
End-to-End Compression for Tabular Foundation Models

Anonymous Authors¹

Abstract

The long-standing dominance of gradient-boosted decision trees for tabular data has recently been challenged by in-context learning tabular foundation models. In-context learning methods fit and predict in one forward pass without parameter updates by leveraging the training data as context for predicting on query test points. While recent tabular foundation models achieve state-of-the-art performance, their transformer architecture based on the attention mechanism has quadratic complexity regarding dataset size, which in turn increases the overhead on training and inference time, and limits the capacity of the models to handle large-scale datasets. In this work, we propose TACO, an end-to-end tabular compression model that compresses the training dataset in a latent space. We test our method on the TabArena benchmark, where our proposed method is up to 94x faster in inference time, while consuming up to 97% less memory compared to the state-of-the-art tabular transformer architecture, all while retaining performance without significant degradation. Lastly, our method not only scales better with increased dataset sizes, but it also achieves better performance compared to other baselines.

1. Introduction

Tabular data has ever since been a prevailing data structure in a plethora of domains due to its flexibility towards storing heterogeneous feature types in relational databases (Borisov et al., 2021; Peleska & Sír, 2024; Jiang et al., 2025). In the last decade the overwhelming success of foundation models catalyzed a parallel agenda for tabular learning (Van Breugel & Van Der Schaar, 2024) next to traditional learning approaches like gradient-boosted decision trees (Chen & Guestrin, 2016; Ke et al., 2017; Prokhorenkova

et al., 2018), random forests (Breiman, 2001), or feed-forward models (Kadra et al., 2021; Gorishniy et al., 2021; Holzmüller et al., 2025; Gorishniy et al., 2025). The main essence of tabular foundation models relies in learning across a broad distribution of synthetic and real datasets, which enables a fast adaptation to an unseen dataset at inference time without task-specific retraining (Hollmann et al., 2023; Qu et al., 2025; Hollmann et al., 2025). A compelling paradigm for prediction on tabular data is in-context learning (Hollmann et al., 2023), where at inference time, the model conditions on a dataset’s training split and produces predictions for test points via a single forwards pass, which effectively amortizes the learning procedure itself into the model’s weights.

Recent tabular foundation models have demonstrated that in-context predictors are on par or even surpass strong classical baselines on tabular benchmark datasets. In particular, recent benchmark efforts (Zabërgja et al., 2025; Erickson et al., 2025) have positioned tabular foundation models as state-of-the-art for tabular classification. Prominent representatives of this paradigm are, e.g., TabPFN (Hollmann et al., 2025), TabICL (Qu et al., 2025), and TabDPT (Ma et al., 2024). These models share a common architectural pattern based on decoder-only transformers that attend over a tokenized representation of the training split and predict labels for test points conditioned on that context in a single forward pass. However, the architectural choice also induces a fundamental scaling bottleneck. The inference costs scale quadratically with the size of the conditional context, resulting in $\mathcal{O}(N^2)$ inference-time bottleneck in the number of training instances N . If Key-Value (KV) caching is used, the inference-time bottleneck can be reduced to $\mathcal{O}(N)$. More concretely, when the training split contains N rows with M features, many implementations effectively induce memory and compute that scale with the number of cells within the tabular data, i.e., $N \times M$, which quickly reaches hardware budget restrictions in practical scenarios (Zhong et al., 2024). Common workarounds are to train only on feasible context sizes that force models to include a subsampling of rows/features (Ye et al., 2025a; Eisenschlos et al., 2021) or to restrict the training only to small and mid-size tables (Hollmann et al., 2023; 2025). Consequently, the predictive performance degenerates on large, information-enriched datasets (Qu et al., 2025). An-

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

other drawback stems from the inference latency and cost itself, where tabular foundation models are only attractive and impactful in practical scenarios if the forward pass is tractable at realistic N rows and M features.

To address this central bottleneck, we propose TACO, a novel modular architecture for fast context compression that can be integrated with any decoder-only tabular foundation model to linearly reduce its inference complexity regarding N . Our approach introduces a learned compressor that reduces the training context before prediction, thereby decoupling predictive quality from the raw size of the conditional context table. Under the hood, TACO maps an input dataset into a compact representation, and performs in-context prediction using only the compressed context. Intuitively, our compressor module learns a condensed embedding capturing $K \ll N$ prototypical patterns of the tabular data, and the predictor module produces test-point predictions conditioned on the compressed representation. Therefore, TACO directly tackles the attention scaling problem as the compressor reduces the effective context size by a factor N/K , denoting the compression rate, resulting in a proportionally reduction of memory and compute costs. TACO reframes tabular in-context inference by learning *how to compress*, i.e., by learning data representations that remain predictive under strict inference-time budgets. Notably, our end-to-end pipeline enables a principled accuracy-latency trade-off via the compression rate. This allows tuning resource use under deployment constraints and allows in-context tabular predictors to scale to datasets from real-world domains.

In a thorough evaluation, we demonstrate that it is possible to compress the context down to 4% while incurring no statistically significant deterioration of predictive performance relative to the uncompressed foundation model baseline.

Overall, our contributions can be summarized as follows:

- We introduce TACO, an end-to-end context compression method for tabular in-context learning.
- We demonstrate empirically that TACO offers a $53\times$ speedup during inference and 94% memory savings with no significant performance degradation at only 4% of the original context size.
- Lastly, we derive a simple chunking-and-stitching strategy that enables ingestion of datasets with over one million rows.

2. Related Work

Deep Learning for Tabular Data. Tabular data has traditionally been dominated by tree-based methods (Chen & Guestrin, 2016; Ke et al., 2017; Prokhorenkova et al., 2018), while numerous novel deep learning architectures (Popov

et al., 2020; Arik & Pfister, 2021; Gorishniy et al., 2021; Ye et al., 2025b) have been proposed, they often fall short in achieving better predictive performance (Grinsztajn et al., 2022; Shwartz-Ziv & Armon, 2022; McElfresh et al., 2023). However, recently simple feed-forward networks (Kadra et al., 2021; Holzmüller et al., 2025) and efficient ensembles of feed-forward neural networks (Gorishniy et al., 2025) have managed to outperform traditional methods such as gradient-boosted decision trees, as verified by multiple studies (Erickson et al., 2025; Zabërgja et al., 2025).

Tabular Foundation Models (TFMs). The recent success of foundation models in the natural language (Radford et al., 2019; Brown et al., 2020), and computer vision domain (Siméoni et al., 2025), has shifted the attention of practitioners in developing tabular foundation models to replicate the success in the tabular data domain (Van Breugel & Van Der Schaar, 2024). TabPFN (Hollmann et al., 2023) is among the initial works that use a transformer-based model trained on synthetic data that achieves state-of-the-art prediction results in small to medium-sized datasets. Later versions of the work enhance the initial architecture by alternating column and row-wise attentions (Hollmann et al., 2025). Alternative approaches (Qu et al., 2025), use a more efficient architecture to handle datasets of larger sizes. Recently, TabDPT (Ma et al., 2025) and ConTextTab (Spinaci et al., 2025) have shown that it is possible to build a tabular foundation model relying solely on real data and achieve competitive performance. Lastly, training on a mixture of synthetic and real data achieves state-of-the-art predictive performance (Grinsztajn et al., 2025).

Compression Methods. While transformer-based architectures achieve state-of-the-art results compared to traditional methods, they suffer from a squared complexity with regards to the number of input instances, which results in a higher latency overhead for training and inference. To that end, there exists prior work that aim to improve the efficiency of the current tabular foundation models. MotherNet (Müller et al., 2023) trains a meta-learned transformer hypernetwork that predicts the weights of a simple feed-forward neural network to reduce the inference time, effectively distilling the parameters of the large transformer architecture in a simple per-dataset neural network. Additionally, TabFlex (Zeng et al., 2025) incorporates linear attention to reduce training and inference overhead, and effectively scale to larger datasets. To the best of our knowledge, there is no published methodology for compressing the training context of tabular foundation models in an end-to-end manner.

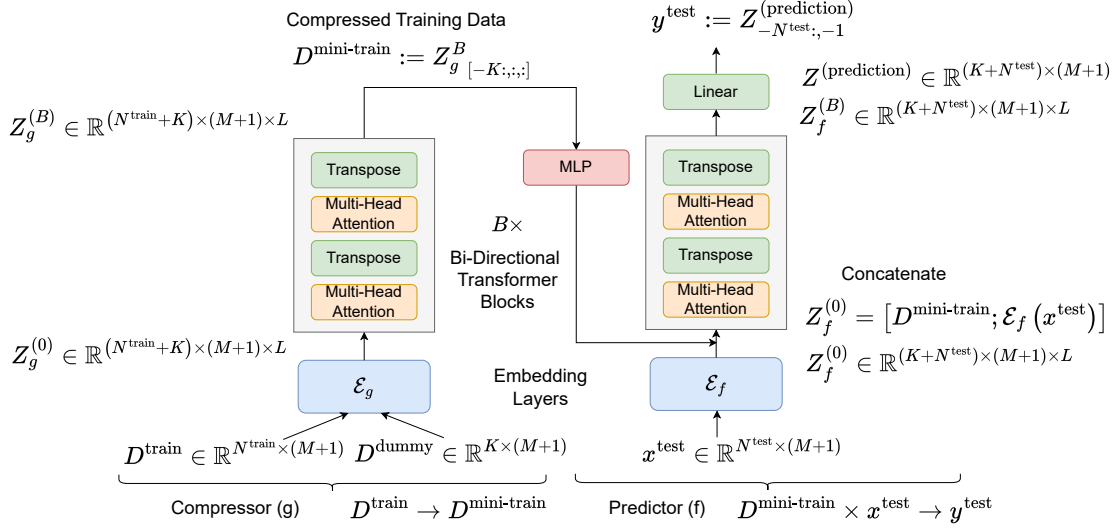


Figure 1. The architecture for training a joint end-to-end compressor and a predictor, with each model having its own parameters.

3. TACO: Tabular Compression

3.1. Tabular Foundation Models

Our problem is tabular prediction, which consists of predicting the target variable $y \in \mathbb{R}^N$ of a feature vector $x \in \mathbb{R}^{N \times M}$, where N is the number of training data points, and M is the number of features. We define a tabular dataset $D := \{(x_i, y_i)\}_{i=1}^N$ to be a matrix $D \in \mathbb{R}^{N \times (M+1)}$ by concatenating the features and the target variables into $(M+1)$ -dimensional vectors. A prior over the datasets, a.k.a. the meta-training collection of datasets, is defined as $(D^{\text{train}}, D^{\text{test}}) \sim p(D)$. A batch of test data points is consequently defined as $(x^{\text{test}}, y^{\text{test}}) \in D^{\text{test}}$. A foundation model with parameters θ is designed as an in-context prediction model $f(x^{\text{test}}, D^{\text{train}}; \theta) : \mathbb{R}^M \times \mathbb{R}^{N^{\text{train}} \times (M+1)} \rightarrow \mathbb{R}$.

We train the foundation model’s parameters through the following log-likelihood minimization over the prior of datasets $p(D)$, by minimizing a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$, such as cross-entropy or least-squares error.

$$\arg \min_{\theta} \mathbb{E}_{(D^{\text{train}}, D^{\text{test}}) \sim p(D)} \mathcal{L}(y^{\text{test}}, f(x^{\text{test}}, D^{\text{train}}; \theta)) \quad (1)$$

$$(x^{\text{test}}, y^{\text{test}}) \sim p_{D^{\text{test}}}$$

The architecture of state-of-the-art tabular foundation models is a bi-directional decoder-only transformer, first proposed by (Lorch et al., 2022) and later popularized by recent models (Hollmann et al., 2023; Qu et al., 2025; Hollmann et al., 2025). First of all, each cell of a table $D \in \mathbb{R}^{N \times (M+1)}$ with N rows, M features, and a target variable, is projected to an L -dimensional space with an encoder layer $\mathcal{E}_f(D) : \mathbb{R}^{N \times (M+1)} \rightarrow \mathbb{R}^{N \times (M+1) \times L}$, yielding a latent representation $Z_f^{(0)}$. Then, a sequence of two multi-head

attention blocks (Vaswani et al., 2017) is applied, in the form of a row-wise attention across N cells for each column, and a column-wise attention across $M+1$ cells for each row (Hollmann et al., 2025). The latent representation of a dataset after each transformer block is denoted by $Z_f^{(b)} \in \mathbb{R}^{N \times (M+1) \times L}$, $b \in \{1, \dots, B\}$.

3.2. End-to-End Compressor and Predictor Networks

The decoder-only transformer architectures suffer from an $\mathcal{O}((N^{\text{train}})^2 \times M + M^2 \times N^{\text{train}}) = \mathcal{O}((N^{\text{train}})^2 \times M)$ inference complexity in the default variant, given that $M \ll N^{\text{train}}$ in most practical scenarios. If KV caching is used, the inference complexity is reduced to $\mathcal{O}(N^{\text{train}} \times M)$.

In this paper, we speed up the inference of tabular foundation models by compressing the training set D^{train} into a $D^{\text{mini-train}}$, where the latter has $K \ll N^{\text{train}}$ rows. Therefore, the inference cost for the attention layers is reduced by a factor of $(\frac{N^{\text{train}}}{K})^2$ in the default transformer variant, and by $(\frac{N^{\text{train}}}{K})$ in the case of KV caching.

Our methodology consists of two end-to-end transformer architectures, a compressor $g : D^{\text{train}} \rightarrow D^{\text{mini-train}}$ with parameters ϕ that computes a lower rank representation of the training data, and a predictor $f : D^{\text{mini-train}} \times x^{\text{test}} \rightarrow y^{\text{test}}$ with parameters θ that takes as an input the compressed training set and the feature matrix of a test batch, to predict the target variable y^{test} of the test features.

We feed the compressor network g with the original training data $D^{\text{train}} \in \mathbb{R}^{N^{\text{train}} \times (M+1)}$, and a dummy table $D^{\text{dummy}} \in$

$\mathbb{R}^{K \times (M+1)}$, where $K \ll N^{\text{train}}$. Such dummy rows are initialized randomly from a subset of D^{train} . Throughout the transformer block, the network is trained to compress D^{train} into the latent representation of the dummy rows, yielding a final representation $D^{\text{mini-train}} \in K \times (M+1) \times L$. To differentiate between the real and dummy cells in the attention and feed-forward layers of the compressor network, we mask the $(M+1)$ -th column of the input dummy matrix, corresponding to the target variable, with a special placeholder value.

The output of the compressor is fed into a predictor network f , which is a standard tabular transformer model (Hollmann et al., 2025). We concatenate the output of the compressor $D^{\text{mini-train}}$ with the embedding of a test batch $\mathcal{E}_f(x^{\text{test}})$ as the initial tensor to the predictor’s attention blocks $Z_f^{(0)}$. The predictor’s transformer network fuses the latent representation of the compressed training set and the embedded test batch to predict the target variable of the test batch. A visual representation of our novel tabular compression architecture is shown in Figure 1.

The compressor g and the predictor f are meta-learned jointly in an end-to-end manner as follows:

$$\arg \min_{\theta, \phi} \mathbb{E}_{\substack{(D^{\text{train}}, D^{\text{test}}) \sim p(D) \\ (x^{\text{test}}, y^{\text{test}}) \sim p_{D^{\text{test}}}}} \mathcal{L}(y^{\text{test}}, f(x^{\text{test}}, g(D^{\text{train}}; \phi)); \theta) \quad (2)$$

As a result of the end-to-end learning, the compressor is incentivized to produce a low-rank training set representation $D^{\text{mini-train}}$ that maximally improves the performance of the predictor. At the same time, the predictor is trained to fuse the L -dimensional cells of the compressed training data and the test batch embedding into a joint representational space.

4. Experimental protocol

This section describes the experimental protocol used to train and evaluate our models. We first detail the training setup, including model architecture, optimization strategy, and the synthetic prior used during pretraining. We then introduce the evaluation setup, covering the datasets, metrics, caching mechanisms, and terminology that will be used throughout the empirical analysis. Based on this setup, we subsequently formulate a set of research hypotheses that structure the results presented in the following section.

4.1. Training Setup

We begin by describing the training configuration shared across all experiments. Unless stated otherwise, all models are trained using the same protocol to ensure stable optimization and fair comparison.

Model Architecture and Hyperparameters. For both the compressor and predictor modules in TACO, we use the

TabPFNv2 (Hollmann et al., 2025) architecture with 2D attention over rows and columns. Each module comprises 12 attention layers, 6 attention heads, and an embedding dimension of 192, for roughly 7M parameters per module. We additionally insert a two-layer residual MLP between the compressor and predictor.

Synthetic Prior. Following prior work (Hollmann et al., 2023; Qu et al., 2025), we train our models exclusively on synthetically generated datasets produced by structural causal models (SCMs), including variants where the causal mechanisms are tree-based. This synthetic prior is designed to expose the model to a broad range of feature types, dependency structures, and label-generating mechanisms that are representative of real-world tabular data. We adopt the synthetic prior released with TabICL (Qu et al., 2025) and use its open-source implementation throughout.

Optimization and Training Length. Following prior work, during training, we fix the sequence length to 1024 and sample the number of features uniformly at random between 2 and 100. We train for 80K optimization steps with a global batch size of 1024, yielding approximately 82M synthetic datasets. To alleviate GPU memory constraints, we use gradient accumulation with a micro-batch size of 16. We optimize using AdamW and train with mixed-precision arithmetic. We warm up the learning rate to 1×10^{-4} and subsequently apply a cosine annealing schedule (Loshchilov & Hutter, 2017). We additionally employ gradient clipping with a maximum norm of 1.0 and use weight decay of 1×10^{-2} . For each of our models, we employ distributed training across 8 NVIDIA H100 GPUs with 95 GB of VRAM each. Pretraining took approximately 15 days to complete. For reproducibility, we open-source our code at: <https://anonymous.4open.science/r/taco-AAA1/README.md>.

4.2. Evaluation Setup

In this subsection, we describe the evaluation protocol used to assess the scalability and predictive performance of TFMs. This setup defines the datasets, metrics, and conventions referenced throughout the results section.

Evaluation Datasets and Metrics. We evaluate all models and baselines on the classification subset of TabArena (Erickson et al., 2025). For most experiments, we restrict this evaluation to the same set of 26 datasets used in the TabPFN v2.0 evaluation (Hollmann et al., 2023). Among these datasets, 20 are binary classification tasks, and 6 are multiclass. For the experiments corresponding to Hypothesis 6, we use all 36 classification datasets with fewer than 10 classes in TabArena, of which 8 are multiclass.

We follow the same data preprocessing pipeline as TabPFN v2.0 (Hollmann et al., 2023) and use the predefined cross-

validation folds from TabArena (Erickson et al., 2025). Model performance is evaluated using ROC-AUC for binary classification tasks and ROC-AUC (one-vs-one) for multi-class tasks (Junge & Dettori, 2018). For simplicity, we refer to both metrics as ROC-AUC throughout this section.

Terminology. Throughout this section, we use the term *compression rate* r to denote the ratio between the number of rows K in the compressed training representation produced by the compressor and the number of rows N^{train} in the original training set, i.e., $r = K/N^{\text{train}}$. We refer to *POT* as the *predictor-only transformer*, which consists solely of the predictor architecture used in TACO and is trained under the same data, optimization, and hyperparameter setup, but without any compression component; the predictor therefore operates directly on the full, uncompressed training data. *TACO* denotes the full model that combines a learned compressor with a predictor transformer of the same architecture as POT, trained jointly in an end-to-end manner. We use *TACO* ($r\%$) to indicate a configuration in which the compressor produces a compressed training representation with a compression rate of $r\%$.

4.3. Research Hypotheses and Experiments

Hypothesis 1: TACO achieves up to $53\times$ speedup during inference, without significant degradation in performance. Initially, we construct a grid of synthetic datasets with $N \in [5000, 50000]$ rows and $M \in [100, 1000]$ features. For each grid point, we evaluate both TACO with a compression rates of $r = 4\%$ and the plain counterpart predictor-only transformer (POT). We evaluate both methods with and without KV caching, and measure fit and prediction time. Where, fit time represents the preprocessing time.

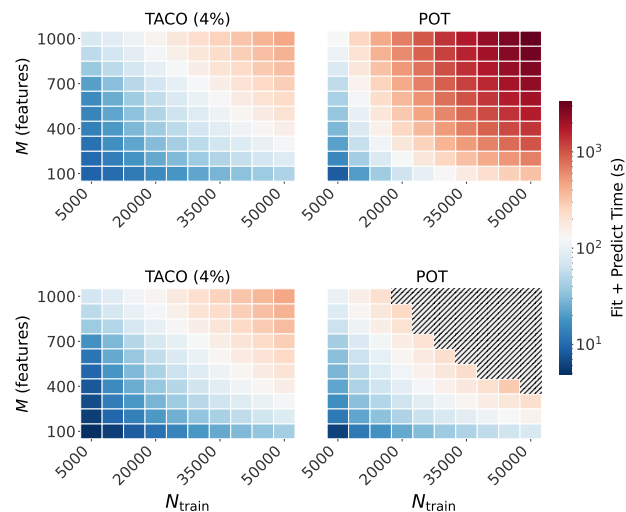


Figure 2. Fit + Predict time heatmap. **Top:** Predict times for TACO and POT, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.

We consider three inference scenarios: (i) providing the full test set to the model in a single call, (ii) splitting the test set into 10 batches, and (iii) splitting the test set into 100 batches. For consistency, we fix the test set to 1,000 rows. Results for scenario (ii) are shown in Figure 2, while the remaining scenarios are reported in Appendix A.1.

In all experiments, KV caching applies exclusively to the predictor transformer module and is used identically for both TACO and the predictor-only transformer (POT). KV caching requires materializing and storing the key-value tensors of all attention layers, which introduces additional memory overhead compared to standard inference, where attention maps can be recomputed on the fly. In contrast, caching the output of the compressor in TACO incurs no additional memory cost, since the compressed context must be materialized regardless in order to be consumed by the predictor. Consequently, for TACO, KV caching refers solely to caching the predictor’s attention states over the compressed context, while the compressor itself is executed only once and never recomputed. This distinction explains why KV caching is often infeasible for POT under tight memory budgets, yet remains practical when combined with TACO.

Figure 2 (Top) presents the heatmap without KV caching. Across all grid points, TACO is consistently faster, as indicated by the blue and light orange color throughout the heatmap. Likewise, Figure 2 (Bottom) shows that TACO remains the faster method under KV caching, with performance advantages that are consistent across the entire grid. In contrast, POT with KV caching runs out of memory on 32 out of 100 datasets in the grid. These results also suggest that TACO can serve as an enabler of KV caching by making cached inference feasible under tighter memory budgets.

To further highlight the speedup of our approach, we simulate an “infinite test-batches” setting, which is common in industrial deployments (e.g., real-time fraud detection, ad click-through prediction, predictive maintenance) where predictions must be produced continuously for a stream of incoming data. We generate a synthetic dataset with 15,000 rows and 500 features, and evaluate inference over 100 sequential test batches of 50 rows each.

Figure 3 reports the cumulative prediction time (y-axis) as a function of the number of processed batches (x-axis). By compressing the context once and reusing it across batches, TACO delivers substantial gains after the first batch, which incurs a one-time cost to generate the compressed context (or to perform the initial fit in the KV caching setting), achieving speedups of up to $\sim 53\times$ at a 4% compression rate for the no KV caching setting, and up to $\sim 33\times$ for the KV caching setting. The gap in time between TACO and POT increases with dataset size.

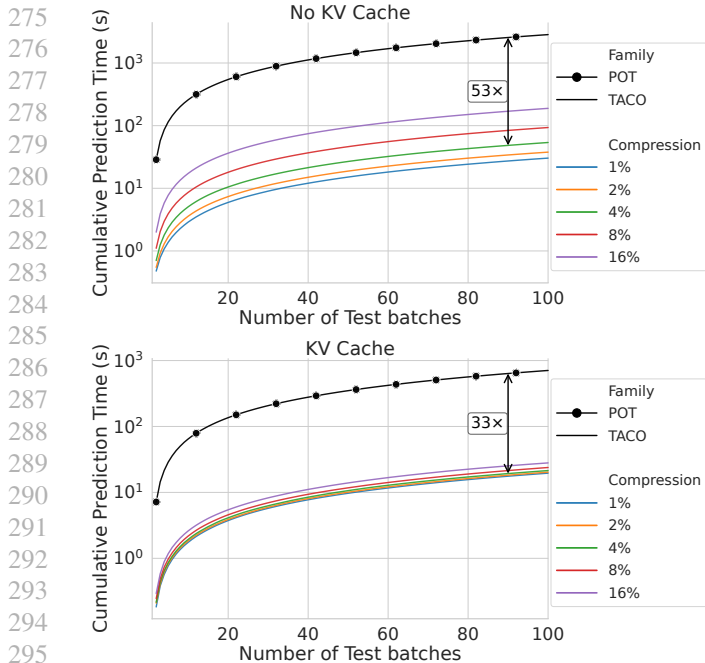


Figure 3. Cumulative prediction time for TACO at different compression rates r compared to POT for 100 test batches. **Top:** Without using KV Caching, **Bottom:** Using KV Caching.




Figure 4. Critical difference (CD) diagram comparing predictor-only transformer (POT) and TACO. Average ranks are computed using Friedman–Nemenyi tests (Demšar, 2006) via autorank (Herbold, 2020b) ($\alpha = 0.05$). Bars connect methods without significant differences.

Next, it is important to validate that the speedup gains and memory consumption reduction that TACO obtains do not come at the cost of significant degraded predictive performance. To assess this trade-off, we compare TACO against POT on the TabArena benchmark (Erickson et al., 2025).

In Figure 4, we provide critical difference (CD) diagrams to investigate the method ranks and the statistical significance of the results. To build the CD diagrams, we use the autorank (Herbold, 2020a) package that runs a Friedman test with a Nemenyi post-hoc test, and a 0.05 significance level. The results indicate that **up to 4% compression rate, TACO provides a 53 \times speedup in inference time with no significant degradation in performance.** Lastly, in Table 1

Table 1. Mean and standard deviation of TACO and POT over the TabArena datasets compared to recent state-of-the-art methods.

Model	Mean ROC AUC (\uparrow)
TabICL	0.866 ± 0.103
TabPFN V2.0	0.866 ± 0.103
POT	0.859 ± 0.102
TACO (r=1%)	0.854 ± 0.099
TACO (r=2%)	0.854 ± 0.099
TACO (r=4%)	0.855 ± 0.099
TACO (r=8%)	0.855 ± 0.099
TACO (r=16%)	0.855 ± 0.099

Note: Highlighted methods use open-weight checkpoints.

we provide the mean performance and standard deviation of TACO with various compression rates, and POT against open-source checkpoints of state-of-the-art architectures, where, as observed, although POT and TACO have not been trained for as long as competitor methods, they achieve comparable performance. Further reports of the performance distribution across datasets can be found in Appendix C.

Hypothesis 2: TACO consumes 94% less memory compared to the predictor-only transformer during prediction. The compressor component of TACO inherits the same asymptotic complexity as the predictor-only transformer architecture; however, this is a one-time cost incurred during compression. While the predictor component of TACO also scales linearly as the predictor-only transformer, it scales with respect to the retained context size K rather than the full dataset size N . Since $K \ll N$, its memory complexity is $\mathcal{O}(K \times M)$, yielding substantially lower memory usage in practice. We measure memory consumption for the same streaming-batch experiment shown in Figure 3 and report the results in Table 2 for the no KV cache setting.

The results show that using a K -sized compressed context yields substantial memory savings during prediction. Specifically, TACO achieves a 78.2% memory reduction for TACO (16%) and up to 97.6% for TACO (1%) in the no KV caching setting. While our fit-time memory footprint is $\sim 2.5\times$ higher than that of the predictor-only transformer in the no-KV-cache setting, **this is typically not a practical concern because fitting without KV caching is comparatively memory-friendly compared to the predict phase.** In contrast, fit-time memory becomes a bottleneck when KV caching is enabled, since the model must store key and value tensors for every token, across all attention layers and heads. In this formulation, where each table cell corresponds to a token, this storage requirement can scale rapidly and become prohibitive. **In the KV caching setting, our model consumes $\sim 3\times$ less memory compared to the**

Table 2. Inference time and memory consumption for the predictor-only transformer (POT), versus TACO with different compression rates, without KV caching.

Method	Fit Time	First Predict Time	Subsequent Predict Time	Fit Mem	Predict Mem	
					Mean	Median
POT	8.41s	28.55s	28.67s±0.05s	36MB	22.45GB	22.45GB
TACO 1%	8.22s	29.68s	306ms±18ms (×93.6)	92MB	776MB (−96.6%)	549MB (−97.6%)
TACO 2%	8.30s	29.80s	382ms±18ms (×75.2)	92MB	1.05GB (−95.3%)	845MB (−96.3%)
TACO 4%	9.86s	30.71s	544ms±17ms (×52.7)	92MB	1.63GB (−92.8%)	1.41GB (−93.7%)
TACO 8%	10.02s	32.47s	943ms±17ms (×30.4)	92MB	2.78GB (−87.6%)	2.56GB (−88.6%)
TACO 16%	9.20s	35.36s	1.91s±0.01s (×15)	92MB	5.10GB (−77.3%)	4.89GB (−78.2%)

Parentheses: × speed-up, −% memory reduction (vs. POT).

predictor-only transformer. For the KV cache results we refer the reader to Table 4 in Appendix A.1.

Hypothesis 3: Joint training of the compressor and predictor is necessary. To test whether the joint training of compressor and predictor is necessary, we train a variant of TACO in which the predictor is initialized with the TabPFN v2.0 (Hollmann et al., 2025) weights and kept frozen throughout meta-learning, while only the compressor parameters are optimized. We compare this variant to the standard TACO model under identical experimental conditions across a range of compression rates.




Figure 5. Distribution of test performances on the TabArena classification tasks, for the: **Left:** joint training of the compressor and predictor, **Right:** training only compressor, while keeping the predictor frozen.

The results in Figure 5 verify that **jointly optimizing the predictor and compressor consistently leads to better performance compared to optimizing the compressor alone with a frozen predictor.** This highlights the importance of joint training: folding compressed representations into a predictor’s representation space allows the predictor itself to adapt, rather than requiring the compressor to align with a fixed predictor.

Hypothesis 4: Meta-learning over various compression rates does not incur a performance loss compared to meta-learning at a fixed compression rate. Multi-rate compression rate training enables a single TACO model to

operate across multiple compression rates by sampling the compression rate uniformly from $\{1, 2, 4, 8, 16\}$ for each synthetic dataset during training. To assess whether this flexibility leads to performance degradation, we compare this *multi-rate model* to *rate-specific* TACO models, each trained using a single fixed compression rate r .

All models use the same architecture, training procedure, prior, and number of training datasets, and are evaluated at their corresponding compression rates.

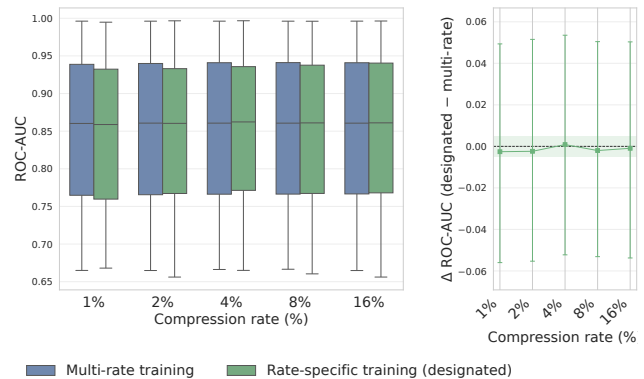


Figure 6. Comparison of dynamic compression rate training for the compressor, versus specific compression rate training. **Left:** Distribution of ROC-AUC test performances over all TabArena classification datasets, **Right:** ROC-AUC differences between specific and dynamic compression rate training; green-hued lines denote effects with **95% confidence intervals.**

As shown in Figure 6, performance differences between the multi-rate and rate-specific models are centered near zero across all compression rates, with no statistically significant advantage for rate-specific training at the 95% confidence level. This indicates that multi-rate compression rate training does not result in a measurable performance loss.

Hypothesis 5: TACO outperforms additional compression baselines. We further compare TACO to two compression approaches on the TabArena benchmark: **kNN-based sampling**, where we form a training context subset by selecting k -nearest neighbors for each test point, increasing k until the union reaches the target size and uniformly




Figure 7. Comparison of TACO and the two compression baselines (kNN and Random). **Left:** ROC-AUC distributions across different compression rates. **Right:** Mean performance difference Δ ROC-AUC (TACO – baseline) across compression rates with **95% bootstrap confidence intervals**. Positive values indicate consistent performance improvements of TACO over the corresponding baseline.

subsampling if needed to ensure the correct compression rate and roughly equal test-point contribution; and **random sampling**, where we subsample uniformly from the training set until we reach the compression rate.

We evaluate these compression baselines with the predictor-only transformer, comparing their performance across a range of compression rates. As shown in Figure 7, **TACO consistently and significantly outperforms both baseline methods across different compression rates ranging from 1% to 16%**. As the compression rate increases, the performance gap gradually narrows, since using more rows from the training set brings all methods closer to using the full training set. Moreover, TACO introduces no additional runtime and only a marginal memory overhead compared to the baselines, which is negligible relative to its memory savings over POT. Additional details on runtime and memory costs are provided in the Appendix A.2.

Hypothesis 6: TACO enables efficient chunking of large tabular datasets while preserving strong learned embeddings under extreme compression. Current tabular foundation models often struggle to scale to large training sets because the memory cost of attention grows quickly with the number of rows. A key advantage of our end-to-end compression pipeline is that it does not require loading the full dataset into the compressor at once. Instead, we use an efficient *chunking-and-stitching* procedure that enables scalable compression and inference.

For example, given a dataset with $N = 10^6$ rows, we partition it into N_C disjoint chunks of size $C = 10^4$, yielding $N_C = N/C = 100$ chunks. We then process chunks sequentially, by applying the compressor with a certain compression rate e.g. 1%, producing a compressed summary with K_C rows (e.g., 100 rows per chunk). Finally, we *stitch* the compressed outputs by concatenating the per-chunk summaries into a single global context, which is provided as context to the predictor at test time.

Using this approach, we successfully perform prediction

on the MetroPT-3 (Davari et al., 2021; Veloso et al., 2022) dataset, which contains approximately 1.5M rows and 15 features. For this dataset we use a simple 80%/20% time-based train/test split and a compression rate of 0.1%, compressing the training set from ~ 1.2 M rows to only 1214 retained compressed rows. We compare against both the random and kNN baselines, and report results in Table 3.

Table 3. Results on MetroPT-3

Method	Accuracy	LogLoss	ROC-AUC	AUPRC
TACO (chunk)	0.9993	0.0162	0.9702	0.8894
Random	0.9981	0.0440	0.9652	0.7956
kNN	0.9809	0.0706	0.9388	0.2316

The results indicate that TACO consistently outperforms compression-related baselines across all the evaluation metrics, validating the chunk and stitch approach.

Moreover, to validate the generality of chunking, we run our method on all 36 TabArena classification datasets with 10 classes or fewer, including regimes where the predictor-only transformer is unable to operate due to the requirement of fitting the full dataset in memory. For more details, we refer the reader to Appendix B.

5. Conclusion

Tabular Foundation Models have emerged as state-of-the-art for tabular classification. However, their inference cost scales with the conditional context. In this work, we tackle the quadratic complexity of transformer attention by introducing TACO, a novel modular architecture for fast context compression. TACO integrates a compressor module that learns a mapping from the input domain to a compact representation. A predictor module, then, is fed with the compressed data to produce test-point predictions. Our evaluation highlights that our model enables a 53x speedup and 94% memory consumption reduction without degrading performance significantly compared to the predictor-only transformer.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Arik, S. Ö. and Pfister, T. Tabnet: Attentive interpretable tabular learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 6679–6687, 2021.
- Borisov, V., Leemann, T., Sessler, K., Haug, J., Pawelczyk, M., and Kasneci, G. Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35:7499–7519, 2021. doi: 10.1109/tnnls.2022.3229161.
- Breiman, L. Random forests. *Mach. Learn.*, 45(1):5–32, October 2001. ISSN 0885-6125. doi: 10.1023/A:1010933404324.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners, 2020.
- Chen, T. and Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pp. 785–794, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939785.
- Davari, N., Veloso, B., Ribeiro, R. P., Pereira, P. M., and Gama, J. Predictive maintenance based on anomaly detection using deep learning for air production unit in the railway industry. In *2021 IEEE 8th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 1–10. IEEE, 2021.
- Demšar, J. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7(1):1–30, 2006.
- Eisenschlos, J. M., Gor, M., Müller, T., and Cohen, W. W. Mate: Multi-view attention for table transformer efficiency. pp. 7606–7619, 2021. doi: 10.18653/v1/2021.emnlp-main.600.
- Erickson, N., Purucker, L., Tschalzev, A., Holzmüller, D., Desai, P. M., Salinas, D., and Hutter, F. Tabarena: A living benchmark for machine learning on tabular data. In

The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track, 2025.

- Gorishniy, Y., Rubachev, I., Khrulkov, V., and Babenko, A. Revisiting deep learning models for tabular data. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 18932–18943. Curran Associates, Inc., 2021.
- Gorishniy, Y., Kotelnikov, A., and Babenko, A. Tabm: Advancing tabular deep learning with parameter-efficient ensembling. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Grinsztajn, L., Oyallon, E., and Varoquaux, G. Why do tree-based models still outperform deep learning on typical tabular data? In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 507–520. Curran Associates, Inc., 2022.
- Grinsztajn, L., Flöge, K., Key, O., Birkel, F., Jund, P., Roof, B., Jäger, B., Safaric, D., Alessi, S., Hayler, A., et al. TabPFN-2.5: Advancing the state of the art in tabular foundation models. *arXiv preprint arXiv:2511.08667*, 2025.
- Herbold, S. Autorank: A python package for automated ranking of classifiers. *Journal of Open Source Software*, 5(48):2173, 2020a. doi: 10.21105/joss.02173.
- Herbold, S. Autorank: A python package for automated ranking of classifiers. *Journal of Open Source Software*, 5(48):2173, 2020b. doi: 10.21105/joss.02173.
- Hollmann, N., Müller, S., Eggensperger, K., and Hutter, F. TabPFN: A transformer that solves small tabular classification problems in a second. In *International Conference on Learning Representations 2023*, 2023.
- Hollmann, N., Müller, S., Purucker, L., Krishnakumar, A., Körfer, M., Hoo, S. B., Schirmermeister, R. T., and Hutter, F. Accurate predictions on small data with a tabular foundation model. *Nature*, 01 2025. doi: 10.1038/s41586-024-08328-6.
- Holzmüller, D., Grinsztajn, L., and Steinwart, I. RealMLP: Advancing MLPs and default parameters for tabular data. In *ELLIS workshop on Representation Learning and Generative Models for Structured Data*, 2025.
- Jiang, J.-P., Liu, S.-Y., Cai, H.-R., Zhou, Q.-L., and Ye, H.-J. Representation learning for tabular data: A comprehensive survey. *ArXiv*, abs/2504.16109, 2025. doi: 10.48550/arxiv.2504.16109.

- 495 Junge, M. R. J. and Dettori, J. R. Roc solid: Receiver
496 operator characteristic (roc) curves as a foundation for
497 better diagnostic tests. *Global Spine Journal*, 8(4):424–
498 429, 2018. doi: 10.1177/2192568218778294. PMID:
499 29977728.
- 500 Kadra, A., Lindauer, M., Hutter, F., and Grabocka, J. Well-
501 tuned simple nets excel on tabular datasets. In Ranzato,
502 M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan,
503 J. W. (eds.), *Advances in Neural Information Process-*
504 *ing Systems*, volume 34, pp. 23928–23941. Curran Asso-
505 ciates, Inc., 2021.
- 507 Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma,
508 W., Ye, Q., and Liu, T.-Y. Lightgbm: A highly efficient
509 gradient boosting decision tree. *Advances in neural infor-*
510 *mation processing systems*, 30, 2017.
- 511 Lorch, L., Sussex, S., Rothfuss, J., Krause, A., and
512 Schölkopf, B. Amortized inference for causal structure
513 learning. In Oh, A. H., Agarwal, A., Belgrave, D., and
514 Cho, K. (eds.), *Advances in Neural Information Process-*
515 *ing Systems*, 2022.
- 517 Loshchilov, I. and Hutter, F. SGDR: Stochastic gradient
518 descent with warm restarts. In *International Conference*
519 *on Learning Representations*, 2017.
- 520 Ma, J., Thomas, V., Hosseinzadeh, R., Kamkari, H., Labach,
521 A., Cresswell, J. C., Golestan, K., Yu, G., Caterini, A. L.,
522 and Volkovs, M. Tabdpt: Scaling tabular foundation
523 models on real data. *arXiv preprint arXiv:2410.18164*,
524 2024.
- 526 Ma, J., Thomas, V., Hosseinzadeh, R., Labach, A., Cress-
527 well, J. C., Golestan, K., Yu, G., Caterini, A. L., and
528 Volkovs, M. TabDPT: Scaling tabular foundation models
529 on real data. In *The Thirty-ninth Annual Conference on*
530 *Neural Information Processing Systems*, 2025.
- 531 McElfresh, D., Khandagale, S., Valverde, J., Prasad C, V.,
532 Ramakrishnan, G., Goldblum, M., and White, C. When
533 do neural nets outperform boosted trees on tabular data?
534 *Advances in Neural Information Processing Systems*, 36:
535 76336–76369, 2023.
- 537 Müller, A., Curino, C., and Ramakrishnan, R. Mothernet:
538 A foundational hypernetwork for tabular classification.
539 *ArXiv*, abs/2312.08598, 2023. doi: 10.48550/arXiv.2312.
540 08598.
- 541 Peleska, J. and Sír, G. Tabular transformers meet relational
542 databases. *ACM Transactions on Intelligent Systems and*
543 *Technology*, 16:1 – 24, 2024. doi: 10.1145/3749991.
- 545 Popov, S., Morozov, S., and Babenko, A. Neural oblivious
546 decision ensembles for deep learning on tabular data. In
547 *International Conference on Learning Representations*,
548 2020.
- 549 Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V.,
and Gulin, A. Catboost: unbiased boosting with categor-
ical features. In *Proceedings of the 32nd International*
Conference on Neural Information Processing Systems,
NIPS’18, pp. 6639–6649, Red Hook, NY, USA, 2018.
Curran Associates Inc.
- Qu, J., Holzmüller, D., Varoquaux, G., and Morvan, M. L.
Tabicl: A tabular foundation model for in-context learn-
ing on large data. *arXiv preprint arXiv:2502.05564*, 2025.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and
Sutskever, I. Language models are unsupervised multitask
learners. 2019.
- Shwartz-Ziv, R. and Armon, A. Tabular data: Deep learning
is not all you need. *Information Fusion*, 81:84–90, 2022.
ISSN 1566-2535. doi: [https://doi.org/10.1016/j.inffus.](https://doi.org/10.1016/j.inffus.2021.11.011)
2021.11.011.
- Siméoni, O., Vo, H. V., Seitzer, M., Baldassarre, F., Oquab,
M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Rama-
monjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang,
J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C.,
Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J.,
Jégou, H., Labatut, P., and Bojanowski, P. Dinov3, 2025.
- Spinaci, M., Polewczyk, M., Schambach, M., and Thelin,
S. Contexttab: A semantics-aware tabular in-context
learner. In *The Thirty-ninth Annual Conference on Neural*
Information Processing Systems, 2025.
- Van Breugel, B. and Van Der Schaar, M. Position: Why
tabular foundation models should be a research priority.
In *International Conference on Machine Learning*, pp.
48976–48993. PMLR, 2024.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones,
L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. At-
tention is all you need. In Guyon, I., Luxburg, U. V.,
Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S.,
and Garnett, R. (eds.), *Advances in Neural Information*
Processing Systems, volume 30. Curran Associates, Inc.,
2017.
- Veloso, B., Ribeiro, R. P., Gama, J., and Pereira, P. M. The
metropt dataset for predictive maintenance. *Scientific*
Data, 9(1):764, 2022.
- Ye, H.-J., Liu, S.-Y., and Chao, W.-L. A closer look at
tabpfn v2: Strength, limitation, and extension. *arXiv*
preprint arXiv:2502.17361, 2025a.
- Ye, H.-J., Yin, H.-H., Zhan, D.-C., and Chao, W.-L. Re-
visiting nearest neighbor for tabular data: A deep tabular
baseline two decades later. In Yue, Y., Garg, A., Peng,
N., Sha, F., and Yu, R. (eds.), *International Conference*
on Representation Learning, volume 2025, pp. 94917–
94936, 2025b.

550 Zabërgja, G., Kadra, A., Frey, C. M. M., and Grabocka, J.
551 Tabular data: Is deep learning all you need?, 2025.
552
553 Zeng, Y., Dinh, T., Kang, W., and Mueller, A. C. Tabflex:
554 Scaling tabular learning to millions with linear attention.
555 In *Forty-second International Conference on Machine*
556 *Learning*, 2025.
557
558 Zhong, H., Li, T., Chen, Y., Tang, W., Wu, J., Yang, H., and
559 Li, X. Nand-tree: A 3d nand flash based processing in
560 memory accelerator for tree-based models on large-scale
561 tabular data. *2024 ACM/IEEE International Conference*
562 *On Computer Aided Design (ICCAD)*, pp. 01–09, 2024.
563 doi: 10.1145/3676536.3676666.
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604

A. Predict Time and Memory Consumption

A.1. Time and Memory Consumption of TACO vs POT

In Hypothesis 1, Figure 2 reports the combined fit-and-predict runtime on the synthetic grid when the test set is split into 10 batches. Here, we additionally report results for processing the full test set as a single batch in Figure 8, and for splitting it into 100 batches in Figure 9.

In Figure 8, the upper plot shows the setting without KV caching, and the lower plot shows the setting using KV caching. In both cases, the overall runtime is comparable. However, POT runs out of memory on 32 out of 100 datasets when KV Caching is used, reflecting the substantial memory required to fit the largest problem instances when the test set is processed as a single batch.

In contrast, Figure 9 reveals a marked shift. Without KV caching (upper plot), TACO dominates across the grid, as indicated by the blue color at essentially all grid points. POT, on the other hand, becomes prohibitively slow on the largest datasets, requiring up to ~ 9 hours to produce predictions. A similar pattern holds with KV caching enabled (lower plot): TACO remains consistently faster at fit + predict time and is more memory efficient, whereas POT again runs out of memory on 32 out of 100 datasets.

Moreover, we observe that even if POT was not constrained by memory, its inference time increases much more rapidly than that of TACO as the dataset size grows. This behavior is already evident at the 45K-row / 400-feature setting, where POT reaches the darkest red region, indicating prohibitively long runtimes. Additionally, we report in Figure 10, Figure 11, Figure 12, and Figure 13 separate heatmaps for fit and predict times for each test scenario.

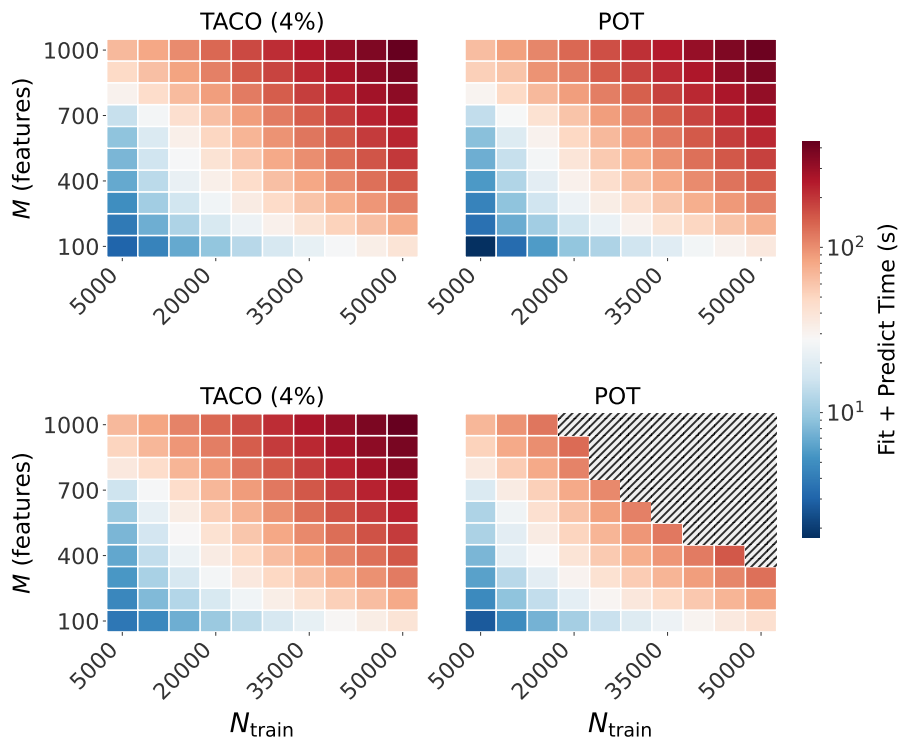


Figure 8. Fit+Predict Time Heatmap for the full test batch scenario. **Top:** Fit+Predict times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.

Lastly, we show empirically, in Table 4 that on a medium-scale dataset with 15,000 rows and 500 features, the predictor-only transformer requires 30.45 GB of GPU VRAM with KV caching. Compressing the context makes KV caching far more tractable: TACO (16%) reduces VRAM usage to 9.94 GB, and TACO (1%) further reduces it to 8.69 GB.

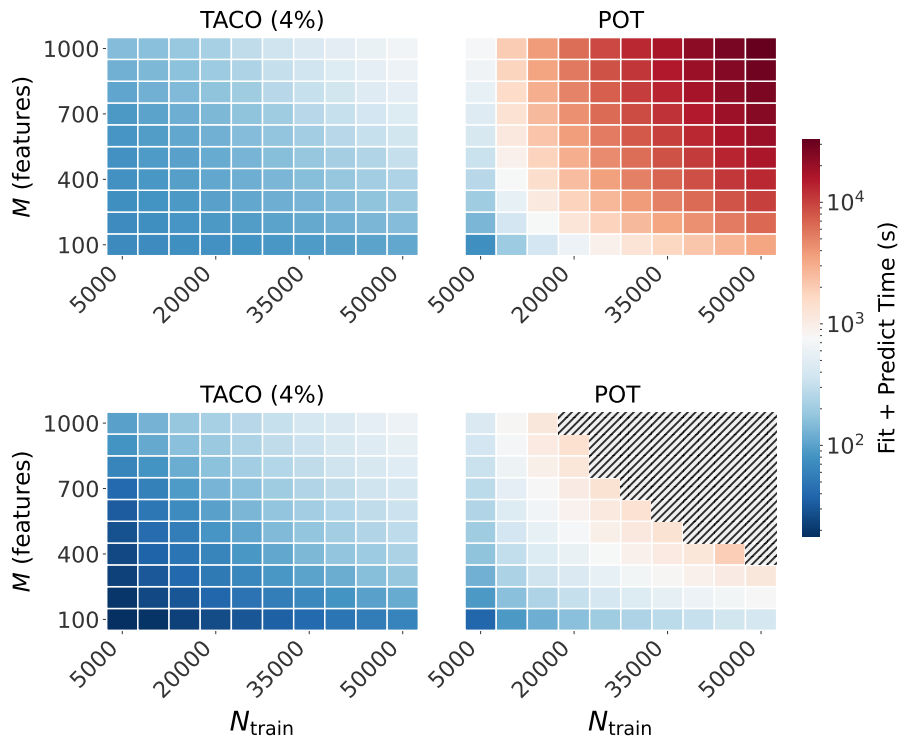


Figure 9. Fit+Predict Time Heatmap for the 100 test batches scenario. **Top:** Fit+Predict times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.

Table 4. Inference time and memory consumption for the predictor-only transformer (POT), versus TACO with different compression rates, with KV-caching.

Method	Fit Time	First Predict Time	Subsequent Predict Time	Fit Mem	Predict Mem	
			Mean \pm Std		Mean	Median
POT	40.60s	7.30s	7.14s \pm 0.08s	30.45GB	8.24GB	8.24GB
TACO 1%	39.00s	279ms	198ms \pm 4ms (\times 36.1)	8.69GB	678MB (-92%)	678MB (-92%)
TACO 2%	39.59s	305ms	205ms \pm 4ms (\times 34.8)	8.77GB	975MB (-88.4%)	976MB (-88.4%)
TACO 4%	39.84s	384ms	217ms \pm 6ms (\times 33)	8.94GB	1.53GB (-81.4%)	1.53GB (-81.4%)
TACO 8%	41.65s	487ms	241ms \pm 5ms (\times 29.6)	9.27GB	2.69GB (-67.3%)	2.69GB (-67.3%)
TACO 16%	46.09s	731ms	283ms \pm 7ms (\times 25.2)	9.94GB	5.02GB (-39.1%)	5.02GB (-39.1%)

Parentheses: \times speed-up, $-\%$ memory reduction (vs. POT).



Figure 10. Fit Time Heatmap. **Top:** Fit times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Fit times with KV caching. Black shaded regions indicate out-of-memory errors.



Figure 11. Predict time heatmap in the full test batch scenario. **Top:** Predict times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.



Figure 12. Predict time heatmap in the 10 test batches scenario. **Top:** Predict times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.

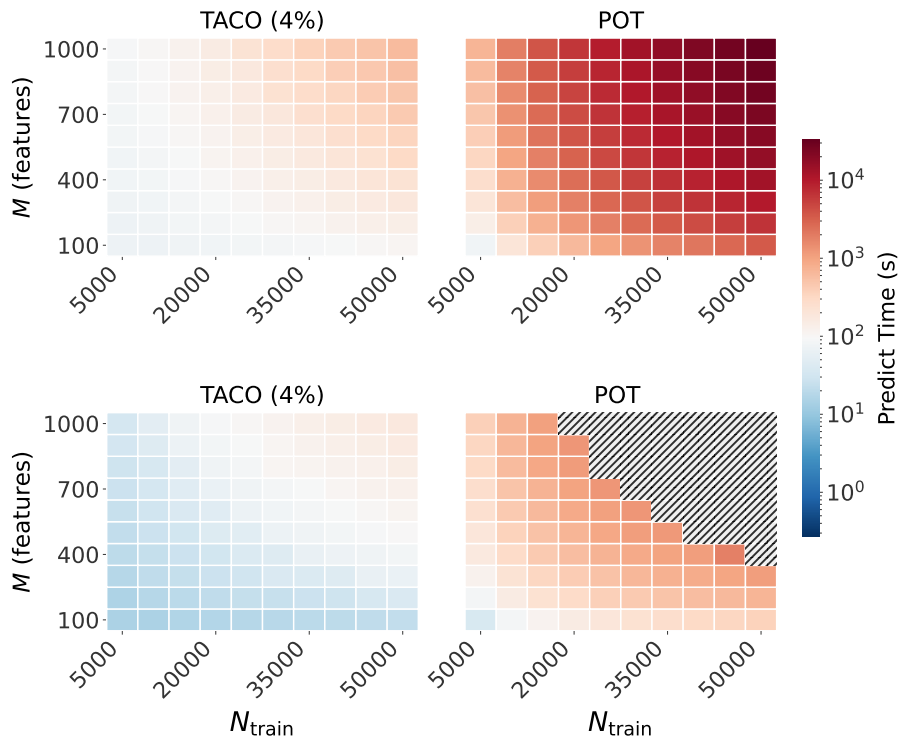


Figure 13. Predict time heatmap in the 100 test batches scenario. **Top:** Predict times for TACO and the Predictor Transformer, without using KV caching. **Bottom:** Predict times with KV caching. Black shaded regions indicate out-of-memory errors.

A.2. Time and Memory Consumption of TACO vs compression baselines

Table 5 reports inference time and memory usage for TACO and the compression baselines. For TACO without KV-cache, inference consists of two phases: (i) a one-time context compression step performed during the first prediction, and (ii) subsequent predictions that reuse the compressed context. Since the compressed context is already embedded, subsequent predictions bypass the predictor’s embedding layer and only execute the remaining inference stages.

As a result, the first prediction of TACO has a runtime comparable to the full-context predictor-only transformer and is therefore slower than inference with reduced-context variants, whereas subsequent predictions are substantially faster than those obtained even using the compression baselines.

At equal retained context ratios, TACO achieves between $1.2\times$ and $1.4\times$ lower mean prediction latency than kNN and random sampling, with the gap widening as the retained context fraction decreases (e.g., 306 ms vs. 408 ms at 1%). This advantage stems from the fact that kNN and random sampling remain stateless and must re-embed the selected context for every prediction, whereas TACO amortizes the compression cost across queries.

Compared to kNN and random sampling, TACO requires higher prediction-time memory at identical context ratios, as it retains a compressed *embedded* context to amortize embedding costs across predictions. However, this memory footprint remains negligible relative to the full predictor, yielding a 78–98% reduction in prediction-time memory while supporting faster subsequent inference. Importantly, in contrast to kNN and random sampling, which exhibit pronounced performance degradation under context reduction, TACO maintains performance close to the full predictor as discussed in the main paper.




Figure 14. Cumulative prediction time for TACO at different compression rates r compared to POT, random sampling and kNN for 100 test batches, without using KV Caching.

B. TabArena Chunking Results

Since we use the TabPFNv2 architecture as our backbone for both the compressor and predictor, we first run TACO on the same TabArena subset as the original method. Results are presented in Table 1. Because TACO can employ chunking (Hypothesis 6), we can then scale this evaluation to the full TabArena classification suite at an extreme compression rate of 0.1% to probe how informative the compressor’s latent context rows remain. We use a fixed compression rate of 0.1% for TACO and the baselines across all datasets. While chunking and stitching is only necessary for large datasets, we nevertheless apply it to smaller datasets strictly to isolate and demonstrate its effectiveness. The TabArena classification datasets span a wide range in size, from $N = 748$ to $N = 150,000$ rows. We therefore adopt a simple chunk-size policy as a function of N :

$$C(N) = \begin{cases} 500, & N < 2,000, \\ 1,000, & 2,000 \leq N < 10,000, \\ 5,000, & 10,000 \leq N < 20,000, \\ 10,000, & N \geq 20,000. \end{cases}$$

End-to-End Compression for Tabular Foundation Models

Table 5. Inference Time and Memory Benchmark of TACO and Compression Baselines without KV-caching

Method	Fit Time	First Predict Time	Subsequent Predict Time	Fit Mem	Predict Mem	
			Mean \pm Std		Mean	Median
Predictor	8.41s	28.55s	28.67s \pm 0.05s	36MB	22.45GB	22.45GB
TACO 1%	8.22s	29.68s	306ms \pm 18ms (\times 93.6)	92MB	776MB (-96.6%)	549MB (-97.6%)
TACO 2%	8.30s	29.80s	382ms \pm 18ms (\times 75.2)	92MB	1.05GB (-95.3%)	845MB (-96.3%)
TACO 4%	9.86s	30.71s	544ms \pm 17ms (\times 52.7)	92MB	1.63GB (-92.8%)	1.41GB (-93.7%)
TACO 8%	10.02s	32.47s	943ms \pm 17ms (\times 30.4)	92MB	2.78GB (-87.6%)	2.56GB (-88.6%)
TACO 16%	9.20s	35.36s	1.91s \pm 0.01s (\times 15)	92MB	5.10GB (-77.3%)	4.89GB (-78.2%)
kNN 1%	8.38s	431ms	408ms \pm 5ms (\times 70.3)	36MB	395MB (-98.3%)	395MB (-98.3%)
kNN 2%	7.95s	495ms	479ms \pm 5ms (\times 59.9)	36MB	625MB (-97.3%)	625MB (-97.3%)
kNN 4%	8.06s	665ms	645ms \pm 5ms (\times 44.5)	36MB	1.06GB (-95.3%)	1.06GB (-95.3%)
kNN 8%	8.57s	1.08s	1.05s \pm 0.01s (\times 27.4)	36MB	1.95GB (-91.3%)	1.95GB (-91.3%)
kNN 16%	7.96s	2.04s	2.02s \pm 0.01s (\times 14.2)	36MB	3.74GB (-83.4%)	3.74GB (-83.4%)
Random 1%	8.61s	422ms	398ms \pm 6ms (\times 72.1)	36MB	396MB (-98.3%)	396MB (-98.3%)
Random 2%	10.53s	515ms	483ms \pm 6ms (\times 59.4)	36MB	625MB (-97.3%)	625MB (-97.3%)
Random 4%	8.93s	664ms	643ms \pm 6ms (\times 44.6)	36MB	1.06GB (-95.3%)	1.06GB (-95.3%)
Random 8%	10.83s	1.08s	1.04s \pm 0.01s (\times 27.7)	36MB	1.95GB (-91.3%)	1.95GB (-91.3%)
Random 16%	8.83s	2.05s	2.01s \pm 0.01s (\times 14.3)	36MB	3.74GB (-83.4%)	3.74GB (-83.4%)

Parentheses: \times speed-up, $-$ % memory reduction (vs. POT).

With a compression rate 0.1%, this policy yields between 2 and 120 retained context rows per dataset after chunking and stitching, depending on N . We summarize performance across all datasets using boxplots in Figure 15. Overall, the results show that even at these extreme compression rates, the embeddings learned by the compressor remain highly informative, and TACO consistently outperforms both Random and kNN by a large margin.

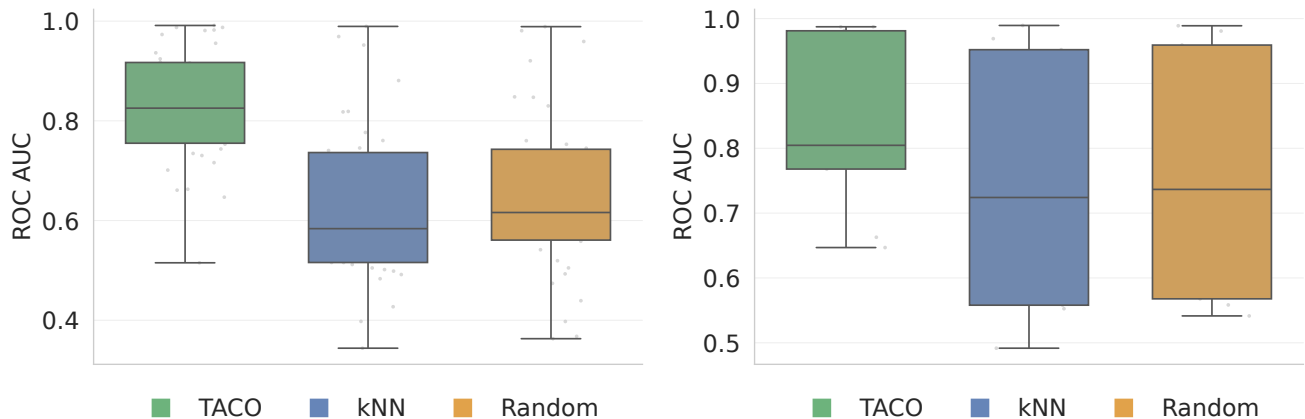


Figure 15. **Left:** Distribution of test performances on all TabArena classification tasks with chunking. **Right:** Distribution of test performances on TabArena classification tasks with $N \geq 20K$ rows, with chunking.

Table 6 reports the full results on all TabArena datasets. Here, K denotes the number of latent context rows retained after compression, ranging from 2 for the smallest dataset to 120 for the largest. We show that, even with as few as $K = 2$ latent context rows, TACO achieves strong performance and outperforms the random sampling and kNN baselines by a large margin.

Table 6. Performance on the full TabArena suite (ROC-AUC; higher is better). K indicates number of latent context rows retained after compression.

Dataset	#Rows	K	TACO	Random	kNN
blood-transfusion-service-center	748	2	0.730	0.573	0.512
diabetes	768	2	0.800	0.398	0.344
anneal	898	2	0.973	0.474	0.502
credit-g	1000	2	0.753	0.439	0.556
maternal_health_risk	1014	2	0.841	0.575	0.427
qsar-biodeg	1054	2	0.924	0.368	0.519
website_phishing	1353	3	0.913	0.505	0.499
Fitness_Club	1500	3	0.810	0.619	0.584
MIC	1699	3	0.661	0.493	0.583
Is-this-a-good-customer	1723	3	0.701	0.588	0.505
Marketing_Campaign	2240	2	0.914	0.724	0.721
hazelnut-spread-contaminant-detection	2400	2	0.956	0.703	0.717
seismic-bumps	2584	3	0.793	0.760	0.700
splice	3190	3	0.991	0.728	0.516
Bioresponse	3751	3	0.860	0.663	0.564
hiva_agnostic	3845	4	0.515	0.519	0.521
students_dropout_and_academic_success	4424	4	0.850	0.599	0.626
churn	5000	4	0.916	0.591	0.570
polish_companies_bankruptcy	5910	5	0.911	0.363	0.483
taiwanese_bankruptcy_prediction	6819	6	0.936	0.830	0.881
NATICUSdroid	7491	6	0.982	0.920	0.746
coil2000_insurance_policies	9822	8	0.716	0.595	0.398
Bank_Customer_Churn	10000	8	0.856	0.753	0.777
heloc	10459	9	0.797	0.745	0.761
jm1	10885	9	0.735	0.697	0.672
E-CommereShippingData	10999	9	0.744	0.613	0.698
online_shoppers_intention	12330	10	0.917	0.848	0.818
in_vehicle_coupon_recommendation	12684	11	0.761	0.592	0.516
HR_Analytics_Job_Change_of_Data_Scientists	19158	16	0.803	0.713	0.740
credit_card_clients_default	30000	24	0.786	0.737	0.724
Amazon_employee_access	32769	27	0.663	0.558	0.558
bank-marketing	45211	37	0.768	0.680	0.708
kddcup09_appetency	50000	40	0.805	0.568	0.492
Diabetes130US	71518	58	0.647	0.541	0.553
APSFailure	76000	61	0.987	0.981	0.969
SDSS17	78053	63	0.987	0.989	0.989
customer_satisfaction_in_airline	129880	104	0.981	0.959	0.952
GiveMeSomeCredit	150000	120	0.860	0.847	0.819

C. Detailed Results Across Datasets

This section presents additional details comparing TACO to the predictor-only transformer (POT) baseline across 26 classification datasets from TabArena (Erickson et al., 2025). Table 7 provides the full test ROC-AUC scores under different compression rates, along with the corresponding differences relative to POT. Across datasets, TACO yields performance that is very similar to POT, with only minor variations. This is also reflected in Figure 16, where the ROC-AUC distributions for POT and the TACO variants largely overlap.

Table 7. Test ROC-AUC performance of POT and TACO across datasets under varying compression rates (1%–16%). Parentheses report the absolute ROC-AUC difference of each TACO variant relative to POT (TACO - POT).

Dataset	#Rows	POT	TACO (1%)	TACO (2%)	TACO (4%)	TACO (8%)	TACO (16%)
blood-transfusion-service-center	748	0.720	0.724 (+0.004)	0.719 (-0.001)	0.714 (-0.007)	0.712 (-0.008)	0.710 (-0.010)
diabetes	768	0.846	0.844 (-0.002)	0.845 (-0.001)	0.846 (+0.000)	0.847 (+0.001)	0.848 (+0.002)
anneal	898	0.991	0.977 (-0.013)	0.976 (-0.015)	0.979 (-0.012)	0.979 (-0.012)	0.981 (-0.009)
credit-g	1000	0.779	0.771 (-0.008)	0.772 (-0.006)	0.775 (-0.004)	0.775 (-0.004)	0.776 (-0.003)
maternal_health_risk	1014	0.929	0.871 (-0.059)	0.874 (-0.055)	0.881 (-0.048)	0.883 (-0.046)	0.887 (-0.043)
qsar-biodeg	1054	0.936	0.934 (-0.002)	0.935 (-0.002)	0.935 (-0.001)	0.935 (-0.001)	0.935 (-0.001)
website_phishing	1353	0.970	0.946 (-0.024)	0.954 (-0.017)	0.955 (-0.015)	0.955 (-0.015)	0.955 (-0.015)
Fitness_Club	1500	0.810	0.810 (+0.000)	0.809 (-0.001)	0.809 (-0.001)	0.809 (-0.001)	0.809 (-0.001)
MIC	1699	0.658	0.665 (+0.007)	0.665 (+0.007)	0.666 (+0.008)	0.667 (+0.009)	0.665 (+0.007)
Is-this-a-good-customer	1723	0.732	0.738 (+0.006)	0.739 (+0.007)	0.739 (+0.007)	0.738 (+0.006)	0.739 (+0.007)
Marketing_Campaign	2240	0.933	0.932 (-0.001)	0.933 (-0.000)	0.933 (+0.000)	0.933 (+0.001)	0.933 (+0.000)
hazelnut-spread-contaminant-detection	2400	0.977	0.968 (-0.008)	0.970 (-0.007)	0.970 (-0.006)	0.971 (-0.006)	0.971 (-0.006)
seismic-bumps	2584	0.787	0.790 (+0.004)	0.790 (+0.007)	0.794 (+0.008)	0.796 (+0.009)	0.796 (+0.009)
splice	3190	0.996	0.996 (-0.000)	0.996 (-0.000)	0.996 (-0.000)	0.996 (-0.000)	0.996 (-0.000)
students_dropout_and_academic_success	4424	0.861	0.859 (-0.002)	0.860 (-0.001)	0.860 (-0.001)	0.860 (-0.001)	0.860 (-0.001)
churn	5000	0.935	0.933 (-0.002)	0.932 (-0.002)	0.932 (-0.003)	0.932 (-0.002)	0.932 (-0.003)
polish_companies_bankruptcy	5910	0.955	0.940 (-0.015)	0.942 (-0.013)	0.943 (-0.012)	0.943 (-0.012)	0.943 (-0.012)
taiwanese_bankruptcy_prediction	6819	0.942	0.945 (+0.002)	0.945 (+0.003)	0.945 (+0.003)	0.945 (+0.003)	0.945 (+0.003)
NATICUSdroid	7491	0.984	0.984 (-0.000)	0.984 (-0.000)	0.984 (-0.000)	0.984 (-0.000)	0.984 (-0.000)
coil2000_insurance_policies	9822	0.748	0.746 (-0.002)	0.746 (-0.002)	0.746 (-0.002)	0.747 (-0.001)	0.746 (-0.001)
Bank_Customer_Churn	10000	0.862	0.862 (-0.001)	0.862 (-0.001)	0.862 (-0.001)	0.862 (-0.001)	0.862 (-0.001)
heloc	10459	0.797	0.799 (+0.002)	0.799 (+0.002)	0.799 (+0.002)	0.799 (+0.002)	0.799 (+0.002)
jm1	10885	0.745	0.736 (-0.009)	0.736 (-0.009)	0.737 (-0.009)	0.737 (-0.008)	0.737 (-0.009)
E-CommereShippingData	10999	0.743	0.743 (-0.001)	0.743 (-0.001)	0.742 (-0.001)	0.742 (-0.001)	0.742 (-0.001)
online_shoppers_intention	12330	0.923	0.921 (-0.002)	0.921 (-0.002)	0.921 (-0.002)	0.921 (-0.002)	0.921 (-0.002)
in_vehicle_coupon_recommendation	12684	0.775	0.763 (-0.012)	0.763 (-0.012)	0.763 (-0.011)	0.763 (-0.011)	0.763 (-0.011)




Figure 16. Boxplots of test ROC-AUC across datasets for the predictor-only transformer (POT) and TACO under varying compression rates (1%–16%). Boxes indicate the interquartile range with the median, while whiskers show the overall spread of performance.