COMPLEMENTING DOMAIN LABELS WITH WAVEENERGY SIGNATURES FOR TIME SERIES HETEROGENEITY

Seongbeom Park[1]*, Hyunwoo Seo[1]*, Yichen Ma[2]*, Hojin Cho[1], Chiehyeon Lim[1,3], Jianjun Shi[2]
[1] Ulsan National Institute of Science and Technology, Ulsan, Republic of Korea
[2] Georgia Institute of Technology, Atlanta, Georgia, USA
[3] POSCO N.EX.T Hub, POSCO Holdings Inc., Seoul, Republic of Korea
{sbpark1021, hseo75, chjin314, chlim}@unist.ac.kr
{yma447, jianjun.shi}@gatech.edu

ABSTRACT

Time series data are ubiquitous across real-world applications, yet they are highly heterogeneous in temporal dynamics, data formats, and acquisition sources, posing a major challenge for pre-training large-scale time series foundation models (TSFMs). Existing TSFMs often exploit human-assigned domain labels, which are dataset-level annotations derived from external metadata, as implicit or explicit supervision to learn domain-specific representations. However, because a single domain label is shared by all windows within a dataset, it provides a coarse and potentially ambiguous signal that fails to reflect window-level heterogeneity, where similar local dynamics can appear across domains and diverse dynamics can coexist within the same domain. This window-level heterogeneity makes domain-label supervision under-specified for learning dynamics-aware representations in TSFMs. To address this limitation, we introduce WaveEnergy, which uses wavelet decomposition to represent time series windows as multi-scale components and derives a dynamics-aware signature from the energy of coefficients that complements domain labels. Experimentally, WaveEnergy provides a stronger alignment with TSFM embeddings than domain labels and enables finer-grained characterization of datasets by capturing within-domain differences in temporal dynamics.

Track: Research

1 INTRODUCTION

Time series data are ubiquitous in various applications such as energy, traffic, and healthcare (Yang et al., 2025; Lawhern et al., 2018). As real-world time series are collected from various environments, they are highly heterogeneous in temporal dynamics, data formats, or source systems (Goswami et al., 2024). This heterogeneity makes it challenging to pre-train large-scale time series foundation models (TSFMs) that learn unified and generalizable representations (Goswami et al., 2024; Woo et al., 2024). To address this challenge, many TSFMs implicitly or explicitly assume that time series characteristics differ primarily across domains, and thus leverage domain labels as supervising signals to learn domain-specific information from multi-domain datasets (Woo et al., 2024; Zhang et al., 2025; Wang et al., 2024).

Domain is a dataset-level annotation assigned by humans based on external metadata, such as the source system and application context (Napoli et al., 2024). Because a single label is used to represent an entire dataset, domain labels cannot capture the diverse temporal dynamics that emerge across different windows; a phenomenon we term window-level heterogeneity in time series. Although TSFMs typically process time series as sliding windows (sub-series), using a single domain label for all windows provides a coarse-grained and potentially ambiguous training signal. Hence, models can become biased toward dataset-level signals, hindering their ability to learn accurate semantic information of local temporal patterns. In practice, our analysis suggests that domain labels

*Equal contribution.

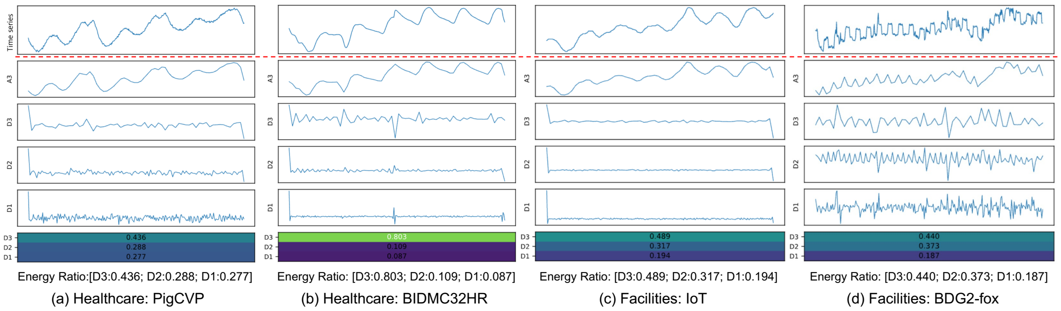

Figure 1: Raw time series and detail coefficients from a 3-level wavelet decomposition. Domain labels fail to capture window-level heterogeneity, whereas the energy distributions of detail coefficients characterize this heterogeneity effectively.

are insufficient for window-level heterogeneity: windows from different domains can exhibit similar temporal patterns (see Figures 1 (b), (c)), while windows within the same domain can be heterogeneous (see Figures 1 (a), (b) and (c), (d), respectively).

We formalize window-level heterogeneity as emerging from interactions among multi-scale components that encode distinct temporal dynamics. A time series comprises multiple dynamic patterns (e.g., trends, events, and spikes) operating at different temporal scales and exhibiting varying strengths over time. To effectively disentangle time series components, the wavelet transform provides a principled multi-scale decomposition in which a time series can be represented as an approximation coefficient and a hierarchy of residual detail coefficients providing multi-scale information (Zhang, 2019). Wavelet transform can effectively characterize window-level heterogeneity: even windows with similar time-domain shapes may exhibit different approximation and detail coefficients, and distinct energy (i.e., the variance of coefficients) distributions (see Figure 1). Thus, we argue that wavelet-derived information complements domain labels by characterizing datasets through their varying temporal dynamics.

In this paper, we introduce WaveEnergy, a wavelet-derived dynamics signature for time series windows that captures interactions among multi-scale temporal patterns. We compute WaveEnergy at the window level and aggregate it across windows to obtain a dataset-level descriptor that complements domain labels. These window- and dataset-level signatures provide fine-grained characterizations of temporal dynamics and enable large-scale dataset categorization via clustering in the WaveEnergy space. Experimentally, our clustering analysis reveals that WaveEnergy aligns with TSFM embeddings more consistently than domain labels. In addition, our analysis indicates that WaveEnergy complements domain labels by resolving within-domain variability in temporal dynamics and offering finer-grained dataset-level annotations.

2 Methodology: WaveEnergy

Notations. Let $\mathcal{D} = \{d_1, \ldots, d_N\}$ be the set of datasets used in our paper, where N is the number of datasets. Each dataset $d \in \mathcal{D}$ provides a collection of time series windows $\{x_{d,i}\}_{i=1}^{n_d}$, where $x_{d,i} \in \mathbb{R}^{C \times T}$ denotes a C-channel window of length T, and i is the index of the window. Also, dataset d is assigned with a domain label $\ell_d \in \{1, \ldots, K_{\text{dom}}\}$. In this paper, we input each time series in a univariate setting by averaging over channels, resulting in $x_{d,i} \in \mathbb{R}^T$.

WaveEnergy Formulation. Given a wavelet basis b, level L, and boundary mode q, we compute the L-level discrete wavelet transform (DWT) of a univariate sequence $x = \text{DWT}(x; b, L, q) \to \{cA_L, cD_L, cD_{L-1}, \ldots, cD_1\}$, where cD_ℓ denotes the detail coefficients at scale ℓ. Then, we define the detail energy at level ℓ as $E_\ell(x) = \|cD_\ell\|_2^2$ and the normalized energy ratio vector as:

$$r_\ell(x) = \frac{E_\ell(x)}{\sum_{j=1}^{L} E_j(x) + \varepsilon}. \tag{1}$$

Parseval's theorem states that the total energy of a time series in the time domain equals the total energy in an orthonormal transform domain, such as a wavelet transform (Mallat, 1999). In the time domain, this energy reflects the aggregate effect of all temporal dynamics entangled in the

Table 1: Clustering analysis on TSFMs embedding space with domain and WaveEnergy labels.

Metrics	k=6	k=9	k=13	k=15	k=18
Intrinsic embedding geometry					
Silhouette (MOMENT embedding)	**0.2800 ± 0.0033**	0.2632 ± 0.0002	0.2616 ± 0.0213	0.2293 ± 0.0120	0.2365 ± 0.0107
Silhouette (LPTM embedding)	**0.2757 ± 0.0159**	0.2495 ± 0.0296	0.2198 ± 0.0134	0.2093 ± 0.0093	0.2085 ± 0.0072
Alignment with MOMENT embeddings					
NMI (Domain labels)	0.2431 ± 0.0185	0.2711 ± 0.0010	0.3059 ± 0.0110	0.3027 ± 0.0057	0.3224 ± 0.0159
NMI (WaveEnergy labels)	**0.2601 ± 0.0142**	**0.2756 ± 0.0018**	**0.3441 ± 0.0058**	**0.3522 ± 0.0089**	**0.3766 ± 0.0167**
ARI (Domain labels)	0.1127 ± 0.0088	0.1565 ± 0.0008	0.1719 ± 0.0184	0.1366 ± 0.0119	0.1449 ± 0.0232
ARI (WaveEnergy labels)	**0.1921 ± 0.0040**	**0.1589 ± 0.0014**	**0.1918 ± 0.0070**	**0.1743 ± 0.0156**	**0.1868 ± 0.0275**
Alignment with LPTM embeddings					
NMI (Domain labels)	0.1049 ± 0.0116	0.1624 ± 0.0104	0.2051 ± 0.0145	0.2220 ± 0.0098	0.2532 ± 0.0052
NMI (WaveEnergy labels)	**0.1480 ± 0.0089**	**0.1893 ± 0.0077**	**0.2736 ± 0.0157**	**0.3149 ± 0.0090**	**0.3707 ± 0.0088**
ARI (Domain labels)	0.0401 ± 0.0024	0.0512 ± 0.0051	0.0544 ± 0.0138	0.0523 ± 0.0092	0.0546 ± 0.0080
ARI (WaveEnergy labels)	**0.0556 ± 0.0084**	**0.0564 ± 0.0132**	**0.0712 ± 0.0073**	**0.0809 ± 0.0073**	**0.0921 ± 0.0106**

signal. In contrast, the wavelet transform decomposes the signal into scale-specific components (approximation and details), so the energy of each coefficient band reflects the contribution of a particular temporal pattern at a specific scale. Accordingly, a wavelet energy ratio vector summarizes how different multi-scale temporal patterns contribute to a window's signal energy, providing a dynamics-aware profile of local behavior and enabling characterization of window-level heterogeneity via the distribution of energy across scales. In practice, approximation coefficients often contain most of the signal energy, dominating the energy profile and masking fine-grained dynamics in the detail bands. We therefore exclude approximation energy and define our vector using only detail-band energies, yielding a signature that better discriminates window-level temporal patterns. Details of the wavelet transform are provided in Appendix B.

We define the energy ratio vector of a window x as $r(x) = [r_1, \ldots, r_L] \in \mathbb{R}^L$. Since $r(x)$ characterizes the window's multi-scale dynamics, we quantize each ratio to construct a dataset-level signature. Detailed explanation on quantization is provided in Appendix C. For each dataset d and level ℓ, we accumulate a histogram $h_{d,\ell} \in \mathbb{R}^B$ over all windows $\{x_{d,i}\}$, and normalize it as:

$$\bar{h}_{d,\ell} = \frac{h_{d,\ell}}{\sum_{k=1}^{B} h_{d,\ell}[k] + \varepsilon}. \tag{2}$$

We represent dataset d by concatenating the normalized histograms across levels:

$$w_d = [\bar{h}_{d,1}; \bar{h}_{d,2}; \ldots; \bar{h}_{d,L}] \in \mathbb{R}^{LB}. \tag{3}$$

Given a large-scale pre-training corpus, we compute w_d for each dataset and perform clustering in the WaveEnergy space. This yields dynamics-aware dataset groupings, enabling the construction of fine-grained dataset labels that reflect detailed temporal patterns beyond coarse domain annotations.

3 Experiments

Experimental Setup. We run all analyses on the Time Series Pile (Goswami et al., 2024) collection, yielding $N = 1049$ datasets with domain annotations spanning $K_{\text{dom}} = 13$ domains (see Table 2). For each dataset $d \in \mathcal{D}$, we extract up to a maximum of $n_d \leq n_{\max}$ windows (with $n_{\max} = 20000$), where the window length and dataset stride are set to 512. For WaveEnergy, we compute wavelet energy signatures using wavelet basis as `sym4` and level $L = 3$ while keeping q as periodization and B fixed to 10. To examine whether WaveEnergy and domain labels align with the TSFM embedding space, we use MOMENT-Large and LPTM as the backbone TSFMs in our analysis (Goswami et al., 2024; Prabhakar Kamarthi & Prakash, 2024). Specifically, we compute a dataset embedding for each dataset d by mean-pooling the embeddings of its windows and then applying ℓ_2 normalization.

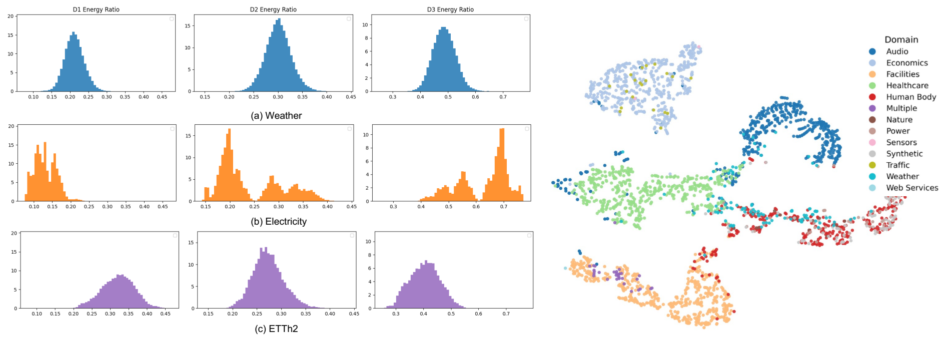

(a) (b)

Figure 2: (a) Distributions of detail-band energy ratios across datasets. (b) t-SNE visualization of window-level WaveEnergy embeddings colored by domain.

Clustering and Evaluation Metrics. To examine whether WaveEnergy and domain labels align with the TSFM embedding space, we conduct clustering-based analyses on TSFM embeddings. Specifically, to probe the intrinsic structure of time series datasets, we apply K-means clustering in the TSFM embedding space and compute the silhouette score using the resulting cluster assignments. We then evaluate Normalized Mutual Information (NMI) and Adjusted Rand Index (ARI) between the TSFM-induced clusters and the labels given by either WaveEnergy or domain annotations across multiple values of k. Higher NMI and ARI indicate that the utilized labels are more consistent with the partitioning implied by the TSFM embedding geometry, suggesting stronger alignment between the label space and the TSFM representations. For robustness, all results are averaged over five random seeds.

3.1 Experimental Results

Table 1 summarizes the clustering analysis of the TSFM embedding space with respect to domain labels and WaveEnergy signatures. We perform K-means clustering with $k \in 6, 9, 13, 15, 18$. Although the Time Series Pile contains 13 domain labels, the embedding spaces of both MOMENT and LPTM achieve the highest silhouette scores at $k = 6$. **This discrepancy suggests that domain labels do not align with the intrinsic granularity of the TSFM embedding structure.** Because TSFMs are trained to capture temporal dependencies through masked patch reconstruction or next-patch prediction, we hypothesize that their embedding spaces preserve the temporal dynamics of input windows. This suggests that domain labels do not fully capture the temporal-dynamics-aware topology of the pre-training corpus. In contrast, WaveEnergy-based labels consistently show higher agreement with TSFM embeddings than domain labels, as measured by NMI and ARI. **This suggests that WaveEnergy signatures more effectively capture the latent grouping structure encoded in TSFM embeddings.**

4 Visualization Analyses on Window-level Heterogeneity

Figure 2 (a) shows histograms of energy ratios for the three detail levels across windows from the Weather, Electricity, and ETTh2 datasets. Overall, the distributions of detail-band energy ratios exhibit significant within-dataset diversity and clear cross-dataset differences, **demonstrating substantial window-level heterogeneity.** Figure 2(b) presents a t-SNE visualization of WaveEnergy embeddings from the pre-training corpus, colored by domain. In each domain, time series windows are widely dispersed in the WaveEnergy space, indicating substantial within-domain variability in temporal dynamics. In addition, although most windows are well separated, noticeable overlap remains, indicating that similar temporal patterns often occur across domains. This cross-domain overlap suggests that **domain labels provide an incomplete description of window-level heterogeneity and WaveEnergy can offer finer-grained dataset-level annotations, capturing various temporal dynamics across domains.**

5 Potential Use and Future Works

Our experiments demonstrate window-level heterogeneity in time series datasets and show that WaveEnergy captures it more effectively than domain labels, as evidenced by stronger alignment with TSFM embeddings. While this workshop paper focuses on analyzing WaveEnergy, as future work, we will investigate how WaveEnergy can be leveraged to improve TSFM pre-training.

First, WaveEnergy enables dynamics-aware construction and curation of pre-training corpora. Existing TSFM corpora are often assembled coarsely to span diverse domains based on researchers' judgment and metadata. Recently, BLAST (Shao et al., 2025) introduced a balanced sampling strategy to improve corpus diversity. In contrast, WaveEnergy is a principled descriptor of dataset-specific temporal dynamics beyond domain tags. We hypothesize that leveraging WaveEnergy for dataset selection and sampling can promote more balanced coverage of diverse dynamics while reducing redundancy from near-duplicate datasets that may be indistinguishable under coarse domain labels.

Second, we will explore TSFM pre-training with WaveEnergy as a self-supervisory signal. While TimesBERT (Zhang et al., 2025) explicitly predicts domain labels during pre-training, we will instead train models to predict or align with window-level WaveEnergy signatures. Because WaveEnergy reflects multi-scale temporal dynamics and varies within a domain, it provides richer supervision than a single dataset-level label. We expect WaveEnergy-supervised pre-training to yield more dynamics-aware and transferable representations than domain-label supervision.

REFERENCES

Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski. Moment: a family of open time-series foundation models. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 16115–16152, 2024.

Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer interfaces. *Journal of neural engineering*, 15(5):056013, 2018.

Xu Liu, Juncheng Liu, Gerald Woo, Taha Aksu, Yuxuan Liang, Roger Zimmermann, Chenghao Liu, Junnan Li, Silvio Savarese, Caiming Xiong, and Doyen Sahoo. Moirai-moe: Empowering time series foundation models with sparse mixture of experts. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=SrEOUSyJcR.

Yong Liu, Haoran Zhang, Chenyu Li, Xiangdong Huang, Jianmin Wang, and Mingsheng Long. Timer: generative pre-trained transformers are large time series models. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 32369–32399, 2024.

Stephane Mallat. *A wavelet tour of signal processing*. Academic Press, 1999.

Otávio Napoli, Dami Duarte, Patrick Alves, Darlinne Hubert Palo Soto, Henrique Evangelista de Oliveira, Anderson Rocha, Levy Boccato, and Edson Borin. A benchmark for domain adaptation and generalization in smartphone-based human activity recognition. *Scientific Data*, 11(1): 1192, 2024.

Harshavardhan Prabhakar Kamarthi and B Aditya Prakash. Large pre-trained time series models for cross-domain time series analysis tasks. *Advances in Neural Information Processing Systems*, 37: 56190–56214, 2024.

Zezhi Shao, Yujie Li, Fei Wang, Chengqing Yu, Yisong Fu, Tangwen Qian, Bin Xu, Boyu Diao, Yongjun Xu, and Xueqi Cheng. Blast: Balanced sampling time series corpus for universal forecasting models. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 2502–2513, 2025.

Yihang Wang, Yuying Qiu, Peng Chen, Kai Zhao, Yang Shu, Zhongwen Rao, Lujia Pan, Bin Yang, and Chenjuan Guo. Rose: Register assisted general time series forecasting with decomposed frequency learning. In *Proceedings of the 42st International Conference on Machine Learning*, 2024.

Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Unified training of universal time series forecasting transformers. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 53140–53164, 2024.

Yang Yang, Weilong Ding, Yuwei Gu, Hao Zhang, Wei Wei, and Honghao Gao. Beyond data heterogeneity: A multivariate time series forecasting for energy systems through enhanced channel fusion in frequency domain. *Information Fusion*, pp. 103870, 2025.

Dengsheng Zhang. Wavelet transform. In *Fundamentals of image data mining: Analysis, Features, Classification and Retrieval*, pp. 35–44. Springer, 2019.

Haoran Zhang, Yong Liu, Yunzhong Qiu, Haixuan Liu, Zhongyi Pei, Jianmin Wang, and Mingsheng Long. Timesbert: A bert-style foundation model for time series understanding. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 10975–10983, 2025.

A Time Series Foundation Model Pre-training with Domain Labels

Time series foundation models (TSFMs) aim to learn generalizable representations that require minimal adaptation across downstream tasks. A key challenge is the substantial heterogeneity of time series windows within and across datasets, which limits the effective transfer of pre-trained models. To organize and supervise learning under such heterogeneity, prior work has often adopted domain labels as a convenient and consistent dataset-level discriminator. Because time series collected from different domains frequently differ in temporal dynamics and even measurement configurations (e.g., the number of channels) due to their acquisition sources, existing TSFMs commonly leverage domain information during pre-training—implicitly or explicitly—to learn domain-specific representations from multi-domain corpora.

In the broader literature, domain labels have been widely used to curate and structure pre-training corpora. Large-scale corpora such as Time Series Pile (Goswami et al., 2024), LOTSA (Woo et al., 2024), and UTSD (Liu et al., 2024) are constructed to increase diversity in terms of the domains. By curating these corpora with domain-specific temporal patterns, TSFMs implicitly leverage domain labels during pre-training to learn domain-discriminative representations and improve cross-domain generalization. Recently, several studies have attempted to explicitly incorporate domain information within a pre-training framework. MOIRAI (Woo et al., 2024) proposed multi-patch size projection layers to activate different patch projections depending on the characteristic timescales of each dataset. The appropriate patch size for each dataset is predefined based on the domain-specific characteristics. TimesBERT (Zhang et al., 2025) explicitly leveraged domain information in pre-training through functional token prediction tasks, including domain classification and variate-domain discrimination. ROSE (Wang et al., 2024) and MOIRAI-MOE (Liu et al., 2025) selectively activate different layers to capture domain-specific information. As such, domain labels have become a central organizing standard and supervisory signal in TSFM pre-training. Therefore, addressing their coarse granularity by complementing them with dynamics-aware supervision constitutes an important contribution.

B Wavelet Transform

Wavelet Transform. Wavelet transform represents a finite-length signal through a hierarchy of scale-dependent components, separating slow-varying structure from progressively finer fluctuations (Mallat, 1999). Given an input window $x_{d,i} \in \mathbb{R}^T$, we apply an L-level discrete wavelet transform $\mathrm{DWT}(x_{d,i}; b, L, q)$, where b is the wavelet basis, L is the level, and q is the boundary mode, and obtain the coefficient streams $\{\mathbf{c}A_L, \mathbf{c}D_L, \mathbf{c}D_{L-1}, \ldots, \mathbf{c}D_1\}$. The approximation coefficient $\mathbf{c}A_L$ summarizes the coarsest (low-frequency) behavior, while each detail coefficient $\mathbf{c}D_\ell$ isolates residual variations at scale ℓ (with ℓ=1 being the finest scale). In this work, we characterize each window by the distribution of scale-wise residual energy across the detail streams, computed as $E_\ell(x_{d,i}) = \|\mathbf{c}D_\ell\|_2^2$ and normalized as $r_\ell(x_{d,i}) = \frac{E_\ell(x_{d,i})}{\sum_{j=1}^{L} E_j(x_{d,i})+\varepsilon}$, yielding the multi-scale signature $r(x_{d,i}) = [r_1(x_{d,i}), \ldots, r_L(x_{d,i})]$ used throughout our analysis.

Choice of Wavelet Basis and Boundary Mode. In our implementation, we adopt the symlet-4 (`sym4`) basis as the default choice and apply the L=3 level DWT with an explicit boundary-handling mode on finite-length windows. The symlet family provides a practical balance between temporal localization and frequency selectivity through moderately long, near-symmetric filters, which helps mitigate phase distortion and reduces sensitivity to edge effects when analyzing discrete-time signals. Relative to very short bases (e.g., haar), `sym4` typically yields smoother approximation behavior and more stable detail coefficient streams, while avoiding the increased computational cost and potential over-smoothing that can arise from substantially higher-order wavelets. For boundary handling, since convolution requires samples beyond a finite window, we adopt periodization as the boundary-handling mode q, treating each window as circularly wrapped at the boundaries during filtering.

C Energy ratio vector quantization

For each window $x_{d,i}$, we compute the energy ratio vector $r(x_{d,i}) = [r_1(x_{d,i}), \ldots, r_L(x_{d,i})] \in \mathbb{R}^L$, where each component $r_\ell(x_{d,i}) \in [0, 1]$ summarizes the relative residual energy at scale ℓ. To obtain a dataset-level signature that is comparable across datasets with different numbers of windows, we quantize each ratio component into B uniform bins and aggregate their empirical distributions. Concretely, for a dataset d with windows $\{x_{d,i}\}_{i=1}^{n_d}$ and for each level $\ell \in \{1, \ldots, L\}$, we discretize each ratio component $r_\ell(x_{d,i}) \in [0, 1]$ into B bins and accumulate a per-level histogram $h_{d,\ell} \in \mathbb{R}^B$ across windows; specifically, we form $h_{d,\ell}$ by counting how many windows fall into each bin. We then normalize it as $\bar{h}_{d,\ell} = \frac{h_{d,\ell}}{\sum_{k=1}^{B} h_{d,\ell}[k]+\varepsilon}$. Each dataset is represented as $w_d = [\bar{h}_{d,1}; \bar{h}_{d,2}; \ldots; \bar{h}_{d,L}] \in \mathbb{R}^{LB}$. Stacking all dataset vectors yields the matrix $W = [w_{d_1}; \ldots; w_{d_N}] \in \mathbb{R}^{N \times LB}$ in the WaveEnergy space.

D Datasets used in WaveEnergy

Table 2: **Domains and examples of datasets from Time Series Pile.** This table is reproduced from the original MOMENT paper (Goswami et al., 2024).

Domains	**Examples of Datasets**
Healthcare	ECG, EEG, Hospital
Human Body	Tongue movement, Finger movement, Muscle Signals
Nature	Fish outlines, Flower outlines, River flow
Audio	Arabic speech, Japanese Speech, Phonetics
Power	Power consumption, Electricity, Home appliance usage
Economics	Exchange Rate, Bitcoin, Tourism
Traffic	Road Traffic, Pedestrian cross, Line occupancy rate
Weather	Temperature, Rain, Wind
Facilities	Machine Status, Spacecraft Status, Web traffic
Web Services	IOPS, NAB
Synthetic	MGAB
Sensors	NASA MSL, NASA SMAP
Gait	Daphnet

We adopt the Time Series Pile curated in MOMENT (Goswami et al., 2024) as our primary corpus for analyzing heterogeneity and representations across diverse time series domains. As shown in Table 2, Time Series Pile aggregates real-world and synthetic time series spanning 13 curated domains (e.g., healthcare, human body, and power). We treat these domain labels as a coarse, dataset-level reference to examine whether TSFM embeddings exhibit domain-level structure and to quantify deviations from the domain taxonomy. We further compare TSFM embeddings with WaveEnergy, a dynamics-aware dataset-level descriptor that complements domain labels.

E ACKNOWLEDGMENTS

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (RS-2024-00443780, Development of Foundation Models for Bioelectrical Signal Data and Validation of Their Clinical Applications: A Noise-and-Variability Robust, Generalizable Self-Supervised Learning Approach; RS-2024-00439932, SW Starlab), Korean Institute for Advancement of Technology(KIAT) grant funded by the Korea Government(MOTIE) (RS-2024-00435815, Human Resource Development Program for Industrial Innovation(Global)), the Basic Science Research Program through the National Research Foundation of Korea (NRF) grant funded by the Ministry of Education (RS-2024-00463188), and the Korea Planning & Evaluation Institute of Industrial Technology (KEIT) funded by the Ministry of Trade, Industry and Resources (RS-2025-25458052, Development of Core Technologies for Manufacturing Foundation Models).