TabDPT: Scaling Tabular Foundation Models on Real Data

Junwei Ma*, Valentin Thomas*,1, Rasa Hosseinzadeh, Alex Labach, Hamidreza Kamkari, Jesse C. Cresswell, Keyvan Golestan, Guangwei Yu, Anthony L. Caterini¹, Maksims Volkovs Layer 6 AI, Toronto

Abstract

Tabular data is one of the most ubiquitous sources of information worldwide, spanning a wide variety of domains. This inherent heterogeneity has slowed the development of Tabular Foundation Models (TFMs) capable of fast generalization to unseen datasets. In-Context Learning (ICL) has recently emerged as a promising solution for TFMs, enabling dynamic adaptation to new tasks without additional tuning. While many studies have attempted to re-purpose large language models for tabular ICL, they have had limited success, so recent works have focused on developing tabular-specific foundation models. In this work, we propose an approach to combine ICL-based retrieval with self supervised learning to train tabular foundation models. We also investigate the utility of real vs. synthetic data for model pre-training, and show that real data can contain useful signal not easily captured in synthetic training. Specifically, we show that incorporating real data during the pre-training phase can lead to significantly faster training and better downstream generalization to unseen data. Our resulting model, TabDPT, achieves strong performance on both regression (CTR23) and classification (CC18) benchmarks. Importantly, we also demonstrate that with our pre-training procedure, scaling both model and data size leads to consistent performance improvements that follow power laws. This echoes scaling laws in LLMs and other foundation models, and suggests that large-scale TFMs can be achievable. We open-source our full pipeline: inference code including trained model weights can be found here, and the training code to reproduce experiments can be found here.

1 Introduction

Tabular data constitutes the backbone of most real-world applications, from finance, health-care, and e-commerce, to many others [56]. However, building a *Tabular Foundation Model* (TFM) – a single model that generalizes across the enormous heterogeneity of tabular datasets – remains a key challenge. The alternative, traditional approach [9, 47] is to train individual models for each new task, which may yield strong results but requires costly model selection and hyperparameter tuning on a per-dataset basis. For deep learning methods, this procedure results in even greater computational overhead,

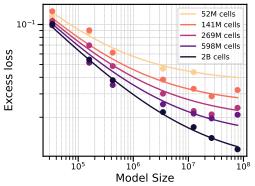


Figure 1: Scaling behavior for our foundation tabular models. Increasing model or pre-training data size (number of cells) leads to consistent improvements predictable by power laws (fitted solid lines).

which has hindered the adoption of neural networks as a universal solution in the tabular domain.

^{*}Equal Contribution; ¹Correspondence to {valentin.t, anthony}@layer6.ai

In-context learning (ICL) offers an appealing alternative by enabling a model to adapt to new tasks by simply modifying the context, obviating the need for per-dataset fine-tuning or hyperparameter selection [8]. Beyond lowering the cost of model deployment, ICL facilitates rapid prototyping and provides a natural mechanism for handling distribution shifts, as the model can efficiently adapt to new data at inference time using only in-context examples, mimicking the effect of conventional training with less overhead.

Recent attempts to repurpose large language models (LLMs) for ICL on tabular data faced several fundamental obstacles [15, 25, 17, 49]. Chief among them is the highly inefficient tokenization of numerical tabular data into text, which quickly saturates the LLM's context window even for moderately sized tables. LLMs' results also vary based on the specific prompt format [52, 60] and are sensitive to the order of the given examples [40], whereas tabular data is inherently unordered. These limitations degrade performance, leading to prompt-tuning reliance and difficulty handling even moderately sized tables. Alternative ICL solutions for tabular data, such as TabPFN [44, 29], are architecturally designed for tabular data, enabling them to handle tables of practical size more effectively. This direction is gaining popularity but is still in early stages with relatively few tabular-based TFMs developed [30, 48]. Further investigation is needed into architecture design choices and training procedures that lead to strong downstream generalization, and in this work we make a major step in this direction.

We show that ICL retrieval combined with self-supervised learning (SSL) based on column masking [11, 26, 45] leads to a robust pre-training procedure for TFMs. We investigate the utility of real data, showing that using it for this pre-training procedure leads to faster convergence and improved downstream accuracy on unseen datasets compared to training exclusively on competing open-source synthetic data generators. Since existing TFMs are predominately trained with synthetic data, our findings suggest that further investigation should be conducted into the benefits of curating real datasets for pre-training. Our pre-training process produces a TFM capable of both classification and regression with leading accuracy on new, unseen datasets without any further training or tuning. We name this new model the Tabular Discriminative Pre-trained Transformer, or **TabDPT** for short.

Comprehensive evaluations on the OpenML-CC18 [5] and OpenML-CTR23 [16] benchmarks confirm the effectiveness of TabDPT. It consistently matches or surpasses the performance of specialized models that undergo extensive per-dataset hyperparameter optimization at a fraction of the deployment time and cost. Furthermore, we show strong results in the few-shot regime, where, with minimal semi-supervised modifications, TabDPT outperforms specialized baselines on 10-shot classification tasks, highlighting its versatility. Finally, we demonstrate that TabDPT scales predictably with both model size and quantity of real pre-training data (Figure 1), underscoring the viability of large-scale foundation models for tabular domains. We summarize our contributions as follows:

- 1. We develop a procedure for pre-training of TFMs based on ICL retrieval and SSL that leads to robust downstream generalization to unseen data without explicit fine-tuning.
- 2. We show that applying our pre-training procedure to real data leads to faster convergence and better downstream accuracy than using purely open-source synthetic data generators. We also demonstrate scaling with this procedure where more data and larger models continue to yield consistent gains akin to scaling laws in LLMs.
- 3. We release TabDPT including open model weights, code, and pre-training datasets to encourage further research and reproducibility. A lightweight inference interface is available here, and the full training code here.

2 Related Work

Tabular Foundation Models Although foundation models in other domains [11, 8, 12] have shown tremendous progress in recent years, foundation models for tabular data have lagged behind [56]. Several attempts have tried to bridge this gap. However, many of these methods require additional training when applied to a novel task [37, 57, 38], hindering widespread adoption in practice due to the high costs associated with fine-tuning. Meanwhile, ICL-based TFMs have started gaining traction. Among them, large language model (LLM) based approaches [17, 27, 54, 61] initially appear to be a natural fit. However, LLMs cannot easily handle numerical content of tables since their tokenization is not tailor-made for numerical data. As a result, LLM-based ICL methods suffer from high memory costs and low performance, as we discuss in Sections 3.1 and 4.3. Furthermore, natural language

follows a sequential left-to-right structure, whereas tabular data is inherently order-invariant with respect to rows. Indeed, LLMs are sensitive to the order of examples in context [40].

On the other hand, tabular-specific ICL methods such as TabPFN [29] can naturally handle tabular data with numerical entries. However, they are completely reliant on synthetic data generators; ensuring that this mechanism captures the full diversity and nuances of real-world data is challenging, and making meaningful improvements to it is difficult. Notable concurrent work by Hollmann et al. [30] does not include open-sourced pre-training and synthetic generators, further complicating direct improvements to their model; another concurrent method requires a complex, three-stage procedure to learn from synthetic generators [48]. In this paper, we hypothesize that real tabular data contains much more information than existing heavily engineered synthetic tabular generators, thus allowing more straightforward improvements by scaling model and data size; this is supported by experiments in Section 4.5.

Scaling Laws Neural scaling laws have been studied extensively in various data modalities [8, 35]. In natural language processing (NLP), scaling laws were first identified in language models, where performance improves predictably with larger models, training corpora, and compute. Similar trends have been observed in computer vision [65]. Recently, Schambach et al. [51] demonstrated preliminary evidence of scaling laws for tabular data with very small-scale experiments. In this paper, we follow the developments from NLP, conducting thorough experiments to show that TabDPT follows scaling laws. Our novel analysis of scaling in the tabular domain paves the way for TFMs to scale and improve, much like foundation models in other domains.

Tabular Self-supervised Learning SSL has proven to be successful for text and images [11, 12], but has not achieved similar success on tabular data. Many tabular SSL methods cannot generalize beyond the dataset on which they were pre-trained [34, 64, 42, 53], raising the question of their potential to benefit from cross-task training. The answer is likely to be "yes", as recent work shows even tree-based methods benefit from hyperparameter tuning across tasks [31], and basic MLPs can be competitive in predictive tabular tasks when leveraging SSL [50]. Consequently, tabular SSL methods have begun to show generalization across tasks and competitive performance [66, 63]. However, they still require task specific fine-tuning and hyperparameter selection, which can be time- and resource-intensive. The only other tabular SSL method we are aware of that generalizes across tasks without per-task tuning is from Gardner et al. [17]. However, despite having 8 billion parameters (several orders of magnitude larger than TabDPT), its performance remains uncompetitive as its LLM-based design limits its context size to only 32 data points. To our knowledge, we are the first to demonstrate competitive performance and generalization of tabular SSL across tasks without task-specific training or hyperparameter tuning.

3 TabDPT Methodology

We now describe TabDPT, our approach for building a TFM, which employs (i) a row-based transformer encoder for in-context learning, (ii) self-supervised learning for *augmenting* the pre-training set, and (iii) retrieval-based context selection for both training and inference. These components are combined in a novel fashion to produce a single foundation model that generalizes to a diverse array of unseen classification and regression tasks without dataset-specific fine-tuning.

3.1 Tabular Transformer Encoder

We use a row-based transformer encoder similar to TabPFN v1 [29], where each table row serves as a "token." Specifically, for an input table with N rows and F features, we standardize its feature dimension to $F_{\rm max}$ via padding ($F < F_{\rm max}$) or dimensionality reduction ($F > F_{\rm max}$), then embed each row into a d-dimensional vector. Rows attend to each other through stacked transformer layers.

A key motivation behind row-based encoding is memory and compute efficiency. In cell- or text-based tokenizations [21, 57, 30], each cell in an $N \times F$ table must be split into multiple tokens (e.g., subwords), resulting in $\mathcal{O}(N \times F \times \langle N_{\text{tok}} \rangle)$ tokens, where $\langle N_{\text{tok}} \rangle$ is the average number of tokens per cell. Even modest-sized tables can inflate the input sequence well beyond typical transformer context limits. By contrast, encoding *entire rows* as tokens reduces the sequence length to N, allowing us to process many more rows with significantly lower memory overhead. It also affords invariance to row ordering (no positional encoding), which matches the perceived structure of tabular data. While cell-based tokenization additionally introduces an appealing invariance to

column ordering, we can achieve it approximately by permuting columns at training time; other works [30, 48] have also found that strict column order invariance can introduce training instability.

Finally, TabDPT uses a shared architecture for both regression and classification, which has been concurrently explored in TabPFN v2 [30]. In our case, this is realized through two separate MLP heads atop one single shared transformer: one head used for classification (supporting up to $C_{\rm max}$ classes) and another for regression. The shared backbone facilitates better parameter sharing across regression and classification tasks. Additional implementation details are provided in Appendix C.

3.2 Self-Supervised Learning on Tabular Data

Real-world tabular datasets are typically structured as $\mathcal{D}=\{X,y\}$, where $X\in\mathbb{R}^{N\times F}$ is the input table containing N rows and F features, and $y\in\mathbb{R}^{N\times 1}$ is the target (class index or regression value). In some instances y can contain multiple targets but that is relatively rare. As the number of high quality publicly available tabular datasets is also relatively low, training TFMs in a supervised fashion to predict y quickly saturates and leads to overfitting. To circumvent this, current methods predominantly leverage synthetic data that is continuously generated throughout pre-training [29, 7, 30, 48]. However, this approach has its own set of challenges where priors that generate synthetic data need to be extensively engineered to approximate the distribution of highly heterogeneous real-world data.

In this work we take a different approach and aim to maximize the value of real data in TFM pretraining. To this end, we leverage self-supervised learning (SSL) to extract maximal information from each table and regularize the model. Inspired by masked modeling in language [11] and vision [26], as well as some initial efforts in tabular domain [45, 17], we randomly designate one column as the "target" to be predicted from the others. Concretely, we randomly pick a task to be either regression or classification, then pick a column c with sufficient unique values as the target. We remove c from the table and standardize its values for regression or bin them into classes for classification. The model then has to predict this auxiliary target y = c from the resulting table $X \setminus c$.

We also shuffle and drop other columns, forcing the model to learn from varying feature combinations. Without these augmentations, the number of prediction tasks would grow only linearly with the number of features. In contrast, our approach scales task count combinatorially, compelling the model to capture richer inter-feature relationships. This provides a stronger training signal and serves simultaneously as a regularizer. Pseudo-code for the SSL procedure is provided in Appendix D.

3.3 Retrieval-Based Pre-Training

In ICL, training data rows are passed as context to the transformer together with a target test row to generate a prediction. Although row-based embeddings allow for larger sample sizes, using full training tables as context still quickly becomes prohibitively large. Retrieval-based techniques, where only the top K most similar training rows are selected as context, have been shown to mitigate this inherent

Algorithm 1 One Training Step of TabDPT

Select B random datasets $\{\mathcal{D}^{(i)}\}_{i=1}^{B}$

for each dataset $\mathcal{D}^{(i)}$ do

Randomly set task as regression or classification Generate target $y^{(i)}$ from a random column $c^{(i)}$

Sample K "close" rows from $X^{(i)} \setminus c^{(i)}$

Split rows into context $\{X_{\text{ctx}}^{(i)}, y_{\text{ctx}}^{(i)}\}$ and query $\{X_{\text{qy}}^{(i)}, y_{\text{qy}}^{(i)}\}$

Shuffle and/or drop columns from $X_{
m ctx}^{(i)}$ and $X_{
m qy}^{(i)}$

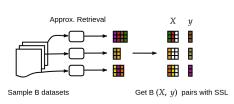
Get transformer predictions $\{\hat{y}_{qy}^{(i)}\}_{i=1}^{B}$ (Equation 1) Calculate loss and perform model update

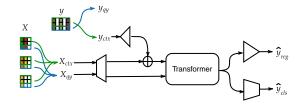
limitation at *inference time* [55, 62], significantly improving the accuracy and scalability of ICL.

We propose to take this one step further and align training with inference by also leveraging retrieval during training batch construction. Formally, after obtaining y=c and $X \setminus c$ through SSL, we sample a set of K rows from $X \setminus c$ that are close to each other in the feature space. These rows are partitioned into two groups uniformly at random: "context" $\{X_{\rm ctx},y_{\rm ctx}\}$ and "query" $\{X_{\rm qy},y_{\rm qy}\}$. The context $\{X_{\rm ctx},y_{\rm ctx}\}$, together with the query features $X_{\rm qy}$, is fed into the model to predict query targets $y_{\rm qy}$. To form model inputs, we pass the context and query through the appropriate row embedding functions $(\phi_x$ or $\phi_y)$, then sum the embeddings of $X_{\rm ctx}$ and $y_{\rm ctx}$, and concatenate with $X_{\rm qy}$:

$$\hat{y}_{qy} = \mathbf{Transformer} \left[\phi_x(X_{\mathsf{ctx}}) \oplus \phi_y(y_{\mathsf{ctx}}), \phi_x(X_{\mathsf{qy}}) \right]. \tag{1}$$

Here, \oplus is element-wise addition, and \hat{y}_{qy} emerges from the classification or regression head depending on the target task. The notation $[\cdot,\cdot]$ indicates the cross-attention split: query points attend to all context points, while context points attend to each other but not to queries – that is, only context serves as keys in the attention mechanism.





(a) Selecting a training batch

(b) Overview of the architecture

Figure 2: (a) We sample B tables from different datasets to construct $X \in \mathbb{R}^{B \times N \times F_{\max}}$ and $y \in \mathbb{R}^{B \times N}$. (b) X and y are partitioned into context $\{X_{\text{ctx}}, y_{\text{ctx}}\}$ and query X_{qy} inputs and passed through embedding functions (indicated by rectangle/triangle). Embeddings of X_{ctx} and y_{ctx} are summed together, concatenated with context embedding of X_{qy} , and passed through a transformer encoder to get classification \hat{y}_{cls} or regression \hat{y}_{reg} prediction for the query. Loss between this prediction and query targets y_{qy} is used to update the model.

Using "similar" rows as context during training and limiting context size naturally aligns it with inference, improving model generalization. It also maintains efficient batch sizes, as context size no longer scales with table size, while still exposing the model to relevant context information. We consistently observe that training with this procedure speeds up convergence and leads to better downstream accuracy. An illustration of our model architecture is shown in Figure 2, with a full training step described in Algorithm 1.

3.4 Inference on New Data

At inference, given a new dataset, we follow the same retrieval protocol. For each test query row $x_{\rm qy}$ we retrieve the top K closest rows from the training set to get context $\{X_{\rm ctx},y_{\rm ctx}\}$. Context and query inputs are then embedded and passed through the transformer encoder (see Equation 1) to get classification $\hat{y}_{\rm cls}$ or regression $\hat{y}_{\rm reg}$ predictions for the query row, depending on the target task. Note that other than context retrieval, no dataset-specific tuning is done and we only run forward passes through the model. Retrieval can add additional latency, however, modern nearest neighbor libraries such as FAISS [13] deliver millisecond responses and scale to billion-row indices. Our backbone imposes a pre-defined maximum number of classes $C_{\rm max}$ and features $F_{\rm max}$; we now discuss how to overcome these limitations with inference-time techniques.

Classes If a dataset contains $C > C_{\max}$ classes, we cannot perform classification in a single forward pass. Naive binary one-versus-all classification would require C forward passes that can significantly impact inference speed as some datasets have hundreds of classes. A more computationally efficient approach is to represent C in base C_{\max} and perform classification on each base- C_{\max} digit as a separate, well-defined prediction task. This approach reduces the required number of forward passes to $\lceil \log_{C_{\max}}(C) \rceil$ and is fully compatible with the TFM setting.

Features When the number of features in a table exceeds $F_{\rm max}$, we can reduce the dimensionality of the table using Principal Component Analysis (PCA) to $F_{\rm max}$, effectively compressing the features to fit the model requirement, while preserving the most salient information.

4 Experiments

In this section, we evaluate TabDPT against leading baselines on standard benchmarks for TFMs, provide a detailed analysis of runtime, and ablate key components.

4.1 Data

Training Data Our training data was collected from OpenML [58] and consists of a wide range of public tabular datasets across diverse domains, all available under the CC-BY licence. To find appropriate datasets, we considered those specified in the Grinsztajn et al. [24], TabZilla [43], and AMLB [18] benchmarks, as well as additional datasets found individually. The full set of pre-training data contains 123 datasets, with a total of 32M rows and 2B cells (individual values within each table) from a diverse set of domains such as biology, finance, industrial applications, and medicine.

The scale of this data is comparable to related work such as Tabula-8B [17] that fine-tuned the LLaMA 3-8B [23] language model for the tabular domain using real-world data. We conjecture that the diversity of domains present in our pre-training data can provide a salient signal and improve downstream generalization. Further details, including the complete list of training datasets and breakdown by size and domain are provided in Appendix B.

Evaluation Data For evaluation, we consider two commonly used public benchmarks containing a total of 107 datasets: CC18 [5] for classification tasks and CTR23 [16] for regression tasks. CC18 is a suite of 72 classification datasets originally sourced from OpenML. These datasets contain between 500 and 100,000 instances, fewer than 5,000 features, and originate from diverse domains such as finance, biology, games, banking, industrial applications, or natural signals such as vision or sound. Datasets were selected according to curation criteria that included removing synthetic data, requiring source information, and removing datasets where a simple algorithm achieves 100% accuracy. CC18 is a common benchmark for evaluating tabular learning on classification tasks [2, 29, 43]. CTR23 is a benchmark suite of 35 datasets also curated from OpenML. It follows most of the design choices of CC18 but contains only regression tasks. In particular, it uses the same restrictions on the number of samples and features as CC18, but replaces the accuracy restriction with a requirement that a linear model must not achieve $R^2 = 1$.

4.2 Baselines

We compare our method against leading baselines that are tuned for each dataset, including tree-based methods such as XGBoost [9], LightGBM [36], and CatBoost [47], and deep learning methods such as TabR [21], TabM [22] and MLP-PLR [20], as well as MLP. For XGBoost, CatBoost, and LightGBM, we use results reported in the TabZilla benchmark [43]. Some datasets are missing results, so we conduct hyperparameter optimization and train models following the TabZilla protocol using the code repository from [21]. For TabR, TabM, MLP-PLR, and MLP, we use the same code repository with the predefined search space and 30 search rounds for both CC18 and CTR23. We choose the best hyperparameters for each dataset fold individually based on the validation performance.

We also compare to ICL baselines including the LLM-based Tabula-8B [17], and tabular-specific foundation models TabPFN v2 [30], and TabPFN (kNN) [55], which retrieves neighbours of each query at inference time. Note that using retrieval with a cell-based method like TabPFN v2 adds prohibitive computational overhead. We run all methods on at least two different splits of the data and report 95% confidence intervals using bootstrapping [1]. For TabDPT, we use the 78M-parameter variant, with 16 transformer layers pre-trained for 600K steps. All training and inference is done on Nvidia A100 GPUs with 40 GB of memory. Further training details are provided in Appendix C

4.3 Results on CC18 and CTR23

Our main results comparing models on the evaluation data are shown in Table 1. TabDPT shows the best overall performance across all models. It is competitive with TabPFN v2 across metrics, and significantly outperforms it on accuracy on CC18. TabDPT also significantly outperforms both deep learning and tree-based methods that are trained for each dataset. These results indicate that real data can be effectively utilized with SSL to train robust TFMs with leading performance. We provide a breakdown of results in Appendix F.1, examining each algorithm's performance under varying dataset sizes, numbers of features, categorical fraction, and percent missing. Results indicate that TabDPT is robust to dataset variations in all of these categories. For Table 1, we use TabDPT with 2,048 context size and 8 ensemble members, where randomness for a given ensemble member emerges by permuting the features and classes (if applicable).

Our foundation models require negligible hyperparameter tuning. Rather than reserving a validation set, we merge it with the training set. In contrast, several baselines perform hyperparameter tuning on a separate validation set (instead of cross-validation with a final refit on the combined train and validation set). To benchmark under the same constraint, we report a variant in which TabDPT receives only the training set, denoted by (no val.). This setting should be compared to classical methods; the default TabDPT setting is comparable to other foundation models that also access the validation set. Since TabDPT (no val.) never uses validation data, it is disadvantaged compared to

¹https://github.com/yandex-research/tabular-dl-tabr

Algorithm	CO	C18	CTR23			
Aigoriumi	AUC	Accuracy	Correlation	R^2		
TabDPT	0.933 ± [0.929, 0.937]	0.884 ± [0.882, 0.887]	0.837 ± [0.826, 0.848]	0.742 ± [0.731, 0.754]		
TabPFN v2	0.932 ± [0.928, 0.936]	$0.872 \pm [0.869, 0.875]$	$0.835 \pm [0.825, 0.845]$	$0.740 \pm [0.729, 0.751]$		
TabPFN (kNN)	0.918 ± [0.915, 0.921]	$0.850 \pm [0.847, 0.853]$	N/A	N/A		
TabPFN	0.898 ± [0.895, 0.901]	$0.812 \pm [0.810, 0.814]$	N/A	N/A		
TabDPT (no val.)	$0.928 \pm [0.924, 0.932]$	$0.874 \pm [0.871, 0.877]$	$0.835 \pm [0.826, 0.844]$	$\underline{0.739}$ ± [0.729, 0.749]		
TabM	0.917 ± [0.913, 0.921]	$0.878 \pm [0.876, 0.881]$	0.824 ± [0.815, 0.834]	0.732 ± [0.722, 0.742]		
TabR	0.925 ± [0.922, 0.928]	$0.874 \pm [0.871, 0.877]$	0.828 ± [0.814, 0.842]	$0.714 \pm [0.693, 0.735]$		
MLP-PLR	0.912 ± [0.905, 0.919]	$0.869 \pm [0.865, 0.872]$	0.829 ± [0.821, 0.838]	$0.716 \pm [0.699, 0.732]$		
MLP	$0.872 \pm [0.867, 0.876]$	$0.809 \pm [0.805, 0.813]$	N/A	N/A		
XGBoost	0.926 ± [0.922, 0.929]	0.869 ± [0.866, 0.872]	0.827 ± [0.818, 0.837]	0.711 ± [0.698, 0.725]		
LightGBM	0.924 ± [0.920, 0.927]	$0.862 \pm [0.859, 0.866]$	$0.825 \pm [0.816, 0.834]$	$0.713 \pm [0.697, 0.729]$		
CatBoost	0.926 ± [0.923, 0.930]	$0.864 \pm [0.860, 0.867]$	0.822 ± [0.808, 0.837]	$0.703 \pm [0.683, 0.723]$		

Table 1: Main results comparing models on evaluation data. We report the average performance on four metrics and their 95% confidence intervals. The best algorithm for each metric is bolded. The best model that does not append the validation set to the training set is underlined. Tabula-8B [17] only reports results on a subset of datasets in CC18 so we conduct pairwise comparison against it on reported datasets in Figure 3a.

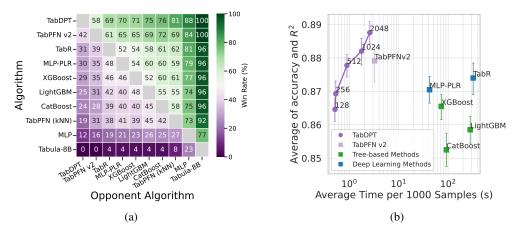
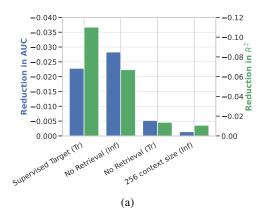


Figure 3: (a) Pairwise win-rate comparison. A win is counted for the method that achieves the higher classification/regression accuracy/ R^2 on a given dataset. (b) Inference runtime vs performance. Tab-DPT models are ordered by context size. Non-TFM baseline runtimes are the total of hyperparameter optimization and inference.

methods that tune on validation data. Even so, TabDPT (no val.) remains competitive and is only outperformed by TabM [22] (using 30 hyperparameter tuning rounds) on CC18 accuracy.

Win-Rate Comparison To get a more direct view of how various methods perform against each other, we compute pairwise win-rate statistics. We assign a "win" to the method if it achieves a higher accuracy score on CC18 or R^2 score on CTR23. This also allows us to compare against Tabula-8B on CC18 as it only reports results on 61 of 72 datasets. Figure 3a shows the win-rate matrix for all methods, painting a similar picture to Table 1: TabDPT performs best overall, followed by the strong tabular-based foundation model TabPFN v2. Tabula-8B with 32 shots – the leading LLM-based tabular foundation model – is not competitive with the other techniques across our benchmarks, indicating that current LLMs are not well adapted to the tabular domain and further techniques are needed. We extend the pairwise comparison between methods in Appendix E using both Elo [14] and Glicko2 [19] metrics, drawing similar conclusions to above.



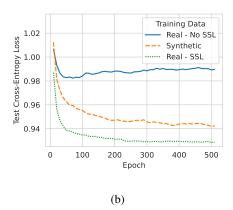


Figure 4: (a) Ablation of key components in training (Tr) and inference (Inf). A higher blue bar and a higher green bar indicate greater reduction in AUC and R^2 respectively. (b) Test loss curves on CC18 when training with and without SSL on real data as well as synthetic data only. We note that "Synthetic" and "Real - SSL" both saturate at 500 epochs, which we confirmed by training for 1024 epochs.

4.4 Ablation Study

In this section, we discuss the ablation of key components in our training and inference strategies, with results visualized in Figure 4.

Training Ablation First, we assess the importance of SSL during training. To ablate SSL, we only use the original target for each table during training and observe a large loss in performance, as shown under "Supervised Target (Tr)" in Figure 4a. The impact on training is further illustrated by the "Real-No SSL" curve in Figure 4b. Training without SSL starts to overfit after around 50 epochs whereas with SSL model continues to improve even after 500 epochs ("Real - SSL" curve). These results demonstrate the critical role of SSL when training TFMs with real data where only one target is typically available per dataset.

Second, we assess the benefit of using retrieval during training to form the context versus random subsampling, the results are shown under "No Retrieval (Tr)" in Figure 4a. We see that removing training retrieval leads to a consistent drop in both classification and regression test accuracy. Although smaller in magnitude than drop from removing SSL, these results confirm that aligning training and inference procedures is beneficial.

Finally, we benchmark the impact of training on real ("Real - SSL" curve) vs. synthetic data ("Synthetic" curve) in Figure 4b. We use the TabPFN [29] synthetic data generator for this experiment. We see that using real data with SSL consistently outperforms training on synthetic data across all epochs, achieving lower test loss under the same compute budget. This further highlights the effectiveness of real data when paired with SSL on our TFM architecture.

Inference Ablation Similarly to [55], we find that using subsampling instead of retrieval during inference significantly decreases performance as indicated by "No Retrieval (Inf)" in Figure 4a. Using a smaller context of 256 size also decreases performance as expected, although it does not decrease nearly as much as the other important components discussed above.

4.5 Scaling Laws

Although preliminary results on tabular scaling have been reported [51], this work provides the first analysis of scaling laws for TFMs that are not restricted to any particular domain. We focus on measuring scaling with pre-training on real data only, and evaluate performance as a function of training data amount and model size, systematically varying both. Model size is varied by changing both the number of layers and their dimensions. Our models range from 33K to 78M parameters, trained on data subsets spanning from 52M cells (104K rows) to 2B cells (32M rows). Following Hoffmann et al. [28], we adopt the joint power-law model:

$$\hat{\ell}(P,D) = AP^{-\alpha} + BD^{-\beta} + E \tag{2}$$

Method	cmc	karhunen	optdigit	diabetes	semeion	pixel	dna	Avg.
TabDPT (semi)	44.24	92.08	94.31	72.01	84.89	93.58	61.85	77.56
STUNT [45]	42.01	86.95	89.91	72.82	74.74	89.90	80.96	76.76
CACTUs [33]	42.14	85.48	87.92	70.75	68.22	87.21	84.40	75.16
VIME + LR[64]	37.92	86.63	89.63	66.56	77.66	88.71	74.73	74.55
kNN (STUNT) [45]	41.07	85.63	87.44	71.32	74.64	87.52	71.15	74.11
ICT [59]	38.00	88.25	90.84	67.63	74.67	89.13	69.55	74.01

Table 2: Few-shot accuracy on seven CC18 datasets. Only 10 labeled examples are available in each class of the training set, the rest are unlabeled.

where $\hat{\ell}$ represents the estimated target metric, P and D denote the number of parameters and total cells in the training set, and A,B,α,β,E are the scaling parameters. Despite using a row-based encoding, we measure data size by cell count, as not all rows contribute equally to the model's learning, particularly in the encoder layer that computes the embeddings. Applying the improved methodology of Besiroglu et al. [4], we estimate the scaling exponents as $\alpha=0.42$ and $\beta=0.39$, indicating that improvements can occur in both dimensions.

In Figure 1, we illustrate the scaling behaviour of our models along with the power-law fit. Since we train on equal proportions of classification and regression tasks, the loss on the y-axis represents the mean of the cross-entropy loss for classification and $1-\rho$ for regression, where ρ is the correlation between prediction and target, equivalent to the MSE for normalized vectors. Visualization is done on a log scale for the excess loss $\hat{\ell}(P,D)-E$ instead of the raw loss, debiasing the estimate by E. We observe consistent improvements when both data and model size increase indicating that information contained in real world tabular datasets can be effectively mined with SSL to pre-train robust TFMs.

4.6 Few-Shot Learning

We next assess the performance of TabDPT on few-shot learning tasks. We adopt the protocol from STUNT [45] with the 10-shot set-up where only 10 labeled rows are provided for each class in each training table, the rest are masked and the goal is to leverage the labeled + unlabeled data to accurately predict the test set. This simulates real-world settings where only small subsets of data can have labels. We compare against baseline results from [45] including STUNT [45], CACTUs [33], VIME+LR [64], and ICT [59] available in their paper. All methods are evaluated on seven datasets from CC18 using the accuracy metric.

TabDPT typically leverages a much larger context than 10 instances during both training and inference. To adapt it to the few-shot setting, we increase the context size by first predicting class probabilities for the unlabeled training set using the 10 labeled examples as context. Then, we take the top-1000 points where predicted probability is highest and use them and their predicted labels – along with the original 10 shots – as context. This results in TabDPT (semi), a semi-supervised version of our TFM that leverages pseudo-labels. This method outperforms STUNT, a leading few-shot method, on 5 of the 7 datasets and on average accuracy (over 50 seeds). Furthermore, it requires only forward passes to generalize to new tasks once we have a pre-trained model, while STUNT trains a new model for each task. This experiment demonstrates the potential of TabDPT as a TFM: it can rapidly adapt to new tabular settings without any additional weight updates.

4.7 Inference Speed

In this section, we analyze the inference runtime of TabDPT against baselines on new datasets. For TFMs we measure the cost of computing the context and making forward passes through the models. For deep learning and tree-based methods that are dataset specific we measure the total time to train the model, including hyperparameter search, and to run inference. We repeat each experiment across dataset folds and measure the average times to process 1,000 rows (computed overall on $\approx 2M$ rows). For TabDPT we report runtimes with different context sizes from 128 to 2048 and the corresponding impact on accuracy. Results are shown in Figure 3b, we see that even our largest model with context size 2048 is at least one order of magnitude faster than dataset-specific tree and deep learning baselines. TabDPT runtime is also comparable to the leading TFM TabPFN v2 while achieving higher accuracy. The cost of pre-training TFMs is not included in this comparison.

However, analogous to LLMs, we stipulate that it is a one-time cost that is offset when model is applied across many datasets and use cases.

On the other hand, when training cost is not part of the budget (e.g., fixed models on fixed datasets), classical tree-based pipelines can deliver substantially faster inference (e.g., $\sim 0.02-0.10$ s per 1,000 rows), often about an order of magnitude faster than TabDPT for inference alone.

5 Conclusion and Future Work

We introduce TabDPT, an open tabular foundation model with a demonstrated ability to generalize on a range of unseen tabular datasets without additional tuning or adaptation. Our training approach provides an effective way to leverage real data with SSL to build robust TFMs. Models pre-trained with our procedure exhibit scaling laws with consistent improvements from both data and model size, analogous to foundation models in other domains. Given the practical ease of use and broad applicability of foundation models, we believe that these contributions will advance their adoption as an alternative to individually trained models in tabular domains.

While TabDPT demonstrates strong performance, opportunities remain to further enhance TFMs – including TabDPT – and address current limitations. (i) Preliminary experiments using feature name embeddings led to overfitting. Expanding training data with more diverse tables containing free-form text may mitigate this and improve textual integration. (ii) TabDPT is designed for rectangular datasets with i.i.d. rows and does not explicitly model temporal dependencies, distribution shifts, or hierarchical structures. Methodological improvements could help overcome these constraints; e.g., ideas similar to Hoo et al. [32] can help tackle temporal dependencies. (iii) TabDPT could be enhanced with techniques from works such as Ma et al. [41] and van Breugel et al. [57] to complement its discriminative capabilities with generative modeling and density estimation.

References

- [1] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. In *Advances in Neural Information Processing Systems*, 2021.
- [2] Dara Bahri, Heinrich Jiang, Yi Tay, and Donald Metzler. SCARF: Self-supervised contrastive learning using random feature corruption. In *International Conference on Learning Representations*, 2022.
- [3] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975.
- [4] Tamay Besiroglu, Ege Erdil, Matthew Barnett, and Josh You. Chinchilla scaling: A replication attempt. *arXiv*: 2404.10102, 2024.
- [5] Bernd Bischl, Bernd Bischl, Giuseppe Casalicchio, Matthias Feurer, Pieter Gijsbers, Frank Hutter, Michel Lang, Rafael Gomes Mantovani, Jan van Rijn, and Joaquin Vanschoren. OpenML benchmarking suites. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.
- [6] Meriem Boubdir, Edward Kim, Beyza Ermis, Sara Hooker, and Marzieh Fadaee. Elo uncovered: Robustness and best practices in language model evaluation. In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics*, 2023.
- [7] Felix den Breejen, Sangmin Bae, Stephen Cha, and Se-Young Yun. Fine-tuned in-context learning transformers are excellent tabular data classifiers. *arXiv*: 2405.13396, 2024.
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Advances in Neural Information Processing Systems, 2020.

- [9] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings* of the 22nd ACM SigKDD International Conference on Knowledge Discovery and Data Mining, 2016.
- [10] Aaron Defazio, Harsh Mehta, Konstantin Mishchenko, Ahmed Khaled, and Ashok Cutkosky. The road less scheduled. In *Advances in Neural Information Processing Systems*, 2024.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [13] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The Faiss library. *arXiv:* 2401.08281, 2024.
- [14] Arpad E Elo. The proposed USCF rating system, its development, theory, and applications. *Chess Life*, 22(8):242–247, 1967.
- [15] Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, Yanjun Qi, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, and Christos Faloutsos. Large language models (LLMs) on tabular data: Prediction, generation, and understanding – A survey. *Transactions on Machine Learning Research*, 2024.
- [16] Sebastian Felix Fischer, Matthias Feurer, and Bernd Bischl. OpenML-CTR23 A curated tabular regression benchmarking suite. In *AutoML Conference (Workshop)*, 2023.
- [17] Josh Gardner, Juan C Perdomo, and Ludwig Schmidt. Large scale transfer learning for tabular data via language modeling. In *Advances in Neural Information Processing Systems*, 2024.
- [18] Pieter Gijsbers, Marcos LP Bueno, Stefan Coors, Erin LeDell, Sébastien Poirier, Janek Thomas, Bernd Bischl, and Joaquin Vanschoren. AMLB: An AutoML benchmark. *Journal of Machine Learning Research*, 25(101):1–65, 2024.
- [19] Mark E Glickman. Example of the Glicko-2 system. *glicko.net*, 2012.
- [20] Yury Gorishniy, Ivan Rubachev, and Artem Babenko. On embeddings for numerical features in tabular deep learning. In *Advances in Neural Information Processing Systems*, 2022.
- [21] Yury Gorishniy, Ivan Rubachev, Nikolay Kartashev, Daniil Shlenskii, Akim Kotelnikov, and Artem Babenko. TabR: Tabular deep learning meets nearest neighbors. In *International Conference on Learning Representations*, 2024.
- [22] Yury Gorishniy, Akim Kotelnikov, and Artem Babenko. TabM: Advancing tabular deep learning with parameter-efficient ensembling. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [23] Aaron Grattafiori et al. The Llama 3 herd of models. arXiv: 2407.21783, 2024.
- [24] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? In Advances in Neural Information Processing Systems, 2022.
- [25] Sungwon Han, Jinsung Yoon, Sercan Ö Arik, and Tomas Pfister. Large language models can automatically engineer features for few-shot tabular learning. In *International Conference on Machine Learning*, 2024.
- [26] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

- [27] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. TabLLM: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*, 2023.
- [28] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models. In Advances in Neural Information Processing Systems, 2022.
- [29] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A transformer that solves small tabular classification problems in a second. In *International Conference on Learning Representations*, 2023.
- [30] Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeister, and Frank Hutter. Accurate predictions on small data with a tabular foundation model. *Nature*, 637(8045):319–326, 2025.
- [31] David Holzmüller, Léo Grinsztajn, and Ingo Steinwart. Better by default: Strong pre-tuned MLPs and boosted trees on tabular data. In *Advances in Neural Information Processing Systems*, 2024.
- [32] Shi Bin Hoo, Samuel Müller, David Salinas, and Frank Hutter. The tabular foundation model TabPFN outperforms specialized time series forecasting models based on simple features. In *NeurIPS Workshop on Time Series in the Age of Large Models*, 2024.
- [33] Kyle Hsu, Sergey Levine, and Chelsea Finn. Unsupervised learning via meta-learning. In *International Conference on Learning Representations*, 2019.
- [34] Xin Huang, Ashish Khetan, Milan Cvitkovic, and Zohar Karnin. TabTransformer: Tabular data modeling using contextual embeddings. *arXiv*: 2012.06678, 2020.
- [35] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv*: 2001.08361, 2020.
- [36] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In Advances in Neural Information Processing Systems, volume 30, 2017.
- [37] Myung Jun Kim, Leo Grinsztajn, and Gael Varoquaux. CARTE: Pretraining and transfer for tabular learning. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [38] Xiaofeng Lin, Chenheng Xu, Matthew Yang, and Guang Cheng. CTSyn: A foundation model for cross tabular data generation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- [40] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022.
- [41] Junwei Ma, Apoorv Dankar, George Stein, Guangwei Yu, and Anthony Caterini. TabPFGen Tabular data generation with TabPFN. In *NeurIPS Second Table Representation Learning Workshop*, 2023.
- [42] Kushal Majmundar, Sachin Goyal, Praneeth Netrapalli, and Prateek Jain. MET: Masked encoding for tabular data. In NeurIPS First Table Representation Learning Workshop, 2022.

- [43] Duncan McElfresh, Sujay Khandagale, Jonathan Valverde, Vishak Prasad C, Ganesh Ramakrishnan, Micah Goldblum, and Colin White. When do neural nets outperform boosted trees on tabular data? In *Advances in Neural Information Processing Systems*, 2023.
- [44] Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers can do Bayesian inference. In *International Conference on Learning Representations*, 2022.
- [45] Jaehyun Nam, Jihoon Tack, Kyungmin Lee, Hankook Lee, and Jinwoo Shin. STUNT: Few-shot tabular learning with self-generated tasks from unlabeled tables. In *International Conference on Learning Representations*, 2023.
- [46] 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. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.
- [47] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, 2018.
- [48] Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. TabICL: A tabular foundation model for in-context learning on large data. In *International Conference on Machine Learning*, 2025.
- [49] Shourav B. Rabbani, Ibna Kowsar, and Manar D. Samad. Transfer learning of tabular data by finetuning large language models. In *International Conference on Electrical and Computer Engineering*, 2024.
- [50] Ivan Rubachev, Artem Alekberov, Yury Gorishniy, and Artem Babenko. Revisiting pretraining objectives for tabular deep learning. *arXiv*: 2207.03208, 2022.
- [51] Maximilian Schambach, Dominique Paul, and Johannes Otterbach. Scaling experiments in self-supervised cross-table representation learning. In *NeurIPS Second Table Representation Learning Workshop*, 2023.
- [52] Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models' sensitivity to spurious features in prompt design or: How I learned to start worrying about prompt formatting. In *International Conference on Machine Learning*, 2024.
- [53] Yi Sui, Tongzi Wu, Jesse Cresswell, Ga Wu, George Stein, Xiaoshi Huang, Xiaochen Zhang, and Maksims Volkovs. Self-supervised representation learning from random data projectors. In *International Conference on Learning Representations*, 2024.
- [54] Yiming Sun, Xumeng Wen, Shun Zheng, Xiaowei Jia, and Jiang Bian. Scaling generative tabular learning for large language models. In *NeurIPS 2024 Third Table Representation Learning Workshop*, 2024.
- [55] Valentin Thomas, Junwei Ma, Rasa Hosseinzadeh, Keyvan Golestan, Guangwei Yu, Maksims Volkovs, and Anthony Caterini. Retrieval & fine-tuning for in-context tabular models. In *Advances in Neural Information Processing Systems*, 2024.
- [56] Boris van Breugel and Mihaela van der Schaar. Why tabular foundation models should be a research priority. In *International Conference on Machine Learning*, 2024.
- [57] Boris van Breugel, Jonathan Crabbé, Rob Davis, and Mihaela van der Schaar. LaTable: Towards large tabular models. *arXiv*: 2406.17673, 2024.
- [58] Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. OpenML: Networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014.
- [59] Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Arno Solin, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *Neural Networks*, 145:90–106, 2022.

- [60] Anton Voronov, Lena Wolf, and Max Ryabinin. Mind your format: Towards consistent evaluation of in-context learning improvements. In *Findings of the Association for Computational Linguistics: ACL*, 2024.
- [61] Xumeng Wen, Shun Zheng, Zhen Xu, Yiming Sun, and Jiang Bian. Scalable in-context learning on tabular data via retrieval-augmented large language models. *arXiv*:2502.03147, 2025.
- [62] Derek Qiang Xu, F Olcay Cirit, Reza Asadi, Yizhou Sun, and Wei Wang. Mixture of incontext prompters for tabular PFNs. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [63] Chao Ye, Guoshan Lu, Haobo Wang, Liyao Li, Sai Wu, Gang Chen, and Junbo Zhao. Towards cross-table masked pretraining for web data mining. In *Proceedings of the ACM Web Conference*, 2024.
- [64] Jinsung Yoon, Yao Zhang, James Jordon, and Mihaela Van der Schaar. VIME: Extending the success of self-and semi-supervised learning to tabular domain. In *Advances in Neural Information Processing Systems*, 2020.
- [65] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
- [66] Bingzhao Zhu, Xingjian Shi, Nick Erickson, Mu Li, George Karypis, and Mahsa Shoaran. XTab: Cross-table pretraining for tabular transformers. In *International Conference on Machine Learning*, 2023.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We were careful to accurately specify the scope of the paper in abstract and introduction, and all claims are supported in later sections.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Discussion of limitations is provided in the final section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results that require proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The experimental procedure to produce the main results and model is fully described across Section 3, Section 4, and the appendices. Code and model weights are provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Inference code including trained model weights can be found here, and the training code to reproduce experiments can be found here. All data is publicly available on OpenML.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be
 possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not
 including code, unless this is central to the contribution (e.g., for a new open-source
 benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Yes, full details are given across Section 3, Section 4, and the appendices.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Confidence intervals are provided where applicable.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: We describe the hardware used for generating results and provide practical analysis of the method's runtime, but we did not record the time used to generate all experimental results. Given that our hardware requirements—a single GPU—are very light for a foundation modelling paper, we do not see this as a major limitation.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We have reviewed the Code of Ethics and verified compliance.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Given the non-specific nature of the modelling discussed in the paper and the lack of specific links between tabular discriminative modelling and societal harms, we do not believe there are societal impacts that need discussion.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The data and models under consideration do not have a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Existing assets from OpenML are used, with the licence described and appropriate attribution given.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: New assets are provided in a documented code repository.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The research did not involve crowdsourcing or human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The research did not involve crowdsourcing or human subjects.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

A Bitter Lessons

Throughout the development of TabDPT we tried many variations to improve performance and/or scalability. Some were successful while others did not lead to an improvement. We list some of these variations here to facilitate future research on TabDPT and related architectures. Overall, our findings align with *The Bitter Lesson*² where efficient use of computation and access to high-quality data are much more important for driving performance than architectural manipulations.

- Different pre-processing techniques that were more robust to outliers, or variants of soft clipping, resulted in no improvement. More advanced methods, such as Robust Scaler and Power Transform, only ended up slowing the training process.
- Class embeddings (either through a separate network or by using class "tokens" in the transformer layer) and computing various similarity metrics between query and class embeddings in a proto-network manner, with the aim of adapting to any number of classes, hurt the performance, especially on real data.
- Different embeddings for y_{ctx} , including a dense layer for regression and a dictionary of $C_{\text{max}} \times d$ embeddings, with the rationale of informing the model about the task, did not lead to performance improvements in large models with sufficient data.
- Specialized tokens for NaN encoding did not improve performance compared to replacing NaNs with mean values (which are zero after preprocessing). Additionally, appending binary features to differentiate actual zeros from NaNs (indicating that the cell was replaced), effectively doubling the number of features, also failed to improve performance.
- Architectures encoding cells as "tokens", with vertical and horizontal attention, similar to spatial and temporal attention in videos, proved more memory intensive. While equivariance to feature order is desirable, processing tensors of size (B, N, f, d) where B is batch size, N is the number of rows, f the number of features, and d the embedding dimension uses much more memory. The simpler architecture with tensors of size (B, N, d) permits a higher embedding dimension d. While Hollmann et al. [30] is able to make this architecture work, we suspect that the differences between synthetic and real data are enough to change which architectures are performant.

We would also like to emphasize that these *Bitter Lessons* reflect our own experience in training an ICL-based TFM with real data. Your mileage may vary.

B Training Datasets

Figure B.1 provides a summary of the sizes and domains of the training datasets, and Table B.1 provides a full list of the datasets. Note that 93 datasets have classification targets, 29 datasets have regression targets, and 1 does not have a default target defined. However, we generate both classification and regression targets for each dataset by applying the SSL procedure described in Section 3.

Table B.1: Details for all training datasets: OpenML Dataset ID, name, dimensions (rows, features, cells), percent of missing cells, target type (classification/regression), domain.

OpenML Dataset ID	Name	# rows	# feat.	# cells	% miss.	Target type	Domain
24	mushroom	8124	22	187K	1.4	Class.	Biology/ecology
30	page-blocks	5473	10	60K	0.0	Class.	Vision/audio/text features
184	kropt	28056	6	196K	0.0	Class.	Deterministic and simulated
273	IMDB.drama	120919	1001	121M	0.0	Class.	Other or not provided
312	scene	2407	299	722K	0.0	Class.	Vision/audio/text features
375	JapaneseVowels	9961	14	149K	0.0	Class.	Vision/audio/text features
382	ipums_la_97-small	7019	60	428K	11.4	Class.	Financial/demographic
389	fbis.wc	2463	2000	4.9M	0.0	Class.	Vision/audio/text features
396	la1s.wc	3204	13195	42M	0.0	Class.	Vision/audio/text features
802	pbcseq	1945	18	37K	3.2	Class.	Medical/human sensor
816	puma8NH	8192	8	74K	0.0	Class.	Deterministic and simulated
821	house_16H	22784	16	387K	0.0	Class.	Financial/demographic

²http://www.incompleteideas.net/IncIdeas/BitterLesson.html

OpenML Dataset ID	Name	# rows	# feat.	# cells	% miss.	Target type	Domain
343	house_8L	22784	8	205K	0.0	Class.	Financial/demographic
346	elevators	16599	18	315K	0.0	Class.	Other or not provided
371	pollen	3848	5	23K	0.0	Class.	Biology/ecology
930	colleges_usnews	1302	33	44K	18.2	Class.	Other or not provided
066	analcatdata halloffame	1340	16	23K	0.1	Class.	Other or not provided
81	kdd internet usage	10108	68	697K	0.4	Class.	Financial/demographic
.002	ipums_la_98-small	7485	55	419K	7.9	Class.	Financial/demographic
018	ipums_la_99-small	8844	56	504K	7.0	Class.	Financial/demographic
036	sylva_agnostic	14395	216	3.1M	0.0	Class.	Biology/ecology
037	ada_prior	4562	14	68K	0.1	Class.	Financial/demographic
1043	ada_agnostic	4562	48	224K	0.0	Class.	Financial/demographic
043	eye_movements	10936	27	306K	0.0	Class.	Medical/human sensor
1111		50000	230	12M	61.9	Class.	Human behaviour
1112	KDDCup09_appetency KDDCup09_churn	50000	230	12M	61.9	Class.	Industrial/operational
1112	1 —			1.1M	0.0	Class.	
	musk	6598	167				Other science
118	chess	28056	6	196K	0.0	Class.	Deterministic and simulat
120	MagicTelescope	19020	10	209K	0.0	Class.	Physics/astronomy
130	OVA_Lung	1545	10935	17M	0.0	Class.	Biology/ecology
142	OVA_Endometrium	1545	10935	17M	0.0	Class.	Biology/ecology
169	airlines	539383	7	4.3M	0.0	Class.	Industrial/operational
444	PizzaCutter3	1043	37	40K	0.0	Class.	Other or not provided
453	PieChart3	1077	37	41K	0.0	Class.	Other or not provided
457	amazon-commerce-reviews	1500	10000	15M	0.0	Class.	Vision/audio/text features
459	artificial-characters	10218	7	82K	0.0	Class.	Deterministic and simulat
466	cardiotocography	2126	35	77K	0.0	Class.	Medical/human sensor
471	eeg-eye-state	14980	14	225K	0.0	Class.	Medical/human sensor
476	gas-drift	13910	128	1.8M	0.0	Class.	Other science
477	gas-drift-different- concentrations	13910	129	1.8M	0.0	Class.	Other science
479	hill-valley	1212	100	122K	0.0	Class.	Deterministic and simulat
481	kr-vs-k	28056	6	196K	0.0	Class.	Deterministic and simular
483	ldpa	164860	7	1.3M	0.0	Class.	Medical/human sensor
493	one-hundred-plants-texture	1599	64	104K	0.0	Class.	Biology/ecology
503	spoken-arabic-digit	263256	14	3.9M	0.0	Class.	Vision/audio/text features
507	twonorm	7400	20	155K	0.0	Class.	Deterministic and simulat
509	walking-activity	149332	4	747K	0.0	Class.	Medical/human sensor
.567	poker-hand	1025009	10	11 M	0.0	Class.	Deterministic and simular
568	nursery	12958	8	117K	0.0	Class.	Financial/demographic
596	covertype	581012	54	32M	0.0	Class.	Biology/ecology
8050	QSAR-TID-11	5742	1024	5.9M	0.0	Reg.	Medical/human sensor
3277	QSAR-TID-10980	5766	1024	5.9M	0.0	Reg.	Medical/human sensor
1135	Amazon_employee_access	32769	9	328K	0.0	Class.	Industrial/operational
1535	Census-Income	299285	42	13M	0.0	None	Financial/demographic
1549	Buzzinsocialmedia_Twitter	583250	77	45M	0.0	Reg.	Human behaviour
23380	cjs	2796	33	95K	73.8	Class.	Biology/ecology
23512	higgs	98050	28	2.8M	0.0	Class.	Physics/astronomy
0536	SpeedDating	8378	120	1.0M	1.8	Class.	Human behaviour
0646	GAMETES_Epistasis_2-Way_	1600	20	34K	0.0	Class.	Biology/ecology
10670	20atts_0.1H_EDM-1_1	10020	10	2001/	0.0	C1	Dli/
0679	magic	19020	10	209K	0.0	Class.	Physics/astronomy
0680	mofn-3-7-10	1324	10	15K	0.0	Class.	Other or not provided
0685	shuttle	58000	9	580K	0.0	Class.	Physics/astronomy
0706	parity5_plus_5	1124	10	12K	0.0	Class.	Deterministic and simular
0733	yeast	1269	8	11K	0.0	Class.	Biology/ecology
0900	Satellite	5100	36	189K	0.0	Class.	Physics/astronomy
1138	APSFailure	76000	170	13M	8.3	Class.	Industrial/operational
1142	christine	5418	1636	8.9M	0.0	Class.	Other or not provided
1143	jasmine	2984	144	433K	0.0	Class.	Other or not provided
1144	madeline	3140	259	816K	0.0	Class.	Other or not provided
1145	philippine	5832	308	1.8M	0.0	Class.	Other or not provided
1146	sylvine	5124	20	108K	0.0	Class.	Other or not provided
1147	albert	425240	78	34M	8.2	Class.	Other or not provided
1150	MiniBooNE	130064	50	6.6M	0.0	Class.	Physics/astronomy
1156	ada	4147	48	203K	0.0	Class.	Other or not provided
1159	guillermo	20000	4296	86M	0.0	Class.	Other or not provided
1161	riccardo	20000	4296	86M	0.0	Class.	Other or not provided
1162	kick	72983	32	2.4M	6.4	Class.	Industrial/operational
							Other or not provided
1163	dilbert	10000	2000	20M	0.0	Class.	
1164	fabert	8237	800	6.6M	0.0	Class.	Other or not provided
11165	robert	10000	7200	72M	0.0	Class.	Other or not provided
11166	volkert	58310	180	11M	0.0	Class.	Other or not provided
11167	dionis	416188	60	25M	0.0	Class.	Other or not provided
11168	jannis	83733	54	4.6M	0.0	Class.	Other or not provided
11169	helena	65196	27	1.8M	0.0	Class.	Other or not provided
11434	Click_prediction_small	39948	11	479K	0.0	Class.	Human behaviour
11540	black_friday	166821	9	1.7M	0.0	Reg.	Human behaviour
+1340							
11940 11980	SAT11-HAND-runtime-	4440	116	519K	5.3	Reg.	Computing

OpenML Dataset ID	Name	# rows	# feat.	# cells	% miss.	Target type	Domain
42563	house_prices_nominal	1460	79	117K	6.0	Reg.	Financial/demographic
42572	Santander_transaction_value	4459	4991	22M	0.0	Reg.	Human behaviour
42705	Yolanda	400000	100	40M	0.0	Reg.	Other or not provided
42724	OnlineNewsPopularity	39644	59	2.4M	0.0	Reg.	Human behaviour
42727	colleges	7063	44	318K	33.5	Reg.	Other or not provided
42728	Airlines_DepDelay_10M	10000000	9	100M	0.0	Reg.	Industrial/operational
42730	us_crime	1994	126	253K	15.6	Reg.	Financial/demographic
42732	sf-police-incidents	2215023	8	20M	0.0	Class.	Human behaviour
42734	okcupid-stem	50789	19	1.0M	16.0	Class.	Human behaviour
42742	porto-seguro	595212	57	35M	2.5	Class.	Human behaviour
42746	KDDCup99	4898431	41	206M	0.0	Class.	Computing
43071	MIP-2016-Reg.	1090	144	158K	0.0	Reg.	Computing
43072	KDDCup09-Upselling	50000	14891	745M	2.6	Class.	Human behaviour
44055	analcatdata_supreme	4052	7	32K	0.0	Reg.	Other or not provided
44056	visualizing_soil	8641	4	43K	0.0	Reg.	Biology/ecology
44061	Mercedes_Benz_Greener_ Manufacturing	4209	359	1.5M	0.0	Reg.	Industrial/operational
44063	Bike_Sharing_Demand	17379	11	209K	0.0	Reg.	Human behaviour
44065	nyc-taxi-green-dec-2016	581835	16	9.9M	0.0	Reg.	Human behaviour
44068	particulate-matter-ukair-2017	394299	6	2.8M	0.0	Reg.	Other or not provided
44069	SGEMM_GPU_kernel_ performance	241600	9	2.4M	0.0	Reg.	Computing
44089	credit	16714	10	184K	0.0	Class.	Financial/demographic
44122	pol	10082	26	272K	0.0	Class.	Industrial/operational
44136	wine_quality	6497	11	78K	0.0	Reg.	Human behaviour
44137	Ailerons	13750	33	468K	0.0	Reg.	Other or not provided
44145	sulfur	10081	6	71K	0.0	Reg.	Other science
45020	default-of-credit-card-clients	13272	20	279K	0.0	Class.	Financial/demographic
45022	Diabetes130US	71090	7	569K	0.0	Class.	Medical/human sensor
45026	heloc	10000	22	230K	0.0	Class.	Financial/demographic
45032	yprop_4_1	8885	42	382K	0.0	Reg.	Medical/human sensor
45038	road-safety	111762	32	3.7M	0.0	Class.	Human behaviour
45039	compas-two-years	4966	11	60K	0.0	Class.	Human behaviour
45041	topo_2_1	8885	255	2.3M	0.0	Reg.	Medical/human sensor
45043	seattlecrime6	52031	4	260K	0.0	Reg.	Human behaviour
45045	delays_zurich_transport	5465575	11	66M	0.0	Reg.	Industrial/operational
45046	Allstate_Claims_Severity	188318	124	24M	0.0	Reg.	Industrial/operational
45047	Airlines_DepDelay_1M	1000000	5	6.0M	0.0	Reg.	Industrial/operational

B.1 Contamination Analysis

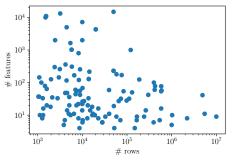
To ensure that the datasets used for training do not contain any information about the evaluation data, we extract a range of metadata from each dataset and compare them across all pairs of training and evaluation datasets. This includes: (i) dataset names, (ii) hashes of dataset files, (iii) numbers of columns and rows, (iv) target mean and variance, (v) mean, variance, skew, and kurtosis of each feature, and (vi) coefficients of a univariate linear fit between each feature and the target if available. To allow for efficient pairwise comparisons between all features in all datasets, we use k-d trees [3] constructed for each dataset that contain the feature statistics. Any dataset pairs with unusual similarities were manually evaluated and those found to be related were removed from training.

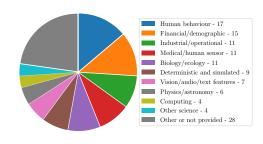
C Model Architecture and Hyperparameters

C.1 Architecture Details

The model architecture (see Figure 2b) is comprised of input embedding functions, multiple transformer encoder layers, and task-specific output heads. The key architectural parameters are summarized in Tables C.1 and C.2.

Preprocessing We deliberately do minimal pre-processing of the data to ensure that our approach has wide applicability. All columns containing non-numerical values are mapped to integers using scikit-learn's [46] LabelEncoder function. The table is then standardized to 0 mean and unit variance, and outliers beyond 10 are clipped. After retrieval, we obtain a local context X_{ctx} and their labels y_{ctx} . X_{ctx} is standardized before the forward pass to avoid distribution shifts, and y_{ctx} is also standardized for the same reason if it is a regression target.





- (a) Sizes of training datasets.
- (b) Breakdown of training dataset domains.

Figure B.1: Breakdown of training datasets by size and domain.

Table C.1: Architectural Parameters

Parameter	Value
Number of Attention Heads	4
Feedforward Network Factor	2
Maximum Number of Classes	10
Maximum Number of Features	100
Normalization First	Yes
Dropout Rate	0.0

Retrieval We use the faiss library³ for fast retrieval. All retrieval is done in the raw feature space after preprocessing, as in [55].

Missing Value Encoding We experimented with several strategies for missing value handling, including concatenating a binary missing-or-not mask, however the improvement was minimal at nearly double the compute cost. Hence, we opt for a simple strategy to zero out missing values and let the model learn how to deal with incomplete inputs. Note that zeroing out is done post normalization, meaning missing values are replaced with the mean.

Optimizer We use the Schedule Free optimizer from Defazio et al. [10] with AdamW [39]. We observed significant increase in performance and optimization speed compared to a cosine scheduler. Label smoothing and weight decay are applied throughout training and are important for smooth convergence. By default we set a learning rate of 5×10^{-4} and weight decay of 5×10^{-2} with label smoothing of 0.1. The batch size is set to 256 and both context and query lengths are set to 1024. Model parameters are kept in brain float 16-bit (bfloat16) format.

D Pseudo-Code for Training Algorithms

In this section, we list the pseudo-code for our training procedure. In Code Block 1, we show the PyTorch Dataloader component. In the initialization phase, we first process the downloaded data and features by filling in missing values with the mean column values and creating a faiss index for fast retrieval. Next, in each worker within the getitem() function, we sample a random dataset, then we sample a random query within the dataset. After that, we mask out the target column and retrieve its approximate neighbours. Then we process the features and targets by random sub-sampling and random partitioning.

In Code Block 2, within each training step, we partition both the data X and targets y into context and query points by sampling an integer uniformly from 10 to its total length (inclusive of start point but exclusive of endpoint). We call this random evaluation position eval_pos in the code block. The points to the left of the evaluation position are then taken as context (i.e., y_ctx), and the points to the right of the evaluation position are taken as queries (i.e., y_qy). Finally, we calculate the appropriate loss depending on the task and optimize the network.

³https://github.com/facebookresearch/faiss

Table C.2: Number of Layers and Transformer Dimensions

Number of Layers	Transformer Dimension
3	32
4	64
5	96
6	256
10	384
12	512
16	768

Code Block 1: Pytorch Dataloader

```
from torch.utils.data import Dataset
   import numpy as np
2
3
   import random
5
   class TrainingDataset(Dataset):
       def __init__(self, dataset_ids):
6
           self.datasets = []
7
           for dataset_id in dataset_ids:
8
                X <- download dataset using dataset_id</pre>
10
                X <- process features of X (handle missing values,
                   scale)
                knn_index <- compute knn index using FAISS
11
12
                self.dataset.append([X, knn_index])
13
       # Random column subsample and shuffling
14
15
       def create_random_columns(self, X):
           N, F = X.shape
16
           num_features_sampled = random.randint(F // 2, F)
17
18
           random_features_indices = np.random.choice(F,
               num_features_sampled, replace=False)
           return X[:, random_features_indices]
19
20
       # Generate a random classification or regression target for
21
           training
       def generate_random_target(self, y, cls_threshold=10):
22
           if len(np.unique(y)) > cls_threshold:
23
                # if there are more than 10 unique values in the
                   target, we keep it as regression 70% of the time
                if np.random.rand() > 0.3:
25
                    return y, "regression"
26
                else:
27
                    # sample a random number of classes by binning
28
                       and divide into classes
                    num_class = np.random.randint(2, cls_threshold)
29
                    cls_boundary = np.random.choice(sorted(np.unique))
30
                        (y))[1:-1], num_class-1, replace=False)
                    y = (y[:, None] > cls_boundary[None, :]).sum(1)
31
32
                    y <- label encode, shuffle y
                    return y, "classification'
33
           else:
34
                assert len(np.unique(y)) > 1
35
                y <- label encode, shuffle y
36
                return y, "classification"
37
38
39
       # Generate a sample for retrieval
       def __getitem__(_):
40
           # sample a random dataset
41
           sample_id = np.random.choice(len(self.dataset), 1)[0]
42
           X_sample, knn_index_sample = self.dataset[sample_id]
```

```
N, F = X_sample.shape
44
45
            # sample a random query from the dataset
46
           x_q = X_sample[random.randint(0, N-1)].copy()
47
48
           # sample a random column to be the target
            target_idx = random.randint(0, F-1)
50
51
            # retrieve approximate neighbours using x_q with
52
                target_idx masked
53
            x_q[:, target_idx] = 0
            X_nn <- find k neighbours using knn_index_sample with</pre>
54
               x_q as query
            y_nn = X_nn[:, target_idx]
55
            X_nn = np.delete(X_nn, target_idx, axis=1)
56
57
           # subsample and shuffle features
58
           X_nn = self.create_random_columns(X_nn)
59
60
61
            # generate random target and task
           y_nn, task = self.generate_random_target(y_nn)
62
63
           return X_nn, y_nn, task
64
```

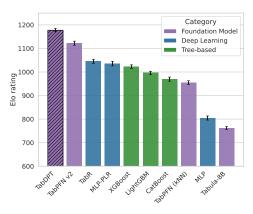
Code Block 2: Training Loop

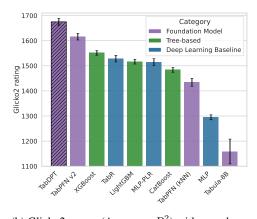
```
model = Transformer()
2
3
   optimizer = schedulerfree.AdamWScheduleFree()
   for epoch in range(num_epochs):
6
     model.train()
     for X, y, task in train_loader:
7
8
        eval_pos = random.randint(10, len(y))
        y_ctx, y_qy = y[:eval_pos], y[eval_pos:]
9
        y_ctx = zero_pad(y_ctx, N_qy, dim=1)
10
11
        output = model(torch.cat(X, y_ctx))
12
13
        if task == "classification":
14
           loss = cross_entropy_loss(output, y_qy)
15
        elif task == "regression":
16
           loss = mse_loss(ouput, y_qy)
17
18
19
        opitmizer.zero_grad()
        loss.backward()
20
        optimizer.step()
21
```

E Elo and Glicko2 Ratings

We expand the pairwise method comparison with Elo [14] and Glicko2 [19] ratings. For the Elo calculation, we estimate uncertainty by bootstrapping over match order permutations [6]. Glicko2, on the other hand, provides uncertainty by design and is less sensitive to match order in our experiments.

In Figures E.1a and E.1b, we report the Elo and Glicko2 scores respectively. The results are consistent between the two plots, with TabDPT performing best on both metrics, followed by the leading TFM baseline TabPFN v2.





(a) Elo scores (Accuracy, \mathbb{R}^2) with error bars.

(b) Glicko2 scores (Accuracy, \mathbb{R}^2) with error bars.

Figure E.1: Duel-based metrics computed on accuracy and \mathbb{R}^2 scores. (a) Elo ratings. (b) Glicko2 ratings.

F Additional Results

F.1 Additional Results by Dataset Statistics

In this section, we analyze the performance of all methods, bucketed by different characteristics of the benchmark datasets. In particular, we analyze performance by number of rows, number of columns, categorical fraction, and missing fraction. We see that TabDPT is robust across various dataset characteristics, with a very slight relative decrease in performance for very large CC18 datasets; this can be mitigated by fine-tuning as suggested in Thomas et al. [55].

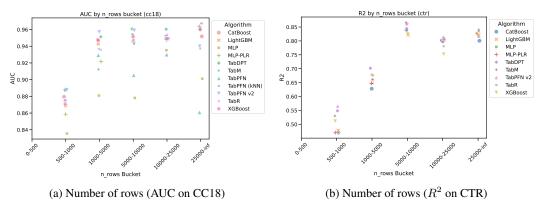


Figure F.1: Comparison for Number of Rows.

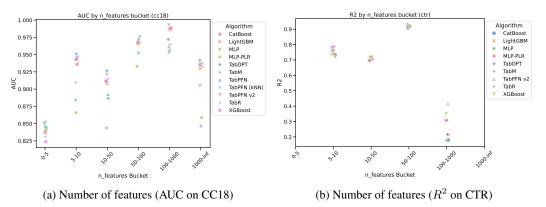


Figure F.2: Comparison for Number of Features.

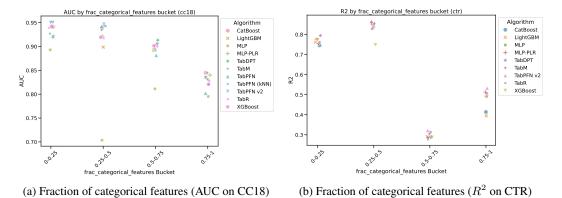


Figure F.3: Comparison for Fraction of Categorical Features.

G Critical Difference Diagrams

Critical difference diagrams computed over CC18 and CTR23 fold results are given in Figures G.1, G.2, G.3, and G.4. The difference in confidence intervals compared to Table 1 are due both to how we compute the ranks and their confidence (using IQM, based on the recommendations from [1]) which computes uncertainty over different realizations of the suite. In contrast, the critical diagrams compare the uncertainty, with multiple hypothesis testing adjustments, over individual dataset realizations.

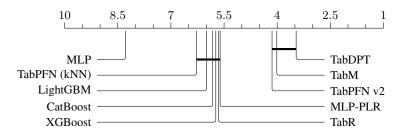


Figure G.1: AUC critical difference diagram

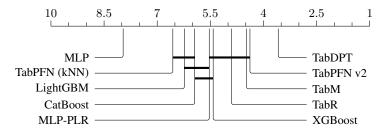


Figure G.2: Accuracy critical difference diagram

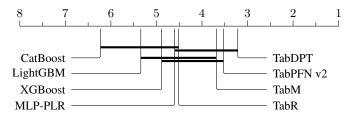


Figure G.3: Correlation critical difference diagram

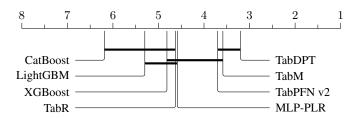


Figure G.4: R2 critical difference diagram