

Prompting Underestimates LLM Capability for Time Series Classification

Anonymous ACL submission

Abstract

Prompt-based evaluations suggest that large language models (LLMs) perform poorly on time-series classification, raising doubts about whether they encode meaningful temporal structure. We show that this conclusion reflects limitations of prompt-based generation rather than the model’s representational capacity by directly comparing prompt outputs with linear probes over the same internal representations. While zero-shot prompting performs near chance, linear probes improve average F1 from 0.15–0.26 to 0.61–0.67, often matching or exceeding specialized time-series models. Layer-wise analyses further show that class-discriminative time-series information emerges in early transformer layers and is amplified by visual and multimodal inputs. Together, these results demonstrate a systematic mismatch between what LLMs internally represent and what prompt-based evaluation reveals, leading current evaluations to underestimate their time-series understanding.

1 Introduction

Time series analysis focuses on patterns in ordered sequences of events and is central to domains such as healthcare (Ayala Solares et al., 2020; Rajkumar et al., 2018; Choi et al., 2016), finance (Fama, 1970; Sirignano and Cont, 2019), and climate science (Reichstein et al., 2019). Time series classification plays a key role in these areas by assigning each sequence to a label based on its temporal structure and dynamics. Common applications include arrhythmia detection, human activity recognition, industrial fault detection, and environmental event identification.

LLMs demonstrate broad generalization across natural language tasks, supporting flexible pattern matching, contextual reasoning, and symbolic manipulation. These capabilities have motivated efforts to extend LLMs to non-linguistic data (Fons et al., 2024; Mai et al., 2024). Recent work has




Figure 1: Does an LLM’s *observed* ability for time series classification match its underlying *potential*?

explored LLMs for time-series forecasting, classification, and analysis (Jain et al., 2023; Gruver et al., 2024; Merrill et al., 2024b; Morales-García et al., 2024), often relying on major modifications and retraining. *However, strong end-to-end performance in these settings does not necessarily imply that LLMs encode discriminative temporal structure, as results may instead reflect task formulation, additional fine-tuning, or prompting choices.*

Merrill et al. (2024a) show that, despite some success in zero-shot forecasting, LLMs struggle to reliably extract discriminative information from time series, performing near-randomly on reasoning and question-answering tasks that require interpreting temporal structure. Others similarly show that removing the LLM component from time-series pipelines often does not degrade performance and can even improve it, calling into question the necessity of language models for these tasks (Tan et al., 2024). In contrast, a growing body of work argues that LLM-centric approaches may enable more general and interactive forms of time series understanding that extend beyond task-specific prediction (Jin et al., 2024b). Together, these mixed findings raise a central question: *do observed limitations reflect a lack of time-series un-*

derstanding in LLMs, or a mismatch between what these models represent and how that information is extracted?

Prior research that has attempted to use LLMs for time-series understanding has not directly assessed the intrinsic representational capacity of LLMs for time-series classification. Instead, performance is often conflated with three factors that obscure the underlying LLM’s contributions. First, *results can depend heavily on prompt engineering choices*, making it difficult to distinguish genuine temporal understanding from prompt-specific behavior. Second, *many approaches rely on task-specific fine-tuning*, which alters the original pre-trained model and prevents evaluation of the LLM’s native capabilities. Third, LLMs are frequently combined with additional learned components such as auxiliary MLPs, transformer layers, or external embedding models, so *improvements cannot be attributed to the language model itself*. As a result, existing evaluations typically treat LLMs as components within end-to-end systems, offering limited insight into whether out-of-the-box LLMs encode temporal patterns that are directly useful for classification.

In this paper, we evaluate prompting and hidden representation probing side by side to assess an LLM’s *observed* ability for time series classification versus its underlying *potential* (see Figure 1). We examine multiple input modalities to determine which are most effective for prompting and which yield representations that are most easily separable by a classifier. Compared with state-of-the-art methods as an upper bound, we find that simple linear probes are competitive with, and in some cases outperform, more complex LLM-based pipelines. Our results show that *LLMs possess substantially greater time series classification representational capacity than is suggested by prompt-only evaluations*. Moreover, *we use linear probing as a diagnostic tool, restricting probe capacity and validating results against a random-weight model (Chrupala et al., 2020; Tenney et al., 2019) baseline to rule out trivial memorization effects*.

To summarize, this work makes three primary contributions. (1) To the best of our knowledge, we are the first to explicitly disentangle prompt-level performance from representational capacity in time series classification. We do this by evaluating LLMs using both direct prompting and linear probing of their hidden states, which provides a clearer view of what temporal information is en-

coded versus what is accessible through generation. (2) We show that simple linear probes over hidden representations consistently match or outperform prompt-based and more complex LLM classification pipelines, indicating that much of the relevant time-series information is present but underused by prompting. (3) Through layer-wise analysis, we find that discriminative temporal structure emerges early in the model, suggesting that current limitations in LLM-based time series classification stem primarily from evaluation interfaces rather than a lack of temporal representations. Together, these results clarify when LLMs are best viewed as reasoning agents, feature extractors, or unnecessary components for time series classification.

2 Related Work

Probing and Hidden Representations in LLMs.

Early work on Layer-wise Relevance Propagation (LRP) introduced methods for attributing model predictions to intermediate representations (Bach et al., 2015) which later extended to more complex architectures (Binder et al., 2016; Arras et al., 2017). Building on this work, probing methods emerged to explicitly measure what information these representations encode by testing their linear separability with lightweight classifiers (Alain and Bengio, 2018).

In recent years, probing methods have increasingly been used as to extract features for a wide range of downstream applications. Zhu et al. (2025) and Wang et al. (2025), show that hidden states can be leveraged to estimate the difficulty an LLM is anticipated to have on a query prior to generation. More broadly, prior work has demonstrated that hidden representations encode rich and actionable signals that can be exploited across diverse settings, including alignment Kong et al. (2024); Li et al. (2023); Zhang et al. (2025), safety Wang et al. (2024); Lu et al. (2025), interpretability Ghahdharlou et al. (2024); Jacobi and Niv (2025), and robustness Lad et al. (2025); Yan et al. (2024).

The most closely related work to ours is TiViT (Roschmann et al., 2025), which repurposes frozen vision and vision–language transformers for time-series tasks by first mapping time series into novel 2D image representations. This input transformation is central to the TiViT framework. In contrast, we study out-of-the-box LLMs and vision–language models (vLLMs), without relying on handcrafted time-series-to-image encodings. In

addition, TiViT does not consider language models or prompt-based evaluation. Our work focuses on how different evaluation interfaces shape conclusions about model capability. We compare prompting and representation probing. For time series understanding, we study what representational power is gained through pretraining (relative to untrained models).

Time-Series Research. Time series research spans a wide range of objectives, with forecasting and imputation forming the historical foundation of the field. Classical statistical methods, such as ARIMA and Kalman filters, model linear temporal dependencies (Box et al., 2015; Hamilton, 1994; Hyndman and Athanasopoulos, 2018), while neural architectures, including LSTMs, convolutional networks, and transformers, enable richer nonlinear and long-range temporal modeling (Hochreiter and Schmidhuber, 1997; Bai et al., 2018; Zhou et al., 2021; Wu et al., 2021; Lim et al., 2021).

Building on this progression, recent work explores how large language models can be adapted for time series analysis. Existing approaches largely follow a small number of recurring patterns (Zhang et al., 2024): direct prompting that treats time series as text, often with modified numerical tokenization (Gruver et al., 2024; Xue and Salim, 2024; Liu et al., 2023; Xie et al., 2023); discretization or quantization schemes that map signals to token sequences (Chung et al., 2023; Duan et al., 2023; Yang et al., 2024; Rubenstein et al., 2023; Yu et al., 2023); and architectures that introduce dedicated time-series encoders aligned with language model representations (King et al., 2023; Sun et al., 2024; Zhou et al., 2023c; Qiu et al., 2023; Zhou et al., 2023a; Chang et al., 2025; Jin et al., 2024a). Related work also explores visual time-series representations processed by vision-language models or uses LLMs as controllers for external analysis tools (Girdhar et al., 2023; Su et al., 2023; Moon et al., 2023; Zhang et al., 2023).

In contrast, we do not propose a new end-to-end system or optimize for task performance. Instead, we use probing as a diagnostic tool to isolate the representational capacity of pretrained LLMs and vLLMs for time series classification. By decoupling model representations from prompt engineering, fine-tuning, and auxiliary components, we provide a more direct analysis of what time-series structure these models encode.

3 Preliminaries and Problem Setup

We study time series classification under LLM and vLLM interfaces. This section introduces the problem definition, the representations used to encode time series data, and the prediction paradigms evaluated in our study.

3.1 Problem Definition

Each time-series dataset is represented as a tensor $X \in \mathbb{R}^{N \times V \times T}$, where N denotes the number of sequences, V the number of variables or channels, and T the temporal length of each sequence. Time series classification is defined as follows: given a training set $\{(X_i, y_i)\}_{i=1}^{N_{\text{train}}}$ and a test set $\{X_j\}_{j=1}^{N_{\text{test}}}$, the objective is to predict class labels $y_j \in \{1, \dots, C\}$ for each observation in the test set by identifying temporal patterns that distinguish between classes.

3.2 Time Series Representations

We consider how LLMs and vLLMs perform with multiple representations of time series data. These representations are used consistently across both prompting and probing methods. The representations are described below:

Text Modality: Digit-Based Tokenization (d). Directly serializing numerical values can lead to unstable tokenization due to byte-pair encoding (BPE), which may fragment or merge digits unpredictably. To mitigate this issue, we adopt a digit-based preprocessing strategy from Gruver et al. (2024). Each numerical value is converted into a digit-space format with fixed precision, where commas separate time steps and decimal points are removed. For example, the time series $[1.0, 20, 0.33]$ is transformed into "1 0 0 , 2 0 0 0 , 0 3 3". This representation ensures that each digit maps to a consistent token and preserves numerical structure without BPE-induced fragmentation.

Visual Representation (v). Each time series is also rendered as an image, using either line plots or spectrograms, depending on the dataset. Plots include appropriate titles, axis labels, and legends for multivariate data. These visualizations provide an alternative modality that leverages the strengths of vLLMs for pattern recognition. Additional details can be found in Appendix D.1.

Multimodal Representation ($d + v$). Finally, we evaluate a multimodal representation that combines digit-based text input with visual input, enabling




Figure 2: Overview of time series representations and prediction paradigms. Raw time series are transformed into digit-based text representations and visualizations. These representations are incorporated into prompts for direct prediction or used to extract hidden representations for layer-wise probing with linear classifiers.

simultaneous access to numerical and visual structure.

3.3 Prediction Paradigms

We study two complementary approaches to time series classification: direct prompt-based prediction and representation probing. Figure 2 provides an overview of the representations and prediction pipelines evaluated in this work.

Prompt-Based Classification. For a given test time series X_j , we construct a prompt $p_j = \tau + q + m$, where τ denotes the task description, q is a multiple-choice question, and $m \in \{d, v, d+v\}$ specifies the input modality. A model f (either an LLM or vLLM) then produces a predicted label $y_j \in \{1, \dots, C\}$ via text generation. Formally, $y_j = f(p_j)$, where $p_j = \phi(X_j, \tau, q, m)$. Please see the Appendix D for a listing of all prompts.

Representation Probing. Using the same prompt p_j , we additionally evaluate whether class-relevant information is encoded in the model’s internal representations. Rather than using the generated output, we extract hidden states from the model. For each layer $\ell \in \{1, \dots, L\}$, we obtain an embedding $h_j^{(\ell)} \in \mathbb{R}^D$, taken from the final-token hidden state.

For each layer, we train an independent logistic regression classifier using the corresponding embeddings. Probes are trained with 5-fold cross validation and a maximum of 1000 iterations to select the regularization parameter C . Formally, $h_j^{(\ell)} = \psi_\ell(f, p_j)$, $\ell \in \{1, \dots, L\}$, and predictions are given by $\hat{y}_j^{(\ell)} = g_\ell(h_j^{(\ell)})$, where g_ℓ denotes a linear probe trained on $\{(h_i^{(\ell)}, y_i)\}_{i=1}^{N_{\text{train}}}$ and evaluated

on $\{h_j^{(\ell)}\}_{j=1}^{N_{\text{test}}}$.

Probing classifiers’ interpretation are sensitive to probe complexity, baselines, and the correlational nature of the analysis (Belinkov, 2022); we address these concerns with additional controls and analyses in Appendix B.

4 Experiments

We evaluate prompting and probing on multiple time-series classification datasets (described in Appendix A.3, descriptive statistics found in Appendix C) and compare them against a range of heuristic, random, and state-of-the-art baselines (Described in Appendix A.1).

Main Findings. Table 1 summarizes the performance of probing and prompting methods relative to state-of-the-art and heuristic baselines across all evaluation datasets. The Majority, Prior, and Uniform baselines achieve average F1 scores of .195, .282, and .296, respectively. These results define absolute lower bounds for random classification. The Random-Probe baseline (More details in Appendix B) achieves an average F1 of .583. This score represents a practical floor for probing-based approaches and reflects how well a linear classifier can extract signal from features without meaningful representations. All probe-based models exceed the Random-Probe baseline, with average F1 scores of .611, .649, and .672. This result alleviates concerns that performance gains are driven solely by supervised learning. Additionally, Table 2 shows that the performance gap between Random-Probe and vLLM-Probe is substantially larger in other modalities.

vLLMs encode useful time-series features but struggle to exploit them during prompting. Across

Model	Method	CTU	EMG	AD	HAR	RWC	TEE	Avg
Attend	TS-Encoder	.848	.985	.732	.933	.726	.524	.791
TS2Vec	TS-Encoder	.631	.933	.667	.865	.672	.794	.761
Moment	TS-Encoder	.660	1.00	.671	.878	.778	.564	.758
OneFitsAll	Encoder Alignment	.587	.296	.681	.885	.790	.475	.615
InstructTime	TS-Quant + Prompt	.627	.167	.278	.527	.434	.235	.378
Random	Probe	.573	.933	.369	.592	.597	.436	.583
	Majority	.333	.167	.015	.048	.434	.036	.195
	Prior	.439	.265	.088	.166	.504	.072	.282
	Uniform	.442	.333	.084	.164	.463	.141	.296
Llama	Probe	.640	1.00	.406	.723	.672	.223	.611
	Prompt	.356	.175	.033	.111	.190	.036	.150
Mistral	Probe	.684	.933	.438	.730	.645	.462	.649
	Prompt	.548	.393	.021	.127	.190	.288	.261
Qwen	Probe	.676	1.00	.374	.737	.656	.592	.672
	Prompt	.324	.413	.016	.174	.222	.304	.242

Table 1: F1 scores across all datasets. Prompt and Probe results pertain only to the $d + v$ modality. Results are organized into three groups: (top) established time-series baselines; (middle) random and heuristic baselines; and (bottom) vLLM Probe and Prompting. Established baselines indicate strong upper-bound performance, while random methods provide a lower-bound reference.

all models, prompting performs substantially worse than probing. For Llama, average F1 increases from .150 with prompting to .611 with probing. Mistral improves from .261 to .649. Qwen improves from .242 to .672. This pattern holds consistently across every dataset and every model. For each case, the probing score exceeds the corresponding prompting score. This effect is most pronounced on the EMG dataset. EMG appears to be an easy classification benchmark. Attend achieves an F1 of .985, and Moment, Llama, and Qwen all achieve perfect separation when probed. Despite this, prompting severely underestimates model capability on this dataset. Llama improves from .175 with prompting to 1.00 with probing. Mistral improves from .393 to .933. Qwen improves from .413 to 1.00. These gaps indicate that the learned representations are strong, but current prompting strategies fail to reliably access them.

vLLMs sometimes outperform state-of-the-art baselines. vLLM probing results are competitive with, and sometimes exceed, established time-series baselines. Qwen (.672) and Mistral (.649) outperform OneFitsAll (.615). Llama narrowly trails at .611. All three vLLMs outperform InstructTime. On individual datasets, vLLMs achieve state-of-the-art performance in several cases. On EMG, Llama and Qwen achieve perfect F1 scores of 1.00, exceeding Attend (.985) and TS2Vec (.933), and matching Moment. On CTU, Llama (.640) and Qwen (.676) outperform all state-of-the-art baselines except Attend. Figure 5 shows two t-SNE

plots comparing Llama v embeddings with Moment embeddings on the HAD dataset. This visualization highlights differences in feature separability and illustrates Llama’s usefulness for classification compared to an established baseline.

Probing and prompting benefit from visual information. Table 2 reports performance across input modalities for both probing and prompting. Probing performance generally improves with the inclusion of v . Llama performs best when using v , achieving an average F1 of .678. Mistral and Qwen perform best with $d + v$, achieving average F1 scores of .631 and .672, respectively. Prompting performance also improves with the inclusion of visual input, particularly for Mistral and Qwen. For Mistral, the v modality achieves the highest average F1 at .354, followed by d at .273. For Qwen, v again performs best with an average F1 of .326, followed by d at .259. This effect is most pronounced on the EMG dataset. For Mistral, prompting with v achieves an F1 of .663, compared to .280 for the other modalities. For Qwen, prompting with v achieves .618, compared to .372 with d and .459 with $d + v$. This subset of results (Mistral & Qwen + EMG + v) is the closest prompting comes to matching probing performance bound by the same modality (Mistral v at .796, Qwen v at .866).

Random-Probe is a strong baseline. The Random-Probe baseline with $d + v$ achieves an average F1 of .583. This score exceeds the probing performance of Llama with d (.547) and Mistral with d (.579). This result indicates that a substantial

Model	Method	Modality	CTU	EMG	HAD	HAR	RWC	TEE	Avg
Llama	Probe	d	.616	.866	.313	.621	.626	.335	.547
		v	.718	.933	.468	.799	.690	.512	.678
		$d + v$	.654	1.00	.376	.708	.672	.223	.597
	Prompt	d	.333	.171	.032	.054	.104	.036	.132
		v	.333	.159	.024	.077	.190	.036	.132
		$d + v$	.344	.171	.033	.111	.190	.036	.160
Mistral	Probe	d	.682	.712	.386	.735	.628	.265	.579
		v	.696	.796	.352	.684	.664	.439	.596
		$d + v$	.698	.933	.393	.729	.645	.462	.631
	Prompt	d	.402	.280	.030	.153	.497	.260	.272
		v	.445	.663	.055	.225	.460	.302	.354
		$d + v$	.511	.280	.021	.173	.190	.200	.254
Qwen	Probe	d	.672	.933	.361	.764	.647	.419	.633
		v	.688	.866	.304	.655	.640	.575	.621
		$d + v$	.676	1.00	.374	.737	.656	.592	.672
	Prompt	d	.393	.372	.039	.146	.432	.159	.259
		v	.341	.618	.028	.220	.296	.246	.326
		$d + v$	.324	.459	.016	.174	.222	.194	.255
Random	Probe	d	.587	.861	.366	.583	.605	.433	.572
		v	.333	.167	.015	.048	.434	.036	.172
		$d + v$	.573	.933	.369	.592	.597	.436	.583

Table 2: F1 for each dataset split by model, method, and modality. All probing results are *much higher* than their corresponding prompt results.

portion of probing performance can be attributed to feature structure alone, independent of semantic alignment. It reinforces that probing gains are not solely an artifact of supervised learning. At the same time, it raises questions about how much semantic understanding is required to achieve competitive performance.

Text-only input is more robust than expected. Text-only input d performs surprisingly well across models and methods. For Qwen, probing with d (.633) is close to its best-performing modality (.672). For Mistral, probing with d (.579) remains competitive with v (.596). Even in prompting, d frequently outperforms the combined $d + v$ modality. This suggests that the textual representation captures meaningful time-series structure.

Layer Analysis. *Information is encoded in early layers.* In both plots of Figure 4, we see a quick improvement. Qwen encodes useful information in layers 1 and 2. Mistral is more gradual but encodes around layer 12 for d and $d + v$ or layer 20 for v . On the Llama side (Figure 3), we observe Llama also peaking early, roughly at layer 5 on $d + v$ and at layer 10 on d alone. We observe that the Random baselines in the same figure immediately return to their mean level at layer 1.

Presence of v boosts over random baseline. In

Figure 3 we see in the d modality that llama and random are intermixed. This compares with the clean separation of $d + v$, the modality in which, starting around layer 5, llama is consistently stronger across layers than the random baseline.

Ablations. Prompting is commonly augmented with in-context examples and chain-of-thought reasoning. We therefore include an ablation to assess the impact of these components. Table 4 reports F1 scores averaged across modalities for different prompting strategies.

Chain-of-thought does not have much effect on representation quality. For Llama and Qwen, average F1 decreases when using CoT, from .609 to .598 and from .667 to .654, respectively. Mistral shows a small improvement, increasing from .638 to .648. Overall, these results indicate that CoT prompting does not reliably enhance the quality of the extracted representations. Similar trends are observed when using two-shot probing, which we report in the appendix.

Similarly, CoT does not consistently improve prompting. Llama benefits from CoT, improving from .174 with direct prompting to .220 with CoT. However, Mistral performs worse with CoT (.296) than with direct zero-shot prompting (.340). Qwen shows no meaningful difference, with CoT




Figure 3: Layer-wise macro F1 of linear probes for Llama versus a random baseline under digit-only (d) and multimodal (d+v) inputs, showing early emergence of discriminative time-series information.




Figure 4: Layer-wise probe performance for Mistral and Qwen across digit (d), visual (v), and multimodal (d+v) inputs, with visual and multimodal representations yielding stronger separability.

(.289) nearly matching direct prompting (.290). These results suggest that CoT benefits are model-dependent and not universal.

In-context examples do not improve prompting performance. Across models, in-context examples either degrade performance or provide negligible gains. For Llama, average F1 decreases from .174 in the zero-shot setting to .163 with two-shot prompting. For Qwen, performance drops from .290 to .280. Mistral shows a small decrease from .340 to .321. These results align with prior findings that few-shot prompting is often ineffective for classification-style tasks.

Taken together, these results indicate that zero-shot direct prompting provides a reasonable and representative evaluation setting. Neither few-shot examples nor in-context learning consistently improve representation quality or downstream prompting performance in this task.

Results extend to text-only LLMs. To evaluate whether our findings depend on multimodal architectures, we repeat our experiments using text-only LLMs. Specifically, we evaluate Gemma and GPT-oss using the *d*-only modality. Table 5 reports the corresponding probe-based and prompt-based results.

Across both models, probing substantially outperforms prompting. Gemma achieves an average F1 of .596 with probing, compared to .206 with prompting. GPT-oss shows a similar pattern, achieving .466 with probing and .202 with prompting. This gap is consistent across all datasets.

These results indicate that the advantages of probing over prompting are not specific to visual inputs or multimodal architectures. Instead, they reflect a more general limitation of prompting as an interface for extracting discriminative representations in time-series classification.




Figure 5: t-SNE visualizations of probe embeddings for Llama (left) and a time-series baseline (Moment, right) on the HAD dataset, illustrating comparable class separation.

Prompt Stability. We hypothesize that performance in LLM-based time series classification is highly sensitive to prompt engineering choices. We evaluate this hypothesis in Table 3, Figure 6, which reports the results from two experiments.

First (orange-shaded columns), we assess *sensitivity to prompt wording* by using gpt-4.1-mini to generate ten meaning-preserving variations of the original prompt. A larger Δ (max—min macro F1) indicates greater variability in performance across prompt variants, and thus higher sensitivity to prompt phrasing. Across datasets and modalities, we observe an average Δ of .088 between the best- and worst-performing prompt variants, indicating substantial sensitivity even under minor prompt modifications. Examples of the prompt-variants can be found in Figures 13, 14, & 15 and the prompts used to generate them in Figures 11 & 12.

Second (blue-shaded columns), we evaluate *output consistency* by prompting the model twenty times with the same prompt. The resulting P@K scores (more information found in Appendix A.4) measure the probability that at least one of K sampled completions is correct. The corresponding $\Delta P@K$ ($P@20 - P@1$) quantifies the benefit of re-

dataset	modality	min	max	mean	median	Δ	P@1	P@20	$\Delta P@K$
CTU	<i>d</i>	.410	.509	.463	.470	.099	.506	.568	.098
	<i>d + v</i>	.335	.444	.380	.389	.109	.471	.548	.000
	<i>v</i>	.357	.434	.394	.395	.077	.457	.556	.052
EMG	<i>d</i>	.167	.395	.202	.167	.228	.343	.400	.057
	<i>d + v</i>	.140	.376	.197	.167	.236	.333	.333	.000
	<i>v</i>	.167	.167	.167	.167	.000	.333	.333	.000
HAD	<i>d</i>	.013	.048	.031	.031	.035	.043	.100	.062
	<i>d + v</i>	.016	.025	.020	.021	.009	.059	.074	.099
	<i>v</i>	.037	.056	.049	.050	.019	.163	.198	.077
HAR	<i>d</i>	.122	.151	.138	.138	.029	.175	.278	.103
	<i>d + v</i>	.088	.145	.130	.137	.057	.274	.314	.034
	<i>v</i>	.080	.157	.143	.152	.077	.310	.344	.040
TEE	<i>d</i>	.000	.081	.027	.016	.081	.140	.238	.057
	<i>d + v</i>	.052	.188	.102	.097	.136	.400	.452	.035
	<i>v</i>	.038	.167	.110	.102	.129	.357	.357	.015
Avg (<i>d</i>)	<i>d</i>	.142	.237	.172	.164	.094	.242	.317	.075
Avg (<i>v</i>)	<i>v</i>	.136	.196	.173	.173	.060	.324	.358	.034
Avg (<i>d+v</i>)	<i>d + v</i>	.126	.236	.166	.162	.109	.308	.344	.037
Total Avg	–	.135	.223	.170	.166	.088	.291	.340	.048

Table 3: Prompt stability results across datasets and modalities. Orange columns reflect sensitivity to prompt wording, and blue columns reflect variability from repeated sampling. Larger Δ (max - min F1) and $\Delta P@K$ indicate less stable prompt-based evaluation.




Figure 6: Prompt-based macro F1 distributions input modalities (*d*, *v*, *d+v*), showing a bit more stable performance for visual-only prompting.

peated sampling. On average, allowing twenty attempts increases accuracy by .048.

Notably, *visual-only prompting (v) is the most stable modality under both criteria*. It exhibits the lowest average prompt-variant sensitivity ($\Delta = .060$) and the smallest sampling gain ($\Delta P@K = .034$), suggesting more consistent behavior both across prompt wording and repeated generations.

Together, these results support our claim that prompt-based evaluation in time series classification is highly sensitive to prompt engineering choices, and that reported performance may reflect prompt design and sampling strategy as much as underlying model capability.

5 Conclusion

This research provides a rigorous reassessment of LLMs’ ability to represent and classify functional time-series data. By contrasting direct prompting with linear probing across diverse architectures and modalities, we identify a persistent gap in representation performance. Our evidence suggests that, although current prompting interfaces fail to capture temporal dynamics, the internal hidden states of these models contain highly discriminative, class-relevant information. The results indicate that the perceived limitations of LLMs in time series tasks are primarily an artifact of the decoding interface rather than a fundamental deficiency in representational capacity. The early emergence of class discriminative signals, often within the first five transformer layers, suggests that these models function as effective general-purpose sequence encoders before higher-level reasoning stages. Furthermore, the competitive performance of simple linear probes against specialized, state-of-the-art time series baselines underscores the quality of pretrained LLMs as robust feature extractors for non-linguistic modalities.

Limitations

This paper is primarily diagnostic. We focus on measuring and explaining the gap between prompt-level performance and representational capacity, rather than proposing a new prompting method to close this gap. Developing more robust prompting strategies is an important direction for future work. While we take care to address standard probing concerns and follow established best practices (see Appendix B), probing remains an imperfect but informative diagnostic.

In particular, strong probe performance should not be interpreted as evidence of high-level temporal reasoning or semantic understanding. Given the nature of the evaluated datasets, it is plausible that probes primarily exploit relatively low-level temporal cues such as amplitude ranges, monotonic trends, periodicity, or short-range autocorrelation. Importantly, these cues are precisely the signals required for accurate classification in many standard time-series benchmarks. The fact that linear probes can reliably extract such task-relevant information, while prompt-based generation cannot, highlights a limitation of prompting as an evaluation interface rather than a deficiency in the underlying representations.

Our random-weight control is limited. We include a single random-weight model to rule out trivial effects from supervised probe training, but extending this analysis to additional random initializations and model families would further strengthen conclusions about architectural separability. Importantly, however, the qualitative gap between probing and prompting is large, consistent across datasets, modalities, and model families, and substantially exceeds the variance introduced by probe training itself, making it unlikely that our core findings are driven by idiosyncrasies of a particular random initialization.

Finally, our evaluation is necessarily scoped. We study six time-series datasets, five language and vision-language models, and a broad set of state-of-the-art and heuristic baselines. In fact, the consistency of the observed trends across these diverse settings suggests that the central result—that prompt-based evaluation systematically underestimates representational capacity—is not sensitive to any single dataset or model choice. Exploring additional datasets, representations, and model scales remains an important direction for future work.

References

2024. [The Llama 3 Herd of Models](#). *arXiv preprint*. ArXiv:2407.21783 [cs].
- Alireza Abedin, Mahsa Ehsanpour, Qinfeng Shi, Hamid RezaTofighi, and Damith C. Ranasinghe. 2021. [Attend and Discriminate: Beyond the State-of-the-Art for Human Activity Recognition Using Wearable Sensors](#). *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(1):1–22.
- Kingma DP Ba J Adam and 1 others. 2014. A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 1412(6).
- Mistral AI. 2025. Mistral-small-3.1-24b-instruct-2503 model card. <https://huggingface.co/mistralai/Mistral-Small-3.1-24B-Instruct-2503>. Model card describing architecture, capabilities, and benchmarks.
- Guillaume Alain and Yoshua Bengio. 2018. [Understanding intermediate layers using linear classifier probes](#). *Preprint*, arXiv:1610.01644.
- D. Anguita, Alessandro Ghio, L. Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. 2013. [A public domain dataset for human activity recognition using smartphones](#). In *The European Symposium on Artificial Neural Networks*.
- Leila Arras, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. [Explaining recurrent neural network predictions in sentiment analysis](#). In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 159–168, Copenhagen, Denmark. Association for Computational Linguistics.
- Jose Roberto Ayala Solares, Francesca Elisa Diletta Raimondi, Yajie Zhu, and 1 others. 2020. [Deep learning for electronic health records: A comparative review of multiple deep neural architectures](#). *Journal of Biomedical Informatics*, 101:103337.
- Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. [On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation](#). *PLOS ONE*, 10(7):1–46.
- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. In *International Conference on Learning Representations*.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhi-fang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025. [Qwen3-VL Technical Report](#). *arXiv preprint*. ArXiv:2511.21631 [cs].

660	Yonatan Belinkov. 2022. Probing classifiers: Promises, shortcomings, and advances . <i>Computational Linguistics</i> , 48(1):207–219.	715
661		716
662		717
663	Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017a. What do neural machine translation models learn about morphology?	718
664	In <i>Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 861–872, Vancouver, Canada. Association for Computational Linguistics.	719
665		720
666		721
667		722
668		723
669		
670	Yonatan Belinkov and James Glass. 2019. Analysis methods in neural language processing: A survey . Preprint, arXiv:1812.08951.	724
671		725
672		726
673	Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017b. Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks . In <i>Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 1–10, Taipei, Taiwan. Asian Federation of Natural Language Processing.	727
674		728
675		729
676		730
677		
678		731
679		732
680		733
681	Alexander Binder, Grégoire Montavon, Sebastian Bach, Klaus-Robert Müller, and Wojciech Samek. 2016. Layer-wise relevance propagation for neural networks with local renormalization layers . Preprint, arXiv:1604.00825.	734
682		735
683		736
684		737
685		738
686	George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. 2015. <i>Time Series Analysis: Forecasting and Control</i> . John Wiley & Sons.	739
687		740
688		741
689	Ching Chang, Wei-Yao Wang, Wen-Chih Peng, and Tien-Fu Chen. 2025. Llm4ts: Aligning pre-trained llms as data-efficient time-series forecasters . <i>ACM Trans. Intell. Syst. Technol.</i> , 16(3).	742
690		743
691		744
692		745
693	Yanping Chen, Eamonn Keogh, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, and Gustavo Batista. 2015. The ucr time series classification archive. www.cs.ucr.edu/~eamonn/time_series_data/ .	746
694		747
695		748
696		749
697		750
698	Mingyue Cheng, Yiheng Chen, Qi Liu, Zhiding Liu, Yucong Luo, and Enhong Chen. 2025. InstrucTime: Advancing Time Series Classification with Multimodal Language Modeling . In <i>Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining, WSDM '25</i> , pages 792–800, New York, NY, USA. Association for Computing Machinery.	751
699		752
700		753
701		754
702		755
703		756
704		757
705		758
706	Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. 2016. Doctor AI: Predicting clinical events via recurrent neural networks. In <i>Proceedings of the 1st Machine Learning for Healthcare Conference (JMLR W&CP 56)</i> , pages 301–318.	759
707		760
708		761
709		762
710		763
711		764
712	Grzegorz Chrupała, Bertrand Higy, and Afra Alishahi. 2020. Analyzing analytical methods: The case of phonology in neural models of spoken language. In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 4146–4156.	765
713		766
714		767
	Hyunseung Chung, Jiho Kim, Joon-Myoung Kwon, Ki-Hyun Jeon, Min Sung Lee, and Edward Choi. 2023. Text-to-ecg: 12-lead electrocardiogram synthesis conditioned on clinical text reports . In <i>ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 1–5.	768
		769
	Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. What you can cram into a single \$&!#* vector: Probing sentence embeddings for linguistic properties. In <i>ACL 2018-56th Annual Meeting of the Association for Computational Linguistics</i> , volume 1, pages 2126–2136. Association for Computational Linguistics.	769
	Yiqun Duan, Jinzhao Zhou, Zhen Wang, Yu-Kai Wang, and Chin-Teng Lin. 2023. Dewave: discrete eeg waves encoding for brain dynamics to text translation. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23</i> , Red Hook, NY, USA. Curran Associates Inc.	770
	Damian R. Eads, Daniel Hill, Sean Davis, Simon J. Perkins, Junshui Ma, Reid B. Porter, and James P. Theiler. 2002. Genetic algorithms and support vector machines for time series classification . In <i>Proceedings of SPIE</i> , volume 4787 of <i>Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation V</i> .	771
	Eugene F. Fama. 1970. Efficient capital markets: A review of theory and empirical work. <i>The Journal of Finance</i> , 25(2):383–417.	772
	Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso, and Svitlana Vyetenko. 2024. Evaluating large language models on time series feature understanding: A comprehensive taxonomy and benchmark . In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> . Association for Computational Linguistics.	773
	Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. Patchscopes: a unifying framework for inspecting hidden representations of language models. In <i>Proceedings of the 41st International Conference on Machine Learning, ICML'24</i> . JMLR.org.	774
	Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Manat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all . In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 15180–15190.	775
	Ary L. Goldberger and 1 others. 2000. Emg database. https://physionet.org/content/emgdb/1.0.0/ . Version 1.0.0, PhysioNet.	776

770	Mononito Goswami, Konrad Szafer, Arjun Choudhry,	Systems, NIPS '23, Red Hook, NY, USA. Curran	824
771	Yifu Cai, Shuo Li, and Artur Dubrawski. 2024. Mo-	Associates Inc.	825
772	ment: A family of open time-series foundation mod-		
773	els . In <i>Proceedings of the Forty-first International</i>	Bryan Lim, Sercan O. Arik, Nicolas Loeff, and Tomas	826
774	<i>Conference on Machine Learning (ICML)</i> .	Pfister. 2021. Temporal fusion transformers for in-	827
		interpretable multi-horizon time series forecasting. In	828
775	Nate Gruver, Ananya Bansal, Somdeb Dey, and 1 others.	<i>International Journal of Forecasting</i> .	829
776	2024. Large language models are zero-shot time		
777	series forecasters. <i>arXiv preprint arXiv:2402.04825</i> .		
778	James D. Hamilton. 1994. <i>Time Series Analysis</i> . Prince-	Haoxin Liu, Chenghao Liu, and B. Aditya Prakash.	830
779	ton University Press.	2025. A Picture is Worth A Thousand Numbers:	831
		Enabling LLMs Reason about Time Series via Vi-	832
780	Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long	sualization . In <i>Proceedings of the 2025 Conference</i>	833
781	short-term memory . <i>Neural Computation</i> , 9(8):1735–	<i>of the Nations of the Americas Chapter of the Asso-</i>	834
782	1780.	<i>ciation for Computational Linguistics: Human Lan-</i>	835
		<i>guage Technologies (Volume 1: Long Papers)</i> , pages	836
783	Rob J. Hyndman and George Athanasopoulos. 2018.	7486–7518, Albuquerque, New Mexico. Association	837
784	<i>Forecasting: Principles and Practice</i> . OTexts.	for Computational Linguistics.	838
785	Jonathan Jacobi and Gal Niv. 2025. Superscopes: Am-	Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu	839
786	plifying internal feature representations for language	Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han.	840
787	model interpretation . <i>Preprint</i> , arXiv:2503.02078.	2020. On the variance of the adaptive learning rate	841
		and beyond. In <i>Proceedings of the Eighth Inter-</i>	842
788	Ayush Jain, Arnav Singh, Rajiv Ratn Shah, and 1 others.	<i>national Conference on Learning Representations</i>	843
789	2023. Time-llm: Time series forecasting by repro-	<i>(ICLR 2020)</i> .	844
790	gramming large language models. <i>arXiv preprint</i>	Nelson F. Liu, Matt Gardner, Yonatan Belinkov,	845
791	<i>arXiv:2310.03589</i> .	Matthew E. Peters, and Noah A. Smith. 2019. Lin-	846
		guistic knowledge and transferability of contextual	847
792	Ming Jin, Shiyang Wang, Lihan Ma, Wei Xue, Jingbo	representations . In <i>Proceedings of the 2019 Confer-</i>	848
793	Liu, and Qingsong Wen. 2024a. Time-llm: Time	<i>ence of the North American Chapter of the Associ-</i>	849
794	series forecasting by reprogramming large language	<i>ation for Computational Linguistics: Human Lan-</i>	850
795	models. In <i>Proceedings of the International Confer-</i>	<i>guage Technologies, Volume 1 (Long and Short Pa-</i>	851
796	<i>ence on Learning Representations (ICLR)</i> .	<i>pers)</i> , pages 1073–1094, Minneapolis, Minnesota.	852
		Association for Computational Linguistics.	853
797	Ming Jin, Yifan Zhang, Wei Chen, Kexin Zhang, Yux-	Xin Liu, Daniel McDuff, Geza Kovacs, Isaac Galatzer-	854
798	uan Liang, Bin Yang, Jindong Wang, Shirui Pan, and	Levy, Jacob Sunshine, Jiening Zhan, Ming-Zher Poh,	855
799	Qingsong Wen. 2024b. Position: what can large lan-	Shun Liao, Paolo Di Achille, and Shwetak Patel.	856
800	guage models tell us about time series analysis. In	2023. Large language models are few-shot health	857
801	<i>Proceedings of the 41st International Conference on</i>	learners . <i>Preprint</i> , arXiv:2305.15525.	858
802	<i>Machine Learning, ICML'24</i> . JMLR.org.		
803	Ryan King, Tianbao Yang, and Bobak J. Mortazavi.	Ilya Loshchilov and Frank Hutter. 2016. Sgdr: Stochas-	859
804	2023. Multimodal pretraining of medical time series	tic gradient descent with warm restarts. <i>arXiv</i>	860
805	and notes. In <i>Proceedings of the Machine Learn-</i>	<i>preprint arXiv:1608.03983</i> .	861
806	<i>ing for Health (ML4H) Workshop</i> , pages 244–255.		
807	PMLR.	Ilya Loshchilov and Frank Hutter. 2017. Decou-	862
		pled weight decay regularization. <i>arXiv preprint</i>	863
808	Lingkai Kong, Haorui Wang, Wenhao Mu, Yuanqi Du,	<i>arXiv:1711.05101</i> .	864
809	Yuchen Zhuang, Yifei Zhou, Yue Song, Rongzhi	Xiaoya Lu, Dongrui Liu, Yi Yu, Luxin Xu, and Jing	865
810	Zhang, Kai Wang, and Chao Zhang. 2024. Aligning	Shao. 2025. X-boundary: Establishing exact safety	866
811	large language models with representation editing:	boundary to shield LLMs from jailbreak attacks with-	867
812	a control perspective. In <i>Proceedings of the 38th</i>	out compromising usability . In <i>Findings of the Asso-</i>	868
813	<i>International Conference on Neural Information Pro-</i>	<i>ciation for Computational Linguistics: EMNLP 2025</i> ,	869
814	<i>cessing Systems, NIPS '24</i> , Red Hook, NY, USA.	pages 5247–5272, Suzhou, China. Association for	870
815	Curran Associates Inc.	Computational Linguistics.	871
816	Vedang Lad, Jin Hwa Lee, Wes Gurnee, and Max	Gengchen Mai, Weiming Huang, Jin Sun, and 1 others.	872
817	Tegmark. 2025. The remarkable robustness of llms:	2024. On the opportunities and challenges of founda-	873
818	Stages of inference? <i>Preprint</i> , arXiv:2406.19384.	tion models for geoai . <i>ACM Transactions on Spatial</i>	874
		<i>Algorithms and Systems</i> , 10(2).	875
819	Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter	Mike A Merrill, Mingtian Tan, Vinayak Gupta, Thomas	876
820	Pfister, and Martin Wattenberg. 2023. Inference-time	Hartvigsen, and Tim Althoff. 2024a. Language mod-	877
821	intervention: eliciting truthful answers from a lan-	els still struggle to zero-shot reason about time series .	878
822	guage model. In <i>Proceedings of the 37th Interna-</i>	In <i>Findings of the Association for Computational</i>	879
823	<i>tional Conference on Neural Information Processing</i>		

880	<i>Linguistics: EMNLP 2024</i> , pages 3512–3533, Miami, Florida, USA. Association for Computational Linguistics.	935
881		936
882		937
883	William Merrill, Haotian Bai, Ananya Das, and Gail Weiss. 2024b. Are large language models good time series forecasters? <i>arXiv preprint arXiv:2402.02487</i> .	938
884		939
885		940
886	Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and 1 others. 2017. Mixed precision training. <i>arXiv preprint arXiv:1710.03740</i> .	941
887		942
888		943
889		944
890		945
891	Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Aparajita Saraf, Amy Bearman, and Babak Damavandi. 2023. <i>IMU2CLIP: Language-grounded motion sensor translation with multimodal contrastive learning</i> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 13246–13253, Singapore. Association for Computational Linguistics.	946
892		947
893		948
894		949
895		950
896		951
897		952
898		953
899	Juan Morales-García, Antonio Llanes, Francisco Arcas-Túnez, and Fernando Terroso-Sáenz. 2024. <i>Developing time series forecasting models with generative large language models</i> . <i>ACM Transactions on Intelligent Systems and Technology</i> .	954
900		955
901		956
902		957
903		958
904	OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, and 108 others. 2025. <i>gpt-oss-120b gpt-oss-20b model card</i> . <i>Preprint</i> , arXiv:2508.10925.	959
905		960
906		961
907		962
908		963
909		964
910		965
911	Jielin Qiu, William Han, Jiacheng Zhu, Mengdi Xu, Michael Rosenberg, Emerson Liu, Douglas Weber, and Ding Zhao. 2023. <i>Transfer knowledge from natural language to electrocardiography: Can we detect cardiovascular disease through language models?</i> In <i>Findings of the Association for Computational Linguistics: EACL 2023</i> , pages 442–453, Dubrovnik, Croatia. Association for Computational Linguistics.	966
912		967
913		968
914		969
915		970
916		971
917		972
918		973
919	Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. <i>OpenAI blog</i> , 1(8):9.	974
920		975
921		976
922		977
923	Alvin Rajkomar, Eyal Oren, Kai Chen, and 1 others. 2018. <i>Scalable and accurate deep learning with electronic health records</i> . <i>npj Digital Medicine</i> , 1(1):18.	978
924		979
925		980
926	Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, and 1 others. 2019. <i>Deep learning and process understanding for data-driven earth system science</i> . <i>Nature</i> , 566(7743):195–204.	981
927		982
928		983
929		984
930	Simon Roschmann, Quentin Bouniot, Vasilii Feofanov, Ievgen Redko, and Zeynep Akata. 2025. Time series representations for classification lie hidden in pretrained vision transformers. <i>arXiv preprint arXiv:2506.08641</i> .	985
931		986
932		987
933		988
934		989
	Paul K. Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, Hannah Muckenhirn, Dirk Padfield, James Qin, Danny Rozenberg, Tara Sainath, Johan Schalkwyk, Matt Sharifi, Michelle Tadmor Ramanovich, Marco Tagliasacchi, and 11 others. 2023. <i>Audiopalm: A large language model that can speak and listen</i> . <i>Preprint</i> , arXiv:2306.12925.	990
	Justin Sirignano and Rama Cont. 2019. <i>Universal features of price formation in financial markets: Perspectives from deep learning</i> . <i>Quantitative Finance</i> , 19(9):1449–1459.	991
	Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. 2023. <i>PandaGPT: One model to instruction-follow them all</i> . In <i>Proceedings of the 1st Workshop on Taming Large Language Models: Controllability in the era of Interactive Assistants!</i> , pages 11–23, Prague, Czech Republic. Association for Computational Linguistics.	
	Chenxi Sun, Hongyan Li, Yaliang Li, and Shenda Hong. 2024. <i>Test: Text prototype aligned embedding to activate llm’s ability for time series</i> . In <i>Proceedings of the International Conference on Learning Representations (ICLR)</i> .	
	Mingtian Tan, Mike A. Merrill, Vinayak Gupta, Tim Althoff, and Thomas Hartvigsen. 2024. Are language models actually useful for time series forecasting? In <i>Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24</i> , Red Hook, NY, USA. Curran Associates Inc.	
	Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, and 179 others. 2024. <i>Gemma 2: Improving open language models at a practical size</i> . <i>Preprint</i> , arXiv:2408.00118.	
	Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, and 1 others. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. <i>arXiv preprint arXiv:1905.06316</i> .	
	Aaron Van Den Oord, Oriol Vinyals, and 1 others. 2017. Neural discrete representation learning. <i>Advances in neural information processing systems</i> , 30.	
	Eva Vanmassenhove, Jinhua Du, and Andy Way. 2017. <i>Investigating ‘aspect’ in nmt and smt: Translating the english simple past and present perfect</i> . <i>Computational Linguistics in the Netherlands Journal</i> , 7:109–128.	
	Will Vickers and Matthew Middlehurst. Rightwhalecalls dataset. https://www.	

992	timeseriesclassification.com/description.php?Dataset=RightWhaleCalls . Accessed: <date>.	
993		
994		
995	Haoyu Wang, Bingzhe Wu, Yatao Bian, Yongzhe Chang, Xueqian Wang, and Peilin Zhao. 2024. Probing the safety response boundary of large language models via unsafe decoding path generation . <i>Preprint</i> , arXiv:2408.10668.	
996		
997		
998		
999		
1000	Yanling Wang, Haoyang Li, Hao Zou, Jing Zhang, Xinlei He, Qi Li, and Ke Xu. 2025. FacLens: Transferable Probe for Foreseeing Non-Factuality in Fact-Seeking Question Answering of Large Language Models . In <i>Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing</i> , pages 18574–18593, Suzhou, China. Association for Computational Linguistics.	
1001		
1002		
1003		
1004		
1005		
1006		
1007		
1008	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, and 1 others. 2020. Transformers: State-of-the-art natural language processing. In <i>Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations</i> , pages 38–45.	
1009		
1010		
1011		
1012		
1013		
1014		
1015		
1016	Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In <i>Advances in Neural Information Processing Systems</i> .	
1017		
1018		
1019		
1020		
1021	Qianqian Xie, Weiguang Han, Yanzhao Lai, Min Peng, and Jimin Huang. 2023. The wall street neophyte: A zero-shot analysis of chatgpt over multimodal stock movement prediction challenges . <i>Preprint</i> , arXiv:2304.05351.	
1022		
1023		
1024		
1025		
1026	Hao Xue and Flora D. Salim. 2024. Promptcast: A new prompt-based learning paradigm for time series forecasting . <i>IEEE Trans. on Knowl. and Data Eng.</i> , 36(11):6851–6864.	
1027		
1028		
1029		
1030	Tianyi Yan, Fei Wang, James Y. Huang, Wenxuan Zhou, Fan Yin, Aram Galstyan, Wenpeng Yin, and Muhao Chen. 2024. Contrastive instruction tuning . In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10288–10302, Bangkok, Thailand. Association for Computational Linguistics.	
1031		
1032		
1033		
1034		
1035		
1036	Dongchao Yang, Jinchuan Tian, Xu Tan, Rongjie Huang, Songxiang Liu, Haohan Guo, Xuankai Chang, Jiaotong Shi, Sheng Zhao, Jiang Bian, Zhou Zhao, Xixin Wu, and Helen M. Meng. 2024. UniAudio: Towards universal audio generation with large language models . In <i>Proceedings of the 41st International Conference on Machine Learning</i> , volume 235 of <i>Proceedings of Machine Learning Research</i> , pages 56422–56447. PMLR.	
1037		
1038		
1039		
1040		
1041		
1042		
1043		
1044		
1045	Xinli Yu, Zheng Chen, Yuan Ling, Shujing Dong, Zongyi Liu, and Yanbin Lu. 2023. Temporal data meets llm – explainable financial time series forecasting . <i>Preprint</i> , arXiv:2306.11025.	
1046		
1047		
1048		
	Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. 2022. TS2Vec: Towards Universal Representation of Time Series . <i>arXiv preprint</i> . ArXiv:2106.10466 [cs].	1049
		1050
		1051
		1052
	Honggen Zhang, Xufeng Zhao, Igor Molybog, and June Zhang. 2025. Real: Response embedding-based alignment for llms . <i>Preprint</i> , arXiv:2409.17169.	1053
		1054
		1055
	Mi Zhang and Alexander A Sawchuk. 2012. Usc-had: A daily activity dataset for ubiquitous activity recognition using wearable sensors. In <i>Proceedings of the 2012 ACM conference on ubiquitous computing</i> , pages 1036–1043.	1056
		1057
		1058
		1059
		1060
	Xiyuan Zhang, Ranak Roy Chowdhury, Rajesh K. Gupta, and Jingbo Shang. 2024. Large language models for time series: A survey . In <i>Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24</i> , pages 8335–8343. International Joint Conferences on Artificial Intelligence Organization. Survey Track.	1061
		1062
		1063
		1064
		1065
		1066
		1067
	Yunkai Zhang, Yawen Zhang, Ming Zheng, Kezhen Chen, Chongyang Gao, Ruian Ge, Siyuan Teng, Amine Jelloul, Jinmeng Rao, Xiaoyuan Guo, and 1 others. 2023. Insight miner: A time series analysis dataset for cross-domain alignment with natural language. In <i>NeurIPS Workshop on AI for Science</i> .	1068
		1069
		1070
		1071
		1072
		1073
	Haoyi Zhou, Shanghang Zhang, Jieqi Peng, and 1 others. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 35, pages 11106–11115.	1074
		1075
		1076
		1077
		1078
	Tian Zhou, Peisong Niu, Xue Wang, Liang Sun, and Rong Jin. 2023a. One fits all: power general time series analysis by pretrained lm. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23</i> , Red Hook, NY, USA. Curran Associates Inc.	1079
		1080
		1081
		1082
		1083
		1084
	Tian Zhou, PeiSong Niu, Xue Wang, Liang Sun, and Rong Jin. 2023b. One Fits All: Power General Time Series Analysis by Pretrained LM . <i>arXiv preprint</i> . ArXiv:2302.11939 [cs].	1085
		1086
		1087
		1088
	Yunjiao Zhou, Jianfei Yang, Han Zou, and Lihua Xie. 2023c. Tent: Connect language models with iot sensors for zero-shot activity recognition . <i>Preprint</i> , arXiv:2311.08245.	1089
		1090
		1091
		1092
	Yubo Zhu, Dongrui Liu, Zecheng Lin, Wei Tong, Sheng Zhong, and Jing Shao. 2025. The LLM Already Knows: Estimating LLM-Perceived Question Difficulty via Hidden Representations . In <i>Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing</i> , pages 1160–1176, Suzhou, China. Association for Computational Linguistics.	1093
		1094
		1095
		1096
		1097
		1098
		1099
	A Baselines, Datasets, Models & Metrics	1100
	A.1 Baselines	1101
	Attend follows the dedicated time-series encoder family. It uses self-attention to model cross-	1102
		1103

Method	Model	Shots	Style	ctu	emg	har	rwc	tee	AVG
Probe	Llama	0	CoT	.651	.789	.629	.566	.351	.598
		0	Direct	.650	.784	.645	.608	.355	.609
	Mistral	0	CoT	.688	.899	.653	.581	.418	.648
		0	Direct	.674	.871	.655	.585	.408	.638
	Qwen	0	CoT	.697	.920	.677	.595	.442	.654
		0	Direct	.702	.938	.704	.592	.473	.667
Prompt	Llama	0	CoT	.434	.233	.126	.231	.077	.220
		0	Direct	.333	.185	.083	.231	.036	.174
		2	Direct	.333	.173	.041	.231	.036	.163
	Mistral	0	CoT	.544	.324	.168	.339	.105	.296
		0	Direct	.420	.393	.199	.398	.288	.340
		2	Direct	.423	.445	.230	.320	.188	.321
	Qwen	0	CoT	.498	.127	.225	.401	.195	.289
		0	Direct	.359	.393	.127	.288	.283	.290
		2	Direct	.472	.360	.228	.234	.107	.280

Table 4: F1 results between prompting and probing. We explore whether Chain-of-Thought (CoT) prompting and few-shot examples improve prompting performance, as well as whether CoT improves the representations during probing. *Shots* represent the number of in-line examples per class in the dataset. For example *har* has 6 possible classes, there are 2 examples per class totaling to 12 examples overall. These results come from a smaller sample using 500 training rows and 100 test rows.

Model	Method	ctu	emg	had	har	rwc	tee	AVG
Gemma	Probe	.680	1.00	.145	.739	.556	.458	.596
GPT-oss		.579	.792	.112	.498	.518	.295	.466
Gemma	Prompt	.286	.314	.023	.075	.346	.100	.206
GPT-oss		.282	.358	.031	.142	.272	.132	.202

Table 5: F1 scores for probe-based and prompt-based inference on text-only large language models.

channel and temporal dependencies in multivariate sensor data and learns discriminative representations directly from raw time series, without adapting language-model token spaces (Abedin et al., 2021). **InstructTime** belongs to the time-series quantization and instruction-based prompting family. It discretizes time series into vector-quantized tokens, aligns them with domain-specific textual instructions, and uses generative instruction tuning to produce label text via next-token prediction (Cheng et al., 2025). **OneFitsAll** follows the encoder-alignment family. It reuses a frozen pre-trained transformer from language or vision domains and learns lightweight input projections and normalization layers to map time-series patches into the model’s embedding space, enabling a single backbone across tasks (Zhou et al., 2023b).

TS2Vec represents dedicated time-series representation learning. It learns multi-scale contextual embeddings through hierarchical contrastive learning over augmented temporal views, which are aggregated and used by simple downstream classifiers without language models (Yue et al., 2022). **Moment** belongs to the TS-Encoder family. It is a high-capacity transformer encoder pre-trained on large-scale, multi-domain time series using masked time-series modeling, and serves as a general-purpose representation learner that can be adapted to classification, forecasting, anomaly detection, and imputation via lightweight task-specific heads or linear probing (Goswami et al., 2024).

A.2 Models

To assess the generality of our findings, we evaluate a diverse set of pretrained models spanning different architectures and scales.¹ Motivated by prior work showing that visualizations provide a strong representation for time series classification (Liu et al., 2025), we primarily focus on vLLMs. Specifically, we evaluate Llama-3.2-11B-Vision-Instruct (gra, 2024), Qwen/Qwen2.5-VL-32B-Instruct (Bai et al., 2025), and Mistral-Small-3.1-24B-Instruct-2503 (AI, 2025). To verify that our conclusions extend beyond multimodal architectures, we also evaluate two text-only LLMs: GPT-oss (OpenAI et al., 2025) and google/gemma-2-9b-it (Team et al., 2024).

A.3 Datasets

Appendix C provides direct links to all source datasets used in our experiments, while Table 6 reports descriptive statistics. We selected datasets to span a wide range of temporal scales, sequence lengths, variable dimensionality, and label complexity. This diversity allows us to evaluate whether LLM-based prompting generalizes across fundamentally different forms of temporal structure. Short descriptions of each dataset can be found below.

CTU records 24-hour household electricity usage patterns and classifies computers as desktops or laptops, capturing differences in device power-consumption cycles (Chen et al., 2015). **EMG** consists of muscle response recordings under electrical stimulation, categorized into healthy, myopathy, and neuropathy classes (Goldberger et al., 2000). **HAR** contains multivariate accelerometer recordings of human activities such as walking, sitting, and stair climbing, and exhibits substantial inter-subject variability (Anguita et al., 2013). **HAD** similarly contains multivariate human activity data; we use both the tri-axis accelerometer data and tri-axis gyroscopic data (Zhang and Sawchuk, 2012). **RWC** comprises underwater acoustic recordings used to detect North Atlantic right whale calls in the presence of ocean noise (Vickers and Middlehurst). **TEE** originates from the FORTE satellite and classifies transient electromagnetic events (e.g., lightning strokes) based on waveform signatures (Eads

¹All models are open source, as access to internal hidden states is required for probing.

et al., 2002).

A.4 Metrics

Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denote a test set of N labeled time series, where $y_i \in \{1, \dots, C\}$ and $\hat{y}_i$ is the predicted class. Let TP_c, FP_c, FN_c denote the counts of true positives, false positives, and false negatives for class c under the one-vs-rest convention.

Macro F1. For each class c , define precision and recall:

$$P_c = \frac{TP_c}{TP_c + FP_c}, \quad R_c = \frac{TP_c}{TP_c + FN_c}.$$

The classwise F1 is

$$F1_c = \frac{2P_cR_c}{P_c + R_c}.$$

Macro F1 averages uniformly across classes:

$$F1_{\text{macro}} = \frac{1}{C} \sum_{c=1}^C F1_c.$$

Pass@K. For each test example i , we obtain n independent generations and mark each generation as correct/incorrect. Let c_i be the number of correct generations among the n samples for example i (in our experiments, $n = 20$). Following the standard unbiased estimator for pass@K, we define:

$$\widehat{\text{Pass@K}}_i = 1 - \frac{\binom{n-c_i}{K}}{\binom{n}{K}},$$

with the conventions $\widehat{\text{Pass@K}}_i = 0$ if $c_i = 0$, and $\widehat{\text{Pass@K}}_i = 1$ if $n - c_i < K$. We report the dataset-level score by averaging across examples:

$$\widehat{\text{Pass@K}} = \frac{1}{N} \sum_{i=1}^N \widehat{\text{Pass@K}}_i.$$

$\Delta\text{Pass@K}$. We quantify the benefit of repeated sampling by comparing multiple attempts against a single attempt:

$$\Delta\widehat{\text{Pass@K}} = \widehat{\text{Pass@K}} - \widehat{\text{Pass@1}}.$$

In particular, we report $\Delta\widehat{\text{Pass@20}} = \widehat{\text{Pass@20}} - \widehat{\text{Pass@1}}$.

B Addressing Probing Concerns

This section provides additional analyses and controls addressing known methodological concerns associated with probing classifiers. Following the recommendations of Belinkov (2022), these experiments are intended to support the interpretation of probing results as diagnostics of representational capacity rather than causal explanations of model behavior.

Probe Complexity and Memorization. A central concern in probing-based analysis is that probe performance may reflect the capacity of the probe itself rather than information encoded in the underlying representations (Belinkov, 2022). To mitigate this risk, all probes used in this work are linear classifiers trained on frozen representations. We use logistic regression, one of the simplest and most interpretable classifiers, commonly employed in probing studies to assess linear extract-ability.

To further guard against probe memorization, we evaluate probe performance across layers and datasets and observe consistent trends rather than isolated improvements at specific layers. All probes are trained using five-fold cross-validation with grid search to select the regularization parameter C , and the maximum number of iterations is fixed to 1000. Together, these design choices reduce the likelihood that observed gains arise from overfitting by the probe and support their interpretation as reflecting systematic representational structure.

Baselines and Controls. Interpreting raw probing accuracy in isolation can be misleading, as even simple or randomized representations may support non-trivial classification performance. Prior work, therefore, emphasizes the importance of contextualizing probing results using appropriate baselines and reference models (Conneau et al., 2018).

As an absolute lower bound, we report random-assignment baselines, including majority, uniform, and prior-based predictors. These baselines characterize dataset properties such as the number of classes and class imbalance, and provide a reference point for interpreting probe performance.

Next, we follow the guidance of Belinkov et al. (2017a) and Conneau et al. (2018), we **train a classifier on an LLM with randomly initialized weights**. Specifically, we chose Llama-3.2-11B-Vision-Instruct as it can be directly compared to its fine-tuned version.

The randomly initialized model serves as a practical floor for probing-based approaches. It reflects how well a linear classifier can extract a signal from features without meaningful representations. This baseline, therefore, controls for the effect of supervised training applied after feature generation.

As an upper-bound reference, established baselines designed specifically to learn discriminative time-series representations, following prior probing work that contrasts frozen representations with task-optimized or state-of-the-art models (Belinkov et al., 2017b; Conneau et al., 2018; Liu et al., 2019).

Correlation Versus Causation. Early probing work often assumed that extractable information reflected model usage, but subsequent studies have demonstrated mismatches between probing performance and downstream task behavior (Vanmassenhove et al., 2017; Belinkov and Glass, 2019). Accordingly, we avoid causal claims about model decision-making and interpret probing results strictly as evidence of representational capacity rather than usage.

Choice of Probed Properties. Another limitation of probing is its reliance on predefined properties, which constrains the scope of conclusions and may bias analysis toward expected features. In our setting, we probe directly for time-series class labels, treating linear separability with respect to the downstream classification task as the property of interest. This choice aligns with our goal of assessing whether LLM representations encode discriminative temporal structure relevant to classification, rather than targeting specific hand-engineered temporal features.

C Dataset Statistics

Below in Table 6 we report the descriptive statistics of each of the datasets used in this study. Links for each dataset can be found below.

Dataset	N-Channels	Series-Length	N-Classes	Train	Test
cpu	1	720	2	250	250
emg	1	1500	3	267	47
had	6	128	12	34779	7929
har	3	206	6	7,352	2,947
rwc	1	4000	2	25,499	4,501
tee	1	319	7	70	73

Table 6: Summary statistics for all time-series datasets used in our classification experiments.

Dataset Links

- CTU: <https://www.>




Figure 7: For the RWC dataset only, we represent the data as a spectrogram.




Figure 8: Example of a univariate time series line plot.

timeseriesclassification.com/description.php?Dataset=Computers

- **HAD:** <http://sipi.usc.edu/HAD/>
- **EMG:** <https://physionet.org/content/emgdb/1.0.0/>
- **HAR:** <https://archive.ics.uci.edu/dataset/240/human+activity+recognition+using+smartphones>
- **RWC:** <https://www.kaggle.com/competitions/whale-detection-challenge/data>
- **TEE:** <https://www.timeseriesclassification.com/description.php?Dataset=Lightning7>

D Prompting Examples

This section illustrates how time series data are visualized in the visual domain and how these representations are incorporated into the prompting structures used for evaluation. We show both line-plot and spectrogram-based visualizations, text-only (*d*) prompts for inference, prompts used to generate prompt-variants, and examples of prompt-variants themselves.

D.1 Example Plots

Figures 7, 8, & 9 demonstrate how time series information we represent time series in the visual domain. For the RWC dataset (Figure 7), following the recommendations of Liu et al. (2025), the time series is rendered as a spectrogram. All other univariate datasets are visualized as line plots (see Figure 8 for an example). For multivariate datasets, each variable is plotted as a separate line with



Figure 9: Example of a multivariate time series line plot.

color-coded legends, expanding to multiple sub-plots when necessary (Figure9). Axis labels, titles, and legends are specified by the user prior to inference.

D.2 Prompting

Below we show the prompting structure for inference as well as the prompt used to make the variants.

Prompts for Inference. Figure 10 shows the instantiated prompt template used to classify the CTU dataset in the text-only (*d*) modality. In the visual (*v*) modality, the numerical time series is replaced with an appropriate visualization (see Section D.1). In the combined (*d+v*) modality, both the numerical sequence and the corresponding visualization are provided. The *Additional information that may help* is specified *a priori* and is adapted directly from Liu et al. (2025), who demonstrated the effectiveness of this prompting strategy.

Prompts for Prompt-Variant Generation. Figures 11 and 12 illustrate the prompts used to generate diverse, meaning-preserving variants. Figure 11 is used to generate variations of the system prompt, while Figure 12 is used to generate variations of the user message. For both system prompt and question generation, we issue two batch requests with batch size $B=5$, yielding a total of ten unique variants.

Examples of Prompt Variants. Figures 13, 14, & 15 present three prompt variants generated for the HAR dataset. Each variant exhibits distinct characteristics. For example, Variant 2 (Figure 14) is substantially more verbose than the others, whereas Variant 1 (Figure 13) is the most concise and omits expanded definitions for acronyms (e.g., using RMS without explicitly stating “Root Mean Square,” as done in the other examples).




Figure 10: System and user prompts for the CTU dataset. The time series is abbreviated for readability. The *Additional information that may help*, *question*, and *options* change depending on the dataset.

Variation Generator: System Prompt Rewriter (System Prompt)

You are a prompt rewriter.

Rewrite the provided SYSTEM PROMPT into N distinct SYSTEM PROMPT variants that preserve meaning, constraints, and all factual task content, while changing phrasing, structure, and formatting.

Hard constraints (must obey):

- Do NOT add, remove, rename, or reorder any class names.
- Do NOT alter the semantic meaning of any class description.
- Preserve answer-choice letters (e.g., [A], [B]) exactly and in order.
- Preserve required output format strings exactly (punctuation/casing), e.g., "The answer is [X] CLASS_NAME".
- Do NOT introduce new task instructions (e.g., reasoning, confidence).
- Maintain an academically appropriate tone.

Diversity requirements:

- Each variant must differ meaningfully in organization and wording.
- Use multiple presentation styles (headings, bullets, numbered steps, etc.).
- Avoid trivial paraphrases.

Output format (strict): return valid JSON only:

```
{
  "variants": [
    {"id": 1, "system_prompt": "..."},
    {"id": 2, "system_prompt": "..."}
  ]
}
```

No extra keys. No markdown. No commentary.

Variation Generator: System Prompt Rewriter (User Message Template)

Per-batch request (batch size = B):

Generate B rewritten variants of the following.

IMPORTANT:

- Do NOT change class names.
- Do NOT change answer letters or their order.
- Do NOT change the required answer format.
- Output must be valid JSON only.

<BASE_SYSTEM_PROMPT>

Batching: if total variants = N and batch size = B , the script makes N/B calls and then renumbers variant ids globally.

Figure 11: System and user prompts for system prompt variation generation. <BASE_SYSTEM_PROMPT> is replaced by the system prompt shown in Figure 10 or by the dataset-specific base system prompt used for evaluation. Two batch requests with batch size $B=5$ are used, producing ten total system prompt variants.

Variation Generator: General Question Rewriter (System Prompt)

You are a question rewriter.

Your task is to rewrite a GENERAL QUESTION into N distinct variants that preserve the task, label space, and decision criteria, while changing wording, structure, and phrasing.

Hard constraints (must obey):

- Do NOT change, rename, remove, or reorder answer choices.
- Preserve answer-choice letters exactly (e.g., [A], [B]) and their order.
- Do NOT introduce new labels, hints, or constraints.
- Do NOT add explanations, reasoning instructions, or output formatting rules.
- The rewritten question must ask the same classification decision.

Diversity requirements:

- Each variant must differ meaningfully in wording and structure.
- Vary tone (instructional vs. role-based), sentence structure, and framing.
- Avoid trivial paraphrases.

Output format (strict): return valid JSON only:

```
{
  "variants": [
    {"id": 1, "question": "..."},
    {"id": 2, "question": "..."}
  ]
}
```

No extra keys. No markdown. No commentary.

Variation Generator: General Question Rewriter (User Message Template)

Per-batch request (batch size = B):

Generate B rewritten variants of the following.

IMPORTANT:

- Do NOT change the answer choice letters, names, or their order.
- Do NOT add output formatting rules.
- Do NOT add reasoning/explanation instructions.
- Output must be valid JSON only.

<BASE_GENERAL_QUESTION>

Batching: if total variants = N and batch size = B , the script makes N/B calls and then rennumbers variant ids globally.

Figure 12: System and user prompts for general question variation generation. <BASE_GENERAL_QUESTION> is replaced by the dataset-specific general question used for evaluation (e.g., *determine whether this computer ...* from Figure 10). As in system prompt variation generation, we issue two batch requests with batch size $B=5$ to obtain ten total question variants.

Example Prompt Variations (Feature list with enumerated structure)

Variant 1.

You are tasked with human activity recognition based on tri-axial accelerometer data (x, y, z) collected over time. Presented with a multiple-choice question and corresponding time series data, your responsibility is to select the correct activity. Provide your answer strictly in the following format: "The answer is [X] CLASS_NAME" (e.g., "The answer is [D] CABBAGE"), without any explanations.

Key features to consider for accurate classification include:

1. Statistical: mean, standard deviation, variance, median, range, IQR
2. Time-domain: RMS, zero crossing rate, signal magnitude area, time between peaks
3. Frequency-domain: FFT, PSD, spectral entropy, peak frequency
4. Correlation-based: correlations between axes
5. Geometrical: angles between axes, magnitude of acceleration vector
6. Entropy-based: signal entropy

Utilize these feature types to interpret the accelerometer signals and determine the human activity. Analyze the sequence of tri-axial accelerometer measurements (x, y, z) collected over time and classify the observed human activity into one of these categories:

[A] WALKING [B] WALKING_UPSTAIRS [C] WALKING_DOWNSTAIRS [D] SITTING [E] STANDING [F] LAYING

Figure 13: Example one of three meaning-preserving prompt variations for human activity recognition.

Example Prompt Variations (Sectioned guidance with descriptive bullets)

Variant 2.

As an expert in human activity recognition, your task is to identify the activity from the tri-axial accelerometer data (x, y, z) recorded over time. You will receive a multiple-choice question along with the time series data. Use the time series to select the correct answer without providing any explanation. Respond strictly in the following format: "The answer is [X] CLASS_NAME". For example: "The answer is [D] CABBAGE".

Additional guidance:

When analyzing accelerometer data for activity recognition, consider extracting the following features:

1. Statistical Features:
 - Mean: Average acceleration per axis.
 - Standard Deviation: Variation in acceleration values.
 - Variance: Spread of data points around the mean.
 - Median: Middle value splitting the data.
 - Range: Difference between max and min values.
 - Interquartile Range (IQR): Spread between 25th and 75th percentiles.
2. Time-domain Features:
 - Root Mean Square (RMS): Magnitude of acceleration.
 - Zero Crossing Rate (ZCR): Frequency of signal sign changes.
 - Signal Magnitude Area (SMA): Energy expenditure over time.
 - Time between Peaks: Duration between local maxima.
3. Frequency-domain Features:
 - Fast Fourier Transform (FFT): Frequency components analysis.
 - Power Spectral Density (PSD): Power distribution across frequencies.
 - Spectral Entropy: Complexity of frequency distribution.
 - Peak Frequency: Dominant frequency component.
4. Correlation-based Features:
 - Correlation between Axes: Coordination of movement.
5. Geometrical Features:
 - Angles between Axes: Body orientation.
 - Magnitude of Acceleration Vector: Overall acceleration magnitude.
6. Entropy-based Features:
 - Signal Entropy: Complexity of the acceleration signal.

Using the tri-axial accelerometer data (x, y, z) recorded over time, identify which human activity is being performed. Choose from the following options:

[A] WALKING [B] WALKING_UPSTAIRS [C] WALKING_DOWNSTAIRS [D] SITTING [E] STANDING [F] LAYING

Figure 14: Example two of three meaning-preserving prompt variations for human activity recognition.

Example Prompt Variations (Compact rubric-style presentation)

Variant 3.

As a professional in human activity recognition, your responsibility is to identify the activity from tri-axial accelerometer data (x, y, z) recorded over time. You will be provided with a multiple-choice question and the time series data. Use the data to select the correct answer without any explanation, adhering strictly to the format: "The answer is [X] CLASS_NAME". Example: "The answer is [D] CABBAGE".

Consider these feature types when analyzing the accelerometer signals:

- 1) Statistical Features: mean, standard deviation, variance, median, range, interquartile range.
- 2) Time-domain Features: root mean square (RMS), zero crossing rate (ZCR), signal magnitude area (SMA), time between peaks.
- 3) Frequency-domain Features: fast Fourier transform (FFT), power spectral density (PSD), spectral entropy, peak frequency.
- 4) Correlation-based Features: correlation between axes.
- 5) Geometrical Features: angles between axes, magnitude of acceleration vector.
- 6) Entropy-based Features: signal entropy.

Acting as a specialist in human activity recognition, use the tri-axial accelerometer readings (x, y, z) collected over time to classify the performed activity. Your options are:

[A] WALKING [B] WALKING_UPSTAIRS [C] WALKING_DOWNSTAIRS [D] SITTING [E] STANDING [F] LAYING

Figure 15: Example three of three meaning-preserving prompt variations for human activity recognition.

Hyperparameter	Value
Model Architecture	GPT-2 (6 layers, pre-trained)
Model Dimension	768
Patch Size	16
Stride	16
Optimizer	RAdam
Learning Rate	0.001
Batch Size	64
Training Epochs	50
Dropout	0.1
Loss Function	Cross-Entropy

Table 7: Hyperparameters and training configuration for OneFitsAll model

E Hyperparameters

For reproducibility, we provide the complete hyperparameter configurations for all baseline methods. These settings were either adopted from the original publications or carefully tuned to ensure fair comparison across all methods on our datasets.

E.1 OneFitsAll

All deep learning models are implemented in PyTorch using the Hugging Face Transformers library (Wolf et al., 2020). We employ a GPT-2-based architecture (Radford et al., 2019) where the pre-trained GPT-2 model serves as the backbone, with most parameters frozen, only the layer normalization and positional embedding layers remain trainable. The input time series are processed using a patching strategy, where patches are embedded into a high-dimensional space before being fed to the transformer layers. Training is performed using the RAdam optimizer (Liu et al., 2020) with cross-entropy loss, and model performance is evaluated using 5-fold stratified cross-validation on the training set. Standardization normalization is applied to input features, and the model is evaluated on a separate held-out test set using accuracy as the key metric. All hyperparameters are summarized in Table 7.

E.2 Attend

We implement our experiments using PyTorch. Our network is trained end-to-end using the Adam optimizer (Adam et al., 2014) with a learning rate that decays periodically. The training data is partitioned into segments using a sliding window approach with 50% overlap between adjacent windows. We employ mixup augmentation as a regularization

Hyperparameter	Value
Epochs	300
Batch size	256
Learning rate	10^{-3}
LR decay factor	0.9
LR decay step	10 epochs
Sliding window (W)	24
Window overlap	50%
Mixup α	0.8
Center loss weight (β)	0.3
Dropout (features)	0.5
Dropout (RNN)	0.25
Dropout (classifier)	0.5
Hidden dimension	128
Filter size	5
Number of filters	64

Table 8: Training hyperparameters for Attend and Discriminate model

technique and incorporate center loss to encourage intra-class compactness in the learned feature space. Dropout is applied at multiple stages of the network, including the RNN layers and the classifier, to prevent overfitting. Model performance is evaluated using 5-fold cross-validation to ensure robust and reliable results. All hyperparameters are summarized in Table 8.

E.3 InstructTime

We implement InstructTime using Python 3.12.7 and PyTorch 2.7.1 on an Nvidia H200 GPU. The framework consists of two stages: (1) a TStokenizer (VQ-VAE) (Van Den Oord et al., 2017) for time series discretization, and (2) an InstructTime model built upon the pre-trained GPT-2 language model (Radford et al., 2019) for classification. For the TStokenizer, we employ standard VQ-VAE training settings with mean pooling. For the InstructTime model, we use AdamW optimizer (Loshchilov and Hutter, 2017) with weight decay and employ a cosine annealing learning rate scheduler with warmup (Loshchilov and Hutter, 2016). Following prior work on pre-trained language models, we use a smaller learning rate and batch size compared to training from scratch, as the PLM has been pre-trained with a large number of parameters. Mixed-precision training (Mickevicus et al., 2017) is enabled using PyTorch’s automatic mixed precision (AMP) to accelerate computation. The detailed hyperparameter settings

Hyperparameter	Value
<i>TTokenizer (VQ-VAE)</i>	
Hidden dimension (d_{model})	64
Codebook size (n_{embed})	256
Wave length	24
Training epochs	60
Batch size (train/test)	32 / 64
Learning rate	5×10^{-4}
Learning rate decay	0.99
Weight decay	$1e^{-5}$
<i>InstructTime</i>	
Batch size	16
Learning rate	$1e^{-5}$
Weight decay	0.01
Warmup ratio	0.05
Training epochs	300
Max sequence length	230
Random seed	2024

Table 9: Hyperparameter settings for TTokenizer and InstructTime models

Hyperparameter	Value
Backbone model	AutonLab/MOMENT-1-large
Task name	classification
Series batch size	64
Window length (W)	512
Window stride	512
Patch length	8

Table 10: Hyperparameters for Moment baseline.

Hyperparameter	Value
Representation Dimension	320
Hidden Channels	64
Residual Blocks	10
Kernel Size	3
Batch Size	8
Crop Length	3000
Normalization	Z-score (per channel)
Learning Rate	.0001
Epochs	10 or 30

Table 11: Hyperparameters and training configuration for TS2Vec.

E.5 TS2Vec

We implement TS2Vec using the authors’ official PyTorch implementation, and, with the exception of epochs, used a fixed set of hyperparameters across all datasets. We used 10 training epochs for larger datasets (HAD, RWC, HAR) and 30 epochs for our smaller datasets (CTU, EMG, TEE). Input time series are standardized prior to training, and representations are extracted from the final encoder layer for downstream evaluation. All hyperparameters are summarized in Table 11.

F Use of AI Assistants

We used AI to help clean up writing, but all thoughts and work are our own.

for both stages are summarized in Table 9.

E.4 Moment Embedding Extraction

We extract fixed-length embeddings using the pretrained AutonLab/MOMENT-1-large pipeline. Each time series is z-normalized using a per-channel StandardScaler fit on the training split and reused for the test split. To handle sequences longer than the model window, each series is segmented into non-overlapping windows of length $W=512$ with left-padding for the final partial window; an input mask marks padded positions. MOMENT outputs patch embeddings with patch length 8, which we average over the valid (non-padded) patches to obtain a window embedding. Window embeddings are then aggregated into a single series embedding using a length-weighted average across windows. All configuration values are summarized in Table 10.