

Temporal Tokenization Strategies for Event Sequence Modeling with Large Language Models

Zefang Liu¹, Nam H. Nguyen¹, Yinzhu Quan², Shi-Xiong Zhang¹

¹Capital One, USA

²Georgia Institute of Technology, USA

Correspondence: zefang.liu@capitalone.com

Abstract

Representing continuous time is a critical and under-explored challenge in modeling temporal event sequences with large language models (LLMs). Various strategies like byte-level representations or calendar tokens have been proposed. However, the optimal approach remains unclear, especially given the diverse statistical distributions of real-world event data, which range from smooth log-normal to discrete, spiky patterns. This paper presents a systematic empirical study of temporal tokenization for modeling event sequences with LLMs, comparing distinct encoding strategies: naive numeric strings, high-precision byte-level representations, human-semantic calendar tokens, classic uniform binning, and adaptive residual scalar quantization. We evaluate these strategies by fine-tuning LLMs on real-world datasets that exemplify these diverse distributions. Our analysis reveals that no single strategy is universally superior; instead, prediction performance depends heavily on aligning the tokenizer with the data’s statistical properties, highlighting temporal tokenization as a critical yet often overlooked design dimension in LLM-based event modeling.

1 Introduction

Large language models (LLMs) have demonstrated remarkable capabilities in processing text, leading to new opportunities in modeling sequences of discrete events. While traditional temporal point processes (TPPs) (Hawkes, 2018; Reinhart, 2018; Shchur et al., 2021; Zhou et al., 2025) provide a framework, recent approaches leveraging LLMs (Liu and Quan, 2024; Kong et al., 2026) uniquely capture the rich textual semantics of events, framing event sequence modeling as a structured generative task that extends the boundaries of language-centric modeling. This integration, however, presents a fundamental challenge: LLMs are inherently discrete text processors, while real-

world events occur in continuous time. The critical step of converting continuous time values into discrete tokens, known as temporal tokenization, is a crucial but under-explored component.

The optimal strategy for temporal tokenization is unclear. While the most naive approach is to represent time as a simple numeric string (Bhatia et al., 2025), other methods include human-semantic calendar tokens (He et al., 2025a,b) and byte-level representations (Kong et al., 2026). This choice is complicated by the diverse nature of real-world event data. Our analysis of several event sequence datasets reveals that temporal distributions vary substantially; they range from smooth log-normal patterns to discrete, spiky distributions and complex mixed-modal structures. A tokenization strategy that excels on one distribution may fail on another.

This paper presents a systematic empirical study of temporal tokenization strategies for event sequence modeling with LLMs¹. We conduct a comprehensive comparison of five distinct encoding strategies: naive numeric strings, high-precision byte-level representations, human-semantic calendar tokens, classic uniform binning in linear or log space, and adaptive, data-driven residual scalar quantization (RSQ). By fine-tuning LLMs on a suite of real-world datasets that exemplify these diverse distributions, we analyze the resulting trade-offs between prediction performance and computational efficiency. Our findings establish a practical guide on data-tokenizer alignment, demonstrating that no single strategy is universally optimal, but that performance is maximized by aligning the tokenization scheme with the data’s specific statistical distribution.

¹GitHub repository: <https://github.com/CapitalOne-Research/temporal-tokenization>

2 Related Work

Integrating LLMs with event sequence modeling presents a fundamental conflict in representing continuous time for discrete text processors. Two distinct strategies have emerged: hybrid frameworks that treat time as a vector embedding (Liu and Quan, 2024, 2025), and pure tokenization approaches that convert continuous time intervals into discrete byte-level tokens (Kong et al., 2026). This methodological divergence, intensified by the challenges posed by benchmarks like EasyTPP (Xue et al., 2024) and DanmakuTPPBench (Jiang et al., 2025), highlights that the need to solve this core representation problem remains unsettled. Research in adjacent domains confirms that varied tokenization strategies are critical for success: naive numeric strings have been shown to be problematic (Bhatia et al., 2025), while specialized methods like hierarchical calendar tokens (He et al., 2025a,b) and patch-wise binning (Han et al., 2025) succeed in mobility and time series forecasting. Our paper addresses this gap by providing a systematic empirical comparison of these temporal tokenization strategies for event sequence modeling across diverse statistical distributions.

3 Temporal Tokenization

An event sequence is a series $\mathcal{S} = \{(t_1, k_1), (t_2, k_2), \dots, (t_N, k_N)\}$, where each event i consists of a continuous time value t_i and a discrete text-based event type k_i . To model such sequences with an LLM, both components must be converted into discrete tokens. While the textual event type k_i is naturally handled by the LLM’s native tokenizer, the continuous time value t_i requires a specialized temporal tokenization strategy to represent either the absolute time t_i or, more commonly, the relative time interval $\Delta t_i = t_i - t_{i-1}$.

3.1 Numeric String Tokenization

The naive numeric string strategy formats the continuous time value v_i (e.g., the time interval) into a plain text string with a fixed precision, such as “0.076”. This string is then passed directly to the LLM’s standard subword tokenizer, a process that often fragments it into semantically meaningless subwords (e.g., “0”, “.”, “07”, “6”). The approach requires no modifications to the model’s vocabulary, instead relying entirely on the LLM’s pre-trained ability to parse and reason about numerical

strings as standard text.

3.2 Byte Tokenization

Byte tokenization, a strategy introduced by Language-TPP (Kong et al., 2026), bypasses the LLM’s subword tokenizer to enable seamless integration of continuous time. The approach first represents the time value v_i as a standard 32-bit float, then deterministically splits its 4-byte representation into four discrete byte values. The LLM’s vocabulary is augmented with 256 new special tokens (e.g., $\langle |byte_0| \rangle \dots \langle |byte_255| \rangle$) to represent each possible byte value. The float v_i is thus always encoded as a fixed-length sequence of four tokens, preserving full float32 precision in a token-efficient manner.

3.3 Calendar Tokenization

The calendar tokenization strategy is human-centric, converting a time value into a sequence of semantically meaningful tokens based on the Gregorian calendar. An absolute timestamp t_i can be encoded into tokens like $\langle |year_2025| \rangle$, $\langle |month_10| \rangle$, and $\langle |day_25| \rangle$. Similarly, a relative interval Δt_i can be encoded into tokens such as $\langle |days_03| \rangle$ and $\langle |hours_05| \rangle$. This approach requires adding calendar-based tokens to the vocabulary and relies on the LLM’s pretrained world knowledge of calendar systems.

3.4 Scale Bin Tokenization

Scale bin tokenization represents a classic TPP-based uniform binning approach. It first transforms all time values v_i into a target space (either linear or log scale). During training, the strategy finds the minimum and maximum of these transformed values and divides this range into K (e.g., 256) uniform bins. Each bin is then assigned a unique token (e.g., $\langle |bin_000| \rangle \dots \langle |bin_255| \rangle$) added to the vocabulary. During decoding, the continuous time value is reconstructed by mapping the token back to the center of its corresponding bin in the transformed space. This method is data-driven but assumes a uniform distribution within the transformed space.

3.5 Residual Scalar Quantization Tokenization

Residual scalar quantization (RSQ) tokenization is an adaptive multi-stage binning strategy analogous to residual vector quantization (RVQ) (Gray, 1984; Makhoul et al., 1985). This method first

transforms all time values v_i into a target space (e.g., linear or log scale), yielding $v'_i = f(v_i)$. Instead of using uniform bins, it employs a hierarchy of quantizers. At the first level, a K-Means model is trained on the transformed data $\{v'_i\}$ to find an initial codebook of K_1 cluster centers, $C^{(1)} = \{c_1^{(1)}, \dots, c_{K_1}^{(1)}\}$. A value v'_i is quantized to its nearest centroid $c_{q_1}^{(1)}$, where q_1 is the token index for the first level. The process then becomes iterative: the residual error from this first quantization, $r_1 = v'_i - c_{q_1}^{(1)}$, is calculated and passed to the second level. A second K-Means model is trained on these first-level residuals $\{r_1\}$ to find a new codebook $C^{(2)}$ with K_2 centers. This second level quantizes r_1 to its nearest centroid $c_{q_2}^{(2)}$, yielding the second token q_2 and a new residual $r_2 = r_1 - c_{q_2}^{(2)}$. This is repeated for N levels, encoding the single scalar v'_i as a fixed-length sequence of N discrete token indices $\langle q_1, q_2, \dots, q_N \rangle$. These are mapped to unique tokens (e.g., $\langle |L0_015| \rangle$, $\langle |L1_122| \rangle$, \dots), creating a fine-grained, compositional representation. Decoding is a simple additive process where the reconstructed value \hat{v}'_i is the sum of the centroids from each level’s codebook, $\hat{v}'_i = c_{q_1}^{(1)} + c_{q_2}^{(2)} + \dots + c_{q_N}^{(N)}$, before applying the inverse transform $v_i = f^{-1}(\hat{v}'_i)$.

4 Experiments

In this section, we empirically evaluate the temporal tokenization strategies discussed previously.

4.1 Data

We conduct our experiments using the five real-world datasets from TPP-LLM² (Liu and Quan, 2024). As detailed in Table 1, these datasets provide sequences of textual event types, raw absolute timestamps, and pre-computed relative time intervals with proper units. Calendar-based tokenizers ingest the absolute timestamps, while the numeric, byte, and binning-based strategies operate on the relative time intervals. We also analyze the statistical properties of the relative time intervals (Δt_i) from each dataset. As shown in Figure 1, the distributions are highly diverse. They range from smooth, log-normal patterns (e.g., Stack Overflow) and multi-modal curves (e.g., NYC Taxi) to extremely discrete, spiky distributions with sharp peaks at specific intervals (e.g., Amazon Review). More dataset distributions and examples can be found in Appendix A and B.

²Datasets: <https://huggingface.co/tppllm>

Table 1: Dataset statistics, including the number of event types, total events, sequences with their train/validation/test splits, average sequence length, and the temporal unit for intervals.

Dataset	Types	Events	Seqs	Train/Val/Test	Seq Len	Unit
Stack Overflow	25	187,836	3,336	2,668/334/334	56.31	Month
Chicago Crime	20	202,867	4,048	3,238/405/405	50.12	Month
NYC Taxi	8	362,370	2,957	2,365/296/296	122.55	Hour
US Earthquake	3	30,218	3,079	2,463/308/308	9.81	Day
Amazon Review	18	127,054	2,245	1,796/224/225	56.59	Week

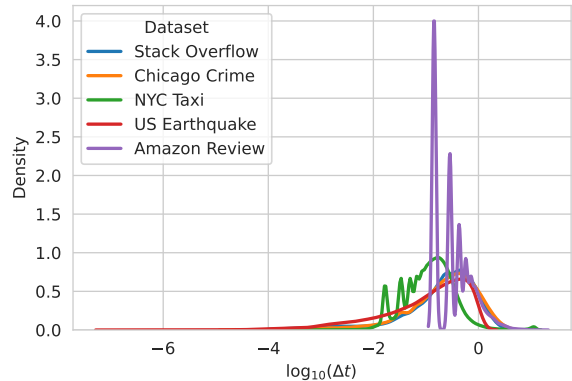


Figure 1: Log-scale distributions of relative time intervals (Δt_i) across all five datasets.

4.2 Experimental Setup

We fine-tune a pretrained LLM backbone Llama-3.2-1B (Llama Team, 2024) with QLoRA (Detmers et al., 2023) through the Hugging Face framework (Wolf et al., 2020) for all experiments and use TPP-LLM (Liu and Quan, 2024) as the baseline. We format each event using a consistent template inspired from Language-TPP (Kong et al., 2026): $\langle |begin_of_event| \rangle \langle |type_prefix| \rangle \langle |type_tokens| \rangle \langle |time_prefix| \rangle \langle |time_tokens| \rangle \langle |end_of_event| \rangle$. We compare multiple temporal tokenization strategies, including numeric strings (6-decimal precision), byte-level tokens, calendar tokens (at second and day resolution), uniform scale binning (256 bins, linear and base-10 logarithm scale), and multi-level RSQ (1 level of 256 bins and 4 levels of 64 bins each, linear and base-10 logarithm scale).

We evaluate model performance on prediction quality, measured by the accuracy of the next event type prediction and the root mean squared error (RMSE) of the predicted time interval, and efficiency, reported by the number of temporal tokens to represent a single time value for each strategy, serving as a hardware-agnostic proxy for inference latency. For strategies predicting absolute timestamps, the RMSE is computed directly between the

Table 2: Main experimental results comparing temporal tokenization strategies. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported across all five datasets. Tokens (\downarrow) measures the number of tokens required to represent a single time value. Results are averaged over 5 runs, with standard deviations reported as subscripts. (Abs. = Absolute; Rel. = Relative; Sec. = Second; L1/L4 = 1 or 4 levels.)

Tokenizer	Tokens	Stack Overflow		Chicago Crime		NYC Taxi		US Earthquake		Amazon Review	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
Numeric String	~4	44.5 \pm .1	0.625 \pm .000	27.1 \pm .1	0.728 \pm .000	92.0 \pm .0	0.888 \pm .002	64.0 \pm .2	0.381 \pm .000	69.9 \pm .1	0.647 \pm .000
Byte	4	44.3 \pm .0	0.516 \pm .004	27.2 \pm .1	0.726 \pm .000	92.0 \pm .0	0.892 \pm .001	64.0 \pm .1	0.348 \pm .006	69.7 \pm .0	0.644 \pm .000
Abs. Calendar (Day)	3	44.7 \pm .1	0.578 \pm .002	27.1 \pm .1	0.698 \pm .001	91.9 \pm .1	16.086 \pm .009	64.4 \pm .3	0.838 \pm .002	70.0 \pm .1	0.633 \pm .002
Abs. Calendar (Sec.)	6	44.8 \pm .1	0.569 \pm .003	27.1 \pm .1	0.692 \pm .001	92.2 \pm .0	0.869 \pm .006	64.2 \pm .3	0.353 \pm .001	70.0 \pm .1	0.634 \pm .001
Rel. Calendar (Day)	3	44.4 \pm .1	0.626 \pm .000	27.1 \pm .1	0.728 \pm .000	91.9 \pm .0	0.912 \pm .000	64.0 \pm .2	0.381 \pm .000	69.8 \pm .1	0.648 \pm .001
Rel. Calendar (Sec.)	6	44.5 \pm .1	0.623 \pm .001	27.2 \pm .1	0.728 \pm .000	92.0 \pm .0	0.891 \pm .002	64.3 \pm .2	0.381 \pm .000	69.7 \pm .1	0.648 \pm .000
Scale Bin (Linear)	1	44.3 \pm .1	0.615 \pm .000	27.1 \pm .1	0.717 \pm .000	92.0 \pm .0	0.894 \pm .007	64.0 \pm .3	0.376 \pm .004	69.7 \pm .1	0.637 \pm .001
Scale Bin (Log)	1	44.3 \pm .0	0.479 \pm .003	27.1 \pm .1	0.562 \pm .004	92.0 \pm .0	1.127 \pm .061	63.9 \pm .2	0.360 \pm .015	69.6 \pm .1	0.632 \pm .001
RSQ (Linear, L1)	1	44.2 \pm .1	0.595 \pm .004	27.1 \pm .1	0.708 \pm .003	92.0 \pm .0	0.911 \pm .017	63.7 \pm .3	0.355 \pm .003	69.7 \pm .2	0.647 \pm .001
RSQ (Linear, L4)	4	44.5 \pm .0	0.603 \pm .002	27.1 \pm .1	0.698 \pm .003	92.0 \pm .0	0.891 \pm .007	64.1 \pm .2	0.363 \pm .002	69.7 \pm .1	0.646 \pm .001
RSQ (Log, L1)	1	44.3 \pm .1	0.497 \pm .005	27.1 \pm .2	0.576 \pm .009	92.0 \pm .0	1.221 \pm .040	63.8 \pm .2	0.368 \pm .010	69.6 \pm .1	0.632 \pm .001
RSQ (Log, L4)	4	44.2 \pm .1	0.474 \pm .003	27.2 \pm .1	0.575 \pm .003	92.0 \pm .0	1.201 \pm .057	64.2 \pm .2	0.366 \pm .010	69.7 \pm .1	0.631 \pm .000
TPP-LLM ¹	-	44.4 \pm .1	0.464 \pm .001	27.3 \pm .1	0.542 \pm .001	91.8 \pm .1	0.826 \pm .002	64.1 \pm .3	0.265 \pm .001	69.5 \pm .1	0.570 \pm .002

predicted and true timestamps (inherently reflecting the relative interval error), with all final RMSEs normalized by the dataset-specific temporal units detailed in Table 1. During inference, if a tokenizer generates malformed or partial sequences that fail to parse, the predicted time interval defaults to zero to ensure a robust error penalty during evaluation. We run 5 independent experiments with different random seeds for each setting, reporting the averaged scores alongside their standard deviations. Additional training and evaluation details are provided in Appendix C.

4.3 Experimental Results

Our main findings, presented in Table 2, reveal a clear and important divergence: the choice of temporal tokenizer has a minimal impact on event type prediction but a significant and data-dependent impact on event time prediction.

For the next event type accuracy, all strategies perform similarly across all five datasets. Conversely, the performance on next event time (RMSE) shows dramatic variation, supporting the observation that a tokenizer must align with the dataset’s underlying temporal distribution. For the datasets with skewed log-normal distributions (Stack Overflow, Chicago Crime) and the spiky log-scale distribution (Amazon Review), the log-based strategies consistently outperform others. Scale Bin

(Log) achieves the best RMSE on Chicago Crime, while RSQ (Log, L4) is the top performer on Stack Overflow and Amazon Review. Intuitively, log-based tokenization allocates a higher density of discrete bins to smaller time intervals, preserving critical precision in the dense regions of these skewed distributions while efficiently compressing the long tail. However, these same log-based strategies perform poorly on the mixed distribution of NYC Taxi, where Absolute Calendar (Second) is the most effective. Unlike the other datasets, NYC Taxi exhibits strong daily and hourly cyclic patterns dictated by human mobility; calendar tokens explicitly capture this intrinsic periodicity, whereas simple interval binning obscures it. For the US Earthquake dataset, which also exhibits a skewed log-normal distribution, Byte achieves the lowest error, suggesting that its high-precision floating-point representation is advantageous for strictly continuous scientific measurements.

This outcome demonstrates that the optimal tokenization strategy is highly dependent on the data’s specific statistical properties, and a similar trade-off applies to efficiency. The single-token log-scale binning and RSQ strategies provide an excellent balance, achieving top-tier RMSE performance on several datasets while being the most token-efficient. In contrast, multi-token strategies like byte and multi-level RSQ explicitly trade token efficiency for higher precision.

While our pure LLM-based approach results in higher RMSEs on several datasets compared to

¹TPP-LLM results are included for indirect comparison, as it uses a distinct framework employing TPP-specific prediction heads and a TPP-based loss function.

TPP-LLM (Liu and Quan, 2024), it offers a simpler and more scalable architecture. Our method treats time prediction as a standard next-token prediction task, entirely avoiding specialized TPP prediction heads, the complex TPP log-likelihood loss function, and substantial computational burden from its integral term. This architectural simplicity, which still maintains competitive event type accuracy, streamlines the modeling process and makes large-scale continual pre-training on event sequences a more feasible and promising direction for future work. More ablation studies can be found in Appendix D.

5 Conclusion

We presented a systematic empirical study comparing temporal tokenization strategies for modeling event sequences with LLMs. Our findings reveal a critical trade-off: while event type prediction accuracy is largely insensitive to the choice of time tokenizer, next event time prediction is highly dependent on the alignment between the tokenization strategy and the dataset’s underlying statistical distribution. We found that log-based strategies (RSQ and scale binning) excel on datasets with log-normal or spiky-log distributions, while calendar tokenization approaches are more robust for mixed-modal distributions. Our work provides a practical guide for this crucial design choice and validates that a pure LLM fine-tuning framework is a simple yet competitive alternative to complex hybrid temporal point process models.

Limitations

Our study, while systematic, has several limitations. The experiments were conducted primarily on Llama-3.2 models up to 3B parameters; these findings may not fully generalize to larger models or different model families. While our five datasets were chosen for their diverse distributions, they do not represent the entire spectrum of real-world event data. Our evaluation is also focused on next-event prediction, and these trade-offs might differ for other tasks, such as long-horizon forecasting or sequence-level generation. Furthermore, the use of PEFT may introduce a representation bottleneck compared to full fine-tuning, potentially limiting the model’s ability to fully integrate specialized temporal tokens into its global attention space. Finally, while we explored a representative set of temporal tokenizers and their configurations, an

exhaustive search across the diverse landscape of tokenization algorithms, binning granularities, and quantization hierarchies remains an area for future investigation.

Ethical Considerations

Our work focuses on the technical trade-offs of tokenization, utilizing established, publicly available benchmarks that are anonymized and contain no personally identifiable information. However, because these datasets model human behaviors, improved predictive accuracy carries potential risks inherited from downstream applications. Crucially, our finding that temporal tokenization strategies are not universal indicates that applying a strategy optimized for one pattern to a mismatched domain may cause performance degradation and result in inconsistent predictive accuracy. Therefore, aligning the temporal tokenization strategy with the underlying data distribution is essential to ensure robust and reliable model performance.

References

- Gagan Bhatia, Maxime Peyrard, and Wei Zhao. 2025. [Date fragments: A hidden bottleneck of tokenization for temporal reasoning](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3201–3219, Suzhou, China. Association for Computational Linguistics.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [QLoRA: Efficient finetuning of quantized LLMs](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 10088–10115.
- Robert Gray. 1984. [Vector quantization](#). *IEEE ASSP Magazine*, 1(2):4–29.
- Xiao Han, Xinfeng Zhang, Yiling Wu, Zhenduo Zhang, Zhiyuan Deng, and Zhe Wu. 2025. [PatchCat: Re-thinking temporal tokenization in time series forecasting](#). ICLR 2026 Conference Withdrawn Submission.
- Alan G. Hawkes. 2018. [Hawkes processes and their applications to finance: A review](#). *Quantitative Finance*, 18(2):193–198.
- Haoyu He, Haozheng Luo, Yan Chen, and Qi R. Wang. 2025a. [Efficient temporal tokenization for mobility prediction with large language models](#). *arXiv preprint arXiv:2507.14017*.
- Haoyu He, Haozheng Luo, Yan Chen, and Qi R. Wang. 2025b. [RHYTHM: Reasoning with hierarchical temporal tokenization for human mobility](#). In *Advances in Neural Information Processing Systems (NeurIPS)*,

- volume 38, pages 84700–84730. Curran Associates, Inc.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations (ICLR)*.
- Yue Jiang, Jichu Li, Yang Liu, Dingkan Yang, Feng Zhou, and Quyu Kong. 2025. [DanmakuTPPBench: A multi-modal benchmark for temporal point process modeling and understanding](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38. Curran Associates, Inc.
- Quyu Kong, Yixuan Zhang, Yang Liu, Panrong Tong, Enqi Liu, and Feng Zhou. 2026. [Byte-token enhanced language models for temporal point processes analysis](#). In *Proceedings of the ACM Web Conference 2026*, pages 7013–7023.
- Zefang Liu and Yinzhu Quan. 2024. [TPP-LLM: Modeling temporal point processes by efficiently fine-tuning large language models](#). *arXiv preprint arXiv:2410.02062*.
- Zefang Liu and Yinzhu Quan. 2025. [Retrieval of temporal event sequences from textual descriptions](#). In *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing (KnowledgeNLP)*, pages 37–49, Albuquerque, New Mexico, USA. Association for Computational Linguistics.
- Llama Team. 2024. [The Llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations (ICLR)*.
- John Makhoul, Salim Roucos, and Herbert Gish. 1985. [Vector quantization in speech coding](#). *Proceedings of the IEEE*, 73(11):1551–1588.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research (JMLR)*, 12:2825–2830.
- Alex Reinhart. 2018. [A review of self-exciting spatio-temporal point processes and their applications](#). *Statistical Science*, 33(3):299–318.
- Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. 2021. [Neural temporal point processes: A review](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 4585–4593. International Joint Conferences on Artificial Intelligence Organization.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP)*, pages 38–45.
- Siqiao Xue, Xiaoming Shi, Zhixuan Chu, Yan Wang, Hongyan Hao, Fan Zhou, Caigao Jiang, Chen Pan, James Y. Zhang, Qingsong Wen, Jun Zhou, and Hongyuan Mei. 2024. [EasyTPP: Towards open benchmarking temporal point processes](#). In *International Conference on Learning Representations (ICLR)*.
- Feng Zhou, Quyu Kong, Jie Qiao, Cheng Wan, Yixuan Zhang, and Ruichu Cai. 2025. [Advances in temporal point processes: Bayesian, neural, and LLM approaches](#). *arXiv preprint arXiv:2501.14291*.

A Data Distributions

This appendix provides the full distributions for the relative time intervals (Δt_i) for all five datasets: Stack Overflow (Figure 2), Chicago Crime (Figure 3), NYC Taxi (Figure 4), US Earthquake (Figure 5), and Amazon Review (Figure 6). Each figure displays the distribution on both a linear scale (top) and a log scale (bottom) to show the overall shape and highlight the behavior of smaller interval values.




Figure 2: Distribution of relative time intervals (Δt_i) for the Stack Overflow dataset, showing linear scale (top) and log scale (bottom).

B Data Examples

This appendix provides tokenization examples for a single event sequence from the Stack Overflow




Figure 3: Distribution of relative time intervals (Δt_i) for the Chicago Crime dataset, showing linear scale (top) and log scale (bottom).




Figure 5: Distribution of relative time intervals (Δt_i) for the US Earthquake dataset, showing linear scale (top) and log scale (bottom).




Figure 4: Distribution of relative time intervals (Δt_i) for the NYC Taxi dataset, showing linear scale (top) and log scale (bottom).




Figure 6: Distribution of relative time intervals (Δt_i) for the Amazon Review dataset, showing linear scale (top) and log scale (bottom).

training set. For brevity, we display only the first two events using the type-time prompt template.

Numeric String Tokenization:

```
'<|begin_of_event|>', '<|type_prefix|>',
'G', 'uru', '<|time_prefix|>', '0',
'.', '000', '000', '<|end_of_event|>',
'<|begin_of_event|>', '<|type_prefix|>',
'Good', 'ĠAnswer', '<|time_prefix|>',
'0', '.', '437', '604',
'<|end_of_event|>', ...
```

Byte Tokenization:

```
'<|begin_of_event|>',
'<|type_prefix|>', 'G', 'uru',
'<|time_prefix|>', '<|byte_000|>',
'<|byte_000|>', '<|byte_000|>',
'<|byte_000|>', '<|end_of_event|>',
'<|begin_of_event|>', '<|type_prefix|>',
'Good', 'ĠAnswer', '<|time_prefix|>',
```

```
'<|byte_147|>', '<|byte_013|>',
'<|byte_224|>', '<|byte_062|>',
'<|end_of_event|>', ...
```

Absolute Calendar Tokenization (Second Resolution):

```
'<|begin_of_event|>',
'<|type_prefix|>', 'G', 'uru',
'<|time_prefix|>', '<|year_2022|>',
'<|month_01|>', '<|day_13|>',
'<|hour_02|>', '<|min_52|>',
'<|sec_08|>', '<|end_of_event|>',
'<|begin_of_event|>', '<|type_prefix|>',
'Good', 'ĠAnswer', '<|time_prefix|>',
'<|year_2022|>', '<|month_01|>',
'<|day_26|>', '<|hour_05|>',
'<|min_56|>', '<|sec_36|>',
'<|end_of_event|>', ...
```

Absolute Calendar Tokenization (Day

Resolution): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|year_2022|>', '<|month_01|>', '<|day_13|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|year_2022|>', '<|month_01|>', '<|day_26|>', '<|end_of_event|>', ...

Relative Calendar Tokenization (Second Resolution): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|year_00|>', '<|month_00|>', '<|day_00|>', '<|hour_00|>', '<|min_00|>', '<|sec_00|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|year_00|>', '<|month_00|>', '<|day_13|>', '<|hour_03|>', '<|min_04|>', '<|sec_28|>', '<|end_of_event|>', ...

Relative Calendar Tokenization (Day Resolution): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|year_00|>', '<|month_00|>', '<|day_00|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|year_00|>', '<|month_00|>', '<|day_13|>', '<|end_of_event|>', ...

Scale Bin Tokenization (Linear Scale): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|bin_000|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|bin_013|>', '<|end_of_event|>', ...

Scale Bin Tokenization (Log Scale): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|bin_000|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|bin_189|>', '<|end_of_event|>', ...

Residual Scalar Quantization Tokenization (Linear Scale, 1 Level): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|L0_059|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>',

'Good', 'ĠAnswer', '<|time_prefix|>', '<|L0_047|>', '<|end_of_event|>', ...

Residual Scalar Quantization Tokenization (Linear Scale, 4 Levels): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|L0_10|>', '<|L1_13|>', '<|L2_21|>', '<|L3_46|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|L0_42|>', '<|L1_42|>', '<|L2_44|>', '<|L3_15|>', '<|end_of_event|>', ...

Residual Scalar Quantization Tokenization (Log Scale, 1 Level): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|L0_238|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|L0_118|>', '<|end_of_event|>', ...

Residual Scalar Quantization Tokenization (Log Scale, 4 Levels): '<|begin_of_event|>', '<|type_prefix|>', 'G', 'uru', '<|time_prefix|>', '<|L0_38|>', '<|L1_35|>', '<|L2_22|>', '<|L3_28|>', '<|end_of_event|>', '<|begin_of_event|>', '<|type_prefix|>', 'Good', 'ĠAnswer', '<|time_prefix|>', '<|L0_41|>', '<|L1_45|>', '<|L2_43|>', '<|L3_57|>', '<|end_of_event|>', ...

C Experimental Settings

In this research, we fine-tune a pretrained LLM backbone Llama-3.2-1B (Llama Team, 2024) for all experiments, employing 4-bit quantization (Dettmers et al., 2023) for efficiency. We use Parameter-Efficient Fine-Tuning (PEFT) via Low-Rank Adaptation (LoRA) (Hu et al., 2022) with a rank $r = 16$ and $\alpha = 32$, applying updates to all attention projection matrices. All models are trained for 5 epochs using the AdamW optimizer (Loshchilov and Hutter, 2019) with a cosine learning rate scheduler, a learning rate of 0.001, and a warmup ratio of 0.1 through the Hugging Face framework (Wolf et al., 2020). We use a per-device train batch size of 4 with 4 gradient accumulation steps, resulting in an effective batch size of 16. All hyper-parameters were determined through preliminary experiments and fixed for the main experiments to avoid exhaustive tuning.

When using tokenization strategies that expand

the LLM’s vocabulary with new temporal tokens, we explicitly set their indices to be trainable within the LoRA configuration to ensure their initialized embeddings are fully optimized alongside the adapters. For the residual scalar quantization (RSQ) tokenizer, we utilize K-Means and default parameter settings for initialization, convergence, and optimization from scikit-learn (Pedregosa et al., 2011). We use TPP-LLM (Liu and Quan, 2024) as our baseline, configured with the same foundation model and QLoRA settings. We follow the original TPP-LLM paper for most hyper-parameters; however, to avoid overfitting, we replace their constant learning rate scheduler with a cosine scheduler and train for 10 epochs, using the same warmup ratio as our experiments. All experiments were conducted on NVIDIA A100 GPUs.

D Ablation Studies

We further investigate the impact of key architectural and design choices through ablation studies on quantization granularity, temporal resolution, prompt template ordering, and the scaling of both the foundation model and the dataset. For these analyses, we selected two representative datasets: Stack Overflow, representing a skewed log-normal distribution, and NYC Taxi, representing a mixed distribution with both log-normal and spiky characteristics.

D.1 Quantization Levels

We analyze the impact of quantization granularity by distributing a fixed budget of 256 newly added tokens across different numbers of RSQ levels, as detailed in Table 3. Consistent with our main results, the choice of base scale (linear vs. log) remains the dominant factor determining performance. However, within the optimal scale for each dataset, we observe a consistent trend where increasing the number of levels improves predictive precision if a suitable scale is selected. For both Stack Overflow (in log scale) and NYC Taxi (in linear scale), the multi-level configurations achieve lower RMSEs compared to the single-level baselines. This demonstrates that a compositional, hierarchical representation, where tokens iteratively refine the residual error, provides superior temporal precision compared to a flat vocabulary of the same size.

Table 3: Ablation study on RSQ levels, comparing strategies that all use a fixed budget of 256 new tokens. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported, with standard deviations provided as subscripts.

Tokenizer	Bins	Stack Overflow		NYC Taxi	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
RSQ (Linear)	256	44.2 \pm .1	0.595 \pm .004	92.0 \pm .0	0.911 \pm .017
RSQ (Linear)	128-128	44.3 \pm .1	0.602 \pm .001	92.0 \pm .0	0.921 \pm .003
RSQ (Linear)	85-85-86	44.4 \pm .1	0.591 \pm .004	92.0 \pm .0	0.904 \pm .004
RSQ (Linear)	64-64-64-64	44.5 \pm .0	0.603 \pm .002	92.0 \pm .0	0.891 \pm .007
RSQ (Log)	256	44.3 \pm .1	0.497 \pm .005	92.0 \pm .0	1.221 \pm .040
RSQ (Log)	128-128	44.3 \pm .1	0.480 \pm .002	92.0 \pm .0	1.191 \pm .041
RSQ (Log)	85-85-86	44.3 \pm .0	0.494 \pm .002	92.0 \pm .0	1.091 \pm .020
RSQ (Log)	64-64-64-64	44.2 \pm .1	0.474 \pm .003	92.0 \pm .0	1.201 \pm .057

D.2 Temporal Resolutions

We investigate the impact of tokenizer granularity by varying the resolution of calendar-based strategies from day down to second, as shown in Table 4. The results highlight a critical dependency between the tokenizer’s temporal resolution and the intrinsic time scale of the dataset. For high-frequency domains like NYC Taxi, where events occur within minutes, a coarse day resolution is insufficient, leading to catastrophically high RMSE scores for the absolute calendar strategy. Increasing the resolution to minute or second drastically corrects this, recovering performance. In contrast, for the Stack Overflow dataset, which operates on a longer time scale, the performance gains from finer granularity are marginal, with the minute resolution offering a slight optimal balance before plateauing. This confirms that while higher resolution generally preserves more information, the optimal choice is dictated by the characteristic granularity of the event sequence itself.

Table 4: Ablation study on temporal resolution. We compare the performance of calendar tokenizers across varying levels of granularity, from day to second. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported, with standard deviations provided as subscripts.

Tokenizer	Resolution	Stack Overflow		NYC Taxi	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
Abs. Calendar	Day	44.7 \pm .1	0.578 \pm .002	91.9 \pm .1	16.086 \pm .009
Abs. Calendar	Hour	44.8 \pm .2	0.568 \pm .003	92.2 \pm .0	1.132 \pm .005
Abs. Calendar	Minute	44.9 \pm .1	0.566 \pm .002	92.2 \pm .0	0.873 \pm .003
Abs. Calendar	Second	44.8 \pm .1	0.569 \pm .003	92.2 \pm .0	0.869 \pm .006
Rel. Calendar	Day	44.4 \pm .1	0.626 \pm .000	91.9 \pm .0	0.912 \pm .000
Rel. Calendar	Hour	44.3 \pm .1	0.624 \pm .001	91.9 \pm .1	0.912 \pm .001
Rel. Calendar	Minute	44.3 \pm .1	0.624 \pm .001	92.0 \pm .0	0.892 \pm .001
Rel. Calendar	Second	44.5 \pm .1	0.623 \pm .001	92.0 \pm .0	0.891 \pm .002

D.3 Template Orders

We explore the sensitivity of model performance to the ordering of information within the prompt template, specifically, whether to place event type tokens before or after event time tokens. As shown in Table 5, our results consistently favor the “Type-Time” ordering (placing the event type before the event time tokens). Notably, on the Stack Overflow dataset, reversing this order to “Time-Type” causes a sharp degradation in event type prediction accuracy for strategies that generate long or semantically opaque sequences, such as numeric string, byte, and relative calendar tokenization. This suggests that intervening complex temporal sequences between event types may dilute the semantic context required for accurate type prediction. Furthermore, in several cases (e.g., RSQ on NYC Taxi), the “Time-Type” order also negatively impacts temporal RMSE. Based on these observations, we standardize on the “Type-Time” template for all main experiments to ensure optimal performance.

Table 5: Ablation study on template order for testing the effect of placing event type tokens or event time tokens first in the sequence. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported, with standard deviations provided as subscripts.

Tokenizer	Template	Stack Overflow		NYC Taxi	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
Numeric String	Type-Time	44.5 \pm .1	0.625 \pm .000	92.0 \pm .0	0.888 \pm .002
Numeric String	Time-Type	33.3 \pm .3.0	0.625 \pm .000	91.8 \pm .0	0.892 \pm .001
Byte	Type-Time	44.3 \pm .0	0.516 \pm .004	92.0 \pm .0	0.892 \pm .001
Byte	Time-Type	39.0 \pm .9	0.540 \pm .004	91.7 \pm .0	0.894 \pm .000
Abs. Calendar (Sec.)	Type-Time	44.8 \pm .1	0.569 \pm .003	92.2 \pm .0	0.869 \pm .006
Abs. Calendar (Sec.)	Time-Type	43.6 \pm .4	3.835 \pm 2.771	91.9 \pm .0	0.874 \pm .006
Rel. Calendar (Sec.)	Type-Time	44.5 \pm .1	0.623 \pm .001	92.0 \pm .0	0.891 \pm .002
Rel. Calendar (Sec.)	Time-Type	39.0 \pm 2.0	0.623 \pm .000	91.8 \pm .1	0.894 \pm .004
Scale Bin (Linear)	Type-Time	44.3 \pm .1	0.615 \pm .000	92.0 \pm .0	0.894 \pm .007
Scale Bin (Linear)	Time-Type	43.8 \pm .1	0.614 \pm .001	91.9 \pm .0	0.906 \pm .007
Scale Bin (Log)	Type-Time	44.3 \pm .0	0.479 \pm .003	92.0 \pm .0	1.127 \pm .061
Scale Bin (Log)	Time-Type	44.1 \pm .1	0.478 \pm .003	91.7 \pm .0	1.312 \pm .016
RSQ (Linear, L4)	Type-Time	44.5 \pm .0	0.603 \pm .002	92.0 \pm .0	0.891 \pm .007
RSQ (Linear, L4)	Time-Type	43.7 \pm .1	0.605 \pm .001	91.9 \pm .0	0.898 \pm .015
RSQ (Log, L4)	Type-Time	44.2 \pm .1	0.474 \pm .003	92.0 \pm .0	1.201 \pm .057
RSQ (Log, L4)	Time-Type	44.1 \pm .1	0.474 \pm .002	91.7 \pm .1	1.479 \pm .104

D.4 Model Sizes

To evaluate the impact of foundation model scale, we compare the performance of our Llama-3.2-1B backbone against the larger Llama-3.2-3B, with results presented in Table 6. The findings show that under our current fine-tuning setting and data size, the 1B model already achieves a highly competitive performance, establishing a strong and efficient baseline. The 3B model performs comparably,

and in some cases, such as with the absolute calendar tokenizer on Stack Overflow, achieves a slight improvement in both accuracy and RMSE. The consistent performance trends observed between these two scales suggest that the fundamental trade-offs are rooted in the tokenization logic itself rather than specific model capacity, providing strong confidence in their generalizability to larger architectures. Thus, while the Llama-3.2-1B model provides an excellent balance of performance and efficiency for this task setup, the framework remains fully compatible with larger models, which may yield further advantages with more extensive training data.

Table 6: Ablation study on foundation model size, comparing Llama-3.2-1B and Llama-3.2-3B backbones. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported, with standard deviations provided as subscripts.

Tokenizer	Model	Stack Overflow		NYC Taxi	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
Numeric String	1B	44.5 \pm .1	0.625 \pm .000	92.0 \pm .0	0.888 \pm .002
Numeric String	3B	43.7 \pm 1.2	0.625 \pm .001	92.1 \pm .0	0.897 \pm .001
Byte	1B	44.3 \pm .0	0.516 \pm .004	92.0 \pm .0	0.892 \pm .001
Byte	3B	43.7 \pm .6	0.552 \pm .047	92.0 \pm .0	0.893 \pm .002
Abs. Calendar (Sec.)	1B	44.8 \pm .1	0.569 \pm .003	92.2 \pm .0	0.869 \pm .006
Abs. Calendar (Sec.)	3B	45.0 \pm .1	0.564 \pm .000	92.2 \pm .1	0.870 \pm .005
Rel. Calendar (Sec.)	1B	44.5 \pm .1	0.623 \pm .001	92.0 \pm .0	0.891 \pm .002
Rel. Calendar (Sec.)	3B	44.3 \pm .1	0.625 \pm .000	92.0 \pm .0	0.895 \pm .006
Scale Bin (Linear)	1B	44.3 \pm .1	0.615 \pm .000	92.0 \pm .0	0.894 \pm .007
Scale Bin (Linear)	3B	44.0 \pm .4	0.614 \pm .001	92.1 \pm .0	0.908 \pm .012
Scale Bin (Log)	1B	44.3 \pm .0	0.479 \pm .003	92.0 \pm .0	1.127 \pm .061
Scale Bin (Log)	3B	44.1 \pm .2	0.481 \pm .002	92.0 \pm .0	1.161 \pm .060
RSQ (Linear, L4)	1B	44.5 \pm .0	0.603 \pm .002	92.0 \pm .0	0.891 \pm .007
RSQ (Linear, L4)	3B	44.4 \pm .1	0.602 \pm .001	92.0 \pm .0	0.906 \pm .004
RSQ (Log, L4)	1B	44.2 \pm .1	0.474 \pm .003	92.0 \pm .0	1.201 \pm .057
RSQ (Log, L4)	3B	44.2 \pm .1	0.473 \pm .001	92.0 \pm .0	1.294 \pm .069

D.5 Data Sizes

To evaluate the impact of dataset scale on performance, we extend the Stack Overflow dataset from 2 years (S) to 2.5 years (M) and 3 years (L), and the NYC Taxi dataset from 7 days (S) to 10 days (M) and 14 days (L), as detailed in Table 7. To ensure a fair comparison, all evaluations are performed on the test set of the smallest data variant (S). The results in Table 8 show mixed outcomes. For NYC Taxi, increasing the data size generally yields slight improvements in accuracy and RMSE across most tokenizers, with the Absolute Calendar (Sec.) achieving the best results on the largest split. However, for Stack Overflow, performance gains are inconsistent or saturate; while some tokeniz-

ers like Scale Bin (Log) improve with more data, others plateau or slightly degrade. This suggests that the capacity of our LoRA-based fine-tuning may be saturating as the dataset complexity grows. While parameter-efficient fine-tuning (PEFT) is effective for smaller datasets, scaling to significantly larger event sequences may require full fine-tuning or continual pre-training (CPT) to fully leverage the additional data, a promising direction for future work.

Table 7: Statistics for the data scale ablation study. We compare small (S), medium (M), and large (L) variants of the Stack Overflow and NYC Taxi datasets, detailing event types, total events, sequence with their train/validation/test splits, average sequence length, and temporal units.

Dataset	Size	Types	Events	Seqs	Train/Val/Test	Seq Len	Unit
Stack Overflow	S	25	187,836	3,336	2,668/334/334	56.31	Month
Stack Overflow	M	27	308,287	5,460	4,368/546/546	56.46	Month
Stack Overflow	L	27	459,269	8,065	6,452/806/807	56.90	Month
NYC Taxi	S	8	362,370	2,957	2,365/296/296	122.55	Hour
NYC Taxi	M	8	444,170	3,641	2,912/364/365	121.99	Hour
NYC Taxi	L	8	680,510	5,579	4,463/558/558	121.98	Hour

Table 8: Ablation study on data scale. We compare the performance of tokenization strategies across three data scale variants (small, medium, and large) on datasets. Next event type accuracy (Acc % \uparrow) and next event time RMSE (\downarrow) are reported, with standard deviations provided as subscripts.

Tokenizer	Data	Stack Overflow		NYC Taxi	
		Acc \uparrow	RMSE \downarrow	Acc \uparrow	RMSE \downarrow
Numeric String	S	44.5 \pm .1	0.625 \pm .000	92.0 \pm .0	0.888 \pm .002
	M	45.0 \pm .1	0.625 \pm .000	<u>92.2</u> \pm .0	0.888 \pm .001
	L	44.4 \pm .5	0.625 \pm .001	<u>92.2</u> \pm .0	0.887 \pm .005
Byte	S	44.3 \pm .0	0.516 \pm .004	92.0 \pm .0	0.892 \pm .001
	M	44.4 \pm .0	0.514 \pm .003	92.1 \pm .0	0.892 \pm .002
	L	44.0 \pm .7	0.556 \pm .052	<u>92.2</u> \pm .0	0.889 \pm .002
Abs. Calendar (Sec.)	S	44.8 \pm .1	0.569 \pm .003	<u>92.2</u> \pm .0	0.869 \pm .006
	M	44.6 \pm .2	1.772 \pm 2.413	92.3 \pm .0	<u>0.860</u> \pm .004
	L	45.0 \pm .1	0.566 \pm .003	92.3 \pm .0	0.859 \pm .003
Rel. Calendar (Sec.)	S	44.5 \pm .1	0.623 \pm .001	92.0 \pm .0	0.891 \pm .002
	M	44.7 \pm .1	0.624 \pm .001	92.1 \pm .0	0.887 \pm .003
	L	44.1 \pm 1.1	0.624 \pm .001	<u>92.2</u> \pm .0	0.885 \pm .003
Scale Bin (Linear)	S	44.3 \pm .1	0.615 \pm .000	92.0 \pm .0	0.894 \pm .007
	M	44.7 \pm .1	0.612 \pm .000	92.1 \pm .1	0.887 \pm .004
	L	44.4 \pm .4	0.598 \pm .000	92.0 \pm .2	0.896 \pm .006
Scale Bin (Log)	S	44.3 \pm .0	0.479 \pm .003	92.0 \pm .0	1.127 \pm .061
	M	44.6 \pm .2	0.481 \pm .001	<u>92.2</u> \pm .0	1.231 \pm .027
	L	<u>44.9</u> \pm .1	<u>0.477</u> \pm .003	<u>92.2</u> \pm .1	1.169 \pm .035
RSQ (Linear, L4)	S	44.5 \pm .0	0.603 \pm .002	92.0 \pm .0	0.891 \pm .007
	M	45.0 \pm .1	0.599 \pm .002	<u>92.2</u> \pm .0	0.897 \pm .004
	L	44.8 \pm .1	0.605 \pm .004	<u>92.2</u> \pm .0	0.910 \pm .008
RSQ (Log, L4)	S	44.2 \pm .1	0.474 \pm .003	92.0 \pm .0	1.201 \pm .057
	M	44.7 \pm .1	0.518 \pm .004	<u>92.2</u> \pm .0	1.293 \pm .054
	L	44.7 \pm .3	0.512 \pm .009	<u>92.2</u> \pm .0	1.158 \pm .015