023; Chang et al., 2023). While these studies demonstrate that models pretrained on massive text corpora can inherit temporal reasoning capabilities, they also highlight persistent "hallucinations" in numeric extrapolation and difficulty handling long-range dependencies (Tang et al., 2025; Merrill et al., 2024; Tan et al., 2024; Cao and Wang, 2024; Zeng et al., 2023). We shift the focus from purely predictive tasks to generative description, exploring how these models translate numeric signals into language.

Related Benchmarks and TSC. Traditional time series archives (Dau et al., 2018; Bagnall et al., 2018; Godahewa et al., 2021) focus on discriminative tasks, while prompt-based forecasting benchmarks like PISA (Xue and Salim, 2023) lack the linguistic pairings and metadata required for descriptive tasks. While recent TSC-specific frameworks have pioneered the area (Fons et al., 2024b;

Table 1: Comparison of TSC benchmarks.

Dataset	Data Nature	Modalities	Domains	Context	Captions	Q&A
TaxoSynth (Fons et al., 2024a)	Synthetic	Numeric + Text	N/A	Minimal	Simplistic	✗
TADACap (Fons et al., 2024b)	Partially Synthetic	Visual	4	Minimal	Simplistic	✗
TRUCE (Jhamtani and Berg-Kirkpatrick, 2021)	Partially Synthetic	Numeric	2	✗	Simplistic	✗
TACO (Dohi et al., 2025)	Mostly Synthetic	Numeric	8	✗	Expressive	✗
BEDTime (Sen et al., 2025)	Partially Synthetic	Numeric + Text	6	✗	Simplistic	✓
CaTS-Bench	HR + Synthetic	Numeric + Text + Visual	11	Rich	Expressive	✓

Jhamtani and Berg-Kirkpatrick, 2021; Fons et al., 2024a; Dohi et al., 2025; Sen et al., 2025), they typically rely on simplistic trend patterns or template-based synthetic generation, failing to capture the descriptive nuance of human discourse (see Table 1). Furthermore, although integrating auxiliary modalities, such as metadata and visual renderings, is vital for grounded interpretability (Zhou and Yu, 2025; Dong et al., 2024; Chen et al., 2024; Wang et al., 2024; Kim et al., 2024; Williams et al., 2024; Liu et al., 2025; Tang et al., 2023), a unified benchmark integrating numeric data, expressive human captions, and multimodal grounding has remained elusive. CaTS-Bench addresses this gap with a unified benchmark for grounded TSC that integrates numeric data, visual representations, contextual metadata, and human language.

3 CaTS-Bench

In this section, we describe the construction of CaTS-Bench, starting from raw numeric series to our dual-stream captioning framework. CaTS-Bench provides an HR gold standard for evaluation, alongside a large-scale synthetic corpus for data augmentation. Moreover, we further enrich the scope of CaTS-Bench by providing an additional suite of Q&A tasks constructed from the same data, enabling a more fine-grained examination of time series and caption reasoning abilities.

3.1 Base Data Curation

We work with 11 diverse real-world source datasets that span climate, public safety, border crossings, demography, health, sales, and agriculture. See Appendix B for details of the source datasets. This variety ensures that models must generalize across different temporal scales, units, and domains.

The curation process begins by generating **base data triplets**. For each entity in a source dataset (e.g., a specific country or product), we apply a random window cropping strategy to extract seg-

ments of varying lengths (See Appendix C). For each segment, we construct a triplet consisting of (i) the **raw numeric time series**, (ii) a structured **metadata JSON** containing domain-specific context (e.g., units, location, temporal bounds) and statistical primitives (mean, dispersion, minimum, and maximum), and (iii) a **line plot** rendering of the series with randomized visual styles (color, width, and size) to promote robustness to visual variation. The overview is shown in Figure 2. To prevent information leakage, we partition the data temporally: the first 80% of each series is reserved for training/augmentation, while the final 20% is used for the actual benchmark data. The sizes of data splits are reported in Table 2.

3.2 Data Generation Streams

Using the base data triplets as a foundation, we employ two distinct streams to generate natural language descriptions: one focused on HR evaluation and the other on scalable training.

Human-Rewritten Gold Standard. To mitigate the potential oracle bias and stylistic homogeneity inherent in single-model benchmarks, we curated a primary evaluation set of 1746 HR captions. To ensure broad linguistic coverage and minimize model-specific artifacts, we utilized a multi-stage curation process: we initially generated candidate drafts from four distinct LLMs (Gemini 2.0 Flash, GPT-4o, Gemma 27B, and Llama 3 70B). Using these diverse drafts as structural scaffolds, we performed extensive rewriting and verification (see Appendix M for details and results of the editing process and Appendix H.4 for diversity analysis).

Across all domains, human rewriting resulted in substantial linguistic revision (average Word Error Rate = 0.60) while maintaining strong numeric overlap (average Numeric Jaccard = 0.74). This approach allowed us to move beyond simple post-editing, resulting in captions that are factually grounded, linguistically diverse, and enriched with

Table 2: Dataset outline by domain. AQ: Air Quality, Border: Border Crossing, Demo: Demography, Injury: Road Injuries, Diet: Calories Consumption, Agri: Agriculture

Metric	Overall	AQ	Border	Crime	Demo	Injury	COVID	CO ₂	Diet	Walmart	Retail	Agri
# Source Time Steps	287M	286M	397k	38k	14k	37k	720k	34k	234k	6k	7k	49k
# Triplets Generated	20k	4.4k	3.2k	764	598	756	5.5k	732	2.1k	544	551	835
# Human Test Samples	1746	0	323	153	120	152	300	100	211	109	111	167
Average TS Length	26.8	-	21.5	76.9	5.0	3.6	71.8	8.8	5.5	11.8	8.1	7.5
# Synth. Train Samples	16k	3.5k	2.6k	611	478	604	4.4k	585	1.7k	435	440	668
Average TS Length	29.1	65.3	21.2	76.8	11.6	5.9	75.8	9.5	12.2	12.2	22.4	7.3
# Synth. Test Samples	4k	886	646	153	120	152	1.1k	147	422	109	111	167
Average TS Length	26.1	66.0	21.2	76.9	5.0	3.6	73.0	8.7	5.5	11.8	8.1	7.5

the interpretive nuances that purely automated systems typically miss. These HR descriptions serve as the gold standard for CaTS-Bench, providing a rigorous test of a model’s ability to generate high-fidelity, context-aware narratives. Beyond evaluation, these captions reflect how humans summarize time series patterns for rapid comprehension.

Synthetic Augmentation Pipeline. To facilitate large-scale training, we utilize a scalable pipeline to generate additional reference captions. We query an oracle LLM (Gemini 2.0 Flash) with a structured prompt comprising the full numeric sequence, enriched metadata with precomputed statistics, and the corresponding line plot. The oracle is strictly instructed to remain within the bounds of the provided data, ensuring that the generated captions are factually grounded (see Appendix N.1 for a prompt example; see Appendix D for an ablation study on the role of metadata). We emphasize that the use of a single oracle model is intended as a proof of concept; different or multiple oracles can be used. While the synthetic data generation primarily serves as a large-scale training resource, we also utilize it to establish a dual-evaluation framework.

We generate a synthetic corpus of approximately 20k captions. Of these, 16k are used for training, while the full synthetic test split contains roughly 4k captions. For evaluation, however, we restrict attention to a subset of 1746 synthetic captions that exactly correspond to the HR evaluation set in the underlying time series. This dual-track allows us to compare synthetic and HR captions on identical samples, revealing where model-generated descriptions diverge from human-rewritten ones.

3.3 Synthetic Data Quality Validation

A key concern when using synthetic captions for training is whether they are factually accurate, lin-

guistically diverse, and free from artifacts that hinder generalization to human-authored text. To validate our synthetic data pipeline, we conduct three complementary analyses, with full protocols and results reported in Appendix H.

First, we perform a manual factual validation over approximately 2.9k synthetic captions, verifying statistical claims and trend descriptions against the underlying numeric series. This analysis shows that over 98.6% of captions are factually correct, indicating minimal numeric hallucination. Second, we conduct a human indistinguishability study with 35 participants, who are asked to distinguish synthetic captions from human-rewritten ones in a blind setting. The resulting detection accuracy is 41.1%, nearing random chance, suggesting that the synthetic captions are stylistically and semantically comparable to human-rewritten descriptions. Finally, we assess linguistic diversity to rule out template-driven generation. Embedding-based similarity analysis shows that near-duplicate captions (cosine similarity > 0.95) occur in only 2.3% of cases, and complementary N-gram and vocabulary analyses confirm broad lexical coverage comparable to human text (Appendix H.4). These results show that synthetic captions provide high-fidelity training signals that closely approximate human descriptions.

3.4 Task 1: Time Series Captioning

The primary objective of TSC is the generation of a linguistically coherent narrative describing the salient temporal characteristics of a numeric series. During evaluation, models are presented with a multimodal prompt comprising four components: i) **numeric series** consisting of raw, time-indexed values that require direct numeric reasoning; ii) **contextual metadata** providing semantic background (e.g., measurement units and location), while ex-

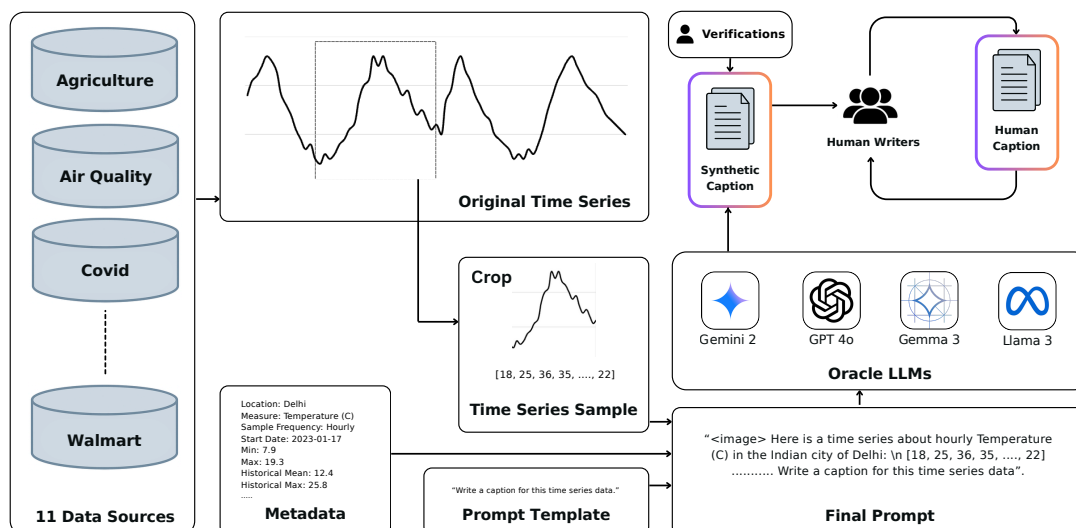


Figure 2: **Overview of CaTS-Bench’s dual-stream data generation pipeline.** A time series window is cropped, metadata is attached, and a human or LLM writes a reference caption. See Appendix Q for examples and Appendix H for the quality verification protocol.

PLICITLY excluding any statistical summaries of that series to ensure models perform genuine statistical inference; iii) a **visual representation** via a line-plot image that enables VLMs to utilize visual trend cues and pattern recognition; and iv) a standardized **instruction template** specifying the desired style and scope of the output (see Appendix N.2). Crucially, the task is not limited to identifying trends, patterns, or summary statistics, but evaluates whether models can synthesize numeric evidence, visual cues, and semantic context into a coherent *free-form* natural language narrative that reflects not only what information is selected, but how it is composed and whether the resulting narrative is informative and correct, rather than producing templated or fragmentary observations.

3.5 Task 2: Diagnostic Q&A

To complement the main captioning task, we introduce a suite of five multiple-choice Q&A tasks designed to isolate and probe specific reasoning bottlenecks. These tasks are automatically derived from our base data triplets, with distractors designed to expose model failures in numeric precision and multimodal alignment. To ensure a high level of difficulty, we initially generated a pool of 4k questions per category and filtered them by removing those correctly answered by Qwen 2.5 Omni, focusing the benchmark on "hard" reasoning cases (see Appendix P.1 for an analysis of this

filtering). The final diagnostic set consists of 910 questions across the tasks below, with examples shown in Appendix P.4:

Time Series Matching. Models must identify the correct time series segment corresponding to a given caption from a set of distractors. These distractors are created via adversarial perturbations (shuffling, temporal reversal, Gaussian noise injection) to ensure that models cannot rely on simple matching and must instead understand the underlying temporal trend (Figure 20; details in P.2).

Caption Matching. Given a numeric series, the model must select the correct caption from a set of semantic and numeric distractors. This isolates the model’s *understanding* of descriptive language from its ability to *generate* fluent text, specifically targeting sensitivities to perturbed statistical claims (Figure 25; prompt in N.5).

Plot Matching. The model is asked to select the correct line plot for a given caption and series. It specifically measures whether the visual modality provides a complementary signal or if the model remains biased toward textual priors (Figure 26).

Comparative Reasoning. Given a pair of time series, the model must select the correct comparative statement (e.g., “Series A exhibits higher variance than Series B”) (Figures 18, 19, 21, 22). This probes the model’s ability to perform relative statistical assessment, a documented area of weakness for generative models (Merrill et al., 2024)

Temporal Matching. This task evaluates a model’s ability to perform exact temporal index alignment over time series data. In the first subtask (text-only), models are given a numeric time series as a sequence of values, along with a start timestamp and a query timestamp, and must select the correct value corresponding to the queried time. In the second subtask (image-based), the numeric series is replaced by a line plot where only the starting time is labeled, requiring the model to infer temporal offsets visually before retrieving the correct value. These subtasks isolate time indexing failures and test whether visual grounding helps or hinders temporal matching (Figures 23, 24).

3.6 Evaluation Metrics

We evaluate models using a multi-faceted framework to assess linguistic fluency, semantic alignment, and numeric grounding. Diagnostic Q&A tasks are evaluated using standard accuracy. Below, we provide an overview of the metrics, with detailed definitions in Appendix G.

Linguistic and Semantic Alignment. We report BLEU (Papineni et al., 2002), ROUGE-L (Chin-Yew, 2004), and METEOR (Banerjee and Lavie, 2005) to measure surface-level lexical overlap. To better capture temporal and semantic nuances, we additionally use embedding-based similarity metrics, namely DEBERTA SCORE (Zhang et al., 2020) and SIMCSE (Gao et al., 2021), which assess semantic consistency with human references independent of exact phrasing.

Numeric Grounding and Hallucination. A critical challenge in TSC is ensuring that language is accurately grounded in numeric reality. Standard NLP metrics are notoriously "number-blind"; a model can achieve a high BLEU score while hallucinating incorrect values. To address this, we introduce two tailored metrics:

1. **Numeric Fidelity Score:** To measure the *coverage* of numeric information, we extract all numeric values from the HR reference (excluding dates) and match them with numbers in the model’s output. We calculate *Accuracy* (precision of the matched numbers) and *Recall* (the proportion of reference numbers successfully captured). The score is a weighted harmonic combination ($\lambda_A = 0.3$, $\lambda_R = 0.7$), reflecting our priority for recall; in a decision-making context, omitting a critical data point is often more detrimental than minor rounding differences.

2. **Statistical Inference Accuracy:** This set of metrics evaluates a model’s ability to correctly infer and verbalize key statistical references (mean, STD, min, max) from raw time series. We parse statistical claims from the generated text with an LLM and verify them against the true precomputed values (details in Appendix N.7). We report the percentage of correctly identified statistics within a 5% relative error to factor out rounding errors (details on the tolerance design are in Appendix G.2). Crucially, this is a *precision-oriented* metric: models are not penalized for omitting a statistic but heavily penalized for reporting an incorrect one, thereby directly measuring hallucination. (see Appendix O.1 for details on attempt analysis)

4 Experiments

We evaluate a number of VLMs on CaTS-Bench, covering both proprietary and open-source models. For TSC, we finetune some open-source models with the synthetic training data, and additionally consider a *program-aided* (PAL) model (Gao et al., 2023). Finetuning details can be found in Appendix E. All models are prompted with the same template-based format to ensure fair comparison, avoiding task- or architecture-specific prompt engineering. Model details and PAL are in Appendix F; the human baseline is in Appendix O.

4.1 Results on Time Series Captioning

To ensure fair comparison across the domains, we report macro-averaged scores for each metric, mitigating sample size imbalances, as some domains contain more data, and preventing any domain from disproportionately influencing the results.

Our TSC results, shown in Table 3, reveal several key insights. First, **proprietary models** consistently outperform pretrained open-source counterparts, with newer models generally establishing a stronger baseline. However, **finetuned models** demonstrate substantial improvements across nearly all metrics, with several open-source models narrowing the gap to proprietary baselines, showcasing the advantage of finetuning on our quality-verified synthetic data. Most notably, **Idfics 2** stands out as the most effective learner: upon finetuning, it achieves the strongest performance among open-source models on several linguistic metrics and substantially narrows the gap to proprietary baselines, achieving, for instance, a DeBERTa

Table 3: Evaluation results of generated captions against HR and synthetic (Synth) ground truths. The "Numeric Metrics" block comprises our tailored metrics (note: *Numeric*: numeric fidelity score, *Mean Inf.* = accuracy of detecting the time series mean). **Bolded** and underlined scores denote first and second places.

	Model	Lexical and Semantic Metrics										Numeric Metrics							
		DeBERTa F1		SimCSE		BLEU		ROUGE-L		METEOR		Numeric		Mean Inf.		Max Inf.		Min Inf.	
		HR	Synth	HR	Synth	HR	Synth	HR	Synth	HR	Synth	HR	Synth	HR	Synth	HR	Synth	HR	Synth
Propri.	Gemini 2.0 Flash	<u>0.694</u>	<u>0.717</u>	0.884	<u>0.907</u>	<u>0.113</u>	<u>0.173</u>	<u>0.283</u>	<u>0.340</u>	0.304	<u>0.378</u>	0.757	0.795	0.894	0.894	0.991	0.991	0.985	0.985
	Claude 3 Haiku	0.676	0.697	0.856	0.873	0.086	0.137	0.261	0.305	0.286	0.349	0.669	0.712	0.842	0.841	0.983	0.983	0.981	0.980
	GPT-4o	0.685	0.702	0.886	0.903	0.090	0.128	0.259	0.300	0.314	0.374	0.739	0.772	<u>0.921</u>	<u>0.921</u>	0.998	0.998	0.995	0.995
Pretrained	InternVL 2.5 38B	0.682	0.704	0.867	0.889	0.085	0.144	0.262	0.314	0.294	0.367	0.740	0.790	0.865	0.865	0.983	0.982	0.978	0.978
	LLaVA v1.6	0.641	0.655	0.796	0.814	0.064	0.088	0.234	0.267	0.248	0.290	0.597	0.618	0.488	0.489	0.792	0.791	0.614	0.614
	Idefics 2	0.615	0.633	0.760	0.787	0.035	0.062	0.213	0.255	0.169	0.206	0.514	0.528	0.460	0.460	0.689	0.688	0.639	0.639
	SmolVLM	0.572	0.586	0.691	0.709	0.029	0.053	0.195	0.230	0.169	0.203	0.508	0.540	0.476	0.494	0.860	0.860	0.647	0.647
	QwenVL	0.638	0.656	0.809	0.832	0.060	0.095	0.229	0.268	0.239	0.282	0.544	0.564	0.496	0.496	0.788	0.789	0.601	0.601
	QwenVL PAL	0.643	0.660	0.814	0.833	0.060	0.108	0.241	0.287	0.224	0.273	0.637	0.656	0.985	0.985	<u>0.992</u>	<u>0.992</u>	<u>0.993</u>	<u>0.993</u>
	Llama 3.2 Vision	0.569	0.574	0.758	0.767	0.039	0.054	0.200	0.222	0.226	0.256	0.665	0.701	0.554	0.551	0.890	0.889	0.781	0.781
	Gemma 3 27B	0.680	0.696	0.883	0.898	0.089	0.111	0.255	0.289	<u>0.307</u>	0.356	0.745	0.778	0.816	0.816	0.984	0.983	0.973	0.973
	Phi-4 Multimodal Instruct	0.650	0.666	0.828	0.849	0.071	0.113	0.247	0.291	0.259	0.310	0.652	0.681	0.617	0.618	0.875	0.876	0.824	0.825
	InternVL 2.5 8B	0.666	0.686	0.856	0.875	0.078	0.123	0.247	0.293	0.264	0.319	0.705	0.732	0.731	0.731	0.985	0.985	0.933	0.934
Gemma 3 12B	0.677	0.695	0.871	0.888	0.093	0.110	0.257	0.289	0.294	0.338	0.715	0.748	0.767	0.765	0.967	0.966	0.949	0.949	
Finetuned	LLaVA v1.6	0.644	0.659	0.802	0.819	0.064	0.090	0.234	0.269	0.250	0.292	0.589	0.613	0.500	0.498	0.827	0.826	0.615	0.615
	Idefics 2	0.713	0.752	<u>0.885</u>	0.914	0.131	0.278	0.325	0.442	0.303	0.427	<u>0.748</u>	<u>0.791</u>	0.859	0.859	0.972	0.972	0.937	0.937
	QwenVL	0.678	0.705	0.863	0.888	0.090	0.166	0.257	0.324	0.272	0.357	0.669	0.707	0.744	0.745	0.977	0.978	0.942	0.943
	SmolVLM	0.601	0.616	0.751	0.771	0.056	0.105	0.224	0.273	0.197	0.243	0.604	0.632	0.433	0.437	0.796	0.796	0.606	0.606
	Llama 3.2 Vision	0.672	0.686	0.851	0.867	0.087	0.130	0.266	0.309	0.295	0.352	0.740	0.761	0.711	0.710	0.942	0.942	0.877	0.876
	Phi-4 Multimodal Instruct	0.664	0.682	0.862	0.880	0.078	0.117	0.243	0.285	0.290	0.349	0.703	0.743	0.651	0.651	0.853	0.855	0.817	0.817

F_1 score of 0.713 and a ROUGE-L of 0.325 against human ground truth. Although its gains in statistical inference are the most pronounced among finetuned models, it still lags behind the strongest proprietary systems in the *Numeric* category. In contrast, **GPT-4o** exhibits exceptional proficiency in statistical inference, followed closely by **QwenVL PAL**. The superior performance of the PAL variant over standard *QwenVL* highlights the substantial benefit of program-aided reasoning and code execution for accurate numerical tasks. Additionally, these results confirm that larger open-source models generally outperform their smaller counterparts.

Crucially, lexical and semantic scores are consistently lower when evaluated against HR ground truth than against synthetic ground truth. This consistent performance gap indicates that HR captions encode richer and less templatic reasoning patterns, and that the gains from finetuning primarily stem from exposure to higher-quality, more human-like data rather than increased data scale. Additionally, we evaluate some newer models only on the human-rewritten set in Section J and provide an information coverage analysis in Section K, showing how finetuning improves the inclusion of core time series reasoning elements.

4.2 Results on Q&A Tasks

Figure 3 summarizes model performance on our Q&A tasks (see Table 24 for detailed results).


Models show no consistent winner across tasks.

Binary time series comparisons are easier, likely due to the narrower decision space, while caption matching is harder, and plot matching is hardest. Proprietary models lead, and Phi-4 M.I. is the strongest open-source model overall. Humans achieve the highest overall scores and are nearly perfect on plot matching, where all models remain near random. Despite the tasks' apparent simplicity, they reveal fundamental limitations in VLMs' temporal reasoning capabilities, which suggests the need to address basic time series understanding before tackling more complex applications.


Temporal matching also unveils significant weaknesses, with results shown in Figure 4. Models perform near chance on long sequences ($\approx 20\text{--}35\%$), with accuracy degrading as the queried time moves further from the series start in most models, while short sequences (length = 6) provide higher accuracy ($\approx 50\text{--}57\%$). These results indicate that current language models struggle with exact index-based temporal reasoning from text, and increased model scale does not reliably mitigate this limitation. See Appendix P.3 for more information about temporal buckets, baseline, and detailed results.

Modality Ablations for Temporal Matching.

Table 4 compares model accuracy on *mid2* Temporal Matching questions when the same underlying time series is provided as numeric text, as a plotted image, or as both image and text. Image-only input yields higher accuracy than text-only input for most models and often reaches the highest performance



(a) Proprietary VLMs



(b) Open-source VLMs

Figure 3: Model accuracy across Q&A sub-tasks (except Temporal Matching). Proprietary models perform best, open-source models lag behind in all tasks.

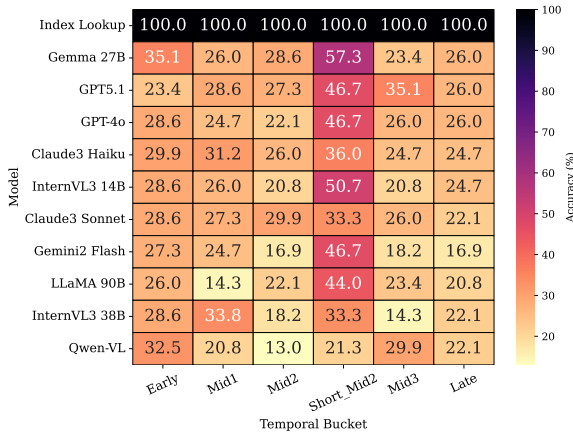


Figure 4: MCQ accuracy (%) for exact temporal value lookup at different query positions. *Short_Mid2* denotes sequences of length 6; other columns use long time series with queries at different temporal offsets. Results are averaged over 75 questions per temporal bucket.

for a given model, suggesting that visual structure strongly supports coarse temporal localization.

However, adding numeric text alongside the image frequently degrades performance, indicating that many models struggle to jointly integrate visual and textual representations of the same time series. This multimodal interference is evident for models such as *Gemini 2.0 Flash*, *GPT-4o*, and *LLaMA-90B*, where image-only accuracy exceeds image-and-text performance.

Notably, only the newer *InternVL3* (Zhu et al., 2025) models, particularly the larger InternVL variants, consistently benefit from multimodal input and achieve the highest accuracies in the combined setting. These results suggest that while vision

Table 4: MCQ accuracy on *Short-Mid2* questions when the time series is provided as numeric text, as a plotted image, or as both image and text. Vision input improves performance for most models, but does not resolve failures in exact temporal indexing.

Model	Text Only	Image Only	Image+Text
Gemini 2.0 Flash	46.7	57.3	34.7
Claude 3 Haiku	33.3	42.7	38.7
GPT-4o	46.7	60.0	57.3
GPT-5.1	46.7	58.7	54.7
InternVL3 14B	50.7	60.0	74.7
InternVL3 38B	33.3	46.7	76.0
LLaMA 90B	44.0	44.0	38.7
Gemma 27B	57.3	58.7	58.7

alone can aid approximate temporal reasoning, effective fusion of redundant visual and numeric signals remains a challenge for most models, with successful multimodal integration emerging only in more recent architectures.

4.3 Role of the Visual Modality

Visual Modality Ablation. We perform a modality removal experiment by stripping away the time series plot and providing only the associated textual metadata and the numeric values of the time series. This quantifies the contribution of the visual channel and enables a better understanding of the model’s captioning performance. We evaluate a selected subset of pretrained baselines to assess their intrinsic reliance on vision.

Our experiments show that visual input contributes little to caption quality for most models. As Figure 5 illustrates, removing the plot image produces only minor changes in both semantic and numeric performance, with numeric accuracy de-




Figure 5: Performance change when adding visual input. Blue indicates improvement; red indicates degradation.

grading modestly and sometimes even improving without the plot. This suggests that models rely mainly on textual priors rather than grounded interpretation of the plotted signal, so caption generation is driven primarily by language-level analysis of typical time series behaviors.

The weak sensitivity to image removal indicates that visual input is only weakly used during generation, for both semantic and numeric metrics, implying that vision plays a limited role in numeric grounding. As discussed in Section I.2, this phenomenon is not restricted to line plots: more expressive encodings such as Gramian Angular Fields and recurrence plots likewise fail to elicit strong visual reasoning from current VLMs in TSC.




Figure 6: Performance deltas on statistical inference with and without annotations. Blue indicates improvement; red degradation.

Figure 6 isolates the effect of numeric annotations on the plot, while keeping the visual modality fixed. In contrast to the full-versus-no-plot comparison, the impact here is substantially larger. When time series values are explicitly annotated, models show clear improvements in fine-grained numeric metrics, particularly Max, Min, and Mean. Removing these annotations while retaining the plotted curve leads to pronounced drops in numeric fidelity. This pattern indicates that models benefit primarily from explicit numeric text embedded in the image, rather than from interpreting the visual shape of the time series curve. The relatively small and inconsistent changes in Std further suggest that dispersion is not reliably inferred from visual structure alone.

Visual Attention Analysis. To better understand how VLMs use plots during caption generation, we examine their attention maps and find minimal visual grounding: models focus mainly on textual elements (e.g., axis labels and titles) with limited attention to the actual line trends. Attention to visual patterns is weak and inconsistent, suggesting that learned parameters largely ignore visual cues in favor of textual priors, highlighting a gap between nominal multimodal input and actual integration (full results in I.1 and Figure 9).

These findings indicate that current VLM pre-training paradigms are ill-suited for structured visual reasoning over time series plots: despite strong performance on natural images, their objectives align poorly with abstract, low-level visual encodings that require relational, temporal, and quantitative interpretation rather than object recognition. Consequently, models default to textual priors and explicit numeric cues, largely ignoring visual structure, even though CaTS-Bench explicitly provides both plots and rich metadata to support multimodal reasoning. That most models fail to leverage visual signals exposes a critical gap in current VLM capabilities and motivates future work on architectures that more effectively integrate plot-based information with textual and numeric cues.

5 Conclusion

We presented **CaTS-Bench**, a multimodal benchmark for TSC and time series reasoning that emphasizes linguistic expressiveness, numeric fidelity, and grounded multimodal alignment. By combining a human-rewritten gold standard with a scalable synthetic augmentation pipeline, CaTS-Bench enables rigorous evaluation of free-form narrative generation over time series while addressing data scarcity. Our results show that proprietary VLMs provide strong zero-shot baselines, while finetuning open-source models on synthetic data yields substantial gains in TSC. However, zero-shot VLMs exhibit systematic failures such as hallucinated statistics in free-form captioning. Diagnostic Q&A tasks further reveal persistent weaknesses in exact temporal indexing and multimodal integration, indicating that the current VLM pretraining paradigm is ill-designed for time series reasoning. Overall, CaTS-Bench exposes key limitations in grounded time series understanding and provides a foundation for developing more reliable multimodal generative models in numeric domains.

Limitations

While CaTS-Bench represents a significant step toward multimodal, context-aware time series understanding, it also has limitations that suggest clear avenues for future work.

First, the training captions are synthetic and generated by a single oracle model (Gemini 2.0 Flash). Our extensive validation studies confirm their factual reliability and linguistic diversity; however, reliance on a single oracle may still introduce subtle modeling biases during finetuning. As discussed earlier, this design choice is intended as a proof of concept to demonstrate the feasibility and effectiveness of scalable augmentation. Extending the pipeline to multiple oracles is straightforward but beyond the scope of this work. Although we release a small subset of paraphrases generated by different LLMs, its limited scale highlights opportunities for future extensions.

Second, although we provide a human-rewritten test set, the scale of expert-verified content remains limited. We acknowledge that captions written by domain experts can offer deeper insights, and future work could incorporate expert-authored annotations across domains such as economics, healthcare, or climate data. At the same time, human authors introduce their own stylistic biases, which should be considered when designing and evaluating benchmarks.

Overall, we view CaTS-Bench as a scalable foundation rather than a fixed resource, with ample room to grow through richer human input, multi-oracle augmentation, and extended coverage of temporal reasoning tasks.

Ethical Considerations

The development of CaTS-Bench was guided by a commitment to ethical research practices. All datasets used in this work are publicly available and do not contain personally identifiable information (PII). The domains, such as climate, public health, and agriculture, were chosen for their public relevance and data accessibility. Our use of an oracle LLM for data augmentation was a deliberate design choice to ensure scalability. We have taken extensive measures to validate the quality of the synthetic captions, gauging their factual accuracy and diversity, as detailed in Section 3.3 and Appendix H. Our human-rewritten test set is also an attempt to further ensure evaluation reliability. For our human Q&A evaluation studies, all participa-

tion was voluntary. We obtained informed consent from all participants, who were university students. The study’s purpose was clearly communicated, and all responses were collected anonymously to protect participant privacy, as shown in an example of a consent form in Appendix O.

LLM Usage Statement

LLMs played a central role in multiple stages of this work.

1. LLMs were employed as **data generators**, producing synthetic captions that mainly serve as training data.
2. LLMs were employed as **data extractors**, for example to parse statistical claims from captions during our evaluation analyses.
3. LLMs and VLMs served as **baselines** in our experiments as captioning models for evaluation.
4. LLMs were employed as a **writing assist tool** to polish the presentation of the paper.

LLMs did not contribute to research ideation or decision-making. All factual claims, analyses, and conclusions are the responsibility of the authors.

Acknowledgments

This work was supported in part by the U.S. Army Research Office under Army-ECASE award W911NF-07-R-0003-03, the U.S. Department Of Energy, Office of Science, IARPA HAYSTAC Program, and NSF Grants #2205093, #2146343, #2134274, #2441832, CDC-RFA-FT-23-0069, DARPA AIE FoundSci and DARPA YFA. We also thank the anonymous reviewers for their constructive feedback, which significantly improved the clarity and presentation of this work.

References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.
- Abdul Fatir Ansari, Lorenzo Stella, Ali Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin

- Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, Jasper Zschiegner, Danielle C. Maddix, Hao Wang, Michael W. Mahoney, Kari Torkkola, Andrew Gordon Wilson, Michael Bohlke-Schneider, and Bernie Wang. 2024. [Chronos: Learning the language of time series](#). *Transactions on Machine Learning Research*. Expert Certification.
- Anthropic. 2024. [The claude 3 model family: Opus, sonnet, haiku](#). Model card, Anthropic.
- Saad Aziz. [Population collapse time series data of the world](#).
- Anthony Bagnall, Hoang Anh Dau, Jason Lines, Michael Flynn, James Large, Aaron Bostrom, Paul Southam, and Eamonn Keogh. 2018. The uea multivariate time series classification archive, 2018. *arXiv preprint arXiv:1811.00075*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Yifu Cai, Arvind Srinivasan, Mononito Goswami, Arjun Choudhry, and Artur Dubrawski. 2024. [Jolt: Jointly learned representations of language and time-series for clinical time-series interpretation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 23447–23448. AAAI Student Abstract and Poster Program.
- California Department of Public Health. n.d. [Road traffic injuries narrative](#).
- Rui Cao and Qiao Wang. 2024. An evaluation of standard statistical models and llms on time series forecasting. *arXiv preprint arXiv:2408.04867*.
- Ching Chang, Wen-Chih Peng, and Tien-Fu Chen. 2023. Llm4ts: Two-stage fine-tuning for time-series forecasting with pre-trained llms. *arXiv preprint arXiv:2308.08469*.
- Georgios Chatzigeorgakidis, Konstantinos Lentzos, and Dimitrios Skoutas. 2024. Multicast: Zero-shot multivariate time series forecasting using llms. In *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*, pages 119–127. IEEE.
- Daqing Chen. 2015. [Online Retail](#). UCI Machine Learning Repository.
- Mouxian Chen, Lefei Shen, Zhuo Li, Xiaoyun Joy Wang, Jianling Sun, and Chenghao Liu. 2024. [Visions: Visual masked autoencoders are free-lunch zero-shot time series forecasters](#). *arXiv preprint arXiv:2408.17253*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yiming Ren, Zixuan Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, and 23 others. 2025. [Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling](#). *Preprint*, arXiv:2412.05271.
- Lin Chin-Yew. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of the Workshop on Text Summarization Branches Out, 2004*.
- City of Los Angeles. n.d. [Crime data from 2020 to present](#).
- Hoang Anh Dau, Eamonn Keogh, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, Yanping, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, Gustavo Batista, and Hexagon-ML. 2018. [The ucr time series classification archive](#).
- Kota Dohi, Aoi Ito, Harsh Purohit, Tomoya Nishida, Takashi Endo, and Yohei Kawaguchi. 2025. Domain-independent automatic generation of descriptive texts for time-series data. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Jiaxiang Dong, Haixu Wu, Yuxuan Wang, Li Zhang, Jianmin Wang, and Mingsheng Long. 2024. [Meta-data matters for time series: Informative forecasting with transformers](#). *arXiv preprint arXiv:2410.03806*.
- European Centre for Disease Prevention and Control. 2024. [Download today's data on the geographic distribution of covid-19 cases worldwide](#).
- Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso, and Svitlana Vyetrenko. 2024a. [Evaluating large language models on time series feature understanding: A comprehensive taxonomy and benchmark](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21598–21634.
- Elizabeth Fons, Rachneet Kaur, Zhen Zeng, Soham Palande, Tucker Balch, Svitlana Vyetrenko, and Manuela Veloso. 2024b. [Tadacap: Time-series adaptive domain-aware captioning](#). In *Proceedings of the 5th ACM International Conference on AI in Finance*, pages 54–62.
- Food and Agriculture Organization of the United Nations. 2024. [Faostat - food balance sheets](#).
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Pal: Program-aided language models](#). *Preprint*, arXiv:2211.10435.

- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple contrastive learning of sentence embeddings. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, and 1332 others. 2025. **Gemini: A family of highly capable multimodal models**. *Preprint*, arXiv:2312.11805.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. **Gemma 3 technical report**. *Preprint*, arXiv:2503.19786.
- Rakshitha Godahewa, Christoph Bergmeir, Geoffrey I. Webb, Rob J. Hyndman, and Pablo Montero-Manso. 2021. Monash time series forecasting archive. In *Neural Information Processing Systems Track on Datasets and Benchmarks*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. **The llama 3 herd of models**. *Preprint*, arXiv:2407.21783.
- Yasser Hassan. 2021. **Walmart dataset**.
- Abhishek S. Jha. 2023. **Time series air quality data of india (2010–2023)**.
- Harsh Jhamtani and Taylor Berg-Kirkpatrick. 2021. Truth-conditional captioning of time series data. In *EMNLP*.
- Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen. 2024. Time-LLM: Time series forecasting by reprogramming large language models. In *International Conference on Learning Representations (ICLR)*.
- Kai Kim, Howard Tsai, Rajat Sen, Abhimanyu Das, Zihao Zhou, Abhishek Tanpure, Mathew Luo, and Rose Yu. 2024. Multi-modal forecaster: Jointly predicting time series and textual data. *arXiv preprint arXiv:2411.06735*.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. What matters when building vision-language models? *Advances in Neural Information Processing Systems*, 37:87874–87907.
- Chen Liu, Shibo He, Qihang Zhou, Shizhong Li, and Wenchao Meng. 2024a. Large language model guided knowledge distillation for time series anomaly detection. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 2162–2170.
- Chenxi Liu, Qianxiong Xu, Hao Miao, Sun Yang, Lingzheng Zhang, Cheng Long, Ziyue Li, and Rui Zhao. 2024b. Timecma: Towards llm-empowered time series forecasting via cross-modality alignment. *CoRR*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning.
- Haoxin Liu, Harshvardhan Kamarthi, Zhiyuan Zhao, Shangqing Xu, Shiyu Wang, Qingsong Wen, Tom Hartvigsen, Fei Wang, and B Aditya Prakash. 2025. How can time series analysis benefit from multiple modalities? a survey and outlook. *arXiv preprint arXiv:2503.11835*.
- Peiyuan Liu, Hang Guo, Tao Dai, Naiqi Li, Jigang Bao, Xudong Ren, Yong Jiang, and Shu-Tao Xia. 2024c. Calf: Aligning llms for time series forecasting via cross-modal fine-tuning. *arXiv preprint arXiv:2403.07300*.
- Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro von Werra, and Thomas Wolf. 2025. Smolvlm: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*.
- Mike Merrill, Mingtian Tan, Vinayak Gupta, Thomas Hartvigsen, and Tim Althoff. 2024. Language models still struggle to zero-shot reason about time series. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3512–3533.
- OpenAI. 2025. **Gpt-5.1 technical report and system card addendum**. Accessed: 2026-01-05.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. **Gpt-4 technical report**. *Preprint*, arXiv:2303.08774.
- Zijie Pan, Yushan Jiang, Sahil Garg, Anderson Schneider, Yuriy Nevmyvaka, and Dongjin Song. 2024. s^2 ip-llm: Semantic space informed prompt learning with llm for time series forecasting. In *Forty-first International Conference on Machine Learning*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

- Hannah Ritchie. 2021. Many countries have decoupled economic growth from co2 emissions, even if we take offshored production into account. *Our World in Data*. <https://ourworldindata.org/co2-gdp-decoupling>.
- Medhasweta Sen, Zachary Gottesman, Jiaying Qiu, C. Bayan Bruss, Nam Nguyen, and Tom Hartvigsen. 2025. *Bedtime: A unified benchmark for automatically describing time series*. *Preprint*, arXiv:2509.05215.
- Chenxi Sun, Hongyan Li, Yaliang Li, and Shenda Hong. 2023. Test: Text prototype aligned embedding to activate llm’s ability for time series. In *The Twelfth International Conference on Learning Representations*.
- Mingtian Tan, Mike A Merrill, Vinayak Gupta, Tim Althoff, and Thomas Hartvigsen. 2024. *Are language models actually useful for time series forecasting?* In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Benny J Tang, Angie Boggust, and Arvind Satyanarayan. 2023. Vistext: A benchmark for semantically rich chart captioning. *arXiv preprint arXiv:2307.05356*.
- Hua Tang, Chong Zhang, Mingyu Jin, Qinkai Yu, Zhen-ting Wang, Xiaobo Jin, Yongfeng Zhang, and Mengnan Du. 2025. Time series forecasting with llms: Understanding and enhancing model capabilities. *ACM SIGKDD Explorations Newsletter*, 26(2):109–118.
- U.S. Department of Transportation, Bureau of Transportation Statistics. n.d. *Border crossing entry data*.
- USDA Economic Research Service. 2024. *International agricultural productivity*.
- Xinlei Wang, Maik Feng, Jing Qiu, Jinjin Gu, and Junhua Zhao. 2024. From news to forecast: Integrating event analysis in llm-based time series forecasting with reflection. *Advances in Neural Information Processing Systems*, 37:58118–58153.
- J. Nan Wilkenfeld, Bei Yan, Jujun Huang, Guirong Luo, and Kristina Algas. 2022. “AI Love You”: Linguistic Convergence in Human-Chatbot Relationship Development. In *Academy of Management Proceedings*, volume 2022, page 17063. Academy of Management.
- Andrew Robert Williams, Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Jithendaraa Subramanian, Roland Riachi, James Requeima, Alexandre Lacoste, Irina Rish, Nicolas Chapados, and Alexandre Drouin. 2024. *Context is key: A benchmark for forecasting with essential textual information*. *Preprint*, arXiv:2410.18959.
- Hao Xue and Flora D Salim. 2023. Promptcast: A new prompt-based learning paradigm for time series forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 36(11):6851–6864.
- Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 11121–11128.
- Jingyang Zhang. 2024. *Visualizing the attention of vision-language models*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.
- Xiyuan Zhang, Ranak Roy Chowdhury, Rajesh K Gupta, and Jingbo Shang. 2024. Large language models for time series: a survey. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 8335–8343.
- Tian Zhou, Peisong Niu, Xue Wang, Liang Sun, and Rong Jin. 2023. One Fits All: Power general time series analysis by pretrained lm. In *NeurIPS*.
- Zihao Zhou and Rose Yu. 2025. *Can LLMs understand time series anomalies?* In *The Thirteenth International Conference on Learning Representations*.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. *Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models*. *Preprint*, arXiv:2504.10479.

Appendix

A Motivation and Utilities of TSC

An interesting question to think about is whether TSC is a meaningful task at all, given that many properties of a time series (e.g., means, extrema, trends) can be computed exactly using deterministic algorithms or statistical tools. While such tools are indeed precise, they fundamentally address a different problem. TSC is not concerned with computing statistics, but with generating *free-form natural language descriptions* that summarize, contextualize, and communicate temporal data in a human-interpretable way.

In many real-world settings, users do not interact directly with raw statistics or code outputs, but instead rely on narrative descriptions. Business reports, financial dashboards, scientific summaries, ECG reports, policy briefs, and automated analytics systems routinely present trends and patterns through textual explanations accompanying plots. In these contexts, the goal is not merely to compute correct values, but to decide *what is salient, how to describe it, and how to express uncertainty, comparison, or change over time* in natural language. These decisions are inherently under-specified and context-dependent, making them unsuitable for purely deterministic pipelines.

LLMs are, therefore, a natural candidate for this task, as they are designed to produce fluent, context-aware descriptions. However, our results demonstrate that current models struggle precisely at this interface between numeric data and free-form generation. We observe frequent hallucination of numeric values, incorrect temporal alignment between values and timestamps, and failures in simple comparative reasoning. Moreover, when strong textual cues are present, models often ignore the underlying plot, generating plausible but weakly grounded descriptions driven by language priors rather than data. These failure modes are not exposed by deterministic statistical tools, which are guaranteed to be correct by construction, but are critical in generative settings where outputs are unconstrained and must be trusted by downstream users. The fact that calculators or statistical algorithms can compute accurate summaries does not obviate the need to evaluate how models *describe* those summaries in natural language. Instead, it highlights the importance of studying TSC as a distinct problem, one that probes a model’s ability to integrate numeric reasoning, temporal structure,

and linguistic generation.

From this perspective, TSC serves as a diagnostic task that reveals the limits of current LLMs and VLMs in grounded free-form generation. Our findings do not argue against deterministic methods, but rather show that replacing human-written narrative descriptions with model-generated captions with direct prompting remains challenging. Understanding these limitations is essential for developing reliable systems that can safely and effectively communicate temporal data in real-world settings.

B Source Datasets

1. **Air Quality** – Hourly air pollution data from 453 Indian cities (2010–2023), covering 30+ parameters including PM_{2.5}, NO_x, CO, and SO₂, compiled from CPCB (Jha, 2023).
2. **Border Crossing** – Monthly inbound border crossing counts at U.S.-Mexico and U.S.-Canada ports, disaggregated by transport mode and collected by U.S. Customs and Border Protection (U.S. Department of Transportation, Bureau of Transportation Statistics, n.d.).
3. **Crime** – Incident-level crime reports in Los Angeles from 2020 onward, provided by LAPD OpenData and updated biweekly, including NIBRS-compliant records (City of Los Angeles, n.d.).
4. **Demography** – Annual global indicators from the UN and World Bank (2000–2021) covering population growth, fertility, life expectancy, death rates, and median age to assess patterns of demographic change and collapse (Aziz).
5. **Injury** – Annual counts of fatal and severe road traffic injuries in California (2002–2010), disaggregated by transport mode and geography, from CDPH’s Healthy Communities Indicators (California Department of Public Health, n.d.).
6. **COVID** – Global daily COVID-19 case and death counts (2020), compiled by ECDC, covering over 200 countries with population-adjusted metrics (European Centre for Disease Prevention and Control, 2024).
7. **CO₂** – National-level per capita CO₂ emissions and GDP trends from Our World in Data, adjusted for trade (consumption-based), spanning 1990–2023 (Ritchie, 2021).

8. **Diet** – Food supply and caloric intake patterns from FAO Food Balance Sheets (Food and Agriculture Organization of the United Nations, 2024).
9. **Walmart** – Weekly sales data from 45 Walmart stores (2010–2012), enriched with features like temperature, fuel price, CPI, unemployment rate, and holiday flags (Hassan, 2021).
10. **Retail** – Transactional records from a UK-based online gift retailer (2010–2011), capturing item-level purchases, cancellations, and customer behavior (Chen, 2015).
11. **Agriculture** – Annual agricultural total factor productivity (TFP) indices from USDA for 1961–2022, covering outputs and inputs like land, labor, capital, and materials across countries (USDA Economic Research Service, 2024).

C Time Series Segment Cropping

Our cropping strategy balances diversity with consistency across datasets. Many source time series (e.g., 50 years of hourly CO₂ emissions) are too long to process directly, so we sample random windows of variable lengths. Each window length is drawn from a dataset-specific range $[min, max]$, with the maximum based on the original series length. This ensures that cropped windows preserve the scale and structure of the data while introducing sufficient variability for training and evaluation. We summarize these rules in Table 5 below.

Table 5: Minimum and maximum segment lengths for each dataset.

Source Dataset	Min Length	Max Length
Air Quality, Crime, Border Crossing, CO ₂ , Walmart, Agriculture	5	$\min(150, 5 + \lfloor \frac{\text{original_length}}{8} \rfloor)$
Demography, Online Retail	5	original_length
Road Injuries	3	$\min(3, 0.2 \times \text{original_length})$
COVID	5	$\min(150, 5 + \lfloor \frac{\text{original_length}}{5} \rfloor)$
Calories	5	$\min(6, 0.2 \times \text{original_length})$

D Why Metadata?

A central challenge in synthetic data generation is ensuring factual fidelity. Free-form synthetic captions produced by unconstrained LLMs are often unreliable as they may hallucinate numeric values, introduce unsupported context, or mischaracterize

extrema and temporal trends, thereby corrupting both supervision and evaluation. Such issues are especially problematic in time series captioning, where even minor numeric inaccuracies can significantly distort semantic interpretation.

To address this, our synthetic captioning pipeline is not based on open-ended prompting. Instead, the oracle is strictly grounded in three sources: (i) the underlying numeric time series, (ii) the corresponding plot, and (iii) enriched metadata capturing domain context and structural attributes. This grounding constrains generation and substantially reduces hallucination.

We validate the importance of metadata grounding through a controlled ablation study. When captions are generated without metadata and manually evaluated for statistical correctness, temporal consistency, and hallucinated values, factual accuracy drops to 10% (5/50). In contrast, captions generated with metadata grounding achieve 98% accuracy (49/50). This performance gap highlights that metadata is not merely auxiliary, but a critical component for ensuring high-fidelity synthetic supervision.

E Hardware and Settings

All experiments were conducted on a high-performance computing node featuring two AMD EPYC 7453 processors, providing a total of 56 logical CPUs, and 125 GB of RAM (with over 117 GB available during runtime). For GPU acceleration, the system includes eight NVIDIA A100 GPUs - six PCIe 80 GB models and two PCIe 40 GB models - along with an ASPEED graphics controller used for display purposes. This configuration offers ample computational and memory resources suitable for mid- to large-scale deep learning training and inference. The models we finetune range in size from 2 billion to 11 billion parameters, with finetuning times spanning from a few hours to a day.

For finetuning, we adopt a unified training strategy guided by best practices in instruction tuning for multimodal inputs (see Table 6). All models are trained using the AdamW optimizer with a cosine learning rate scheduler and a base learning rate of 2×10^{-5} . We apply gradient accumulation to simulate a larger batch size. Mixed precision training and gradient checkpointing are enabled for memory efficiency. Low Rank Adaptation (LoRA) is used to adapt large models by tuning a small subset of parameters, while keeping the rest of the model

Table 6: Finetuning configuration details

Hyperparameter	Value
Batch size	4
Gradient Accumulation Steps	12
Training Epochs	3
Learning Rate	2×10^{-5}
Learning Rate Scheduler	Cosine
Optimizer	AdamW
Numerical Precision	bfloat16
LoRA Rank	8 or 16
Dropout Probability	0.05
Input Image Resolution	224–560

frozen or partially frozen. To ensure deterministic and focused generation, we use a temperature of 0.3 during inference across all evaluated models. Each model is finetuned using a structured JSONL dataset comprising time series plot images and corresponding image-grounded chat-style conversations. We preprocess data with each model’s native processor and apply minimal resizing to maintain fidelity in the visual input. Special care is taken to exclude padding and <image> tokens from loss computation by assigning them an ignore index.

F Baseline Models

For TSC, we evaluate Gemini 2.0 Flash (Gemini Team et al., 2025), Claude 3 Haiku (Anthropic, 2024), GPT-4o (OpenAI et al., 2024), GPT-5.1 (OpenAI, 2025), InternVL 2.5 (8b & 38b) (Chen et al., 2025), LLaVA v1.6 Mistral 7b (default) and 34b (Liu et al., 2023), Phi-4 Multimodal Instruct 5.6b (Abdin et al., 2024), Idefics 2 (8b) (Laurençon et al., 2024), SmolVLM (2b) (Marafioti et al., 2025), QwenVL (7b) (Bai et al., 2023), Llama 3.2 Vision (11b) (Grattafiori et al., 2024), and Gemma 3 (12b & 27b) (Gemma Team et al., 2025) for both TSC and Q&A tasks.

TSC requires precise numeric reasoning alongside text generation, making it suitable for program-aided language (PAL) models (Gao et al., 2023). Hence, we also evaluate QwenVL 32b by prompting it to generate a Python program that outputs the full time series caption. The program is executed in Python, and its return value is taken as the caption. Most (90%) generated programs succeed on the first attempt; if a program fails, we increase the token limit and regenerate until successful. The full prompt example can be found in Appendix N.4.

G Evaluation Metrics

G.1 Linguistic Metrics

DeBERTa Score The DeBERTa SCORE is a contextual similarity metric based on cosine similarity between contextual embeddings of tokens in the candidate (c) and reference (r) captions. Given token embeddings from the DeBERTa encoder, the metric computes token-level precision, recall, and F1:

$$F1_{\text{DeBERTa}} = \frac{2PR}{P+R},$$

$$\text{where } P = \frac{1}{|c|} \sum_{i \in c} \max_{j \in r} \cos(\mathbf{e}_i, \mathbf{e}_j), \quad (1)$$

$$R = \frac{1}{|r|} \sum_{j \in r} \max_{i \in c} \cos(\mathbf{e}_j, \mathbf{e}_i).$$

where \mathbf{e}_i and \mathbf{e}_j are the contextual embeddings of candidate and reference tokens, respectively.

BLEU BLEU evaluates n-gram overlap between a candidate caption and reference using precision with a brevity penalty to discourage short outputs:

$$\text{BLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^N w_n \log p_n\right), \quad (2)$$

$$\text{where } \text{BP} = \begin{cases} 1, & \text{if } c > r, \\ e^{1-r/c}, & \text{if } c \leq r, \end{cases}$$

where p_n is the modified precision for n -grams, w_n are weights (usually uniform), c is candidate length, and r is reference length.

ROUGE-L ROUGE-L measures fluency via the length of the longest common subsequence (LCS) between candidate and reference:

$$\text{ROUGE-L}_{F1} = \frac{(1 + \beta^2) \cdot \text{LCS}}{r + c} \quad (3)$$

where β balances recall and precision (often $\beta = 1$), and r and c are the reference and candidate lengths.

METEOR METEOR aligns unigrams using exact matches, stems, synonyms, and paraphrases. It then computes an F-score and applies a fragmentation penalty:

$$\begin{aligned} \text{METEOR} &= F_{\text{mean}} \cdot (1 - \text{Pen}), \\ \text{where } F_{\text{mean}} &= \frac{10 \cdot P \cdot R}{R + 9P}, \\ \text{Pen} &= 0.5 \left(\frac{\text{chunks}}{\text{matches}} \right)^3. \end{aligned} \quad (4)$$

where P and R are unigram precision and recall, and chunks refers to non-contiguous matched segments.

SimCSE SIMCSE computes semantic similarity at the sentence level using cosine similarity between sentence embeddings:

$$\text{SimCSE}(c, r) = \cos(\mathbf{h}_c, \mathbf{h}_r) = \frac{\mathbf{h}_c \cdot \mathbf{h}_r}{\|\mathbf{h}_c\| \|\mathbf{h}_r\|} \quad (5)$$

where \mathbf{h}_c and \mathbf{h}_r are candidate and reference sentence embeddings, produced by a contrastively trained RoBERTa encoder.

G.2 Tolerance Design in Numeric Metrics

We adopt a 5% relative tolerance for both numeric metrics, as it is a widely accepted threshold in numeric evaluation across data science and time series literature. This value balances sensitivity and robustness: it is tight enough to catch meaningful deviations from the true value, ensuring that significant errors are penalized, yet lenient enough to accommodate minor variations due to rounding, numeric precision, or natural approximations in model-generated captions. By using this standard threshold, our evaluation aligns with common practice while focusing on practically relevant numeric accuracy.

H Quality Validation

H.1 Manual Verification of Synthetic Captions

To assess the reliability of synthetic captions used as ground truth, we manually verified statistical and trend claims in around 2.9k captions of the test set under a three-tier scoring system: *exact* (within ± 0.05 of the true value), *near* (within ± 0.1), and *incorrect*.

Method Each claim was extracted and checked against the underlying time series metadata. We extracted statistical claims and trend patterns from captions using structured keyword clustering across

Table 7: Manual verification of SS captions. Accuracy remains high across all categories.

Task	Feature/Type	Occurrences	Acc.
Statistical	mean / average	570	0.980
	minimum / dip	310	0.994
	standard deviation	123	1.000
	maximum / peak	217	0.994
Trend	Upward	467	0.980
	Downward	328	0.997
	Stability	45	0.970
	Fluctuation	87	0.974
Historical (Trends + Stats)	mean / average	435	0.980
	standard deviation	26	1.000
	maximum / peak	95	0.994
	minimum / dip	12	1.000
	norms	164	0.980
Total / Avg.		2879	0.986

selected statistical categories (mean/average, min/max, standard deviation) and trend keywords (increasing, decreasing, plateau, fluctuation). We (1) identified terms using keyword matching, (2) extracted ground truth values therein, and (3) created verification sheets comparing claims against actual metadata for manual verification. Then, we scored each claim by comparing statements against ground truth data and verifying that trend claims matched actual trajectories.

Accuracy across both statistical descriptors and trend types is consistently high as shown in Table 7, confirming that oracle-generated captions provide factually reliable reference annotations.

H.2 Human Study on Detectability

We conducted a blind study with 35 participants, each reviewing 11 captions (roughly half written by humans, half by Gemini, using the same context information). Participants labeled each as human or AI-written but achieved only 41.1% accuracy, essentially random, suggesting **Gemini’s captions were indistinguishable from human ones**. Participation form with guidelines is in Appendix O.

H.3 Robustness of Evaluation

A legitimate concern when using a single LLM to generate reference captions is the potential for evaluation bias towards the specific linguistic style of that model. To ensure that CaTS-Bench evaluates generalizable time series understanding capabilities rather than the tendency to mimic Gemini’s linguistic style, we conducted a robustness experiment.

We systematically paraphrased a representative subset of our ground truth captions (agriculture,

Table 8: Evaluation of generated captions across paraphrased/original ground truths.

Model	GT Style	Embedding		N-gram			Numeric	Stat. Inference		
		SimCSE	DeBERTa	METEOR	ROUGE-L	BLEU	Numeric	Mean	Max	Min
Gemini 2.0 Flash	GPT-4o Phr.	0.8707	0.6803	0.2313	0.2605	0.0820	0.6715	0.3333	0.9823	0.9377
	Gemma Phr.	0.8635	0.6748	0.2068	0.2443	0.0605	0.6578	0.5357	0.9823	0.9374
	Llama Phr.	0.8680	0.6715	0.2080	0.2445	0.0710	0.6745	0.3750	0.9807	0.9351
	Original	0.8716	0.6860	0.2720	0.3068	0.1315	0.6802	0.3750	0.9823	0.9377
GPT-4o	GPT-4o Phr.	0.8752	0.6740	0.2583	0.2488	0.0845	0.6773	0.8000	0.9921	0.9393
	Gemma Phr.	0.8645	0.6665	0.2250	0.2268	0.0578	0.6673	0.8167	0.9921	0.9393
	Llama Phr.	0.8726	0.6678	0.2295	0.2380	0.0725	0.6673	0.8000	0.9921	0.9379
	Original	0.8773	0.6785	0.2880	0.2785	0.1048	0.6558	0.8000	0.9921	0.9379
Claude 3 Haiku	GPT-4o Phr.	0.8636	0.6720	0.2497	0.2480	0.0693	0.6383	0.7500	0.9797	0.9338
	Gemma Phr.	0.8563	0.6665	0.2263	0.2295	0.0495	0.6223	0.8000	0.9781	0.9321
	Llama Phr.	0.8598	0.6678	0.2320	0.2473	0.0605	0.6285	0.8333	0.9797	0.9339
	Original	0.8683	0.6795	0.2873	0.2870	0.1038	0.6333	0.7500	0.9797	0.9338
Idefics 2	GPT-4o Phr.	0.7952	0.6113	0.1463	0.2035	0.0243	0.3893	0.8056	0.8908	0.8377
	Gemma Phr.	0.7897	0.6058	0.1368	0.1915	0.0158	0.4005	0.8056	0.8908	0.8377
	Llama Phr.	0.7850	0.6045	0.1335	0.1950	0.0183	0.4198	0.8056	0.8908	0.8377
	Original	0.7962	0.6178	0.1623	0.2250	0.0380	0.4580	0.8056	0.8908	0.8377
QwenVL	GPT-4o Phr.	0.8347	0.6323	0.2250	0.2205	0.0483	0.4098	0.4375	0.7948	0.6793
	Gemma Phr.	0.8262	0.6278	0.2043	0.2025	0.0350	0.4160	0.4375	0.7926	0.6793
	Llama Phr.	0.8270	0.6303	0.2113	0.2213	0.0453	0.4470	0.4375	0.7948	0.6776
	Original	0.8427	0.6405	0.2548	0.2488	0.0798	0.4895	0.4375	0.7926	0.6776
Llama 3.2 Vision	GPT-4o Phr.	0.8663	0.6625	0.2663	0.2473	0.0795	0.6890	0.4667	0.9562	0.8937
	Gemma Phr.	0.8591	0.6575	0.2393	0.2315	0.0572	0.6810	0.4722	0.9562	0.8952
	Llama Phr.	0.8647	0.6613	0.2588	0.2545	0.0793	0.6915	0.4722	0.9546	0.8961
	Original	0.8704	0.6680	0.2990	0.2843	0.1060	0.6853	0.4722	0.9562	0.8937

crime, demography, Walmart) using architecturally distinct LLMs (GPT-4o, Gemma 27B, and Llama 90B), resulting in three additional sets of ground truth captions. The paraphrasing prompt was designed to instruct the model to thoroughly alter sentence structure, syntax, and word choice while preserving all factual information, numeric values, and statistical details with absolute fidelity. The prompt used for paraphrasing is shown in N.3. We define this paraphrased caption set as CaTS-Bench-Paraphrased, which contains captions that are semantically equivalent to the original ground truth but differ only in linguistic style.

We then re-evaluated the outputs of all six representative pretrained models against CaTS-Bench-Paraphrased using the full suite of metrics. Results of this analysis are presented in the Table 8, one for each ground-truth generator’s linguistic style. Values represent the average across select domains. We also report the results obtained with our original Gemini captions as ground truth.

Next, we provide an analysis of the rank correlation of model performances between the original and paraphrased evaluation settings. A high

correlation in model rankings would indicate that our benchmark is robust to linguistic variation; the metrics would be consistently measuring the underlying semantic content and numeric accuracy of the captions, not their surface-level similarity to a specific writing style. A low correlation would suggest a non-trivial dependence on the oracle’s particular linguistic patterns. For each linguistic evaluation metric, we provide the model ranking across the four linguistic styles in Figure 7.

Qualitatively, we observe a negligible impact of model-specific linguistic style on the model rankings, suggesting that our evaluations are robust to particular linguistic styles. For each linguistic metric, we measure the average Spearman correlation between the ranking according to Gemini’s style and the ranking according to the other three styles (see Table 9).

In summary, this experiment demonstrates that our evaluation framework is robust to variations in the linguistic style of the reference captions. This conclusion is quantitatively supported by the high Spearman correlation coefficients observed in model rankings between the original and para-




Figure 7: Model ranking heatmaps across metrics under four reference styles. Rankings: 1 (highest) to 6 (lowest). Model mappings: Gemini (Gemini 2.0 Flash), Claude (Claude 3 Haiku), Llama-V (Llama 3.2 Vision).

Table 9: Spearman correlation between Gemini-based rankings and rankings induced by different metrics. A value of 1 indicates identical rankings.

Metric	Correlation	<i>p</i> -value
DeBERTa F1	0.971	0.0018
SimCSE	1.000	0.0000
BLEU	0.814	0.0557
ROUGE-L	0.905	0.0257
METEOR	0.943	0.0079
Average	0.927	0.0182

phrased benchmark sets. These results indicate that **our evaluation framework captures the semantic fidelity and factual quality of the generated content, rather than rewarding models for merely mimicking the stylistic patterns of our oracle model, Gemini 2.0 Flash**. Consequently, the benchmark evaluates fundamental capabilities in time series understanding and description, not superficial stylistic alignment. Furthermore, Gemini 2.0 Flash maintains its superior rank in most metrics regardless of the linguistic style of the ground truth. This consistent dominance validates its selection as a highly capable oracle, reinforcing that its utility stems from its intrinsic ability to generate high-quality descriptions rather than from any benchmark-specific bias.

H.4 Diversity Analysis

H.4.1 Content in the Captions

We analyzed all the 4005 Gemini-generated test set captions using a structured keyword-based approach (see Table 10). Captions across 11 domains were scanned for statistical descriptors (e.g., *mean*,

average, *standard deviation*, *maximum*, *minimum*, *range*) and trend-related expressions (e.g., *increase*, *decrease*, *stability*, *volatility*, *seasonality*). These terms were grouped into clusters such as *Central Tendency*, *Dispersion*, *Extremes*, *Increasing/Decreasing Trends*, *Stability and Volatility*, and *Comparative Trends*. The results in Table 11 show that **captions consistently draw from a diverse mix of descriptors, spanning both statistical features and temporal patterns**. While some categories (e.g., percentiles, distribution-shape terms) were rare, coverage of core descriptors was broad, and every caption included at least one statistical or trend-related element.

We apply the same keyword-based analysis to the HR caption set to examine the types of statistical and temporal descriptors used. Across domains, human captions draw from a wide range of descriptor clusters, including measures of central tendency, extrema, dispersion, and multiple forms of trend characterization. In addition to core statistical terms, human captions frequently include compound trend expressions, contextual qualifiers, and multi-stage temporal descriptions that span increases, decreases, stabilization, and recovery phases within a single caption.

Coverage extends beyond surface-level descriptors to include higher-order patterns such as structural transitions, regime changes, and shape-based trends, which appear with greater regularity in the human set. While all human captions contain at least one statistical or trend-related element, many incorporate several distinct clusters simultaneously, reflecting a richer and more varied composition of quantitative and temporal content.

Table 10: Coverage of statistical and temporal descriptor clusters across human captions. For each cluster, representative surface-form keywords are shown (non-exhaustive).

Category	Cluster (Representative Keywords)	Captions
Statistical	Central Tendency (<i>average, mean, median, mode, count, etc.</i>)	3184
	Maximum Values (<i>max, maximum, peak, highest value, record high, etc.</i>)	1487
	Minimum Values (<i>min, minimum, dip, trough, lowest value, etc.</i>)	1530
	Dispersion (<i>standard deviation, deviation, variance, std, iqr, etc.</i>)	684
	Variability (<i>variability, variation, fluctuate, oscillation, noise, etc.</i>)	665
	Range / Spread (<i>range, spread, difference, gap, extent, etc.</i>)	341
	Extremes (<i>peak, spike, trough, extreme, record high, etc.</i>)	2966
Trend and Patterns	Increasing Trends (<i>increase, rising, growing, upward, gain, etc.</i>)	1207
	Decreasing Trends (<i>decrease, declining, falling, drop, reduction, etc.</i>)	1308
	Stability / Flat Trends (<i>stable, steady, flat, constant, plateau, etc.</i>)	1082
	Sudden Changes (<i>spike, jump, shock, abrupt, rapid increase, etc.</i>)	219
	Volatility (<i>fluctuate, volatile, oscillate, erratic, instability, etc.</i>)	1210
	Comparative Relationships (<i>higher than, lower than, exceeds, below, gap, etc.</i>)	1433
	Structural Changes (<i>inflection, shift, transition, regime change, pivot, etc.</i>)	190
	Recovery / Correction (<i>recovery, rebound, normalize, correction, bounce back, etc.</i>)	362
	Cyclical Patterns (<i>seasonal, cyclical, periodic, recurring, repeats, etc.</i>)	135
Shape-based Patterns (<i>v-shaped, u-shaped, w-shaped, rise-and-fall, hill-shaped, etc.</i>)	75	

Table 11: Coverage of statistical and trend descriptor clusters across synthetic captions. Captions consistently include diverse descriptors capturing both quantitative and temporal aspects of the data.

Category	Cluster and Keywords in the Cluster	Captions
Statistical	Central Tendency (mean, average, median, mode)	3930
	Minimum Values (min, minimum, lowest value)	1528
	Maximum Values (max, maximum, highest value)	1487
	Dispersion (std, variance, deviation, iqr)	1376
	Range / Spread (range, spread)	492
	Extremes (peak, spike, dip, trough)	2966
Trend & Patterns	Peaks and Valleys (peak, dip, spike, trough)	3013
	Increasing Trends (increase, rising, growing)	1987
	Decreasing Trends (decrease, drop, falling)	2197
	Comparative Trends (higher/lower, difference)	2217
	Stability and Volatility (stable, fluctuating)	2276

H.4.2 N-Gram Diversity

We analyze n-gram diversity using the type–token ratio (TTR) to assess whether the generated captions exhibit repetitive or templated phrasing. For the synthetic corpus, computed over approximately 411k tokens, TTR increases steadily with n-gram order, rising from 0.0288 for unigrams to 0.6050 for 5-grams. This monotonic growth indicates rapid diversification at the phrase level rather than reliance on fixed templates. In particular, the presence of over 250k unique 5-grams out of 411k tokens suggests substantially higher-order lexical variability in the synthetic captions.

We perform the same analysis on the human-rewritten gold captions to provide a reference point. As shown in Table 12, HR captions exhibit consistently higher TTR across all n-gram orders, reflecting the greater compositional variability expected from human-authored text. Importantly, the syn-

thetic captions follow a similar upward trajectory with increasing n, differing primarily in magnitude rather than qualitative behavior. This pattern suggests that while synthetic captions are less diverse than their human-revisited counterparts, they remain far from being rigidly templated and capture a broad range of phrase structures.

Table 12: Type–token ratio (TTR) for synthetic and human-rewritten captions across n-gram orders. TTR increases monotonically with n for both corpora, indicating growing phrase-level diversity. Human-rewritten captions exhibit higher TTR across all n, while synthetic captions remain highly non-templated at higher n-grams.

N-gram	Synthetic TTR	Human TTR
1-gram	0.0288	0.0472
2-gram	0.1288	0.2320
3-gram	0.2971	0.5011
4-gram	0.4638	0.7066
5-gram	0.6050	0.8344

H.4.3 Latent Similarities

We further assessed whether Gemini-generated captions introduce systematic stylistic or linguistic bias. Results across multiple embedding models with different dimensions (Dim) show minimal template reliance and high diversity, consistent with prior work on linguistic convergence between AI and humans (Wilkenfeld et al., 2022). We perform $\sim 8M$ pairwise comparisons in the 4005 captions using nine embedding models (Table 13). Intra-domain similarity was consistently

Table 13: Embedding similarity statistics for human-rewritten and synthetic captions across encoder models.

Encoder	Dim	Human Captions					Synthetic Captions						
		Mean	Med	STD	Intra	Inter	% > 0.95	Mean	Med	STD	Intra	Inter	% > 0.95
MiniLM-L12-v2	384	0.2755	0.2437	0.1560	0.6090	0.2308	0.17	0.2932	0.2500	0.1771	0.6202	0.2269	5.26
MiniLM-L6-v2	384	0.3367	0.3113	0.1429	0.6291	0.2975	0.12	0.3339	0.3013	0.1510	0.5999	0.2800	1.47
mpnet-base-v2	768	0.3260	0.2859	0.1679	0.6866	0.2776	0.44	0.3278	0.2791	0.1775	0.6631	0.2599	8.29
bge-large-v1.5	1024	0.5921	0.5753	0.0955	0.7903	0.5655	0.23	0.5807	0.5605	0.0986	0.7600	0.5443	1.67
mxbai-large-v1	1024	0.5344	0.5111	0.1096	0.7691	0.5030	0.24	0.5293	0.5032	0.1116	0.7370	0.4871	2.15
Qwen3-4B	2560	0.3833	0.3613	0.1195	0.6344	0.3496	0.02	0.3816	0.3530	0.1273	0.6098	0.3354	0.20
Qwen3-8B	4096	0.3800	0.3571	0.1188	0.6306	0.3465	0.00	0.3736	0.3465	0.1225	0.5947	0.3288	0.32
e5-mistral-7b	4096	0.6469	0.6354	0.0670	0.7839	0.6285	0.01	0.6507	0.6358	0.0697	0.7770	0.6251	1.33
Average		0.4344	0.4101	0.1222	0.6910	0.3999	0.15	0.4386	0.4092	0.1256	0.6678	0.3922	2.30

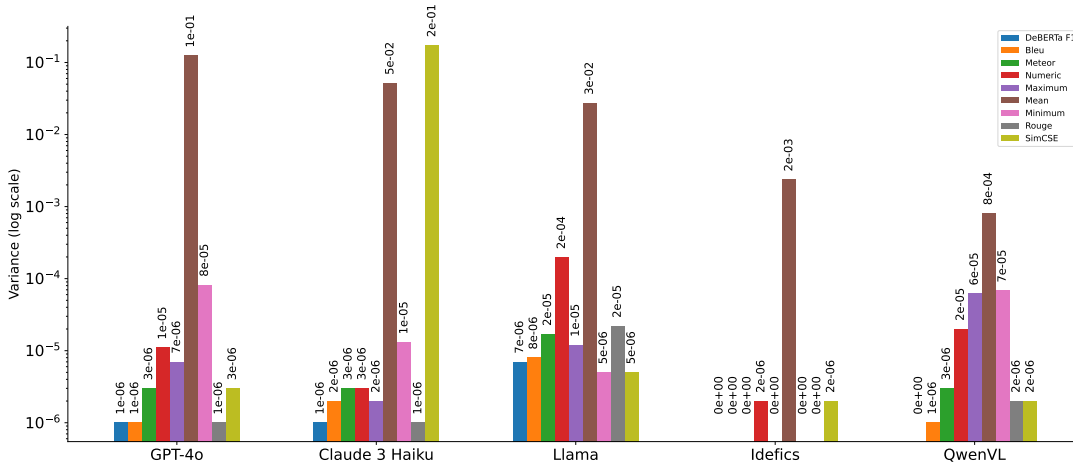


Figure 8: Variance across three independent runs (approximately 500 samples) for each model–metric pair. The logarithmic scale highlights both very small variances (10^{-6}) and moderately larger values.

higher (0.59 – 0.78) than inter-domain similarity (0.23 – 0.54), with large effect sizes (Cohen’s $d > 3.26$). Near-duplicates (cosine > 0.95) were rare, with an average of 2.3% of pairs. Even within domains, similarity showed non-trivial variance, indicating that **captions are not rigid templates but semantically varied**.

H.4.4 Human vs. AI Embedding Similarity

We measure cosine similarity between captions authored by humans, GPT-4o, and Gemini using Qwen-8B embeddings to assess whether AI-generated captions diverge stylistically from human-rewritten references. Across both CaTS-Bench and the STOCK dataset, we observe only modest gaps between human-AI and human-human similarity in embedding space, indicating overlap in semantic space despite substantial surface-level differences. This suggests that while AI captions differ lexically and structurally from human captions, they do not collapse to a distinct semantic mode. For CaTS-Bench, we report results on a balanced subset of 500 human and 500 AI captions,

averaged over five random draws. For STOCK, we evaluate $3k$ human and $3k$ AI captions following prior work. Full results are shown in Table 14.

Table 14: Cosine similarities between human, GPT-4o, and Gemini captions using Qwen-8B embeddings. Human-AI similarity is consistently lower than human-human similarity, but remains semantically comparable.

Comparison	Human	GPT	Gemini
CaTS-Bench (500 vs 500)			
Human	0.389	0.341	0.336
GPT		0.414	0.321
Gemini			0.372
STOCK (3k vs 3k)			
Human	0.658	0.621	0.603
GPT		0.692	0.588
Gemini			0.601

H.5 Reproducibility Validation

We re-ran our evaluation three times on ~ 500 randomly sampled test examples from the synthetic

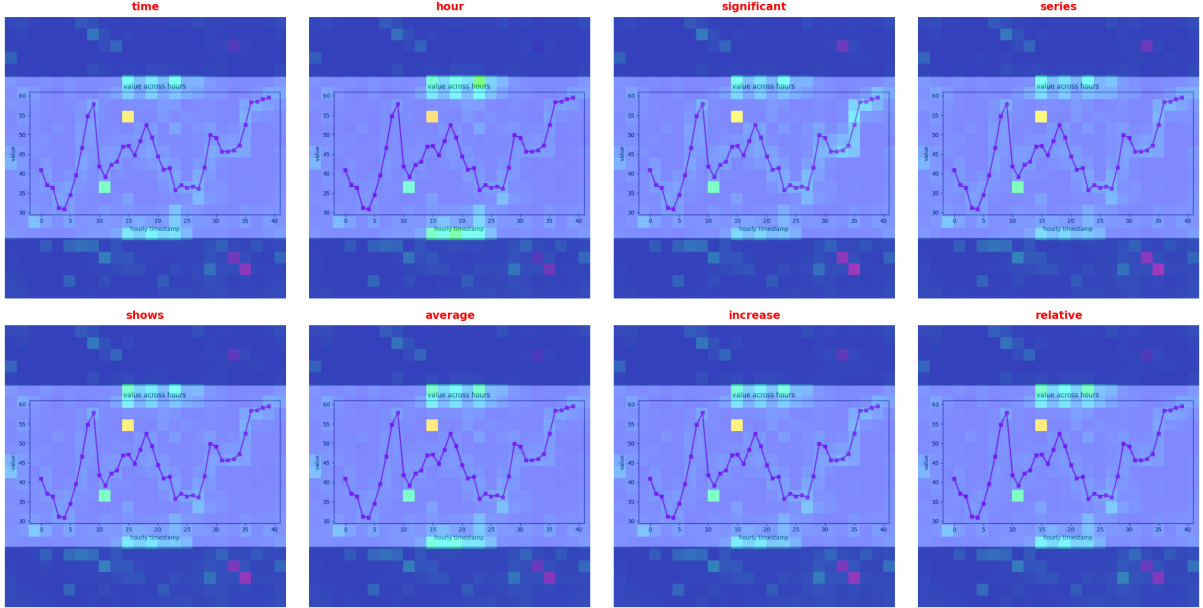


Figure 9: Word-level attention maps for the top 8 tokens from LLaVA 1.6 Mistral overlaid on a time series plot. Despite expectations of alignment with visual trends, attention remains largely diffuse, offering only weak evidence of localized visual grounding in caption generation.

data across five representative models as generators (GPT-4o, Claude 3 Haiku, LLaMA, Idefics, Qwen-VL). Figure 8 shows a log-scale visualization. **Across nearly all metrics, the variance is vanishingly small, often on the order of 10^{-6} , which confirms stability and supports single-run robustness.**

I Role of Vision

I.1 Visual Attention Analysis

Interpreting visual grounding in large multimodal models is non-trivial, as not all of them expose interpretable cross-modal attention mechanisms. We attempt this using the LLaVA model, which provides access to decoder-level cross-attention weights over vision tokens. We adapt the approach in (Zhang, 2024) for the LLaVA 1.6 model. We visualize per-token visual grounding via the following steps. For each generated token, we extract the decoder cross-attention matrix $\mathbf{A}_{\text{llm}} \in \mathbb{R}^{T \times V}$, where T is the number of generated tokens and V is the number of vision tokens.

Next, we zero out the attention to the beginning-of-sequence token and normalize each row:

$$\tilde{\mathbf{A}}_{\text{llm}}[t, v] = \begin{cases} 0, & \text{if } v = \langle \text{bos} \rangle \\ \frac{\mathbf{A}_{\text{llm}}[t, v]}{\sum_{v'} \mathbf{A}_{\text{llm}}[t, v']}, & \text{otherwise} \end{cases} \quad (6)$$

From the CLIP style vision encoder, we extract

attention matrices $\mathbf{A}_{\text{vit}}^{(l)} \in \mathbb{R}^{V \times V}$ from multiple layers and average them:

$$\bar{\mathbf{A}}_{\text{vit}} = \frac{1}{L} \sum_{l=1}^L \mathbf{A}_{\text{vit}}^{(l)} \quad (7)$$

For each token t , we compute its attention-weighted vision token distribution and project it back to the 2D image grid, and the projected map $\hat{\mathbf{H}}_t$ is rendered as a heatmap and overlaid on the original image. This allows inspection of which visual regions contribute to each generated token.

$$\hat{\mathbf{H}}_t = \text{Upsample}(\text{reshape}(\mathbf{H}_t, \text{grid})) \quad (8)$$

I.2 Testing Alternative Visual Encodings

To investigate whether more complex visualizations could help in the modality collapse issue, we experimented with Gramian Angular Fields (GAFs) and recurrence plots (RPs) in addition to standard line plots (Figure 10). Each was generated from the same univariate time series windows, using first-order deltas as input. There are also other encodings, such as multi-series overlays or confidence intervals, but we find these are less applicable in our strictly univariate setting.

Table 15 reports results for three representative models (Idefics2-8B, LLaMA 3.2 Vision, Qwen-VL). Values show the baseline metric with no plot

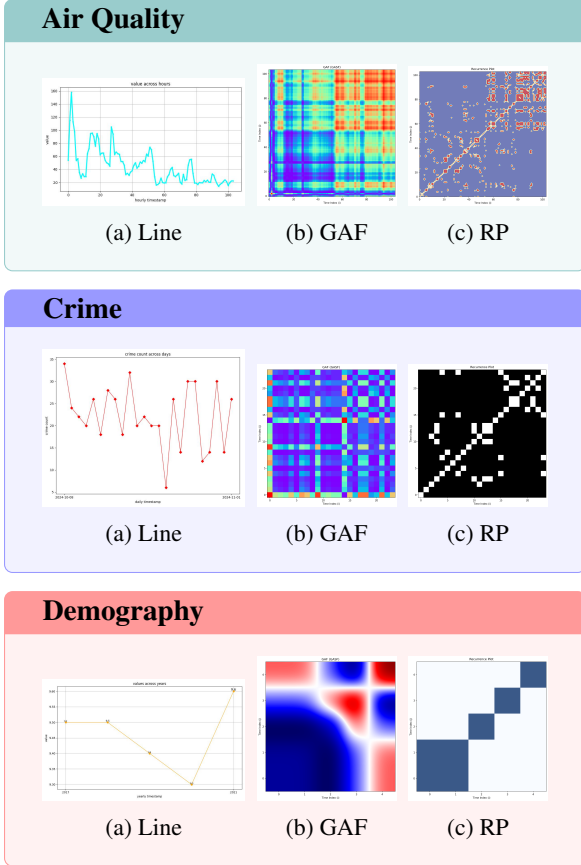


Figure 10: Examples of visual encodings across three domains across three sample domains (Air Quality, Crime, Demography) (a) Line, (b) GAF, and (c) RP.

(TS+Text only), followed by relative gains or losses for each visualization type. Across models and metrics, neither GAFs nor recurrence plots significantly improved performance; in many cases, they degraded results relative to line plots or even no visual input. These findings suggest that the bottleneck lies not in the choice of visualization but in the models’ inability to effectively integrate visual cues. With this, we motivate the development of models and encodings that better exploit the structured information available in visual time series.

J Bonus TSC Evaluations on Newer Models

We additionally evaluate some newer proprietary models (GPT-5.1, Gemini-3 Flash) as well as specialized time series language models (ChatTS 14B, TimeOmni-1 7B) on the human-rewritten set of CaTS-Bench. Their results are provided in Table 16. Comparing them to the models in the main paper (Table 3), the overall takeaways are: i) the newest proprietary models do not consistently out-

Table 15: Effect of alternative visual encodings on captioning performance on a subset of synthetic captions. Values show the baseline score for *No Plot (TS+Text)* and relative differences (Δ) for line plots, Gramian Angular Fields (GAFs), and recurrence plots (RPs).

Model	Metric	No Plot	Line (Δ)	GAF (Δ)	RP (Δ)
IdeFics2 (8B)	DeBERTa F1	0.6255	-0.0780	-0.0525	-0.0525
	BLEU	0.0690	-0.0405	-0.0460	-0.0430
	METEOR	0.2275	-0.0925	-0.0905	-0.0815
	NUMERIC	0.4455	-0.1665	-0.1715	-0.1285
	MAXIMUM	0.8080	+0.0730	-0.1520	-0.1300
	MEAN	0.4195	+0.1285	+0.0945	+0.2995
	MINIMUM	0.7590	-0.0845	-0.2510	-0.3940
	STD	0.5500	-0.1835	-0.1060	-0.3360
	ROUGE	0.2375	-0.0460	-0.0465	-0.0415
	SIMCSE	0.7915	-0.1735	-0.1935	-0.1705
LLaMA 3.2 Vision	DeBERTa F1	0.6550	-0.0020	-0.0030	-0.0080
	BLEU	0.0840	+0.0005	-0.0090	-0.0100
	METEOR	0.2865	-0.0080	-0.0195	-0.0315
	NUMERIC	0.5070	+0.0315	+0.0120	+0.0170
	MAXIMUM	0.8611	+0.0584	-0.0381	-0.0151
	MEAN	0.5395	+0.0690	+0.0755	+0.1215
	MINIMUM	0.6670	+0.0820	+0.0740	+0.0670
	STD	0.8750	-0.3270	-0.2310	-0.1250
	ROUGE	0.2570	+0.0040	0.0000	-0.0040
	SIMCSE	0.8265	-0.0020	-0.0085	-0.0235
Qwen-VL	DeBERTa F1	0.6315	-0.0045	-0.0045	-0.0095
	BLEU	0.0645	-0.0040	+0.0025	-0.0015
	METEOR	0.2385	-0.0050	+0.0025	-0.0025
	NUMERIC	0.3555	+0.0560	-0.0675	-0.0515
	MAXIMUM	0.6610	+0.1135	+0.0200	+0.0430
	MEAN	0.4510	-0.0110	-0.2150	-0.2570
	MINIMUM	0.5350	-0.0055	-0.0650	-0.0620
	STD	0.1070	+0.2095	-0.0620	-0.0070
	ROUGE	0.2310	-0.0045	+0.0130	+0.0030
	SIMCSE	0.7835	+0.0020	-0.0085	-0.0090

perform their older generation counterparts but remain competitive, and ii) time-series-specialized models perform poorly compared to generalist language models, since TSC is still fundamentally a free-form text generation task, though with a taste of numeric and temporal reasoning.

K Information Coverage Analysis

To better understand the gains from finetuning, we conduct an information coverage analysis comparing pretrained and finetuned IDEFICS2 on $\sim 1.3k$ generated captions. Using GPT-4o as an automatic annotator (with manual spot-checking), we measure whether captions include core time series reasoning elements such as central tendency, extrema, trends, and comparative reasoning.

As shown in Table 17, finetuning substantially increases the inclusion of essential reasoning components, particularly for central tendency and comparative reasoning. These improvements help explain the gains observed across evaluation metrics, which suggests that finetuning teaches the model

Table 16: Performance across metrics of additional models on the HR set.

Model	BERT F1	SimCSE	BLEU	ROUGE	METEOR	Numeric	Mean	Max	Min	STD
Gemini 3 Flash	0.690	0.885	0.083	0.259	0.287	0.783	0.967	0.995	0.992	0.718
GPT-5.1	0.657	0.875	0.058	0.227	0.298	0.820	0.926	1.000	0.999	0.523
ChatTS 14B	0.605	0.667	0.040	0.186	0.215	0.587	0.929	0.701	0.667	0.044
TimeOmni-1 7B	0.613	0.798	0.046	0.205	0.225	0.717	0.763	0.974	0.953	0.130

Table 17: Information coverage of key time series reasoning elements in captions generated by pretrained vs. finetuned models.

Category	Pretrained	Finetuned	Gain
Central Tendency	14.0%	95.5%	+81.6%
Peak/Maxima	55.5%	84.7%	+29.2%
Trough/Minima	51.7%	64.1%	+12.4%
Trend Direction	56.0%	74.6%	+18.6%
Temporal References	76.1%	100.0%	+23.9%
Comparative Reasoning	6.5%	95.6%	+89.1%
Variability/Dispersion	31.6%	88.1%	+56.6%
Pattern/Shape	15.3%	67.3%	+52.0%

to produce more structurally complete and informative time series descriptions.

L Inference Examples

L.1 Numeric Time Series Ablation

We compare the generated captions with and without explicitly providing the numeric time series data as input to the model in the prompt. As shown in Figure 11, when a numeric time series is included, the model produces some factual and interpretative errors. Notably, it describes the trend as “increased steadily,” despite the clear dip in 2014 and a decline post-2017. It also incorrectly identifies 2018 as the year of the peak value 102.48, while the actual peak occurs in 2017. Similarly, the slight dip is misattributed to 2016 instead of the correct year 2014.

In contrast, when the numeric time series is removed from the input, the generated caption becomes significantly more erroneous. The model fabricates plausible-sounding but incorrect values, for example, claiming the index reached 90.5 in 2009 and spiked to 105 in 2014, neither of which is present in the actual plot. This suggests that the absence of explicit numeric context forces the model to hallucinate plausible trajectories based solely on the shape of the line plot. While both versions demonstrate limitations in temporal precision, the numeric-aware caption is more grounded and less prone to hallucinating specific values, producing some factual errors but still outperforming the




Figure 11: Comparison of generated captions with and without numeric time series input. Model evaluated is LLaVA 1.6.

numeric-agnostic version, which fabricates values entirely.

L.2 Statistical Inference Failure & Success

In the main paper, we mentioned that finetuned models often become overconfident when inferring statistical properties such as means and standard deviations, despite lacking the capability to compute them accurately. In this section, we present two concrete cases that illustrate this overconfidence in practice, and one case where the finetuned model actually successfully inferred the statistics.

Case 1: Failure The following error case shows the finetuned LLaVA v1.6 Mistral erroneously attempting to infer the time series mean, while its pretrained version does not. We compare the captions generated by both versions of the model against the numerically grounded reference.

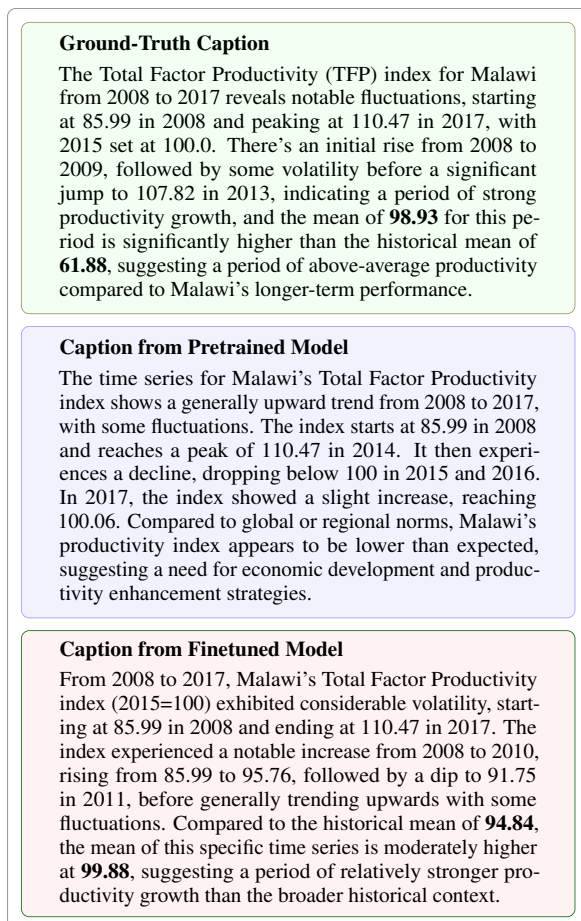


Figure 12: Finetuned model attempts to infer the mean, but it fails. Compare the **bolden** values.

We observe that, while the pretrained model refrains from inferring the mean of the time series, displaying a form of "self-awareness" about its limitations, the finetuned model tends to make confident guesses. This behavior underscores a key limitation of the cross-entropy loss objective. Since ground-truth captions often include the time series mean, the model can reduce its loss by always attempting to predict them, as number tokens are treated like any text tokens. When the model generates captions with sentence structures resembling those in the ground truth, even erroneous guesses of these statistics incur less loss than omitting them entirely.

Case 2: Success The following is a success case where the finetuned Idefics 2 is able to infer the time series mean accurately with a negligible error. We compare the captions generated by the pretrained and finetuned versions of the model against the ground-truth.

Interestingly, the issue of statistical overconfi-

dence appears to be model-specific, as different models exhibit varying behaviors after finetuning. In this case, the finetuned Idefics 2 was able to infer both the mean and the standard deviation with reasonable accuracy, when even the ground-truth caption did not explicitly include the standard deviation. This signals that some models benefit more from finetuning on our training data.

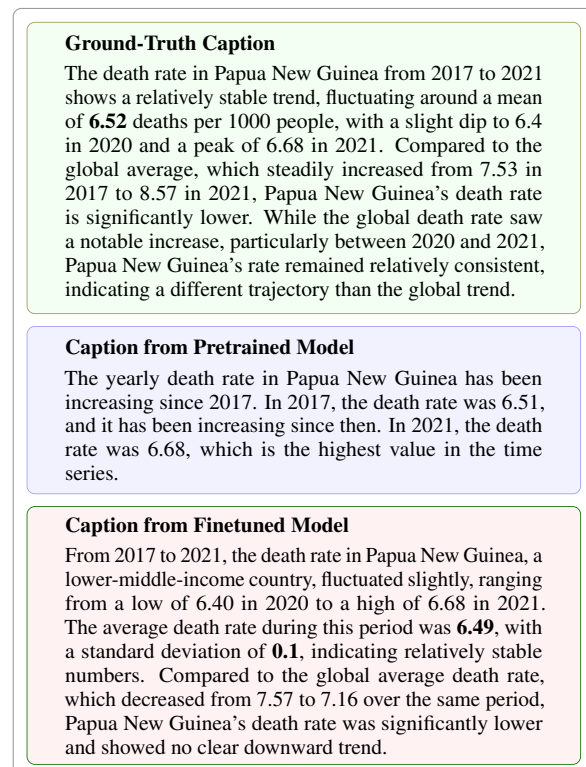


Figure 13: Finetuned model successfully infers the mean and standard deviation with negligible error. Compare the **bolden** values.

M Editing and Review

M.1 Edits made in the captions

During the editing process conducted by two of the authors, captions were systematically refined to ensure accuracy, clarity, and consistency across the dataset. The following key rules were applied:

- 1. Removal of external speculation:** Captions were restricted strictly to information verifiable from the metadata, time series, and plot, avoiding any causal claims or conjecture not grounded in the time series values or provided metadata.
- 2. Variation in phrasing:** To reduce repetitiveness, sentence openings and phrases were varied

rather than uniformly beginning and phrasing the same sentences.

3. **Pattern summarization:** When trends or unique structures (such as V-shaped or monotonic movements, etc.) were clearly visible, they were explicitly noted.
4. **Quantitative grounding:** Values such as maxima, minima, averages, and percentage changes were consistently included when relevant to ensure captions remained data-driven.
5. **Consistency with variation:** While maintaining factual accuracy and grounding in the data, captions were intentionally varied in structure and style to avoid monotony and ensure more natural, human-like phrasing across the dataset.

This systematic review process resulted in captions that were both faithful to the underlying data and stylistically coherent across the dataset.

M.2 Semantic Drift Between Unedited and Edited Captions

In this section, we show how much the captions were edited quantitatively. Table 18 quantifies the extent of human edits across domains using word-level divergence (WER), token-level change, and numeric overlap. The magnitude of rewriting varies substantially by domain. For example, *border crossing*, *diet*, and *Walmart* captions exhibit very large edits, with mean WER values of 0.944, 0.885, and 0.939, respectively, alongside sizeable average token increases of +28.04, +26.88, and +21.24, indicating extensive rewriting and content addition.

Table 18: Domain-wise mean rewrite magnitude (WER), lexical overlap (ROUGE-L), numeric faithfulness (Numeric Jaccard), and caption length change (Δ Tokens) between unedited and human-edited captions.

Domain	WER	ROUGE-L	Numeric Jaccard	Δ Tokens
Agriculture	0.663	0.541	0.892	7.63
Border Crossing	0.944	0.287	0.509	28.04
CO ₂	0.733	0.452	0.674	18.76
COVID	0.232	0.859	0.846	-2.95
Crime	0.593	0.555	0.854	0.88
Demography	0.504	0.637	0.956	4.67
Diet	0.885	0.365	0.618	26.88
Online Retail	0.396	0.730	0.816	-10.76
Road Injuries	0.132	0.921	0.862	2.68
Walmart	0.939	0.288	0.329	21.24
Average	0.602	0.564	0.736	9.71

Other domains reflect more moderate editing. *Agriculture* and *CO₂* show mean WER values of

0.663 and 0.733, with average token increases of +7.63 and +18.76, suggesting partial restructuring rather than complete rewrites. Further, domains such as *COVID*, *road injuries*, and *online retail* exhibit lower rewrite magnitudes, with WER values of 0.232, 0.132, and 0.396, respectively, and small or negative mean token deltas, indicating that many edits involve localized refinements or concision rather than expansion.

Across domains, numeric Jaccard scores remain relatively high in several cases (e.g., 0.956 for *demography* and 0.892 for *agriculture*), showing that even when captions are substantially rewritten, numeric content is often preserved. Overall, human editing spans a wide spectrum of intervention levels, ranging from targeted phrasing adjustments to near-complete caption rewrites, depending on domain characteristics and the needs of individual captions.

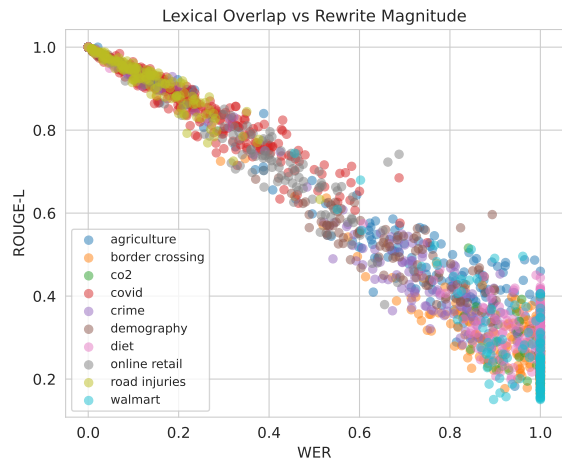


Figure 14: Lexical overlap versus rewrite magnitude across domains.

Figure 14 plots lexical overlap against rewrite magnitude, where each point corresponds to a single caption pair, with word error rate (WER) between the unedited and human-edited captions on the x-axis and ROUGE-L overlap on the y-axis. Colors denote different domains. We observe a strong inverse relationship between rewrite magnitude and lexical overlap: captions with low WER retain high ROUGE-L scores, while captions undergoing substantial rewriting (WER approaching 1.0) exhibit markedly lower lexical overlap. This consistent trend across domains indicates that increased human intervention is primarily expressed through lexical and structural rephrasing. The plot further highlights how human edits progressively

reduce surface-level overlap as rewrite magnitude increases, reflecting varying degrees of paraphrasing and content restructuring.

Figure 15 examines numeric preservation under similar conditions. Each point represents a caption pair, plotted by WER on the x-axis and numeric Jaccard similarity on the y-axis, with colors indicating domains. The figure shows that captions span a wide range of rewrite magnitudes, from minor edits (low WER) to near-complete rewrites (WER close to 1.0). Despite substantial rewriting in many cases, numeric overlap often remains high, with numerous samples exhibiting numeric Jaccard scores above 0.8 even at high WER values. At the same time, certain domains display greater dispersion at high rewrite magnitudes, reflecting more extensive numeric restructuring when captions are heavily rewritten.




Figure 15: Numeric preservation versus rewrite magnitude across domains.

M.3 Interface

In Figure 16, we provide a screenshot of the editing interface we used to edit the human-rewritten test set.

N Template-based Prompts

In this section, we illustrate the prompts used in our sample generation pipeline, evaluation, paraphrasing, PAL, and distractor generation in Q&A. Angular brackets are used as placeholders for the actual values.

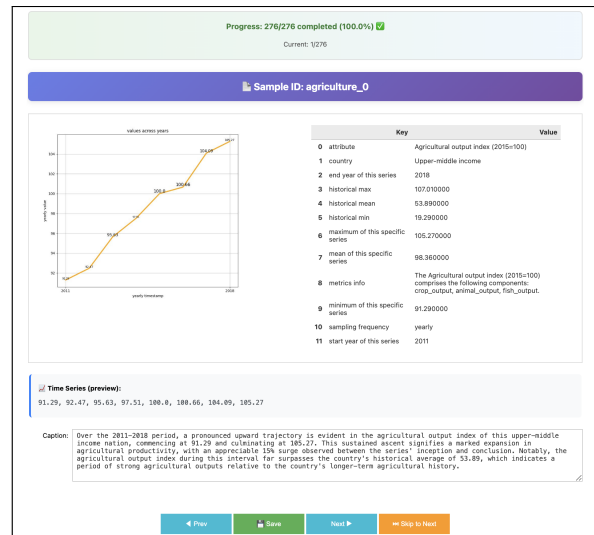


Figure 16: Interface used to edit and verify the captions.

N.1 Ground-Truth Synthetic Caption Generation Prompt

The following is an example of a prompt for generating the ground-truth caption from the source dataset *Crime*.

```

Here is a time series about the number of
<sampling frequency> crimes in <town>, Los
Angeles, from <start date> to <end date>:

<time series>

The all-time statistics of <town> until today
are:
Mean: <general mean of this town>
Standard Deviation: <general standard deviation
of this town>
Minimum: <general minimum of this town>
Maximum: <general maximum of this town>

And the statistics for this specific time
series are:
Mean: <mean of this specific series>
Standard Deviation: <standard deviation of this
specific series>
Minimum: <minimum of this specific series>
Maximum: <maximum of this specific series>

Describe this time series by focusing on trends
and patterns. Discuss concrete numbers you see
and pay attention to the dates.

For numeric values, ensure consistency with the
provided time series. If making percentage
comparisons, round to the nearest whole number.
Report the dates when things happened.

Use the statistics I provided you for comparing
this example to the normalcy.
Do not add any extra information beyond what is
given.
Highlight significant spikes, dips, or patterns.

You don't have to explicitly report the numeric

```

values of general statistics; you just use them for reference.

Compare the trends in this time series to global or regional norms, explaining whether they are higher, lower, or follow expected seasonal patterns.

When making comparisons, clearly state whether differences are minor, moderate, or significant.

Use descriptive language to create engaging, natural-sounding text.

Avoid repetitive phrasing and overused expressions.

Answer in a single paragraph of four sentences at most, without bullet points or any formatting.

N.2 Baseline Caption Generation Prompt

When evaluating the baselines on our benchmark, we provide limited metadata, excluding the pre-computed statistics of the time series, as the models are expected to infer them on their own. An example of the prompt is the following.

```
Here is a time series about <sampling
frequency> <measure> in the Indian city of
<city name>.
```

```
The starting time is <starting time> for the
series. \n <time series> \n
```

Describe this time series following these rules:

1. specify and include all contextual data that is needed.
2. make statistical comparisons (max, min, mean, standard deviation) by providing numerical values.
3. comment on patterns, shapes, and trends of the time series.
4. explain how the period compares to historical or global norms.
5. describe the nature of the movements, patterns, and anomalies if any.
6. walk through the timeline and describe sub-periods, phases, and periodic patterns when present.
7. Report dates when events happened.

****Not all information for the above rules might be present or needed**, focus on information that is actually evident in the data inputs.**

Here is some historical context:

```
- all-time maximum: <all-time maximum>
- all-time minimum: <all-time minimum>
- all-time average value until today:
<all-time average value until today>
- all-time standard deviation until today:
<all-time standard deviation until today>
```

Use descriptive language to create engaging, natural-sounding text.

Avoid repetitive phrasing and overused expressions.

Answer in a single paragraph of four to five sentences at most, without bullet points or any formatting.

N.3 Caption Paraphrasing Prompt

To rephrase a caption into a different linguistic style while preserving semantic and numeric information, we feed the following prompt into a paraphraser model of choice.

```
You are a helpful assistant. Your task is to
rephrase the following paragraph that describes
a time series. You MUST strictly follow these
rules:
```

```
Preserve all factual information: All numeric
values, statistics (min, max, mean, etc.),
trends ('increased', 'peaked'), comparisons
('higher than'), and dates must remain exactly
the same.
```

```
Change the style completely: Use different
sentence structures, synonyms, and grammatical
constructions. Alter the tone (e.g., make it
more formal or more conversational). Do not use
the same phrasing as the original.
```

```
Output only the rephrased paragraph, with no
additional explanation.
```

```
Here is the paragraph to rephrase:
<caption>
```

N.4 PAL Prompt

Here we reproduce the full prompt used for the program-aided context:

```
### Task
<caption_prompt>

### Instructions for the assistant
1. You are an expert coding assistant; think
through the task **step-by-step**.
2. Write **Python 3.12** code (inside one
```python``` block) that computes the final
answer.
 * Use only the Python Standard Library
(e.g., you may use the `math`, `statistics`
libraries).
 * Wrap everything in a `solve()` function
that will be invoked to produce the final
caption.
 * The code **must produce the caption string
itself**. Any numeric values can be computed
in Python and formatted into the caption
string. Make sure to use any values you compute
in the resulting caption string.
3. The `solve()` function you write will be
invoked to produce the final caption.

Output format (exactly; no extra text,
explanations, or formatting)
```python
# code that defines solve() and any desired
strings
```

```
solve()
---
```

The full TSC prompt from N.2 is injected as the `caption_prompt` string.

N.5 Semantic Perturbation Prompt

To perturb a caption so that its semantic meaning is altered while keeping numbers intact, we feed the following prompt into Gemini 2.0 Flash.

```
Your task is to minimally modify a time series description so that its meaning is altered but the numbers are maintained.

For example, you can switch "increase" with "decrease", "upward" to "downward", or something more sophisticated. Keep the description structurally identical to the original text; you don't have to alter too much information. Altering anywhere between 1 and 3 parts is enough. Do not edit the numbers.

Here's the description to modify:
<caption>

Give your answer in a paragraph of text as the given description, without any explanation or formatting.
```

N.6 Numeric Perturbation Prompt

To perturb a caption so that its numbers are altered while its semantic information is preserved, we feed the following prompt into Gemini 2.0 Flash.

```
Your task is to slightly modify the numbers in a time series description so that its semantics remain the same, but the numbers are slightly altered.

For example, you can replace "12" with "12.2", "45%" with "46%". Keep the description structurally and semantically identical to the original text; you don't have to alter all numbers but anywhere between 1 to 3 times is enough. Make sure that the altered number still makes sense and fits the scale of the phenomenon.

Here's the description to modify:

<caption>

Give your answer in a paragraph of text as the given description, without any explanation and formatting.
```

N.7 Numeric Parsing

To compute the statistical inference scores (min, max, mean, STD), we employ an LLM (*Gemini 2.0 Flash*) to extract the statistical claims from the captions. The extracted values will be compared against ground-truth precomputed values from the

Table 19: Inter-rater ranking correlation (Spearman ρ) across three extractors ($p < 0.001$ for all values).

(a) Minimum value inference accuracy			
Extractor	GPT-4o	C3H	G2.0F
GPT-4o	1.0000	0.8643	0.9643
Claude 3 Haiku	0.8643	1.0000	0.8464
Gemini 2.0 Flash	0.9643	0.8464	1.0000
(b) Maximum value inference accuracy			
Extractor	GPT-4o	C3H	G2.0F
GPT-4o	1.0000	0.8500	0.9214
Claude 3 Haiku	0.8500	1.0000	0.8107
Gemini 2.0 Flash	0.9214	0.8107	1.0000
(c) Mean value inference accuracy			
Extractor	GPT-4o	C3H	G2.0F
GPT-4o	1.0000	0.9643	0.9893
Claude 3 Haiku	0.9643	1.0000	0.9750
Gemini 2.0 Flash	0.9893	0.9750	1.0000
(d) Std. Deviation inference accuracy			
Extractor	GPT-4o	C3H	G2.0F
GPT-4o	1.0000	0.9571	0.7536
Claude 3 Haiku	0.9571	1.0000	0.8286
Gemini 2.0 Flash	0.7536	0.8286	1.0000

Note: C3H = Claude 3 Haiku, G2.0F = Gemini 2.0 Flash.

metadata. A statistical value is correct if it's within 5% error compared to the true value. The statistics extraction prompt is shown below.

```
I will provide you with a paragraph of text, which described a time series. Your job is to extract important numbers from it. Specifically, you have to extract minimum, maximum, mean, and standard deviation of this specific time series (not historical statistics) from the text if the text mentions them. You must ignore any statistic about the general history or all-time. For example, you should ignore statistics like "historical maximum", "all-time mean", or something like that. I only need you to extract statistics of the current specific time series.

Here's the text:
\n
{caption}
\n
```

Provide your answer in a JSON format as follows:

```
{  
  "minimum": A,  
  "maximum": B,  
  "mean": X,  
  "std": Y  
}
```

IMPORTANT: For each statistic (minimum, maximum, mean, std):

- If the text DOES NOT mention or attempt to report that statistic at all, set its value to null

- If the text DOES mention or attempt to report that statistic (even if the value is uncertain or approximate), provide the numerical value

The key distinction is whether the caption attempts to report the statistic or not. If it attempts to report it, extract the number. If it doesn't attempt to report it at all, use null.

Provide your answer in a json format and do not say anything more and don't give any explanation. Start your answer directly by opening a brace.

""

Furthermore, to prove that using a single LLM as the extractor is robust enough, we leveraged two additional LLMs, *GPT-4o* and *Claude 3 Haiku*, as extractors. We show that the choice of extraction model does not significantly impact the evaluation outcomes across all four statistical inference metrics. For pairwise ranking correlations, we observed strong to excellent agreement: the minimum metric averaged $\rho = 0.8917$ (range: $0.8464 \sim 0.9643$), the maximum metric averaged $\rho = 0.8607$ (range: $0.8107 \sim 0.9214$), the mean metric averaged $\rho = 0.9762$ (range: $0.9643 \sim 0.9893$), and the standard deviation metric averaged $\rho = 0.8464$ (range: $0.7536 \sim 0.9571$). All pairwise correlations were statistically significant ($p < 0.001$). Critically, when examining the mean metric, which measures accuracy in reporting the average value of time series, all three extractors exhibited 100% agreement on the top-5 performing models (*qwenvl_pal*, *OpenAI GPT-4o*, *Gemini 2.0 Flash*, *InternVL-38B*, and *Claude 3 Haiku*). Similarly, for the standard deviation metric, all three extractors showed 100% agreement on the top-3 models. These results validate the inter-rater reliability of our statistical inference evaluation methodology across all statistical claims and confirm that

our findings are robust and not artifacts of extractor-specific biases. Specific correlations are illustrated in Tables 19.

O Human Baseline

To establish a human performance baseline, we invited university students to complete four Q&A tasks voluntarily. These tasks span a range of reasoning types, including fine-grained statistical comparisons, semantic interpretation, and multimodal alignment. Participants were recruited through academic networks and completed the tasks without the aid of external tools, ensuring a fair comparison with models operating under similar conditions. Participation was entirely voluntary, with no compensation, and individuals could withdraw at any time. Below, we present the instructions given to the volunteers for their participation.

Participant Information and Consent Form for Time Series Q&A Questionnaire

Thank you for considering participation in our study!

This questionnaire is part of a research project evaluating human performance on time series understanding tasks. Your responses will help us establish a baseline for comparing human performance to that of current language models. You will be given a Google Form consisting of 10 to 14 multiple-choice questions of the same type, and you should not use any external tools.

Please read the following information carefully before continuing:

Voluntary Participation: Your participation is entirely voluntary. You may choose not to participate or to withdraw at any time without any consequences.

Duration: The questionnaire is brief and is estimated to take between 3 and 6 minutes to complete.

Anonymity & Data Use: No personal information will be collected or

stored. Your answers will remain anonymous and will be used solely for research purposes, such as evaluating and reporting model performance in academic publications.

No Compensation: There is no monetary or material compensation for participating in this study.

Confidentiality: All collected data will be handled securely. Only aggregated and anonymized results will be published.

By proceeding, you confirm that you understand the above terms and agree to participate in this research study.

Thank you for your collaboration and contribution to our research!

Date: _____

Signature: _____

O.1 Attempt Analysis in Numeric Metrics

We provide a detailed diagnostic analysis of model behavior across the four numeric inference tasks: mean, minimum (Min), maximum (Max), and standard deviation (STD). We analyze each statistic independently to expose distinct reasoning behaviors, robustness patterns, and failure modes. All observations are grounded in successes, attempts, and omissions (Null) of these statistics.

Specifically, **success rate** is defined as the proportion of correctly reported statistical claims among attempts. To illustrate, if the model only attempted to report the time series mean once across the entire test set, and the mean is correct, the success rate would be 1. To account for this flaw, the **attempt rate** is defined as the proportion of times where the model does make the statistical claim, whether it is correct or not. By multiplying the success rate by the attempt rate, we obtain the proportion of correct statistical claims over all evaluation instances. Finally, we count the statistical omissions (Null) where the model did not attempt the statistical claim whatsoever.

Mean Inference. Mean inference is the most stable and reliable numeric operation across all evaluated models. Strong models consistently

Table 20: Mean inference performance across domains and models.

Domain	Model	Success	Attempt	Null
Agriculture	Gemini 2.0 Flash	0.9940	1.0000	0
	GPT-4o	0.9940	1.0000	0
	QwenVL	0.4850	0.9581	7
	QwenVL PAL	0.9701	0.9701	5
Border Crossing	Gemini 2.0 Flash	0.7709	0.9907	2
	GPT-4o	0.8359	0.9938	1
	QwenVL	0.2879	0.9876	4
	QwenVL PAL	0.9659	0.9721	9
COVID	Gemini 2.0 Flash	0.5167	0.9700	9
	GPT-4o	0.5367	1.0000	0
	QwenVL	0.1200	0.9233	23
	QwenVL PAL	0.8400	0.9833	5
Crime	Gemini 2.0 Flash	0.8235	1.0000	0
	GPT-4o	0.9281	1.0000	0
	QwenVL	0.2092	0.9673	5
	QwenVL PAL	0.9804	0.9804	3
Demography	Gemini 2.0 Flash	0.9917	1.0000	0
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.6000	0.6333	44
	QwenVL PAL	0.9500	0.9500	6
Diet	Gemini 2.0 Flash	0.9763	0.9858	3
	GPT-4o	0.9858	0.9858	3
	QwenVL	0.6066	0.9005	21
	QwenVL PAL	0.9242	0.9336	14
Road Injuries	Gemini 2.0 Flash	0.9868	0.9868	2
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.3816	0.9145	13
	QwenVL PAL	0.9474	0.9539	7
Walmart	Gemini 2.0 Flash	1.0000	1.0000	0
	GPT-4o	0.9817	1.0000	0
	QwenVL	0.4128	0.7431	28
	QwenVL PAL	0.9908	0.9908	1

achieve high success rates (≥ 0.97) with near-universal attempt rates (≈ 1.00) and negligible null counts across all settings. For example, *Gemini2.0 Flash* and *GPT-4o* attain success rates between 0.97 and 1.00 with zero or near-zero nulls throughout Table 20. Weaker models exhibit a distinct pattern: while attempt rates remain high (typically 0.90–0.98), success drops substantially (e.g., 0.12–0.61), accompanied by elevated null counts (often > 20). This decoupling of attempt and success indicates that failures arise from imprecise aggregation rather than abstention. Importantly, mean inference performance does not correlate with time series length. Domains with long horizons show comparable success and null rates to short-horizon settings, confirming that temporal span is not a limiting factor.

Min Inference. Minimum inference is substantially more fragile than mean inference. Although

Table 21: Minimum inference performance across domains and models.

Domain	Model	Success	Attempt	Null
Agriculture	Gemini 2.0 Flash	0.9641	0.9641	6
	GPT-4o	0.9940	0.9940	1
	QwenVL	0.3533	0.8503	25
	QwenVL PAL	0.9641	0.9641	6
Border Crossing	Gemini 2.0 Flash	0.9721	0.9907	2
	GPT-4o	0.9876	0.9969	0
	QwenVL	0.3839	0.9505	16
	QwenVL PAL	0.9505	0.9876	4
COVID	Gemini 2.0 Flash	0.9000	0.9600	12
	GPT-4o	0.9567	0.9933	2
	QwenVL	0.3533	0.8133	56
	QwenVL PAL	0.9633	0.9800	6
Crime	Gemini 2.0 Flash	0.9869	0.9935	1
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.1765	0.9804	3
	QwenVL PAL	0.9804	0.9869	2
Demography	Gemini 2.0 Flash	0.9667	0.9667	4
	GPT-4o	0.9917	0.9917	1
	QwenVL	0.5083	0.5083	59
	QwenVL PAL	0.9333	0.9417	7
Diet	Gemini 2.0 Flash	0.9668	0.9763	5
	GPT-4o	0.9763	0.9763	5
	QwenVL	0.5261	0.8863	24
	QwenVL PAL	0.9289	0.9526	10
Road Injuries	Gemini 2.0 Flash	1.0000	1.0000	0
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.6382	0.7829	33
	QwenVL PAL	0.9605	0.9605	6
Walmart	Gemini 2.0 Flash	0.9817	0.9817	2
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.6606	0.9908	1
	QwenVL PAL	0.9908	0.9908	1

attempt rates remain high across models (typically 0.85–1.00), success rates drop sharply for weaker models, and nulling increases significantly. For example, QwenVL exhibits success as low as 0.18–0.53, with null counts exceeding 30–50 in several settings (Table 21). In contrast, stronger models maintain success rates above 0.95 with minimal nulling. Crucially, this fragility does not align with time series length. Short sequences can yield failure rates comparable to or worse than long sequences. Qwen PAL variants consistently reduce nulling (e.g., null counts reduced by 2–4× relative to base Qwen) and improve success, but still fail to match the robustness of stronger closed models.

Max Inference. Maximum inference is the most robust extrema-related operation. Across nearly all models and settings, attempt rates are ≈ 1.00 , omission is minimal, and success rates are consistently high. Strong models achieve success rates of

Table 22: Maximum inference performance across domains and models.

Domain	Model	Success	Attempt	Null
Agriculture	Gemini 2.0 Flash	0.9760	0.9760	4
	GPT-4o	0.9940	0.9940	1
	QwenVL	0.9042	0.9880	2
	QwenVL PAL	0.9820	0.9940	1
Border Crossing	Gemini 2.0 Flash	0.9783	0.9907	2
	GPT-4o	0.9907	0.9969	0
	QwenVL	0.7771	0.9783	7
	QwenVL PAL	0.9536	0.9938	2
COVID	Gemini 2.0 Flash	0.9667	0.9733	8
	GPT-4o	0.9900	0.9933	2
	QwenVL	0.5367	0.8933	32
	QwenVL PAL	0.9800	1.0000	0
Crime	Gemini 2.0 Flash	0.9935	0.9935	1
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.7386	0.9935	1
	QwenVL PAL	0.9804	1.0000	0
Demography	Gemini 2.0 Flash	0.9667	0.9667	4
	GPT-4o	0.9917	0.9917	1
	QwenVL	0.6833	0.6917	37
	QwenVL PAL	0.9333	0.9750	3
Diet	Gemini 2.0 Flash	0.9810	0.9810	4
	GPT-4o	0.9810	0.9810	4
	QwenVL	0.8341	0.9479	11
	QwenVL PAL	0.9336	0.9953	1
Road Injuries	Gemini 2.0 Flash	0.9737	0.9737	4
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.8487	0.9079	14
	QwenVL PAL	0.9605	0.9934	1
Walmart	Gemini 2.0 Flash	0.9725	1.0000	0
	GPT-4o	1.0000	1.0000	0
	QwenVL	0.2569	0.9908	1
	QwenVL PAL	0.9908	1.0000	0

0.97–1.00, while even weaker models often exceed 0.73–0.90 success (Table 22). Notably, high success coincides with aggressive attempt behavior: failures are not driven by abstention. For instance, QwenVL attempts maximum inference in > 97 of cases but still exhibits reduced success in some settings, indicating semantic or conceptual errors rather than omission. Models exhibit a directional asymmetry in extrema reasoning: identifying peaks is substantially easier than identifying troughs, and high maximum scores are not artifacts of selective answering.

STD Inference. Standard deviation inference exposes a **fundamental capability gap** across all evaluated models. Unlike the other statistics, models frequently attempt STD inference (attempt rates often 0.90–1.00), yet produce null or unscorable outputs in the majority of cases. This results in near-zero success across many settings. For ex-

Table 23: Standard deviation inference performance across domains and models.

Domain	Model	Success	Attempt	Null
Agriculture	Gemini 2.0 Flash	0.0000	0.9940	167
	GPT-4o	0.0000	1.0000	167
	QwenVL	0.0000	0.7246	167
	QwenVL PAL	0.0000	0.9581	167
Border Crossing	Gemini 2.0 Flash	0.4149	0.9412	18
	GPT-4o	0.4551	0.9907	2
	QwenVL	0.0248	0.9876	4
	QwenVL PAL	0.7678	0.9659	11
COVID	Gemini 2.0 Flash	0.3633	0.9700	9
	GPT-4o	0.3367	1.0000	0
	QwenVL	0.0100	0.9267	22
	QwenVL PAL	0.8000	0.9567	13
Crime	Gemini 2.0 Flash	0.4771	0.9935	1
	GPT-4o	0.4771	1.0000	0
	QwenVL	0.0458	0.9869	2
	QwenVL PAL	0.9608	0.9804	3
Demography	Gemini 2.0 Flash	0.5167	1.0000	0
	GPT-4o	0.2833	0.9917	1
	QwenVL	0.0417	0.8833	14
	QwenVL PAL	0.0333	0.9333	8
Diet	Gemini 2.0 Flash	0.0000	0.9716	211
	GPT-4o	0.0000	0.9953	211
	QwenVL	0.0000	0.7536	211
	QwenVL PAL	0.0000	0.9336	211
Road Injuries	Gemini 2.0 Flash	0.2566	0.9605	6
	GPT-4o	0.1316	1.0000	0
	QwenVL	0.0395	0.9342	10
	QwenVL PAL	0.0197	0.9013	15
Walmart	Gemini 2.0 Flash	0.0000	0.9725	109
	GPT-4o	0.0000	0.9817	109
	QwenVL	0.0000	0.8073	109
	QwenVL PAL	0.0000	0.9908	109

ample, several domains exhibit success = 0.00 with null counts exceeding 100–200 despite high attempt rates (Table 23). This failure pattern is not driven by abstention, nor by sequence length. Even short-horizon settings show complete failure. PAL increases attempt rates but rarely improves success, demonstrating that the issue is not instruction following but the absence of internalized variance modeling. STD inference requires explicit reasoning about dispersion, which current models do not reliably internalize. Unlike extrema, variance cannot be approximated by semantic heuristics.

P Q&A Tasks

We present accuracy scores for VLMs on the Q&A task in Table 24. An analysis of the highlighted statistics reveals a striking contrast between the finetuned and pretrained models. The finetuned model frequently produces highly confident yet sometimes incorrect predictions, whereas the pre-

trained model demonstrates more caution, acknowledging that the mean is lower than expected without attempting to estimate a specific value. Notably, certain proprietary models are now reaching, and at times even surpassing, human performance on specific subsets of tasks. While this signals exciting progress in the field, it also highlights the nuances of human cognitive performance, particularly under conditions where distraction might occur. It is vital to note, however, that no singular model has consistently achieved near-human proficiency across the entirety of the benchmark’s demands. The plot retrieval task, in particular, stands out as a significant hurdle, robustly affirming the unparalleled human capacity for holistic visual-numeric integration, a critical frontier for time series understanding.

Table 24: Model accuracy for time series Q&A tasks. **Bold** and underlined scores denote first and second places (excluding human). Cap, Plt, and TS refer to caption, plot, and time series matching, whereas the comparative reasoning task is split into Amp, Peak, Mean, and Var, denoting amplitude, peak-earlier, mean, and variance comparison.

Model	Cap	Plt	TS	Amp	Peak	Mean	Var
Proprietary							
Gemini2 F	<u>0.78</u>	0.30	<u>0.61</u>	0.80	0.42	<u>0.70</u>	0.62
Gemini2.5	0.66	0.30	0.31	1.00	1.00	1.00	0.85
Claude3 H	0.68	0.29	0.57	0.65	0.40	0.53	0.33
GPT-4o	0.96	<u>0.31</u>	0.77	0.83	0.73	<u>0.70</u>	0.63
Pretrained							
InternVL 2.5	0.55	0.17	0.49	0.60	0.47	0.45	0.40
LLaVA v1.6	0.39	0.27	0.32	0.45	0.45	0.42	0.45
Phi-4 M.I.	0.62	0.29	0.45	0.70	0.82	0.68	<u>0.70</u>
Idefics 2	0.49	0.25	0.29	0.35	0.40	0.40	0.50
SmolVLM	0.26	0.34	0.28	0.40	0.48	0.44	0.60
QwenVL	0.68	0.27	<u>0.61</u>	0.70	0.50	0.60	0.40
Llama 3.2	0.66	0.24	0.27	0.45	0.63	0.43	0.30
<i>Human</i>	0.81	0.95	0.83	0.93	0.85	0.95	0.90

P.1 Qwen-Based Filtering

To show that questions erroneously answered by Qwen 2.5 Omni are indeed harder, we evaluated a subset of models on both an easy set of 600 questions and the hard set generated by Qwen 2.5 Omni. The questions in the easy set are randomly sampled from those correctly answered by Qwen 2.5 Omni. Table 25 depicts the comparison. All models, regardless of architecture, show a consistent performance gain on the “easy” subset, demonstrating that Qwen-filtered questions are broadly harder, not uniquely harder for Qwen. To ensure a balanced benchmark, we release both the full set (38.4k ques-

tions) and the hard subset (7k questions), enabling evaluation and training across the entire difficulty spectrum.

Table 25: Performance on easy vs. hard questions and corresponding lift.

Model	Easy	Hard	Lift
Idefics 2	65%	46%	+19%
InternVL 2.5	72%	43%	+29%
Phi-4	61%	46%	+15%
SmolVLM	53%	48%	+5%
Llama-3.2	54%	49%	+5%

P.2 Distractor Generation in Q&A Tasks

To increase task difficulty, artificial perturbations are applied in the Time Series Matching and Caption Matching tasks. As shown in the table below, these perturbations significantly impacted model performance in Time Series Matching, increasing task difficulty and forcing models to reason over trends rather than relying on superficial cues.

To illustrate the rationale of this design choice, assume the correct time series option is [1, 2, 3]. Having distractors like [101, 99, 102] makes the question trivial due to its totally different scale and nature. Our distractors are generated by perturbations on the correct time series, resulting in the following distractors.

1. **Shuffled:** [2, 1, 3], avoids answering correctly by exploiting numeric lookup without temporal order awareness.
2. **Reversed:** [3, 2, 1], avoids reasoning without trend awareness.
3. **Gaussian-noised:** [1.03, 1.99, 3.002], forces precise numeric reasoning instead of superficial numeric and trend similarity.

Table 26: Qwen-2.5-Omni-7B accuracy by task and distractor type.

Question Type	Distractor Type	Accuracy
TS Matching	Cross-domain	0.9803
TS Matching	Same-domain	0.9586
TS Matching	Artificially Perturbed	0.6864
Caption Matching	Cross-domain	0.8325
Caption Matching	Same-domain	0.8250
Caption Matching	Artificially Perturbed	0.8532

Table 26 shows that these perturbations indeed augment question difficulty. For Caption Matching,

we do not find a significant difference between artificially perturbed distractors compared to original negative examples drawn from the dataset. We release the set of Q&A questions without artificial perturbations as well. Perturbations are not applied to the other tasks; negative samples are instead drawn directly from other time series in the dataset.

P.3 Temporal Matching Task

Figure 17 illustrates the temporal bucket structure used across our temporal matching diagnostic questions.

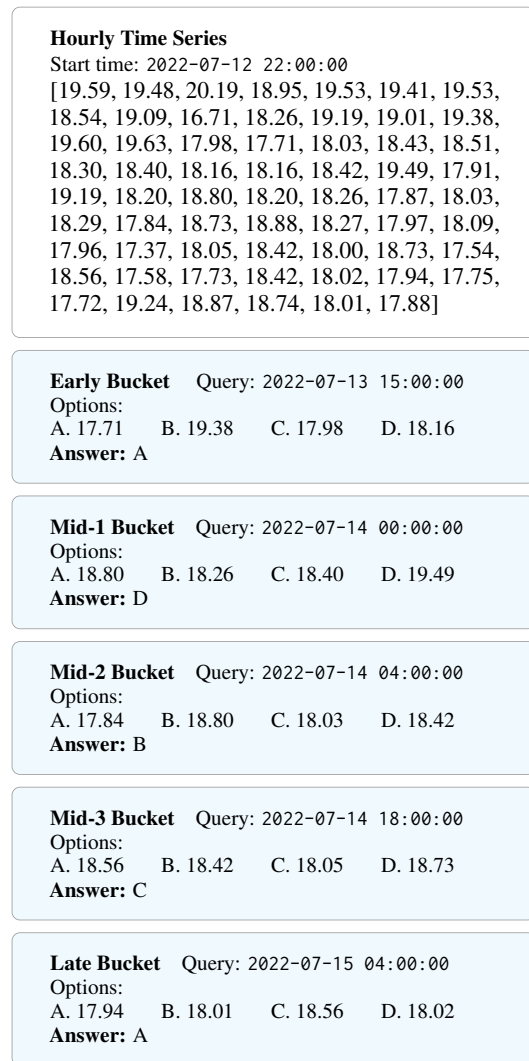


Figure 17: Illustration of temporal buckets used in our diagnostic MCQ task. All questions query the same hourly time series, with increasing temporal distance from the start timestamp, isolating temporal indexing difficulty.

Distractor Generation for Temporal Matching

Distractor options are generated using a temporal locality-driven sampling strategy. Given a time

series $\mathbf{x} = \{x_1, \dots, x_N\}$ and a correct index i^* , we first collect candidate values from a symmetric temporal window centered at i^* . If the number of unique candidates is insufficient, the window is expanded to a larger radius, with a final fallback to sampling from the entire series. This ensures distractors are temporally plausible while avoiding trivial elimination strategies (Algorithm 1).

Algorithm 1 Temporal Locality–Driven Distractor Sampling (TLDS)

Require: Time series \mathbf{x} , correct index i^* , window size w , maximum window W , number of options K

Ensure: $K - 1$ distractor values

- 1: Initialize empty candidate set \mathcal{C}
- 2: **for** $o = -w$ to w **do**
- 3: $i \leftarrow i^* + o$
- 4: **if** $0 \leq i < N$ **and** $i \neq i^*$ **then**
- 5: **if** $x_i \neq x_{i^*}$ **then**
- 6: Add x_i to \mathcal{C}
- 7: **end if**
- 8: **end if**
- 9: **end for**
- 10: **if** $|\mathcal{C}| < K - 1$ **then**
- 11: Repeat the above procedure using window size W
- 12: **end if**
- 13: **if** $|\mathcal{C}| < K - 1$ **then**
- 14: $\mathcal{C} \leftarrow \{x_1, \dots, x_N\} \setminus \{x_{i^*}\}$
- 15: **end if**
- 16: Sample $K - 1$ distinct values uniformly from \mathcal{C}
- 17: **return** sampled distractors

Index Lookup Baseline. We implement an index-based lookup baseline (Algorithm 2) that deterministically retrieves the ground-truth value for a queried timestamp when exact temporal alignment is possible. Given a start time and query time, the baseline computes the hour offset using strict datetime parsing and requires the offset to be an exact integer. If the resulting index is within bounds, the corresponding time series value is retrieved and matched against the multiple-choice options using exact numeric equality, with no rounding or tolerance. The baseline returns an answer only when a unique exact match exists; otherwise, it abstains. This baseline serves as a correctness upper bound for the task and isolates model failures from ambiguities in data representation or parsing.

Algorithm 2 Temporal Index Lookup Baseline

Require: Time series $\mathbf{x} = [x_0, \dots, x_{N-1}]$, start time t_0 , query time t_q , option set \mathcal{O}

Ensure: Predicted option label or failure

- 1: Parse t_0 and t_q using ISO datetime parsing
- 2: Compute time offset $\Delta t \leftarrow t_q - t_0$
- 3: Convert offset to hours: $h \leftarrow \Delta t / 3600$
- 4: **if** h is not an integer **then**
- 5: **return** failure
- 6: **end if**
- 7: $i \leftarrow h$
- 8: **if** $i < 0$ OR $i \geq |\mathbf{x}|$ **then**
- 9: **return** failure
- 10: **end if**
- 11: $v^* \leftarrow x_i$
- 12: Parse option values using numeric parsing
- 13: $\mathcal{M} \leftarrow \{k \mid \mathcal{O}[k] = v^*\}$
- 14: **if** $|\mathcal{M}| = 1$ **then**
- 15: **return** the unique label in \mathcal{M}
- 16: **else**
- 17: **return** failure
- 18: **end if**

P.4 Sample Q&A Questions

We provide examples of Q&A questions in Figures 18, 19, 21, 22, 25, 26, 20, 23, and 24, covering one example per question type.

Question

Given two time series A and B, detect which one has a higher amplitude defined as the maximum - minimum.

A: [1.15, 0.92, 0.85, 0.75, 0.57, 0.62, 0.6, 0.5, 0.68, 0.72, 0.8, 0.67, 0.8, 0.55, 0.55, 0.7, 0.88]

B: [87.0, 83.0, 77.0, 74.0, 84.0]

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:
{"answer": <string>}
<string> must be an answer string containing only A, B.
Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

"answer": "B"

Figure 18: Example of a *time series amplitude comparison* question.

Question

Given two time series A and B, detect which one reaches its maximum earlier.

A: [66.76, 83.06, 85.77, 90.77, 98.81, 90.62, 80.05, 91.36, 89.59, 76.4, 80.1, 85.6, 84.41]
 B: [949.0, 689.0, 561.0, 552.0, 563.0]

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:
 {"answer": <string>}
 <string> must be an answer string containing only A, B.

Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

"answer": "B"

Figure 19: Example of a *time series peak comparison* question.

Question

Given the following two time series A and B, please identify which one has higher overall values.

A: [65.0, 65.0, 64.0, 37.0, 55.0, 51.0]
 B: [6.29, 6.29, 6.29, 7.0, 7.0, 7.0, 7.0, 6.71, 6.71, 6.71, 6.71, 6.717, 7.57, 7.57, 7.14, 7.14, 7.14, 7.14, 7.43]

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:
 {"answer": <string>}
 <string> must be an answer string containing only A, B.

Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

"answer": "A"

Figure 21: Example of a *time series mean comparison* question.

Question

Here is a time series caption:
 From 2014 to 2019, Bulgaria's Agricultural output index (2015=100) generally increased, starting at 103.4 in 2014 and reaching a peak of 109.23 in 2019, with a slight dip to 100.0 in 2015. The average output index during this period was 105.4, notably lower than the historical mean of 126.73, suggesting a period of relatively lower agricultural productivity compared to Bulgaria's long-term performance. The increase from 2015 to 2019 indicates a moderate recovery and growth phase-t within this specific timeframe.

What time series is best described by this caption?
 (A) [109.23, 107.24, 108.45, 104.1, 100.0, 103.4]
 (B) [108.45, 100.0, 104.1, 107.24, 103.4, 109.23]
 (C) [103.9, 99.8, 104.1, 109.2, 106.8, 109.23]
 (D) [103.4, 100.0, 104.1, 108.45, 107.24, 109.23]

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:
 {"answer": <string>}
 <string> must be an answer string containing only A, B, C, or D.

Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

"answer": "D"

Figure 20: Example of a *time series matching* question.

Question

Given the following two time series A and B, please identify which one has higher volatility.

A: [0.14, 0.14, 0.14, 0.29, 0.29, 0.29, 0.29, 0.29, 0.29, 0.29, 0.57, 0.57, 0.57, 0.57, 0.57, 0.57]
 B: [0.21, 0.33, 0.41, 0.39, 0.44, 0.35, 0.35, 0.43, 0.51, 0.65, 0.69, 0.74]

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:
 {"answer": <string>}
 <string> must be an answer string containing only A, B.

Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

"answer": "A"

Figure 22: Example of a *time series variance comparison* question.

Question

You are given an hourly time series.

Time series:
 [19.59, 19.48, 20.19, 18.95, 19.53, 19.41, 19.53, 18.54, 19.09, 16.71, 18.26, 19.19, 19.01, 19.38, 19.6, 19.63, 17.98, 17.71, 18.03, 18.43, 18.51, 18.3, 18.4, 18.16, 18.16, 18.42, 19.49, 17.91, 19.19, 18.2, 18.8, 18.2, 18.26, 17.87, 18.03, 18.29, 17.84, 18.73, 18.88, 18.27, 17.97, 18.09, 17.96, 17.37, 18.05, 18.42, 18.0, 18.73, 17.54, 18.56, 17.58, 17.73, 18.42, 18.02, 17.94, 17.75, 17.72, 19.24, 18.87, 18.74, 18.01, 17.88]

The series starts at:
 2022-07-12 22:00:00

Question:
 What is the value of the time series at the following time?
 2022-07-13 15:00:00

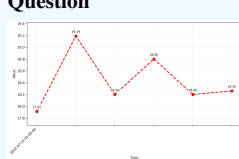
Options:
 A. 17.71 B. 19.38 C. 17.98 D. 18.16

You MUST answer using exactly one capital letter.
 Final answer (one letter only):

Answer
 A

Figure 23: Example of a *time series temporal matching* question.

Question



You are given an hourly time series shown as a plot. Only the starting time is labeled on the x-axis. Each point represents one hour.

Start time: 2022-07-14 01:00:00

Question: What is the value of the time series at the following time?
 2022-07-14 04:00:00

Options:
 A. 18.26 B. 18.8 C. 19.19 D. 18.2

Answer with exactly one capital letter from {A, B, C, D}. Do not explain anything. Just answer in one capital letter. Answer:

Answer
 B

Figure 24: Example of a *plot-based temporal matching* question.

Question

Here is a time series:
 37.00, 37.00, 37.00, 37.00, 37.00, 40.57, 40.57, 40.57, 40.57, 40.57, 40.57

What caption best describes this time series?

(A) From May 1st to July 26th, 2024, the daily COVID-19 deaths in China show a fluctuating pattern, with values generally ranging between 0.29 and 2.86. There are periods of relative stability, such as the initial days of May with a consistent 0.86, interspersed with occasional spikes to 2.86, and dips to 0.29 towards the end of July. Compared to the general daily death statistics for China, where the mean is 73.0 and the maximum reaches 6812.0, this specific time series indicates a period of significantly lower daily deaths, suggesting a substantial improvement in the COVID-19 situation during this timeframe.

(B) From October 24, 2023, to November 3, 2023, the daily COVID-19 cases in Luxembourg show a relatively stable pattern, beginning at 37.3 cases and rising to 40.57 cases by October 29, 2023, where it remains for the rest of the period. Compared to the country's general statistics, where the mean is 236, the daily cases during this period are significantly lower, suggesting a period of reduced viral transmission. This trend does not follow any expected seasonal patterns, as COVID-19 case numbers are known to fluctuate unpredictably.

(C) From October 24, 2023, to November 3, 2023, the daily COVID-19 cases in Luxembourg show a relatively stable pattern, beginning at 37 cases and rising to 40.57 cases by October 29, 2023, where it remains for the rest of the period. Compared to the country's general statistics, where the mean is 236.0, the daily cases during this period are significantly lower, suggesting a period of reduced viral transmission. This trend does not follow any expected seasonal patterns, as COVID-19 case numbers are known to fluctuate unpredictably.

(D) From October 24, 2023, to November 3, 2023, the daily COVID-19 cases in Luxembourg show a relatively unstable pattern, beginning at 37 cases and decreasing to 40.57 cases by October 29, 2023, where it remains for the rest of the period. Compared to the country's general statistics, where the mean is 236.0, the daily cases during this period are significantly lower, suggesting a period of reduced viral transmission. This trend does follow expected seasonal patterns, as COVID-19 case numbers are known to fluctuate unpredictably.

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:

```

{"answer": <string>}
<string> must be an answer string containing only A, B, C, or D.
  
```

 Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer
 "answer": "C"

Figure 25: Example of a *caption matching* question.

Q CaTS-Bench Samples

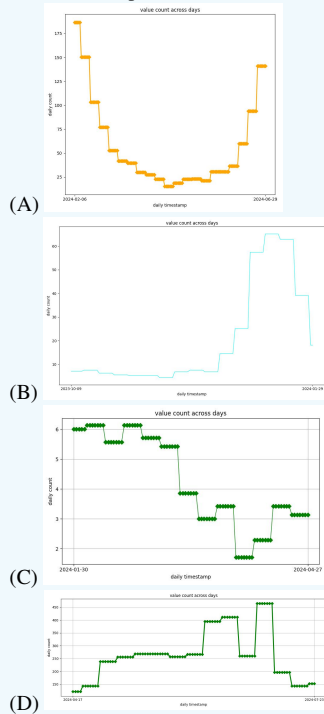
Representative CaTS-Bench samples are shown in Figures 27, 28, 29, 30, and 31.

Question

Here is a time series:

186.57, 186.57, 186.57, 186.57, 186.57, 150.29, 150.29, 150.29, 150.29, 150.29, 103.14, 103.14, 103.14, 103.14, 103.14, 103.14, 77.00, 77.00, 77.00, 77.00, 77.00, 52.71, 52.71, 52.71, 52.71, 41.71, 41.71, 41.71, 41.71, 41.71, 41.71, 41.71, 39.71, 39.71, 39.71, 39.71, 39.71, 39.71, 29.86, 29.86, 29.86, 29.86, 29.86, 29.86, 29.86, 27.43, 27.43, 27.43, 27.43, 27.43, 27.43, 27.43, 22.57, 22.57, 22.57, 22.57, 22.57, 22.57, 22.57, 22.57, 15.14, 15.14, 15.14, 15.14, 15.14, 15.14, 15.14, 18.71, 18.71, 18.71, 18.71, 18.71, 22.71, 22.71, 22.71, 22.71, 22.71, 23.14, 23.14, 23.14, 23.14, 23.14, 23.14, 21.00, 21.00, 21.00, 21.00, 21.00, 21.00, 30.57, 30.57, 30.57, 30.57, 30.57, 30.57, 30.57, 30.57, 30.57, 30.57, 36.29, 36.29, 36.29, 36.29, 36.29, 36.29, 36.29, 59.71, 59.71, 59.71, 59.71, 59.71, 59.71, 59.71, 93.71, 93.71, 93.71, 93.71, 93.71, 140.86, 140.86, 140.86, 140.86, 140.86

Here are four plots of different time series:



Which plot corresponds to the time series provided above?

You must respond only with valid JSON, and no extra text or markdown.

The JSON schema is:

```
{"answer": <string>}
```

<string> must be an answer string containing only A, B, C, or D.

Ensure your output parses as JSON with exactly one top-level object containing the answer field.

Answer

```
"answer": "A"
```

Figure 26: Example of a *plot matching* question.


Time Series Segment

7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 8.0, 25.5, 42.0, 163.33, 258.0, 214.5, 322.5, 354.75, 402.0, 182.33, 141.25, 69.25, 47.0, 12.5, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 11.25, 84.0, 194.5, 338.75, 374.75, 427.25, 272.75, 332.67, 377.75, 232.33, 111.67, 113.25, 23.5, 10.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 7.0, 9.0, 72.25, 91.0, 54.5, 213.0, 299.25, 233.75, 259.0, 390.75, 367.75, 285.25, 207.75, 63.25, 9.5, 7.0, 7.0, 6.75, 6.75, 7.0, 7.0

Metadata JSON

```
{ "all-time maximum": 730.0, "all-time average value until today": 127.62, "all-time minimum": 0.0, "all-time standard deviation until today": 175.79, "average value in this time series": 105.78, "city": "Visakhapatnam", "maximum value in this time series": 427.25, "measure": "SR (W/mt2)", "minimum value in this time series": 6.75, "sampling frequency": "hourly", "standard deviation in this time series": 133.47, "start_month": "July", "start_year": 2017, "starting time": "2022-06-27 22:00:00", "state": "Andhra Pradesh", "station_location": "GVM Corporation, Visakhapatnam " }
```

Line Plot Image



Caption

The hourly solar radiation (sr) data from Visakhapatnam, starting on June 27, 2022, exhibits a clear daily pattern of low values around 7 w/m² during the night and early morning, sharply increasing to peaks during daylight hours, with a maximum value of 427.25 w/m². compared to the city's all-time average of 127.62 w/m², the average value in this time series is 105.78 w/m². The data follows a consistent diurnal cycle, with repeated peaks during the day and low values at night, showing a stable daily pattern.

Figure 27: Sample 1 showing time series data, metadata, plot image, and reference caption.


Time Series Segment

57, 75, 100, 93, 48, 71, 38, 40, 51

Metadata JSON

```
{ "border": "US-Canada Border", "end date of the series": "2020-02-01", "general maximum in the history of this port": 1922, "general mean in the history of this port": 70.79, "general minimum in the history of this port": 0, "general standard deviation in the history of this port": 108.73, "maximum in this specific series": 100, "mean of this specific series": 63.67, "means": "Trucks", "minimum in this specific series": 38, "port": "Del Bonita", "sampling frequency": "monthly", "standard deviation of this specific series": 21.17, "start date of the series": "2019-05-01", "state": "Montana" }
```

Line Plot Image



Caption

From May 2019 to February 2020, the Del Bonita port recorded noticeable variability in truck crossings. The series begins with 57 crossings in May 2019, rises to a maximum of 100 in July 2019, and declines to 38 in November 2019. The mean of 63.67 for this period remains below the historical average of 70.79, indicating slightly reduced throughput relative to the port's long-term behavior. The standard deviation of 21.17, substantially lower than the historical variability of 108.73, reflects constrained fluctuation levels. Overall, this period represents a phase of subdued yet stable truck movement with limited deviations from month to month and a high peak at the start.

Figure 28: Sample 2 showing time series data, metadata, plot image, and reference caption.


Time Series Segment

0.52, 0.32, 0.30, 0.38, 0.41, 0.51, 0.43, 0.41, 0.47

Metadata JSON

```
{ "attribute": "co2_emissions", "country": "Djibouti", "end year of this series": 2018, "maximum of this specific series": 0.52, "mean of this specific series": 0.42, "minimum of this specific series": 0.30, "population at the end year": 1071886.0, "population at the start year": 930251.0, "region": "Middle East & North Africa", "sampling frequency": "yearly", "standard deviation of this specific series": 0.07, "start year of this series": 2010 }
```

Line Plot Image



Caption

Djibouti's CO2 emissions from 2010 to 2018 show fluctuations around the average of 0.42 million metric tons, with a standard deviation of 0.07. Starting at 0.52 million metric tons in 2010, emissions generally decreased to a low of 0.3 million metric tons by 2012, before experiencing some increases and decreases, ending the period at 0.47 million metric tons in 2018.

Figure 29: Sample 3 showing time series data, metadata, plot image, and reference caption.


Time Series Segment

90.0, 90.0, 104.0, 104.0, 104.0, 104.0, 104.0, 104.0, 104.0, 110.57, 110.57, 110.57, 110.57, 110.57, 110.57, 110.57, 121.29, 121.29, 121.29, ..., 268.86, 268.86, 257.29, 257.29, 257.29, 257.29, 257.29, 257.29, 257.29, 257.29, 266.14, 266.14, 266.14, 266.14, 266.14, 266.14, 266.14, 394.57, 394.57, 394.57, 394.57, 394.57, 394.57, 394.57, 394.57, 411.57, 411.57, 411.57, 411.57, ..., 196.71, 196.71, 196.71, 143.43, 143.43, 143.43, 143.43, 143.43, 143.43, 143.43, 152.43

Metadata JSON

```
{ "attribute": "cases", "country": "Thailand", "end date of this series": "2024-07-21", "historical maximum in this country": 26073.0, "historical mean in this country": 2877.0, "historical minimum in this country": 0.0, "historical standard deviation in this country": 5778.0, "income group": "Low & Middle Income", "maximum of this specific series": 465.14, "mean of this specific series": 240.7, "minimum of this specific series": 90.0, "population": 71697024, "region": "East Asia & Pacific", "sampling frequency": "daily", "standard deviation of this specific series": 106.03, "start date of this series": "2024-03-29" }
```

Line Plot Image



Caption

From March 29th to July 21st, 2024, the daily COVID-19 cases in Thailand fluctuated, starting at 90.0 and peaking at 465.14 around mid-June. This peak is significantly lower than the general maximum of 26073.0 daily cases observed in Thailand, indicating a substantial decrease in case numbers during this period. The time series shows an overall pattern of initial stability, followed by increases and decreases, with a final value of 152.43, suggesting a moderate decline towards the end of the observed period.

Figure 30: Sample 4 showing time series data, metadata, plot image, and reference caption.


Time Series Segment

83.65, 81.28, 86.98, 93.55, 96.89, 100.21, 101.88, 100.0, 103.96, 109.67

Metadata JSON

```
{ "attribute": "Total Factor Productivity index (2015=100)",  
  "country": "India", "end year of this series": 2017, "historical  
max": 115.04, "historical mean": 68.79, "historical min": 46.19,  
  "maximum of this specific series": 109.67, "mean of this specific  
series": 95.81, "metrics info": "The Total Factor Productivity  
index (2015=100) is computed as the Output quantity divided by  
aggregated inputs.", "minimum of this specific series": 81.28,  
  "sampling frequency": "yearly", "start year of this series": 2008  
}
```

Line Plot Image



Caption

India's Total Factor Productivity index demonstrates consistent upward momentum from 2008 to 2017. Starting at 83.65 in 2008, the index experiences an initial dip to 81.28 in 2009 before steadily climbing to a peak of 109.67 in 2017, a 35% increase from the initial value. This growth trajectory, with a mean of 95.81, suggests a period of increasing efficiency in India's production processes during this decade, contrasting with the lower historical mean of 68.79 and remaining well above the historical minimum of 46.19. The overall pattern reveals sustained growth with minimal volatility following the initial 2009 recovery.

Figure 31: Sample 5 showing time series data, metadata, plot image, and reference caption.