
TabSTAR: A Tabular Foundation Model for Tabular Data with Text Fields

Alan Arazi Eilam Shapira Roi Reichart
{alanarazi7, eilam.shapira, roireichart}@gmail.com
Technion - IIT

Abstract

While deep learning has achieved remarkable success across many domains, it has historically underperformed on tabular learning tasks, which remain dominated by gradient boosting decision trees. However, recent advancements are paving the way for Tabular Foundation Models, which can leverage real-world knowledge and generalize across diverse datasets, particularly when the data contains free-text. Although incorporating language model capabilities into tabular tasks has been explored, most existing methods utilize static, target-agnostic textual representations, limiting their effectiveness. We introduce TabSTAR: a Tabular Foundation Model with Semantically Target-Aware Representations. TabSTAR is designed to enable transfer learning on tabular data with textual features, with an architecture free of dataset-specific parameters. It unfreezes a pretrained text encoder and takes as input target tokens, which provide the model with the context needed to learn task-specific embeddings. TabSTAR achieves state-of-the-art performance for both medium- and large-sized datasets across known benchmarks of classification tasks with text features, and its pretraining phase exhibits scaling laws in the number of datasets, offering a pathway for further performance improvements.¹

1 Introduction

In recent years, deep learning has profoundly reshaped research and practice in computer vision [50, 66, 33, 17] and natural language processing [56, 5, 79, 16, 13]. This transformation was notably accelerated by the rise of foundation models [8, 88, 4], capable of cross-modal understanding and generalization from massive pretraining across heterogeneous data sources. Importantly, they enabled an end-to-end approach that outperformed previous modular alternatives [71, 1]. Moreover, deep learning models excel at transfer learning [91], generalizing from their pretraining data to new tasks. Their strength, combined with techniques like In-Context Learning (ICL) [13] and Parameter-Efficient Fine-Tuning (PEFT) [41], has enabled rapid adaptation to new tasks with only limited labeled data.

Despite this progress, deep learning has historically lagged behind gradient-boosted decision trees (GBDTs) on tabular data [12, 15, 60], in both classification and regression tasks [65, 11, 31, 55]. The heterogeneity of tabular data, which lacks the spatial locality of images or the sequential order of text, makes it more challenging for deep models to learn. Consequently, GBDTs have remained the de facto standard for tabular learning, offering strong out-of-the-box performance, computational efficiency, and built-in inductive biases (e.g., robustness to skewed feature distributions and automatic feature selection) that make them especially well-suited to heterogeneous datasets [31]. Nonetheless, GBDTs cannot be pretrained to reuse strong representations for downstream tasks. This limitation becomes critical in low-data settings like those often found in healthcare applications [52]. Crucially, they must rely on external embedding models to process unstructured data types like text and images, yielding fixed feature representations that cannot be finetuned for a specific prediction task.

¹Code is available at <https://github.com/alanarazi7/TabSTAR>.

Table 1: A binary classification toy dataset for hospital patient release outcomes. *Decision* is the target variable. *Age* (numerical), *Department* (high-cardinality), and *Report* (textual) are the features.

Age	Department	Report	Decision
45	Cardiology	Mild chest discomfort.	Released
62	Neurology	Complaints of headache and occasional dizziness.	Hospitalized
38	Oncology	Completed treatment cycle without adverse reactions.	Released
55	Neurology	Reports episodes of vertigo and memory lapses.	Hospitalized

The emerging field of Tabular Foundation Models (TFMs) has begun addressing these shortcomings, introducing powerful cross-dataset learning strategies [84, 47, 38]. However, the flagship model TabPFN-v2 [38] still handles text inputs no more flexibly than conventional GBDTs. This design choice is not incidental; historically, tabular benchmarks have prioritized numerical datasets without free-text features, largely for ease of modeling and evaluation. A recent study [48] of mainstream tabular datasets benchmarks [25, 23, 87, 55] found that half of these datasets are more than 20 years old, being a poor representation of modern real-world data.

Real-world tabular datasets often include high-cardinality² and free-text features [14], illustrated by a toy example in Table 1. In such datasets, free-text features (e.g., *Report*) carry rich semantic information critical for tasks like predicting whether a patient will be discharged from the hospital or require continued care. Yet, most models encode them in a target-agnostic manner, delegating to a generic embedding that fails to capture task-specific nuances for predicting *Decision*. Crucially, that same embedding would have been used for a different target variable (e.g., *Treatment Cost*). Similarly, categorical features with dozens of unique values (e.g., *Department*) are difficult to encode efficiently without external knowledge, making naive approaches brittle and limiting generalization. Importantly, the column names, which could guide the model toward more effective representations, are typically ignored. Addressing these limitations is crucial for developing tabular models that leverage semantic information, transfer knowledge from many datasets, and generalize across domains.

In this paper, we introduce **TabSTAR**: a novel **Tabular Foundation Model with Semantically Target-Aware Representations**, designed explicitly for end-to-end handling of purely textual features. By integrating an unfrozen text encoder at its core, TabSTAR can optimize free-text feature representations, demonstrating their clear superiority over alternative frozen embedding approaches. Additionally, it introduces a novel approach of *target-aware tokens*, which inject semantic information about the target variable as part of the input, allowing for efficient parameter sharing and resulting in an architecture with no dataset-specific parameters (see Figure 1). TabSTAR’s training is highly efficient³ and its performance steadily improves with more pretraining data. Empirically, TabSTAR achieves state-of-the-art (SOTA) performance on classification datasets containing textual features, surpassing leading TFMs as well as GBDTs tuned for 4 hours.

2 Related Work

This section reviews prior work in supervised tabular learning. We begin with deep learning methods tailored for tabular data, which were applied to a single dataset. We then discuss cross-dataset transfer learning techniques that improve generalization by leveraging related datasets. Next, we cover the field of TFMs, which aim to generalize across diverse tasks and datasets through large-scale pretraining. Finally, we review recent work on applying large language models (LLMs) to tabular data and elaborate on existing AutoML [34] multimodal solutions. As we focus on supervised learning, we do not cover self-supervised methods [82, 6] and their downstream applications.

Deep Learning on a Single Tabular Dataset Several architectures have been proposed to enhance deep learning for tabular data [70, 45, 81, 85]. *TabNet* [3] and *TabTransformer* [43] introduced attention mechanisms into tabular deep learning, while *FT-Transformer* [26] and its improvement [27] jointly integrated numerical and categorical features into a transformer [79]. Other novel approaches leveraged inter-example information at inference time, with *SAINT* [69] proposing row-

²High-cardinality features are categorical columns with a large number of unique values.

³Pretraining within 48 hours on a single A40 GPU. Finetuning with PEFT for a low memory footprint.

level attention between examples, *Non-Parametric Transformers* [49] processing the entire dataset, including labels, in a single forward pass, and *TabR* [28] combining a k-nearest-neighbor mechanism with a traditional Multi-Layer Perceptron (MLP) architecture. Recent works such as *TabM* [29] and *RealMLP* [39] focused on refining MLPs without an attention component. Despite these innovations, single-dataset deep learning models have not yet convincingly outperformed GBDTs [65, 31, 64]. Furthermore, none of them addressed the challenge of modeling tabular data with rich textual features.

Cross-Dataset Transfer Learning Deep learning has been proven to shine when performing transfer learning in many machine learning domains [91]. Motivated by this success, [52, 89] proved that cross-dataset learning can boost single-dataset performance, but were limited to strict requirements such as partial overlap of feature names. To address this limitation, *TransTab* [83] integrated semantic understanding into feature tokenization, and *XTab* [90] pretrained a transformer backbone with dataset-specific parameters, proving that pretraining contributes to a stronger initialization for a downstream task. Despite their small scale, these studies demonstrated cross-dataset transfer learning’s potential, laying essential groundwork for the rise of TFM.

Tabular Foundation Models TFM represent an emerging paradigm in tabular learning. While the definition is still evolving, we adopt the framing proposed by [77], which identifies key desired characteristics of TFM: large-scale pretraining with adaptability to downstream tasks, mixed-type column support, cross-domain generalization, use of textual metadata,⁴ and column-order invariance.

TabPFN [36] is recognized as the first TFM, and its successor *TabPFN-v2* [38] currently sets the SOTA in tabular learning, becoming a popular approach for TFM [61, 21, 54]. *TabPFN-v2* was the first model to consistently outperform GBDTs on medium-sized datasets, by pretraining Bayesian Prior-Data Fitted Networks (PFNs) [58] on 130 million synthetic datasets. Using ICL at inference time, it accepts up to 10,000 examples as input and predicts without updating its weights. *TabPFN-v2* inspired *TabICL* [61] which improved scalability, and *TabDPT* [54], trained on real data. However, they all rely on off-the-shelf embeddings for text features like GBDTs, limiting their performance.

CM2 [86], *CARTE* [47], and *TP-BERTa* [84] represent a shift toward semantic tabular modeling, leveraging textual signals and external knowledge at a greater scale. Unlike prior methods, these models transfer knowledge via language representations. *CM2* was pretrained on over 2,000 datasets, but did not focus on free-text features and used static word embeddings without further finetuning them. *CARTE* encodes tables as star-shaped graphs, jointly representing features by their names and values, and applies attention over the graph to capture contextual relations. While effective for high-cardinality features, it lacks exposure to longer free-text fields during pretraining and was proven useful mainly for small datasets. *TP-BERTa* adapts *RoBERTa* [53] with intra-feature attention and a tokenization scheme that maps numerical values into discrete relative-magnitude bins, to address the weakness of language models when tokenizing numbers [75]. Although it performs well, its use of dataset-specific output layers limits scalability and complicates multi-task learning. Consequently, they trained two separate models,⁵ wasting potential for better cross-dataset learning. Notably, none of these approaches finetune semantic representations during downstream task training. In our work, we demonstrate that this is critical to align textual and tabular features.

Large Language Models for Tabular Data The remarkable success of LLMs is unprecedented [13, 59]. During the past years, several research attempts have tried to combine LLMs and tabular data. One line of work focuses on using LLMs directly for tabular prediction by converting tabular data into serialized text. *TabLLM* [35] assessed LLMs under few-shot scenarios, while *Tabula-8b* [24] finetuned the Llama 3-8B model extensively on tabular data. Although useful for few-shot learning, these models are computationally expensive,⁶ suboptimal for numerical features [75, 77], and potentially compromised on widely-used benchmarks due to prior exposure during training [9]. While current generations of LLMs weren’t adopted for tabular learning, their emergent knowledge from their pretraining could be crucial when textual features are present [77, 20]. Additionally, LLMs can be used in multiple aspects of tabular learning, as they seem to be promising synthetic data generators [10, 68], useful data cleaners [7], and clever feature engineers [37].

⁴Contextual information such as the dataset description, column names, and category names.

⁵One for classification and one for regression. A joint model for both tasks performed significantly worse.

⁶Llama 3-8b has orders of magnitude more parameters than TP-BERTa, which has roughly 110M parameters.

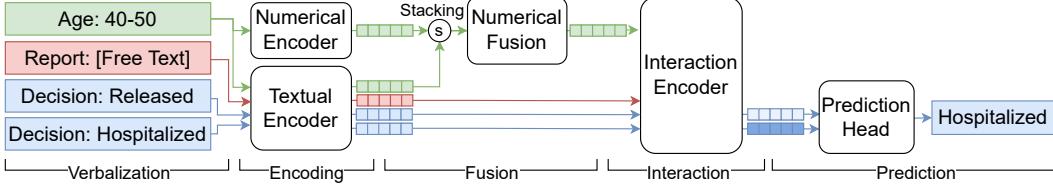


Figure 1: The TabSTAR architecture illustrated with our toy dataset. The model processes **numerical features**, **textual features**, and **all possible target values** for classification.

Multimodal AutoML Historically, textual tabular datasets have been largely overlooked in classical tabular benchmarks. However, the AutoML [34] community has made significant progress in developing multimodal solutions. In particular, *AutoGluon* [19] introduced the *AutoML Multimodal Benchmark* [63], initially focusing on text features and later evolving into *AutoGluon-Multimodal* [73, 72], which incorporates images as well. This powerful AutoML framework can fuse text and image foundation models with tabular models and ensemble multiple models through a meta-learning approach [22], making it one of the few systems able to refine static textual representations via joint learning. Nevertheless, this line of work should not be seen as a single model but rather as a highly optimized, production-ready system. According to the authors, it is "a collection of tricks that significantly enhance performance" [73], establishing itself as a robust approach for multimodal, multi-model tabular learning. However, this line of work remains somewhat orthogonal to the development of novel TFM.

3 TabSTAR

In this section, we introduce TabSTAR: a Tabular Foundation Model with Semantically Target-Aware Representations. Our training framework consists of two stages: (1) **Pretraining**, where the model is pretrained over a corpus of tabular datasets⁷ in a multi-task regime, mixing classification with regression tasks, then (2) **Finetuning**, where the pretrained model is further trained with *LoRA* [42] on a single downstream task. TabSTAR is designed to enable effective cross-dataset learning by applying supervised learning on the target variable in both stages. At its core, it uses an unfrozen encoder-only language model, which can potentially invoke world knowledge acquired during the language model pretraining.⁸ The encoder is combined with a tabular-specific architecture tailored to structured data, mitigating the known limitations of language models in tabular settings [75, 77].

TabSTAR’s architecture comprises five core modules: (1) **Verbalization**, mapping every feature into a textual representation composed of both the column name and value, with a special treatment to numerical features for full numerical precision; (2) **Encoding**, transforming semantic and numerical inputs into meaningful embeddings of the same dimension; (3) **Fusion**, integrating textual and numerical representations on each feature independently; (4) **Interaction**, modeling dependencies and relationships between different features through cross-feature self-attention; and (5) **Prediction**, where outputs are projected into a real value for regression or a probability distribution for classification. Figure 1 illustrates the architecture, Appendix A elaborates it, and Appendix B discusses the training.

A key innovation of TabSTAR is the introduction of *target-aware tokens*, a novel approach that integrates the target variable’s identity as an input to the model. Unlike existing TFM [38, 47, 84, 86, 90, 83], which treat the target value as a mere label, TabSTAR fuses target-awareness from the very beginning. For classification tasks, each target value is verbalized and encoded like any other feature. Then, features and target tokens interact with each other, building representations that are then used for prediction. Crucially, this target-awareness allows parameter sharing between all target tokens, which can later use a shared prediction head that maps tokens to probabilities regardless of the number of classes and their identity. By doing so, TabSTAR eliminates the need for dataset-specific components commonly found in prior work [26, 90, 84]. TabSTAR’s flexible architecture effortlessly scales⁹ to any dataset size, and handles any number of classes in multiclass classification tasks.

⁷Ranging from metadata-rich, text-heavy datasets to numeric-only tables lacking column names.

⁸Note that the language-model pretraining occurs before TabSTAR’s pretraining. Unless specified differently, the term pretraining refers to TabSTAR’s pretraining, which assumes the use of a pretrained language model.

⁹Except when the number of features becomes very large, where memory limitations may arise.

Table 2: An illustrative verbalization of the first patient of Table 1. Each semantic feature is verbalized with its name and value. The numerical *Age* value 45 is standardized (mapped into z-scores, e.g., 0.27) and binned (providing a range to the verbalization, e.g., 40-50, and its quantile). The target variable *Decision* is mapped into its two possible elements, regardless of its original true value.

Name	Value	Semantic	Numerical
Age	45	“Age: 40–50 (Quantile 50–60%)”	0.27
Department	Cardiology	“Department: Cardiology”	–
Report	Mild chest discomfort.	“Report: Mild chest discomfort.”	–
Decision	Hospitalized	“Target. Decision: Hospitalized”	–
Decision	Released	“Target. Decision: Released”	–

Verbalization All the features and each of the target values are processed into a sequence of *elements*. Numerical features are processed into two inputs: a numerical one and a semantic one. The numerical input is standardized using z-scores, with outlier clipping at ± 3 standard deviations. In addition, they are verbalized using quantile discretization into 10 bins, a novel approach to mitigate the precision loss inherent in language models [75]. Appendix A.1 shows a precise example and §6 discusses different verbalization strategies. In contrast, semantic features are directly verbalized by concatenating the feature name and textual value, without any numerical representation. The target variable is also included as part of the input: In classification tasks, each of the C possible values is represented by an element, constant for every example, while the true value remains hidden. For regression tasks, a single element is verbalized, carrying only the target name. Table 1 shows a toy dataset of patient records and outcomes and Table 2 shows the verbalization for the first patient.

Encoding We employ a pretrained *e5-small-v2* [80] embedding model for semantic encoding, chosen for its strong performance on the *MTEB* benchmark [57] with a relatively modest parameter count. By unfreezing the upper half of its layers, the representations are optimized for predicting the target variable, which leads to a significant impact on TabSTAR performance (see §6). Each verbalization element is encoded independently into a semantic representation, with attention applied between tokens within each sequence element. In parallel, we encode standardized numerical values by projecting them into the same dimension using a small MLP. For the patient in Table 2, this results in a numerical embedding for *Age* alongside semantic representations for each of the 5 verbalizations.

Fusion To obtain a unified representation for each sequence element, we apply a fusion block consisting of a single encoder-only transformer layer. Crucially, each numerical feature is fused independently, as the block attends only to its numerical and semantic embeddings. In our running example, the representation of *Age* now jointly captures both its semantic context (the fact that the value represents age) as well as its numerical value (the patient’s age, 45, or 0.27 after standardization).

Interaction The fused, semantically-rich and numerically-grounded representations of all elements interact via a 6-layer Transformer encoder [79]. Each input element is now a token, with feature tokens and target tokens all attending to each other. Unlike standard language models, which integrate positional encoding, the Interaction module’s inputs are order-invariant, a desideratum for TFM’s, as defined by [77]. The encoder produces contextualized representations for each target value. In our example, this yields dedicated embeddings for the *Release* and *Hospitalization* target values. The role of these representations is to carry information about how likely each value is to be the true value.

Prediction TabSTAR is designed for cross-dataset learning, with shared regression and classification heads used during both pretraining and finetuning. For classification, each of the C target tokens is processed independently through the same classification head, which projects them to scores. We then apply a softmax over all the possible values to yield a probability distribution. Crucially, the fact that target tokens for every class in every dataset share the same classification head allows efficient parameter sharing, flexibly supports any number of output classes, and removes any need for dataset-specific parameters. This is not only efficient during pretraining, but also provides a better initialization for finetuning. In our example, both the *Released* and *Hospitalized* tokens go through the same classification head, which maps them from representations to logits. Applying softmax yields predicted probabilities. For regression tasks, a single target token is projected into a real value.

4 Experiments

The TabSTAR Pretraining Corpus While TabSTAR could be pretrained on a massive scale, for this work we limit ourselves to a modest pretraining corpus focusing on classification, as we believe that TabSTAR’s inductive biases are best suited to shine in this task. We manually curate a pretraining corpus of 350 high-quality tabular datasets (253 classification, 97 regression), in a tedious process in which we uncover numerous duplications in the most popular tabular repositories, OpenML [78] and Kaggle,¹⁰ as elaborated by [76]. We begin by sourcing datasets from popular benchmarks [25, 47, 23, 22, 31, 63, 55, 32], but observe that the presence of textual tabular datasets in them is very rare. Thus, we furthermore expand our corpus, focusing on classification datasets with rich semantic content. See Appendix C for more details.

Benchmark Tabular datasets with free-text have seen little prior research, and accordingly, benchmarks are rare. To address this, we compile all available datasets from three sources: (1) the *AutoML Multimodal Benchmark* [63], (2) the CARTE paper [47], and (3) the analysis of free-text and high-cardinality features by [32]. After deduplication, our final benchmark includes 50 datasets. Despite its breadth, this collection has two key limitations: First, the benchmark is heavily skewed towards regression tasks, with 36 datasets. Secondly, 29 out of these 36 datasets were solely contributed by the CARTE benchmark, which focuses more heavily on high-cardinality features, rather than longer texts, as it was pretrained over knowledge graphs. While our main motivation is classification tasks with textual features, we decide nevertheless to evaluate on the full set of 50 datasets, although it is biased toward regression problems and high-cardinality features (see Appendix D).

Baselines We compare TabSTAR against **GBDTs**: *Random Forest* [12], *LightGBM* [46], *XGBoost* [15], and *CatBoost* [60]; **MLPs**: *RealMLP* [39] and *TabM*; and **TFMs**: [29], *CARTE* [47], *TabDPT* [54], *TabPFN-v2* [38], and *TabICL* [61] which only supports classification tasks. For several models¹¹ we consider both default variants as well as tuned ones, where hyperparameters are optimized separately for each task using random search with 5-fold cross-validation under a 4-hour budget on 8 CPU cores. Since the public TabPFN-v2 model does not support text, we use their closed-sourced API client.¹² For models lacking native support for textual features, we embed text using *e5-small-v2* [80], allowing a fair comparison. For more details about the hyperparameters for each baseline as well as exclusion of baselines due to potential leakage concerns, see Appendix E.

Experimental Setup Each of the 50 datasets in the benchmark is evaluated with 10 random train-test splits (90% training, 10% testing), resulting in 500 runs per model. While 30 of the datasets have more than 10,000 examples, most TFM can’t effectively scale beyond it: TabPFN-v2, for example, employs ICL and thus receives as input at most 10,000 examples. While CARTE imposes no strict size cap, its tuning is slow, inefficient, and impractical for larger datasets.¹³ Because of these important limitations, we consider two experiment conditions: (1) **10K**: Each model is trained¹⁴ over at most 10,000 training examples, and (2) **Unlimited**: We evaluate TabSTAR and the most competitive, scalable baselines on the full version of the 30 datasets.¹⁵

The TabSTAR Training To maximize the value of cross-dataset learning, instead of pretraining TabSTAR once, we create five dataset splits. Each variant is pretrained on the 350 pretraining datasets and 40 of the benchmark datasets, while the other 10 serve exclusively as its test set. Crucially, the whole collection was carefully curated to prevent any data leakage from duplicate or overly similar datasets. As a result, each dataset is evaluated by finetuning the single pretrained variant, which excludes it from its pretraining. For finetuning, while dataset-specific hyperparameter tuning can boost performance, we believe that robust out-of-the-box defaults are essential for TFM and their evaluation. Therefore, we use a default hyperparameter configuration that was found robust over a disjoint set of tabular datasets, as detailed in Appendix B.2.

¹⁰<https://www.kaggle.com/datasets>

¹¹All MLPs and GBDTs, except the naive Random Forest baseline, used as a weak baseline.

¹²<https://github.com/PriorLabs/tabpfn-client>

¹³In their own paper, CARTE was evaluated only over up to 2,048 examples, without scaling guarantees.

¹⁴While ICL-based TFM aren’t technically trained, we adopt this term for conciseness.

¹⁵We technically cap the amount of examples to 100,000 for computational efficiency.

5 Results

We evaluate each model using AUROC (classification) and R^2 (regression) as metrics. Following [38], we normalize scores per dataset split to the $[0, 1]$ range, using the best and worst model performance as anchors.¹⁶ The normalized scores are averaged across all runs, with 95% CIs. Performance for all models on both conditions is shown in Figure 2 (classification) and Figure 3 (regression). For certain datasets, we are unable to execute TabPFN-v2, CARTE, and TabDPT due to model-specific implementation issues or scalability constraints. Reported averages for these models are computed only over the datasets where evaluation is feasible. Appendix F.1 elaborates on technical limitations, Appendix F.2 on dataset-level performance, and Appendix F.3 on head-to-head comparisons.

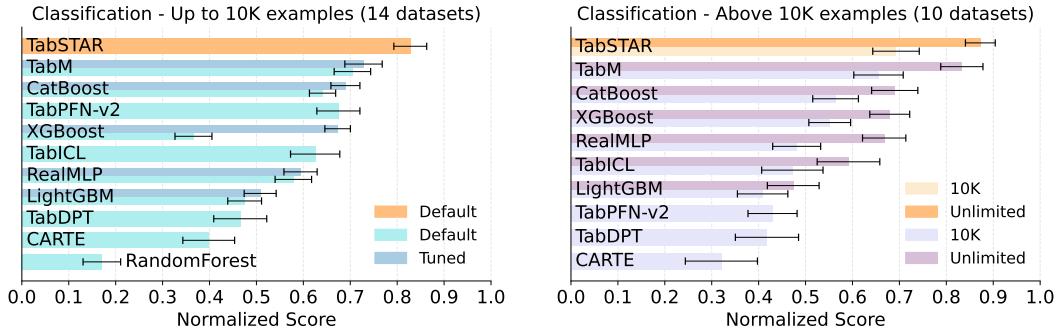


Figure 2: Comparison of normalized scores with 95% CIs between TabSTAR and baseline models in classification tasks, evaluated on up to 10,000 examples (left) and above 10,000 (right).

In classification problems, TabSTAR consistently achieves SOTA performance. This is evident both when restricting the dataset size to 10,000 examples and when using larger datasets in the unlimited condition. For the 10K condition, TabSTAR achieves a 0.83 score, followed by TabM-Tuned (0.73), CatBoost-Tuned (0.69), and TabPFN-v2 (0.67). When analyzing head-to-head comparisons (see Appendix F.3), TabSTAR outperforms TabPFN-v2 (8/11 datasets), TabM-Tuned (10/14), and CatBoost-Tuned (12/14). For the unlimited condition, TabSTAR-Unlimited achieves a 0.84 score, followed by TabM-Tuned-Unlimited (0.79), and much above the rest. Importantly, unlimited variants significantly surpass the 10K ones, emphasizing the importance of models with no scaling restrictions.

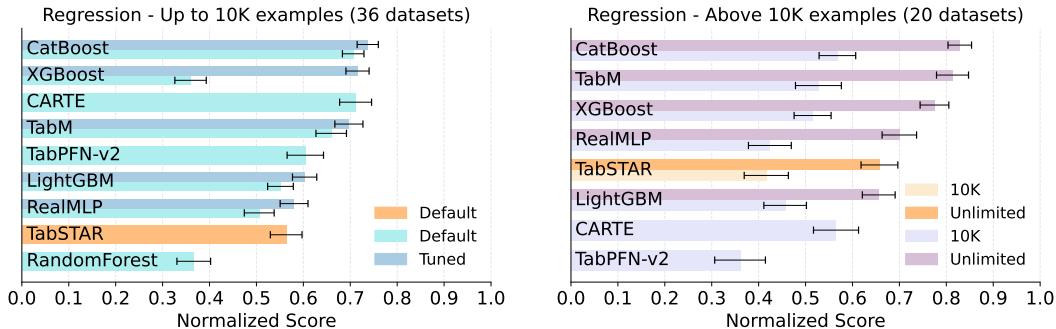


Figure 3: Comparison of normalized scores with 95% CIs between TabSTAR and baseline models in regression tasks, evaluated on up to 10,000 examples (left) and above 10,000 (right).

Although regression is not our main focus, TabSTAR achieves competitive results in the 10K condition, but clearly does not set the SOTA. Surprisingly, while TabPFN-v2 is superior, it significantly underperforms compared to GBDTs which dominate this category. This emphasizes the need for better modeling of textual tabular learning, especially since TabPFN-v2 has shown remarkable performance in non-textual tabular datasets, and CARTE set the SOTA for small datasets. In the unlimited setting, TabSTAR scales well and surpasses other TFM models which cannot scale, but the gap from GBDTs remains significant. §7 discusses this limitation and suggests promising directions for future generations of TabSTAR to achieve SOTA in regression as well.

¹⁶For a single run, the best model gets 1, the worst gets 0, and the rest are linearly scaled accordingly.

Table 3: Cost Analysis. Median training and inference times and peak memory usage on GPU and CPU, aggregated over the 50 datasets of the benchmark with up to 10,000 examples.

Model	Train			Inference		
	Time (s)	GPU (GB)	CPU (GB)	Time (s)	GPU (GB)	CPU (GB)
CatBoost-CPU	360.8	–	2.2	34.0	–	2.1
CatBoost	68.5	1.3	1.5	2.0	1.3	1.5
LightGBM	39.1	1.3	1.5	2.0	1.3	1.5
RandomForest	86.1	1.3	1.5	2.5	1.3	1.5
RealMLP	136.4	1.4	1.7	2.5	1.3	1.7
TabDPT	10.4	2.2	1.8	161.4	33.8	6.0
TabICL	42.6	1.2	1.7	13.4	25.4	2.1
TabM	19.1	1.8	4.4	1.5	1.8	4.3
TabSTAR	493.2	4.7	1.8	3.0	1.3	1.8
XGBoost	38.4	1.3	1.5	2.5	1.3	1.5

Cost Analysis Table 3 reports the computational requirements of TabSTAR and competing baselines, measured by runtime and peak memory usage, for both training and inference. To highlight the importance of GPU acceleration when text embeddings are part of the preprocessing, we include a variant of CatBoost restricted to CPU only. We exclude TabPFN-v2, whose performance is probably comparable to TabICL, since it is only accessible via an API, and CARTE, whose extensive tuning requirements and prohibitive training costs forced evaluation on less capable hardware.

We find that TabSTAR incurs modest inference costs, comparable to GPU-accelerated GBDTs, since GPU processing of text embeddings dominates runtime. This is evident from the CatBoost-CPU variant, whose inference is nearly ten times slower than TabSTAR. Moreover, TabICL and TabDPT are orders of magnitude slower and consume far more memory, emphasizing the scalability limitations for ICL-based models. During training, TabSTAR is considerably slower than the baselines. However, while fitting a single GBDT is much faster, this advantage quickly diminishes in practice when hyperparameter tuning is needed or GPU acceleration for text embeddings is not available. Appendix F.4 provides dataset-level runtimes and F.5 analyzes costs as a function of the feature count.

6 Analysis

We analyze the factors contributing to TabSTAR’s strong performance by addressing three key research questions: **Q1:** How important is the encoder language model unfreezing? **Q2:** Does the number of datasets during pretraining contribute to the downstream task performance? and **Q3:** How do different verbalization methods of numerical features impact performance?

To answer these questions, we pretrain several variants of TabSTAR for each analysis, limiting ourselves to a subset of the tabular datasets used for the main experiment (see §4). Specifically, each variant is pretrained over 256 datasets¹⁷ including 30 datasets from our benchmark, and evaluated over the remaining 20 datasets (12 regression, 8 classification). This reduced setup allows leveraging transfer learning and exploiting our corpus, without the burden of training multiple folds per variant. Appendix G.1 lists the 20 datasets used for evaluation along with per-dataset results.

Q1: The Role of the Encoder Unfreezing We examine the effect of unfreezing different numbers of textual encoder layers during pretraining. In finetuning, we keep all base layers frozen and instead apply LoRA adapters to the same layers that were unfrozen in pretraining.¹⁸ Figure 4 shows the validation loss during TabSTAR pretraining (left) and the normalized score on the downstream tasks (right) as a function of the number of unfrozen encoder layers. Notably, unfreezing even a single encoder layer significantly outperforms using static embeddings. Further substantial improvements are observed as more layers are tuned, with the best results achieved when unfreezing 6 layers. While unfreezing 9 layers shows lower performance, it is plausible that adding more datasets to the pretraining phase will affect this finding. See Appendix G.2 for more details.

¹⁷Except for variants of Q2, which analyze the effect of the number of datasets on pretraining.

¹⁸For simplicity, we will refer to them as unfrozen to differentiate them from layers without LoRA.

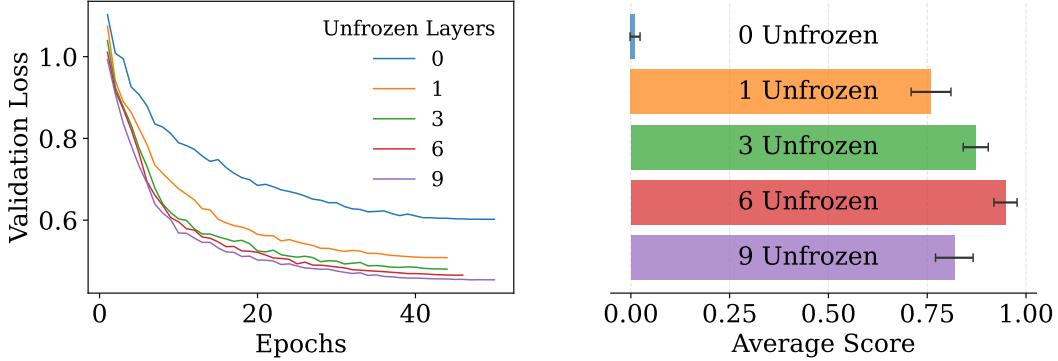


Figure 4: Performance as a function of the number of encoder layers unfrozen: Validation loss during TabSTAR’s pretraining (left) and normalized scores with 95% CIs on the downstream tasks (right). Unfreezing even a single encoder layer significantly improves the performance of TabSTAR.

Q2: The Effect of Pretraining To evaluate the impact of pretraining on TabSTAR’s downstream performance, we compare a pretrained version of TabSTAR with a version that was finetuned from scratch.¹⁹ In line with previous work [90, 86], the pretrained model performs significantly better, highlighting the critical role of transfer learning for TabSTAR’s success. To further investigate the effect of the number of pretraining datasets on downstream task performance, we train two additional versions: one pretrained on 16 datasets and another on 64 datasets.

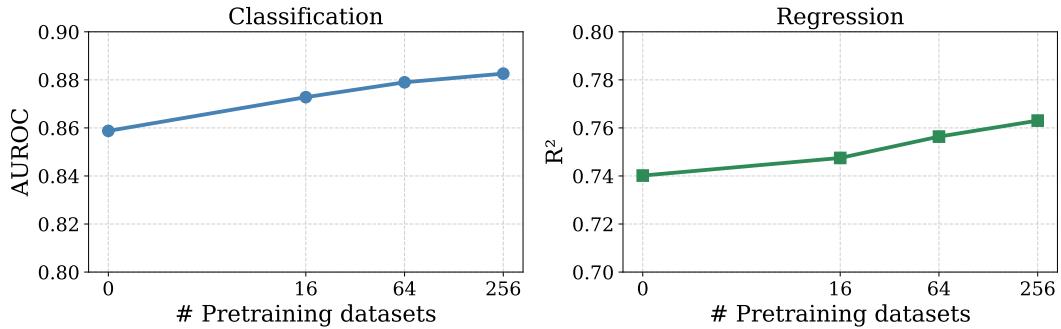


Figure 5: Average performance on downstream tasks as a function of the number of pretraining datasets (in log scale). We use AUROC for classification (left), and R^2 for regression (right).

As shown in Figure 5, increasing the number of pretraining datasets consistently improved performance in both classification and regression tasks. Notably, the substantial gain in regression tasks suggests that TabSTAR’s downstream performance on §5 could improve with more pretraining data, potentially reaching SOTA performance with enough scale. See Appendix G.3 for more details.

Q3: Numerical Verbalization A key challenge in integrating language models with numerical data is determining how to best represent numerical values within a linguistic framework. While some semantic tabular methods omit numerical features from the verbalization [83, 86], TP-BERTA [84] introduced *Relative Magnitude Tokenization* [84], which encodes numerical information through non-semantic special bin tokens. In contrast, TabSTAR injects semantic numerical information into the verbalization of numerical features, as illustrated in Table 2. To quantify the effect of our novel verbalization, we explore two thinner variants: (1) **Name + Bin**, which excludes the quantile information, and (2) **Name**, which omits numeric information entirely and verbalizes the feature name only. Appendix G.4 shows an illustrative example for each variant and presents the full results. As demonstrated in Table 4, our findings reveal that incorporating numerical information significantly enhances performance, highlighting the importance of balancing numerical precision with a representation format that aligns with the language model’s parametric knowledge.

¹⁹Since LoRA underperforms on random weights, we finetune the entire non-pretrained model.

Table 4: Normalized score with 95% CIs by the numerical verbalization method.

Verbalization Method	Name	Name + Bin	TabSTAR
Classification	0.386 ± 0.095	0.544 ± 0.093	0.593 ± 0.097
Regression	0.386 ± 0.081	0.584 ± 0.076	0.596 ± 0.079

7 Discussion and Conclusion

We introduce TabSTAR, a Tabular Foundation Model with Semantically Target-Aware Representations, which integrates textual features through an unfrozen pretrained encoder. In addition, its novel target-aware tokens enable efficient cross-dataset generalization without dataset-specific parameters. Despite limited pretraining data and a relatively small text encoder [80], TabSTAR sets the SOTA in tabular classification with textual features, significantly surpassing GBDTs and leading TFM.

Since scaling laws in data and model size have proven themselves for LLMs [44] and TabSTAR improves with the number of pretraining datasets (see §6), future work should scale TabSTAR across both model and data dimensions. For model scaling, we envision a family of model sizes, common for LLMs [30, 74, 51], that will allow a trade-off between quality and costs. Data scaling might leverage self-supervised learning [62, 82] over large-scale table corpora [18], or realistic synthetic tabular data generators [10], which have proven successful [38, 2]. At scale, it could potentially unlock few-shot learning capabilities and develop automatic feature-engineering skills [37].

Beyond scaling, TabSTAR’s semantic approach has tremendous potential to explicitly include world knowledge by leveraging LLMs, which to date have had a limited impact on tabular learning. As a few motivating examples, LLMs could improve TabSTAR’s numerical verbalization binning approach by providing semantically informed thresholds, or by providing explicit world knowledge that could be injected as a strong prior in small data scenarios. While these directions seem like plausible research paths, they come with a risk of data leakage due to the memorization properties of LLMs [9]. Evaluating TFM fairly while keeping benchmarks uncontaminated would be an important enabler for tabular research. As a step in this direction, we are releasing several TabSTAR variants, each with a different dataset withheld during pretraining, ensuring that for every dataset there is a TabSTAR model that has never seen it. We urge fellow researchers to adopt this approach in their own work.

While TabSTAR sets a new bar in classification, its regression results lag behind GBDTs, which outperform other TFM as well. This gap could be narrowed through additional scaling, by exploring regression-via-classification techniques like [38, 2], and by enriching the numerical encoder with distribution-aware statistics per feature as done by [61]. In addition, similar to other TFM, TabSTAR may encounter memory bottlenecks on datasets with hundreds of features, and its training speed lags behind untuned GBDTs and ICL-based TFM. Yet TabSTAR achieves GBDT-level efficiency during inference, as opposed to TabPFN-v2 and TabICL. Furthermore, TabSTAR has not been extensively evaluated in few-shot scenarios and in purely numerical datasets.

Despite these limitations, TabSTAR offers a promising pathway toward improving performance on tabular datasets with textual fields, common in industries with high social impact (e.g., healthcare, education), or with significant economic value (e.g., banking, manufacturing). In addition, its architecture lends itself to multimodality and could be extended to tabular datasets that combine numerical, textual, and image features. We believe TabSTAR paves the way for a new generation of semantically enriched, Multimodal TFM, and invite the research community to advance this vision.

Acknowledgments and Disclosure of Funding

Roi Reichart and Eilam Shapira have been partially supported by a VATAT grant on data science.

We thank Omri Feldman for brainstorming since the very beginning; Elad Hoffer and Ofir Lindenbaum for consulting and feedback; David Holzmüller, Lennart Purucker, Myung Kim, and Gaël Varoquaux for assisting with evaluations and benchmarks; and Frank Hutter, Noah Hollmann, Léo Grinsztajn, and the rest of the Prior Labs team for providing extensive access to TabPFN-v2’s API version.

References

- [1] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, and et al. Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 173–182. PMLR, June 2016. URL <https://proceedings.mlr.press/v48/amodei16.html>. ISSN: 1938-7228.
- [2] Abdul Fatir Ansari, Lorenzo Stella, Ali Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, Jasper Zschiegner, Danielle C. Maddix, Hao Wang, Michael W. Mahoney, Kari Torkkola, Andrew Gordon Wilson, Michael Bohlke-Schneider, and Bernie Wang. Chronos: Learning the Language of Time Series. *Transactions on Machine Learning Research*, May 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=gerNCVqqtr>.
- [3] Sercan Ö Arik and Tomas Pfister. TabNet: Attentive Interpretable Tabular Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(8):6679–6687, May 2021. ISSN 2374-3468. doi: 10.1609/aaai.v35i8.16826. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16826>. Number: 8.
- [4] Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundation Models Defining a New Era in Vision: A Survey and Outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(4):2245–2264, April 2025. ISSN 1939-3539. doi: 10.1109/TPAMI.2024.3506283. URL <https://ieeexplore.ieee.org/abstract/document/10834497>.
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate, May 2016. URL <http://arxiv.org/abs/1409.0473>. arXiv:1409.0473 [cs].
- [6] Dara Bahri, Heinrich Jiang, Yi Tay, and Donald Metzler. Scarf: Self-Supervised Contrastive Learning using Random Feature Corruption. October 2021. URL https://openreview.net/forum?id=CuV_qYkmKb3.
- [7] Tommaso Bendinelli, Artur Dox, and Christian Holz. Exploring LLM Agents for Cleaning Tabular Machine Learning Datasets. March 2025. URL <https://openreview.net/forum?id=RXnQPYsoun>.
- [8] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, and et al. On the Opportunities and Risks of Foundation Models, August 2021. URL <https://arxiv.org/abs/2108.07258v3>.
- [9] Sebastian Bordt, Harsha Nori, Vanessa Rodrigues, Besmira Nushi, and Rich Caruana. Elephants Never Forget: Memorization and Learning of Tabular Data in Large Language Models. *First Conference on Language Modeling*, 2024.
- [10] Vadim Borisov, Kathrin Sessler, Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. Language Models are Realistic Tabular Data Generators. September 2022. URL <https://openreview.net/forum?id=cEygmQNOeI>.
- [11] Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. Deep Neural Networks and Tabular Data: A Survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35(6):7499–7519, June 2024. ISSN 2162-2388. doi: 10.1109/TNNLS.2022.3229161. URL <https://ieeexplore.ieee.org/abstract/document/9998482>.
- [12] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, October 2001. ISSN 1573-0565. doi: 10.1023/A:1010933404324. URL <https://doi.org/10.1023/A:1010933404324>.
- [13] Tom Brown, Benjamin Mann, Nick Ryder, and et al. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bf8ac142f64a-Abstract.html>.

[14] Patricio Cerdá and Gaël Varoquaux. Encoding High-Cardinality String Categorical Variables. *IEEE Transactions on Knowledge and Data Engineering*, 34(3):1164–1176, March 2022. ISSN 1558-2191. doi: 10.1109/TKDE.2020.2992529. URL <https://ieeexplore.ieee.org/abstract/document/9086128>.

[15] 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*, KDD ’16, pages 785–794, New York, NY, USA, August 2016. Association for Computing Machinery. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939785. URL <https://dl.acm.org/doi/10.1145/2939672.2939785>.

[16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.

[17] 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, June 2021. URL <http://arxiv.org/abs/2010.11929>. arXiv:2010.11929 [cs].

[18] Gus Eggert, Kevin Huo, Mike Biven, and Justin Waugh. TabLib: A Dataset of 627M Tables with Context, October 2023. URL <http://arxiv.org/abs/2310.07875>. arXiv:2310.07875 [cs].

[19] Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data, March 2020. URL <http://arxiv.org/abs/2003.06505>. arXiv:2003.06505 [stat].

[20] Xi Fang, Weijie Xu, Fiona Anting Tan, Ziqing Hu, Jiani Zhang, Yanjun Qi, Srinivasan H. Sengamedu, and Christos Faloutsos. Large Language Models (LLMs) on Tabular Data: Prediction, Generation, and Understanding - A Survey. *Transactions on Machine Learning Research*, March 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=IZnrCGF9WI>.

[21] Benjamin Feuer, Robin T. Schirrmeister, Valeria Cherepanova, Chinmay Hegde, Frank Hutter, Micah Goldblum, Niv Cohen, and Colin White. TuneTables: Context Optimization for Scalable Prior-Data Fitted Networks. *Advances in Neural Information Processing Systems*, 37:83430–83464, December 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/97dc07f1253ab33ee514f395a82fa7cc-Abstract-Conference.html.

[22] Matthias Feurer, Katharina Eggensperger, Stefan Falkner, Marius Lindauer, and Frank Hutter. Auto-Sklearn 2.0: Hands-free AutoML via Meta-Learning. *Journal of Machine Learning Research*, 23(261):1–61, 2022. ISSN 1533-7928. URL <http://jmlr.org/papers/v23/21-0992.html>.

[23] Sebastian Felix Fischer, Matthias Feurer, and Bernd Bischl. OpenML-CTR23 – A curated tabular regression benchmarking suite. August 2023. URL <https://openreview.net/forum?id=HebAOoMm94>.

[24] Josh Gardner, Juan C. Perdomo, and Ludwig Schmidt. Large Scale Transfer Learning for Tabular Data via Language Modeling. *Advances in Neural Information Processing Systems*, 37:45155–45205, December 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/4fd5cf2e31bebbccfa5ffa354c04bdc-Abstract-Conference.html.

[25] Pieter Gijsbers, Marcos L. P. 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. ISSN 1533-7928. URL <http://jmlr.org/papers/v25/22-0493.html>.

[26] Yury Gorishniy, Ivan Rubachev, Valentin Khrulkov, and Artem Babenko. Revisiting deep learning models for tabular data. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, pages 18932–18943, Red Hook, NY, USA, December 2021. Curran Associates Inc. ISBN 978-1-7138-4539-3.

[27] Yury Gorishniy, Ivan Rubachev, and Artem Babenko. On embeddings for numerical features in tabular deep learning. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, pages 24991–25004, Red Hook, NY, USA, November 2022. Curran Associates Inc. ISBN 978-1-7138-7108-8.

[28] Yury Gorishniy, Ivan Rubachev, Nikolay Kartashev, Daniil Shlenskii, Akim Kotelnikov, and Artem Babenko. TabR: Tabular Deep Learning Meets Nearest Neighbors. The Twelfth International Conference on Learning Representations, October 2023. URL <https://openreview.net/forum?id=rhgIgTSSxW>.

[29] Yury Gorishniy, Akim Kotelnikov, and Artem Babenko. TabM: Advancing Tabular Deep Learning with Parameter-Efficient Ensembling, February 2025. URL <http://arxiv.org/abs/2410.24210> [cs].

[30] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, and et al. The Llama 3 Herd of Models, November 2024. URL <http://arxiv.org/abs/2407.21783>. arXiv:2407.21783 [cs].

[31] Leo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, 35:507–520, December 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/0378c7692da36807bdec87ab043cdadc-Abstract-Datasets_and_Benchmarks.html.

[32] Léo Grinsztajn, Edouard Oyallon, Myung Jun Kim, and Gaël Varoquaux. Vectorizing string entries for data processing on tables: when are larger language models better?, December 2023. URL <http://arxiv.org/abs/2312.09634>. arXiv:2312.09634 [stat].

[33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition, December 2015. URL <http://arxiv.org/abs/1512.03385>. arXiv:1512.03385 [cs].

[34] Xin He, Kaiyong Zhao, and Xiaowen Chu. AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622, January 2021. ISSN 0950-7051. doi: 10.1016/j.knosys.2020.106622. URL <https://www.sciencedirect.com/science/article/pii/S0950705120307516>.

[35] 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 *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, pages 5549–5581. PMLR, April 2023. URL <https://proceedings.mlr.press/v206/hegselmann23a.html>. ISSN: 2640-3498.

[36] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. The Eleventh International Conference on Learning Representations, September 2022. URL https://openreview.net/forum?id=cp5Pvc16w8_.

[37] Noah Hollmann, Samuel Müller, and Frank Hutter. Large Language Models for Automated Data Science: Introducing CAAFE for Context-Aware Automated Feature Engineering. *Advances in Neural Information Processing Systems*, 36:44753–44775, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/8c2df4c35cdbee764ebb9e9d0acd5197-Abstract-Conference.html.

[38] 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, January 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08328-6. URL <https://www.nature.com/articles/s41586-024-08328-6>. Publisher: Nature Publishing Group.

[39] David Holzmüller, Léo Grinsztajn, and Ingo Steinwart. Better by default: Strong pre-tuned MLPs and boosted trees on tabular data. *Advances in Neural Information Processing Systems*, 37: 26577–26658, December 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/2ee1c87245956e3eaa71aabaf5753eb-Abstract-Conference.html.

[40] 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. October 2024. URL <https://openreview.net/forum?id=H02X7R030C#discussion>.

[41] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-Efficient Transfer Learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2790–2799. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/houlsby19a.html>. ISSN: 2640-3498.

[42] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. October 2021. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.

[43] Xin Huang, Ashish Khetan, Milan Cvitkovic, and Zohar Karnin. TabTransformer: Tabular Data Modeling Using Contextual Embeddings, December 2020. URL <http://arxiv.org/abs/2012.06678> [cs]. arXiv:2012.06678 [cs].

[44] 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, January 2020. URL <http://arxiv.org/abs/2001.08361>. arXiv:2001.08361 [cs].

[45] Liran Katzir, Gal Elidan, and Ran El-Yaniv. Net-DNF: Effective Deep Modeling of Tabular Data. October 2020. URL <https://openreview.net/forum?id=73WTGs96kho>.

[46] 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. Curran Associates, Inc., 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html.

[47] Myung Jun Kim, Léo Grinsztajn, and Gaël Varoquaux. CARTE: pretraining and transfer for tabular learning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML’24*, pages 23843–23866, Vienna, Austria, July 2024. JMLR.org.

[48] Ravin Kohli, Matthias Feurer, Katharina Eggensperger, Bernd Bischl, and Frank Hutter. Towards Quantifying the Effect of Datasets for Benchmarking: A Look at Tabular Machine Learning.

[49] Jannik Kossen, Neil Band, Clare Lyle, Aidan N Gomez, Thomas Rainforth, and Yarin Gal. Self-Attention Between Datapoints: Going Beyond Individual Input-Output Pairs in Deep Learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 28742–28756. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/f1507aba9fc82ffa7cc7373c58f8a613-Abstract.html>.

[50] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html.

[51] Barak Lenz, Opher Lieber, Alan Arazi, and et al. Jamba: Hybrid Transformer-Mamba Language Models. October 2024. URL <https://openreview.net/forum?id=JFPaD71pBD>.

[52] Roman Levin, Valeriia Cherepanova, Avi Schwarzschild, Arpit Bansal, C. Bayan Bruss, Tom Goldstein, Andrew Gordon Wilson, and Micah Goldblum. Transfer Learning with Deep Tabular Models. The Eleventh International Conference on Learning Representations, September 2022. URL <https://openreview.net/forum?id=b0RuGUYo8pA>.

[53] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach, July 2019. URL <http://arxiv.org/abs/1907.11692>. arXiv:1907.11692 [cs].

[54] Junwei Ma, Valentin Thomas, Rasa Hosseinzadeh, Hamidreza Kamkari, Alex Labach, Jesse C. Cresswell, Keyvan Golestan, Guangwei Yu, Anthony L. Caterini, and Maksims Volkovs. TabDPT: Scaling Tabular Foundation Models on Real Data, July 2025. URL <http://arxiv.org/abs/2410.18164>. arXiv:2410.18164 [cs].

[55] 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? *Advances in Neural Information Processing Systems*, 36:76336–76369, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/f06d5ebd4ff40b40dd97e30cee632123-Abstract-Datasets_and_Benchmarks.html.

[56] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space, September 2013. URL <http://arxiv.org/abs/1301.3781>. arXiv:1301.3781 [cs].

[57] Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. MTEB: Massive Text Embedding Benchmark. In Andreas Vlachos and Isabelle Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.eacl-main.148. URL <https://aclanthology.org/2023.eacl-main.148>.

[58] Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers Can Do Bayesian Inference. October 2021. URL <https://openreview.net/forum?id=KSugKcbNf9>.

[59] OpenAI. GPT-4 Technical Report, March 2024. URL <http://arxiv.org/abs/2303.08774>. arXiv:2303.08774 [cs].

[60] 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*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/14491b756b3a51daac41c24863285549-Abstract.html>.

[61] Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. TabICL: A Tabular Foundation Model for In-Context Learning on Large Data, February 2025. URL <http://arxiv.org/abs/2502.05564>. arXiv:2502.05564 [cs].

[62] Ivan Rubachev, Artem Alekberov, Yury Gorishniy, and Artem Babenko. Revisiting Pretraining Objectives for Tabular Deep Learning, July 2022. URL <http://arxiv.org/abs/2207.03208>. arXiv:2207.03208 [cs].

[63] Xingjian Shi, Jonas Mueller, Nick Erickson, Mu Li, and Alex Smola. Benchmarking Multimodal AutoML for Tabular Data with Text Fields. August 2021. URL <https://openreview.net/forum?id=Q0z0Iaec8HF>.

[64] Assaf Shmuel, Oren Glickman, and Teddy Lazeznik. A Comprehensive Benchmark of Machine and Deep Learning Across Diverse Tabular Datasets, August 2024. URL <http://arxiv.org/abs/2408.14817>. arXiv:2408.14817 [cs].

[65] Ravid Schwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90, May 2022. ISSN 1566-2535. doi: 10.1016/j.inffus.2021.11.011. URL <https://www.sciencedirect.com/science/article/pii/S1566253521002360>.

[66] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition, April 2015. URL <http://arxiv.org/abs/1409.1556>. arXiv:1409.1556 [cs].

[67] Leslie N. Smith and Nicholay Topin. Super-Convergence: Very Fast Training of Neural Networks Using Large Learning Rates, May 2018. URL <http://arxiv.org/abs/1708.07120> [cs].

[68] Aivin V. Solatorio and Olivier Dupriez. REaLTabFormer: Generating Realistic Relational and Tabular Data using Transformers, February 2023. URL <http://arxiv.org/abs/2302.02041> [cs].

[69] Gowthami Somepalli, Micah Goldblum, Avi Schwarzschild, C. Bayan Bruss, and Tom Goldstein. SAINT: Improved Neural Networks for Tabular Data via Row Attention and Contrastive Pre-Training, June 2021. URL <http://arxiv.org/abs/2106.01342>. arXiv:2106.01342 [cs].

[70] Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. AutoInt: Automatic Feature Interaction Learning via Self-Attentive Neural Networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, pages 1161–1170, New York, NY, USA, November 2019. Association for Computing Machinery. ISBN 978-1-4503-6976-3. doi: 10.1145/3357384.3357925. URL <https://dl.acm.org/doi/10.1145/3357384.3357925>.

[71] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, volume 2 of *NIPS'14*, pages 3104–3112, Cambridge, MA, USA, December 2014. MIT Press.

[72] Zhiqiang Tang, Haoyang Fang, Su Zhou, Taojiannan Yang, Zihan Zhong, Cuixiong Hu, Katrin Kirchhoff, and George Karypis. AutoGluon-Multimodal (AutoMM): Supercharging Multimodal AutoML with Foundation Models. AutoML Conference 2024 (ABCD Track), April 2024. URL <https://openreview.net/forum?id=irStSm9waW>.

[73] Zhiqiang Tang, Zihan Zhong, Tong He, and Gerald Friedland. Bag of Tricks for Multimodal AutoML with Image, Text, and Tabular Data, December 2024. URL <http://arxiv.org/abs/2412.16243>. arXiv:2412.16243 [cs].

[74] Gemini Team. Gemini: A Family of Highly Capable Multimodal Models, May 2025. URL <http://arxiv.org/abs/2312.11805>. arXiv:2312.11805 [cs].

[75] Avijit Thawani, Jay Pujara, Pedro A. Szekely, and Filip Ilievski. Representing Numbers in NLP: a Survey and a Vision, March 2021. URL <http://arxiv.org/abs/2103.13136>. arXiv:2103.13136 [cs].

[76] Andrej Tschalzev, Lennart Purucker, Stefan Lüdtke, Frank Hutter, Christian Bartelt, and Heiner Stuckenschmidt. Unreflected Use of Tabular Data Repositories Can Undermine Research Quality, March 2025. URL <http://arxiv.org/abs/2503.09159>. arXiv:2503.09159 [cs].

[77] Boris Van Breugel and Mihaela Van Der Schaar. Position: why tabular foundation models should be a research priority. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML'24*, pages 48976–48993, Vienna, Austria, July 2024. JMLR.org.

[78] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. OpenML: networked science in machine learning. *SIGKDD Explor. Newsl.*, 15(2):49–60, June 2014. ISSN 1931-0145. doi: 10.1145/2641190.2641198. URL <https://doi.org/10.1145/2641190.2641198>.

[79] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

[80] Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Dixin Jiang, Rangan Majumder, and Furu Wei. Text Embeddings by Weakly-Supervised Contrastive Pre-training, February 2024. URL <http://arxiv.org/abs/2212.03533>. arXiv:2212.03533 [cs].

[81] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. DCN V2: Improved Deep & Cross Network and Practical Lessons for Web-scale Learning to Rank Systems. In *Proceedings of the Web Conference 2021*, WWW '21, pages 1785–1797, New York, NY, USA, June 2021. Association for Computing Machinery. ISBN 978-1-4503-8312-7. doi: 10.1145/3442381.3450078. URL <https://dl.acm.org/doi/10.1145/3442381.3450078>.

[82] Wei-Yao Wang, Wei-Wei Du, Derek Xu, Wei Wang, and Wen-Chih Peng. A Survey on Self-Supervised Learning for Non-Sequential Tabular Data. *CoRR*, January 2024. URL <https://openreview.net/forum?id=idHPqbNraV>.

[83] Zifeng Wang and Jimeng Sun. TransTab: Learning Transferable Tabular Transformers Across Tables. October 2022. URL https://openreview.net/forum?id=A1yGs_SWiIi.

[84] Jiahuan Yan, Bo Zheng, Hongxia Xu, Yiheng Zhu, Danny Chen, Jimeng Sun, Jian Wu, and Jintai Chen. Making Pre-trained Language Models Great on Tabular Prediction. The Twelfth International Conference on Learning Representations, October 2023. URL <https://openreview.net/forum?id=anzIzGZuLi>.

[85] Junchen Yang, Ofir Lindenbaum, and Yuval Kluger. Locally Sparse Neural Networks for Tabular Biomedical Data. In *Proceedings of the 39th International Conference on Machine Learning*, pages 25123–25153. PMLR, June 2022. URL <https://proceedings.mlr.press/v162/yang22i.html>. ISSN: 2640-3498.

[86] Chao Ye, Guoshan Lu, Haobo Wang, Liyao Li, Sai Wu, Gang Chen, and Junbo Zhao. Towards Cross-Table Masked Pretraining for Web Data Mining. May 2024. URL <https://openreview.net/forum?id=9jj7cMOXQo>.

[87] Han-Jia Ye, Si-Yang Liu, Hao-Run Cai, Qi-Le Zhou, and De-Chuan Zhan. A Closer Look at Deep Learning on Tabular Data. *CoRR*, January 2024. URL <https://openreview.net/forum?id=eu2cABIHge>.

[88] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, Hao Peng, Jianxin Li, Jia Wu, Ziwei Liu, Pengtao Xie, Caiming Xiong, Jian Pei, Philip S. Yu, and Lichao Sun. A Comprehensive Survey on Pretrained Foundation Models: A History from BERT to ChatGPT, May 2023. URL <http://arxiv.org/abs/2302.09419>. arXiv:2302.09419 [cs].

[89] Qile Zhou, Han-Jia Ye, Leye Wang, and De-Chuan Zhan. Unlocking the Transferability of Tokens in Deep Models for Tabular Data. October 2023. URL <https://openreview.net/forum?id=u20VQ2Xvq1>.

[90] Bingzhao Zhu, Xingjian Shi, Nick Erickson, Mu Li, George Karypis, and Mahsa Shoaran. XTab: Cross-table Pretraining for Tabular Transformers. In *Proceedings of the 40th International Conference on Machine Learning*, pages 43181–43204. PMLR, July 2023. URL <https://proceedings.mlr.press/v202/zhu23k.html>. ISSN: 2640-3498.

[91] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A Comprehensive Survey on Transfer Learning. *Proceedings of the IEEE*, 109(1):43–76, January 2021. ISSN 1558-2256. doi: 10.1109/JPROC.2020.3004555. URL <https://ieeexplore.ieee.org/abstract/document/9134370>.

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 claim "TabSTAR achieves state-of-the-art performance for both medium- and large-sized datasets across known benchmarks" in the abstract and "TabSTAR achieves state-of-the-art (SOTA) performance on classification benchmarks containing substantial textual content" in the introduction - we describe these results in detail in §5, showing TabSTAR is the top performer comparing vs leading gradient boosting decision trees (GBDTs) and leading Tabular Foundation Models (TFMs)
- We claim "...its pretraining phase exhibits scaling laws in the number of datasets" in the abstract and "its performance steadily improves with more pretraining data" in the intro, and show in §6 that TabSTAR's performance improve as we significantly increase the number of pre-training datasets.
- We claim in the introduction "TabSTAR can optimize free-text feature representations, demonstrating their clear superiority over alternative frozen embedding approaches" - we describe these results in detail in §6, by showing that TabSTAR's performance increase greatly when finetuning the embeddings, making this the advantage over other algorithms compared in §5.

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: Yes, in §7 we discuss TabSTAR's limitations, such as its lower performance in regression and its lack of evaluation in other scenarios like datasets without textual content. Some more details about the regression performance are elaborated at §5.

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 doesn't present theoretical proofs, but focuses on algorithmic development and empirical validation.

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 cross-referenced.
- 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: Yes. Our model and method is highly elaborated through the paper (see §3) and the main experiments as well (§4). In addition, multiple supplemental materials cover these points in more details (e.g., Appendix A, Appendix B). Furthermore, we release our code, alongside documentation.

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: See reference from §1 to our anonymous GitHub repository, where we provide our code and running instructions.

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: §4, Appendix B and Appendix E give the full experimental protocol.

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: The results are accompanied by 95% confidence intervals (see Appendix F). This is true also for the analysis section §6, which follows the same standards.

Guidelines:

- 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: [\[Yes\]](#)

Justification: Yes, see Appendix F.4 for compute information and running times, for both TabSTAR and other baselines.

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: The study relies exclusively on publicly available, de-identified datasets (OpenML, Kaggle and published benchmarks) and releases code and models under an open-source licence; no human subjects or sensitive personal data are involved, fully satisfying the NeurIPS Code of Ethics.

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: [Yes]

Justification: We use as a running example introduced in Table 1 an important real use case with high impact for the healthcare industry, motivated by the importance of improving decision making in a small-data regime. In addition, we briefly discuss other industries in §7.

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: TabSTAR is not a generative model, and thus can't be used to create potentially high-risk text and images. In addition, we use publicly available datasets which don't impact individual privacy or security.

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: Yes, we credited them in an appropriate way, adding URLs and package names in multiple occasions throughout the paper and its supplemental materials.

Guidelines:

- 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: Upon acceptance, we will release several TabSTAR pretrained checkpoints, plus complete training and inference code. We believe that this research line is of huge potential and we encourage researchers to continue our line of work, also providing feasible evaluations setups for TFM, such as default hyperparameters.

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 paper doesn't 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: As the study didn't involve any study participants, this section isn't applicable.

Guidelines:

- 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: While TabSTAR uses a language model at its core, this is an encoder-only model, and not a generative LLM. This widely used language model, e5-small-v2 [80], is used as part of the architecture in a standard usage. Thus, this section isn't applicable to our research.

Guidelines:

- 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 Architecture

This appendix provides additional technical details for the architecture introduced in §3. First, we discuss the verbalization module; next, we formally describe the architecture step-by-step; and finally, we present selected experiments on the TabSTAR architecture.

A.1 The Verbalization Module

TabSTAR’s verbalization module standardizes heterogeneous tabular inputs by converting each column, whether predictive feature or target variable, into templated text blocks. We first describe the detection of column types, and then detail the processing steps for each type.

Feature Detection We classify each column as either numerical, referring to quantitative values, or semantic, referring to textual values including categorical and boolean fields encoded as text. We rely on heuristics involving both the primitive data type (e.g., string, float) and human annotation (e.g., OpenML metadata). However, real-world datasets pose challenges, as numerical features can often be stored as strings (e.g., “35 years”, “unknown age”) or may lack inherent order (e.g., country calling codes). Leveraging LLMs for contextualized data cleaning can be a promising direction [7].

A special case is the handling of timestamp and date columns. Similarly to [40], we rely on *skrub*’s *DatetimeEncoder*²⁰ to detect datetime columns and decompose each one of them into a set of new features. Each extracted feature then undergoes its own processing: For example, the weekday is treated as semantic, while the total seconds since the Unix epoch is treated as numerical. Integrating date features more holistically remains an open research question.

Table 5: Illustrative verbalization of a numerical feature (*Age*) with 10 bins. Examples outside the range and missing values are considered as special bins.

Bin	Range	Example Value	Illustrative Verbalization
–	Lower than 18	17	Age: Lower than 18 (Quantile 0%)
1	18–23	20	Age: 18–23 (Quantile 0–10%)
2	23–27	25	Age: 23–27 (Quantile 10–20%)
3	27–31	29	Age: 27–31 (Quantile 20–30%)
4	31–35	33	Age: 31–35 (Quantile 30–40%)
5	35–40	38	Age: 35–40 (Quantile 40–50%)
6	40–45	42	Age: 40–45 (Quantile 50–60%)
7	45–51	48	Age: 45–51 (Quantile 60–70%)
8	51–58	55	Age: 51–58 (Quantile 70–80%)
9	58–67	63	Age: 58–67 (Quantile 80–90%)
10	67–87	83	Age: 67–87 (Quantile 90–100%)
–	Higher than 87	93	Age: Higher than 87 (Quantile 100%)
–	Unknown	–	Age: Unknown Value

Numerical Features Numerical features are represented by both a numerical and a semantic representation. For the numerical representation, given a value x , we compute the clipped z-score $z' = \text{clip}((x - \mu)/\sigma, -3, 3)$ where μ, σ are the training set mean and the standard deviation, and missing values are set to 0. For the semantic representation, we build $B = 10$ quantile bins over the training distribution to map the value accordingly. Table 5 shows an illustrative example for the feature *Age* from our running example in Table 2.

Semantic Features Semantic features are sanitized (e.g., normalizing whitespaces) and verbalized using the template presented in Table 6. Missing values are mapped to “Unknown Value”, just like for numerical features. If a text exceeds the model’s context window (512 tokens for *e5-small-v2*), it is naively truncated to fit it. This limitation is far more pronounced for methods that serialize the entire example into a single textual sequence [84], thereby dramatically reducing the effective context size.

²⁰<https://skrub-data.org/>

Target Variables The verbalization templates for the target values are prepended to every example. For classification tasks, each possible label is verbalized, while for regression we verbalize a single element consisting solely of the feature name. Employing a binning strategy to treat regression as a classification task is a future work direction, as discussed in §7. For regression tasks, target values go through the same standardization with outlier clipping as numerical features, being used solely as the ground truth without going through the input.

Table 6: Verbalization templates for semantic features and target values.

Element Type	Verbalization Template
Predictive feature	"Predictive Feature: {feature_name}" "Feature Value: {feature_value}"
Classification target	"Target Feature: {target_name}" "Feature Value: {target_value}"
Regression target	"Numerical Target Feature: {target_name}"

A.2 The Annotated TabSTAR

Table 7 describes the number of parameters per component in the TabSTAR architecture, when using *e5-small-v2* [80] as the text encoder. It has approximately 47.26M parameters, most of which come from the text encoder. When unfreezing 6 layers of the text encoder, about 24.70M parameters are tuned, with the remaining 11.92M embedding parameters and 10.65M layer ones being kept frozen.

Table 7: Parameter counts for TabSTAR components

Module	# Parameters
Encoding: Semantic	33,360,000
Encoding: Numerical	296,832
Fusion	1,774,464
Interaction	10,646,784
Prediction	1,185,794

To describe the architecture more precisely, we start by defining the dataset formally. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ denote a tabular dataset with n examples. Each example $x_i = [x_{i1}, \dots, x_{im}]$ has m features. The target variable y_i is either continuous (regression) or discrete (classification) taking one of C classes. For simplicity, we describe the architecture at the example level, though all computations are actually carried out on mini-batches of size B . The batches are always drawn from a single dataset in both pretraining and finetuning, removing any need for padding.

Verbalization We denote by t the number of target entries where $t = C$ for classification and $t = 1$ for regression. We then form a raw sequence of length $e = t + m$ by listing the t target values, followed by the m feature entries. Each element j in this sequence is then verbalized into a semantic string s_j and a numerical value n_j , set to be the clipped z-score for numerical non-missing features, and zero otherwise. The example is thus represented by parallel sequences (\mathbf{s}, \mathbf{n}) of length e .

Encoding Each semantic string s_j and numerical value n_j are projected into a d -dimensional vector. Semantic strings are encoded with an encoder-only language model (*e5-small-v2* [80]). Each string is tokenized, passed through the model, and pooled to produce its final embedding. This process is independent between elements, i.e. the attention is at the token level within a single element. In parallel, each numeric value is fed through a two-layer MLP that first projects from 1 to $2d$ dimensions, applies a ReLU and dropout, and then projects back down to d . This produces matching d -dimensional embeddings for each of the e elements, ready to be fused.

Fusion To unify semantic and numerical embeddings into a single representation, we apply a single-layer Transformer Encoder²¹ over each element’s pair of vectors. Concretely, for each element

²¹With 2 attention heads, a feed-forward hidden size of $4d$, dropout 0.1, and ReLU activation.

we stack its d -dimensional text and numeric embeddings and feed them through the encoder layer. For every element, the attention is applied between its two representations, and we average the two outputs to produce one fused d -dimensional embedding. This yields a final sequence of length e and dimension d , which will serve as tokens for the Interaction block.

Interaction The fused sequence of e tokens is processed by a standard Transformer Encoder with model dimension $d = 384$, $L = 6$ layers, 6 attention heads per layer, feed-forward size $4d$, dropout 0.1, ReLU activation and using a pre-norm configuration. Unlike in language modeling, feature ordering is irrelevant, so no positional encodings are used. The encoder produces contextualized embeddings for every position, and we retain the t target embeddings for the prediction.

Prediction We initialize two identical MLP heads, one for regression and one for classification. Each of them consists of a hidden layer of size $4d$ (with ReLU activation) followed by a linear projection to a single output. For each dataset, we choose the relevant head and process the t target token embeddings. For classification, we independently feed each one of the $t = C$ target tokens to the classification head to obtain a score (logit) per class. Notably, the same head is shared across classes and datasets. We apply softmax over these scores, yielding a probability distribution regardless of the number of classes. For regression, the single target token is projected into a real value. Note that the heads is shared between datasets, as regression outputs are always clipped z-scores.

A.3 Architecture Experiments

In this section we explore the effect of different design choices for TabSTAR’s architecture. For each experiment, we only vary the parameter of interest, keeping everything else fixed. We follow the same pretraining regime as in Appendix B.1, except that for computational efficiency we train only 25 epochs (instead of 50) with 128 pretraining datasets (instead of 390). We evaluate each variant relying solely on pretraining performance, as an approximation for downstream task performance. We acknowledge that our conclusions might depend on this limited scale, hence we discuss a subset of the experiments briefly to reflect the depth of our work and inspire future research.

The Fusion Block’s Mechanism For the fusion block, we consider two simpler alternatives to the attention mechanism, both of them underperforming: (1) Concatenation, by concatenating the semantic and numerical d -dimensional vectors into a $2d$ -dimensional vector, and projecting them back via an MLP, and (2) Multiplication, by multiplying the semantic representation directly with the numerical value²² in a straightforward, parameter-free manner as in [83, 86].

The Number of Interaction Layers We experiment with the number of encoder layers, and observe that 3 yields anecdotally worse performance than 6, with lower parameter count. Nevertheless, we prioritize a deeper network as for datasets with very complex relationships we believe that this might be beneficial. Additionally, we try a 9-layer variant which performs significantly worse, while also increasing the parameter count.

Row-Level Attention We experiment with adopting the architecture proposed by SAINT [69], which adds row-level attention to each encoder layer. Similar concepts are also employed by models that get the input the whole dataset, labels included [38, 49]. We run experiments with 2, 4 and 6 layers as they are roughly equivalent in parameter count to 3, 6, and 9 layers without row-attention. We observe no substantial gain, and thus we prioritize the simpler solution, as row-level attention is sensitive to the batch size and adds complexity to inference time.

Verbalization We experiment with various verbalization strategies. For features, including the column name alongside the value consistently outperforms using the value alone. For target-aware tokens, we find that explicitly verbalizing all possible target values yields better performance than omitting them. In contrast, adding a dataset-level description prefix to improve contextualization offers no measurable advantage. Interestingly, TabSTAR maintains high performance even when no semantic metadata is present. We see this robustness as a strength.

²²After rescaling it to be centered around 1 rather than 0, using a learned scaling factor.

B Training

In this section we elaborate on the two stages of TabSTAR’s training: pretraining and finetuning, presented in §3. As in Appendix A.3, we summarize key pretraining experiments.

B.1 Pretraining

TabSTAR is pretrained employing supervised learning in a multi-task regime, jointly learning regression, binary and multiclass classification tasks. The parameters of the architecture are fully-shared, without any need for dataset-specific parameters. Every example during the pretraining updates all the model’s weights, with the sole exception of the prediction head, for which every example uses its respective head depending on the nature of the task (classification or regression).

Sampling For computational efficiency, each dataset is subsampled once before the pretraining. At the example level, we sample up to 300,000 examples from each dataset, stratified by the target variable for classification tasks. Since we only use a fraction of each dataset for each pretraining epoch, this decision has negligible influence. In addition, we randomly sample up to 200 features per dataset. While straightforward, this decision is suboptimal as feature importance isn’t taken into consideration. As this work does not focus on wide-feature datasets, we consider this trade-off acceptable. Importantly, this setup is enabled during finetuning as the TabSTAR architecture is agnostic to the number of features. We split each dataset into train-validation splits (95%-5%),²³ without any need for test splits, and cap the validation set at a maximum of 1,000 examples used for evaluating pretraining performance.

Batching Every epoch, we randomly sample up to 2,048 examples from each dataset in mini-batches of 32, and shuffle all the batches. We conduct gradient accumulation and update the model every 4 steps to reduce the chances of a single update being dominated by a single dataset, so the global batch size is effectively 128. Appendix B.3.1 elaborates on the effect of batch size.

Metrics Our loss function is cross-entropy for classification, and MSE for regression. With standardized targets, $R^2 \approx 1 - \text{MSE}$, although this equivalence is degraded by clipping targets in preprocessing. We train with mixed-precision and apply gradient clipping to stabilize training without task-specific weights, with Appendix B.3.2 discussing the limitations of this approach. We use as metrics AUROC for classification and R^2 , so for each task the optimal metric value is 1. We average performance across all datasets into a single metric that reflects the pretraining performance.

Training We pretrain for 50 epochs with the *OneCycleLR* [67] optimizer, with warmup during the first 5 epochs (10%) and cosine annealing. Early stopping is conducted after 3 epochs without improvement on the pretraining metric. The weight decay is set to 0.001, and a max learning rate of $lr = 5 \times 10^{-5}$ is applied uniformly across all layers. Appendix B.3.3 discusses experiments with differential learning rate. Pretraining running time varies depending on the number of epochs and the included datasets. The full models (390 datasets) reported in §5 train for less than 48 hours on a single NVIDIA A40 GPU (48GB memory), and we believe that this could be optimized much further.

B.2 Finetuning

We finetune downstream tasks using LoRA’s implementation of the *peft* package²⁴. We use a rank of $r = 32$, set $\alpha = 2r = 64$ and $dropout = 0.1$. We employ the same scheduler as in the pretraining phase, with the only difference being that we set $lr = 0.001$, and increase the patience parameter for early stopping to 5. We apply a train-test split of 90%-10% and sample a validation set of 10%. As opposed to the pretraining, all batches are drawn from the same dataset. Therefore, we observe no effect from changing the mini-batch size when keeping the global batch size fixed to 128. We tune 1,597,440 out of TabSTAR’s 47,263,874 parameters (3.4%), spanning all blocks of the architecture. The frozen layers of the text encoder do not receive LoRA adapters.

Finetuning hyperparameters are selected by pretraining TabSTAR over 256 datasets and performing grid-search over a held-out set of 25 downstream tasks disjoint from the 50 datasets in the benchmark

²³We choose only 5% for efficiency, as we use hundreds of pretraining datasets.

²⁴<https://github.com/huggingface/peft>

evaluated in §4. The search space is presented in Table 8, and we observe that average performance is relatively robust across this space. An interesting observation is that decreasing the number of parameters by setting $r = 16$ mildly hurts performance, but it has upsides on memory and latency aspects, allowing a future trade-off exploration. As a final note, we argue that providing a strong default configuration for TFMIs is crucial for evaluating them, but for real-world applications, it is still recommendable to find the best hyperparameters tailored to the downstream task.

Table 8: LoRA hyperparameter tuning grid search for TabSTAR’s finetuning.

Hyper-parameter	Search Space
LoRA rank (r)	16, 32, 64
Learning Rate	0.0005, 0.001, 0.002, 0.005, 0.01
Dropout	0, 0.1

The only experiment in this paper where we employ full finetuning instead of LoRA is for the non-pretrained variant discussed in the analysis in §6. For this variant we fully finetune the pretrained model on each downstream task. Compared to the pretraining, we use $lr = 2.5 \times 10^{-5}$ and increase the patience to 5. These hyperparameters are lightly tuned using the same procedure as for LoRA, and we observe that fully finetuning the model achieves comparable performance, except for small datasets, where training is more prone to overfitting.

B.3 Pretraining Experiments

In this section, we briefly elaborate on some experiments performed over TabSTAR’s pretraining protocol. As in Appendix A.3, we highlight only a subset of them in a high-level manner.

B.3.1 Batch Size

During pretraining, we use a mini-batch size of 32, each of them drawn from a single dataset. Since we train with gradient accumulation and a global batch size of 128, varying the batch size affects the diversity of a single model update: lower batch sizes are likely to be exposed to more datasets. We decrease the batch size to 16 and 8 and observe an improvement at the cost of slower training. An interesting direction for future work is moving to mixed-datasets batches, which require more complex implementation but might benefit from more regularized learning. Such approach, however, goes against row-level attention methods and ICL, as discussed in Appendix A.3.

B.3.2 Loss Weights

Pretraining the model over hundreds of datasets in a multi-task regime presents a key challenge: the loss scale of each dataset can vary substantially, depending on task difficulty or task type. For example, a multiclass classification task with dozens of classes will naturally yield a higher average loss than a binary task. These dynamics can also shift during training. Our default approach naively averages the loss across all datasets, which risks over-weighting tasks for potentially arbitrary reasons.

To address this, we explore two alternative weighting strategies: (1) Assigning a constant weight per task type, accounting for the number of classes in classification tasks, and (2) Normalizing each dataset’s contribution by the best loss achieved by CatBoost [60] when being fitted to that dataset. While these strategies better reflect task-specific characteristics, they hardly impact performance and introduce additional complexity. Notably, adjusting loss weights across tasks impacts metric interpretability, as each weighting scheme implicitly optimizes a different objective.

We do not explore more sophisticated methods such as learning per-dataset weights, as these often require mixed-dataset batches and introduce additional learnable parameters. We believe, however, that multi-task pretraining over tabular datasets remains an open and important research question.

B.3.3 Differential Learning Rate

TabSTAR’s weights initialization is not balanced: the textual encoder is a pretrained embedding model while the rest of the architecture parameters are randomly initialized. To counteract this imbalance, we experiment with using differential learning rates for the textual encoder layers, and experiment

with scaling it by a factor of 0.5 and of 0.75. To our surprise, this decision hurts performance, so we stick to a uniform learning rate across all layers.

C Training Datasets

In this appendix we briefly expand on the pretraining corpus elaborated in §4. It is composed of 350 datasets, spanning all the datasets appearing in *AMLB* [25], *OpenML-CTR23* [23], *TabZilla* [55] and the ones presented by *Grinsztajn* [31]. After deduplication,²⁵ this results in 152 datasets (94 classification, 58 regression). Interestingly, only 6 of these 152 datasets have free-text or high-cardinality features. We manually add datasets from OpenML [78] and Kaggle, as well as from the *AutoML-Benchmark-Train* [22] corpus, and achieve a total of 350 datasets, with 49 textual datasets.

Table 9 details the 253 classification datasets and Table 10 the 97 regression ones. We elaborate the *Dataset* name, the number of examples n , the number of features m , and the number of classes C for classification. In addition, we mark datasets that belong to one of the benchmarks, and the ones that have text features. Importantly, the textual flag is quite permissive, as it includes features with relatively short texts or potentially low predictive power (e.g., people names or addresses).

Table 9: The 253 Classification Datasets of the Pretraining Corpus, with their n examples, m features, C classes, presence in a benchmark (B) and whether they are textual (T).

Dataset	n	m	C	B	T
KDDCup99	4,898,422	40	20	✓	
mimic_extract_los_3	4,155,270	17	68		✓
Online-P2P-Lending	2,875,146	16	5		
sf-police-incidents	2,215,023	8	2	✓	✓
physionet_sepsis	1,552,210	42	2		
poker-hand	1,025,009	10	10	✓	
Higgs	1,000,000	28	2	✓	
BAF_base	1,000,000	30	2		
Credit_Card_Fraud_	1,000,000	7	2		
Harry-Potter-fanfiction-data	648,493	13	4		✓
porto-seguro	595,212	57	2	✓	
covertype	581,012	54	7	✓	
AVIDa-hIL6	573,891	3	2		✓
airlines	539,383	7	2	✓	✓
HolisticBias	472,991	14	4		✓
albert	425,240	78	2	✓	
DBPedia	342,781	3	219		✓
hcdr_main	307,511	120	2		
Mental_Health_Dataset	292,364	16	5		
Kuzushiji-49	270,912	784	49		
spoken-arabic-digit	263,256	14	10		
cdc_diabetes	253,680	21	2		
skin-segmentation	245,057	3	2		
LT-Vehicle-Loan-Default-Prediction	233,154	38	2		
Churn_Telco_Europa	190,776	17	2		
ldpa	164,860	6	11		
Give-Me-Some-Credit	150,000	10	2	✓	
walking-activity	149,332	4	22		
social_bias_frames	144,649	16	3		✓
Wikipedia_Talk_Labels	140,379	12	15		✓
Municipal-Debt-Risk-Analysis	138,509	13	2		
MiniBooNE	130,064	50	2	✓	
nba-shot-logs	128,069	15	2		✓

Continued on next page

²⁵And the exclusion of the *fifa* dataset, which is included in the benchmark.

Table 9: The 253 Classification Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	<i>C</i>	B	T
college_scorecard	124,699	117	2		
drug-directory	120,215	16	7	✓	
TVS_Loan_Default	119,528	29	2		
road-safety	111,762	32	2	✓	
Diabetes130US	101,766	46	3	✓	
fars	100,968	29	8		
Credit_Score_Classification	100,000	26	3		✓
numerai28.6	96,320	21	2	✓	
Run_or_walk_information	88,588	6	2		
jannis	83,733	54	4	✓	
KDD98	82,318	477	2		
APSFailure	76,000	169	2	✓	
kick	72,983	32	2	✓	✓
human-choice-prediction	71,579	20	2		✓
Traffic_violations	70,340	20	3		✓
Fashion-MNIST	70,000	784	10	✓	
Cardiovascular-Disease-dataset	70,000	11	2		
mnist_784	70,000	719	10		
connect-4	67,557	42	3	✓	
mobile_churn	66,469	63	2		
helena	65,196	27	100	✓	
LICD	63,634	413	2		✓
CIFAR_10	60,000	3,072	10		
REASONER	58,497	34	2		✓
volkert	58,310	147	10	✓	
shuttle	58,000	9	7	✓	
GTSRB-HueHist	51,839	256	43		
okcupid-stem	50,789	19	3	✓	✓
KDDCup09-Upselling	50,000	13,419	2	✓	
KDDCup09_appetency	50,000	207	2	✓	
adult	48,842	14	2	✓	
League-of-Legends-Diamond	48,651	14	2		
tamilnadu-electricity	45,781	2	20		
bank-marketing	45,211	16	2	✓	
meta_stream_intervals	45,164	74	11		
jungle_chess	44,819	6	3	✓	
Dynamically-Generated-Hate-Speech-Dataset	41,144	8	2		✓
Breast-cancer-prediction	39,998	11	2		
Click_prediction_small	39,948	11	2	✓	
Hotel-Reviews	38,932	3	2		✓
electricity	38,474	8	2	✓	
nomao	34,465	118	2	✓	
Employee-Turnover-at-TECHCO	34,452	9	2		
Amazon_employee_access	32,769	9	2	✓	
Credit-Risk-Dataset	32,581	11	2		
Default-of-Credit-Card-Clients-Dataset	30,000	23	2	✓	
funpedia	29,819	3	3		✓
credit_risk_china	27,522	27	5		
Insurance	23,548	10	2		
guillermo	20,000	4,281	2	✓	
riccardo	20,000	4,283	2	✓	
insurance_dataset	20,000	26	4		
letter	20,000	16	26		
game-of-thrones-script-all-seasons	16,825	5	43		✓

Continued on next page

Table 9: The 253 Classification Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	<i>C</i>	B	T
NewspaperChurn	15,855	16	2		✓
mozilla4	15,545	5	2		
pol	15,000	26	11	✓	
eeg-eye-state	14,980	14	2		
MagicTelescope	13,376	10	2	✓	
nursery	12,958	8	4		
online-shoppers-intention	12,330	17	2		
Disaster-Tweets	11,370	4	2		✓
mammography	11,183	6	2		
PhishingWebsites	11,055	30	2	✓	
Binary-Dataset-of-Phishing-and-Legitimate-URLs	11,000	14	2		
pendigits	10,992	16	10		
WBCAtt	10,298	11	5		
artificial-characters	10,218	7	10	✓	
internet_usage	10,108	71	46		✓
robert	10,000	7,200	10	✓	
dilbert	10,000	2,000	5	✓	
shrutime	10,000	10	2		
JapaneseVowels	9,961	14	9		
GesturePhaseSegmentationProcessed	9,873	32	5	✓	
FICO-HELOC-cleaned	9,871	23	2	✓	
IBRD_Loans_Classification	9,215	6	10		
Indian_pines	9,144	220	8		
SpeedDating	8,378	120	2	✓	✓
fabert	8,237	795	7	✓	
mushroom	8,124	21	2		
isolet	7,797	617	26		
eye_movements	7,608	23	2	✓	
twonorm	7,400	20	2		
blastchar	7,043	19	2		
musk	6,598	167	2		
first-order-theorem-proving	6,118	51	6	✓	
HMEQ_Data	5,960	12	2		
philippine	5,832	308	2	✓	
optdigits	5,620	62	10		
BachChoralHarmony	5,586	15	68		
page-blocks	5,473	10	5		
wall-robot-navigation	5,456	24	4		
christine	5,418	1,611	2	✓	
phoneme	5,404	5	2	✓	
Is_fraud	5,227	19	2		✓
sylvine	5,124	20	2	✓	
Satellite	5,100	36	2	✓	
Multiclass_Classification_for_Corporate_Credit	5,000	7	10		
Personal-Loan-Modeling	5,000	12	2		
churn	5,000	20	2	✓	
waveform-5000	5,000	40	3		
air-quality-and-pollution-assessment	5,000	9	4		
Heart_Failure_Prediction	5,000	12	2		
compas-two-years	4,966	11	2	✓	
wine-quality-white	4,898	11	7	✓	
wilt	4,839	5	2	✓	
spambase	4,601	57	2		
StackOverflow-polarity	4,423	1	3		✓

Continued on next page

Table 9: The 253 Classification Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	<i>C</i>	B	T
hiva_agnostic	4,229	1,617	2		
Fraud-Detection-Updated	4,156	27	2		
ada	4,147	46	2	✓	
analcatdata_supreme	4,052	7	10		
hypothyroid	3,770	27	3		
Bioresponse	3,751	1,776	2	✓	
Internet-Advertisements	3,279	1,558	2	✓	
led24	3,200	24	10		
kr-vs-kp	3,196	36	2	✓	
splice	3,190	60	3	✓	
dna	3,186	180	3	✓	
gina	3,153	970	2	✓	
madeline	3,140	259	2	✓	
jasmine	2,984	144	2	✓	
cjs	2,796	29	6		
madelon	2,600	500	2		
ozone-level-8hr	2,534	72	2	✓	
segment	2,310	16	7	✓	
cardiotocography	2,126	23	10		
Estimation_of_Obesity_Levels	2,111	16	7		
kc1	2,109	21	2	✓	
Corporate-Credit-Rating	2,026	30	8		✓
mfeat-factors	2,000	216	10	✓	
South_Asian_Churn_dataset	2,000	13	2		
mfeat-zernike	2,000	47	10	✓	
mfeat-fourier	2,000	76	10	✓	
pbcseq	1,945	18	3		
steel-plates-fault	1,941	27	7	✓	
car	1,728	6	4	✓	
GAMETES_Heterogeneity	1,600	20	2		
one-hundred-plants-texture	1,599	64	100	✓	
audit-data	1,552	35	2		
OVA_Breast	1,545	10,935	2		
amazon-commerce-reviews	1,500	10,000	50	✓	
yeast	1,484	8	10	✓	
cmc	1,473	9	3	✓	
ibm-employee-attrition	1,470	31	2		
pc4	1,458	37	2	✓	
Data_Science_Nigeria_Telecoms_Churn	1,400	14	2		
hepatitis_c_virus_hcv_for_egyptian_patients	1,385	28	4		
Bank-Note-Authentication-UCI	1,372	4	2		
baseball	1,340	16	3		
Titanic	1,309	13	2		✓
mental-health-in-tech-survey	1,259	26	2		✓
hill-valley	1,212	100	2		
Heart-Disease-Dataset-(Comprehensive)	1,190	11	2	✓	
volcanoes-e1	1,183	3	5		
Airlines-Tweets-Sentiments	1,097	1	3		✓
MiceProtein	1,080	77	8		
cnae-9	1,080	856	9	✓	
solar_flare	1,058	9	5	✓	
qsar-biodeg	1,055	41	2	✓	
SOCC	1,043	13	4		✓
rmftsa_sleepdata	1,024	2	4		

Continued on next page

Table 9: The 253 Classification Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	<i>C</i>	B	T
autoUniv-au1-1000	1,000	20	2		
collins	1,000	19	30		
credit-g	1,000	20	2	✓	
vowel	990	12	11		
The-Estonia-Disaster-Passenger-List	989	6	2		✓
xd6	973	9	2		
tokyo1	959	42	2		
tic-tac-toe	958	9	2		
Tour-and-Travels-Customer-Churn-Prediction	954	6	2		
acp-breast-cancer	949	1	4		✓
oil_spill	937	48	2		
anneal	898	18	5		
Cervical_Cancer_Risk_Factors	858	30	5		
vehicle	846	18	4	✓	
analcatdata_authorship	841	70	4		
glioma_grading_clinical_and_mutation_features	839	23	2		
analcatdata_dmft	797	4	6		
regensburg_pediatric_appendicitis	780	55	3		
QSAR_Bioconcentration_classification	779	12	3		✓
Diabetes_Dataset	768	8	2		
blood-transfusion-service-center	748	4	2	✓	
eucalyptus	736	19	5	✓	
breast-w	699	9	2		
Australian	690	14	2	✓	
soybean	683	35	19		
profb	672	8	2	✓	
Student_Performance	666	11	4		
balance-scale	625	4	3	✓	
Loan-Predication	614	11	2		
monks-problems-2	601	6	2	✓	
synthetic_control	600	60	6		
ilpd	583	10	2		
micro-mass	571	1,082	20	✓	
wdbc	569	30	2		
arsenic-male-lung	559	4	2		
cylinder-bands	540	34	2		
climate-model-simulation-crashes	540	18	2		
water-treatment	527	36	2		
Early-Stage-Diabetes-Risk-Prediction-Dataset	520	16	2		
dresses-sales	500	12	2		
irish	500	5	2		
arrhythmia	443	262	10		
wholesale-customers	440	7	2		
vote	435	16	2		
cars	406	7	3		
chronic-kidney-disease	400	25	2		
differentiated_thyroid_cancer_recurrence	383	16	2		
colic	368	26	2	✓	
breast-cancer	286	9	2		
qualitative-bankruptcy	250	6	2		
us-2020-presidential-election-speeches	245	5	7		✓
audiology	192	57	8	✓	
bone_marrow_transplant_children	187	36	2		
darwin	174	450	2		

Continued on next page

Table 9: The 253 Classification Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	<i>C</i>	B	T
tae	151	5	3		
EgyptianSkulls	150	4	5		
lymph	148	18	3	✓	
arcene	100	9,920	2	✓	

 Table 10: The 93 Regression Datasets of the Pretraining Corpus, with their *n* examples, *m* features, presence in a benchmark (*B*) and whether they are textual (*T*).

Dataset	<i>n</i>	<i>m</i>	B	T
delays_zurich_transport	5,465,575	14	✓	
New-York-Citi-Bike-Trip	4,500,000	7		
USA-Airport-Dataset	3,606,803	14		✓
New-York-Taxi-Trip	2,083,778	21		
Buzzinsocialmedia_Twitter	583,250	77	✓	
nyc-taxi-green-dec-2016	581,835	18	✓	
515K-Hotel-Reviews-Data-in-Europe	515,738	16		✓
dionis	416,188	54	✓	
Yolanda	400,000	100	✓	
Allstate_Claims_Severity	188,318	130	✓	
Football_players_Fifa_stats	183,142	37		
black_friday	166,821	9	✓	
medical_charges	163,065	3	✓	
football-manager-data	159,541	87		✓
wave_energy	72,000	32	✓	
video_transcoding	68,784	18	✓	
dating_profile	59,946	30		✓
diamonds	53,940	9	✓	
sarcos	48,933	21	✓	
physicochemical_protein	45,730	9	✓	
fried	40,768	10		
2dplanes	40,768	10		
mv	40,768	10		
Perth-House-Prices	33,656	17		✓
cps88wages	28,155	6	✓	
fps_benchmark	24,624	39	✓	
news_popularity2	24,007	4		✓
house_16H	22,784	16	✓	
health_insurance	22,272	11	✓	
house_sales	21,613	21	✓	
superconductivity	21,263	81	✓	
california_housing	20,640	8	✓	
avocado-sales	18,249	11		
Bike_Sharing_Demand	17,379	12	✓	
elevators	16,599	18	✓	
FIFA20-Players	14,999	72		✓
miami_housing	13,932	15	✓	
naval_propulsion_plant	11,934	14	✓	
Brazilian_houses	10,692	11	✓	
German-House-Prices	10,552	24		✓
sulfur	10,081	5	✓	
climate_change_impact	10,000	14		
grid_stability	10,000	12	✓	

Continued on next page

Table 10: The 93 Regression Datasets of the Pretraining Corpus.

Dataset	<i>n</i>	<i>m</i>	B	T
Credit-Card-Dataset-for-Clustering	8,949	16		
topo_2_1	8,885	261	✓	
yprop_4_1	8,885	212	✓	
seoul_bike_sharing_demand_cat	8,760	13		
pumadyn32nh	8,192	32	✓	
kin8nm	8,192	8	✓	
cpu_activity	8,192	21	✓	
bank32nh	8,192	32		
Pollen-Luxembourg-1992-2018	7,784	36		
colleges	7,063	44	✓	✓
wind	6,574	14		
QSAR-TID-10980	5,766	1,024	✓	
QSAR-TID-11	5,742	1,024	✓	
Myanmar-Air-Quality	5,122	10		
Santander_transaction_value	4,459	4,735	✓	
SAT11-HAND-runtime-regression	4,440	114	✓	
Mercedes_Benz_Greener_Manufacturing	4,209	364	✓	
abalone	4,177	8	✓	
pollen	3,848	4		
space_ga	3,107	6	✓	
scotch-whiskey-reviews-update-2020	2,247	4		✓
quake	2,178	3	✓	
auction_verification	2,043	7	✓	
us_crime	1,994	126	✓	
airfoil_self_noise	1,503	5	✓	
house_prices	1,460	80		
house_prices_nominal	1,460	79	✓	
NBA-PLAYERS-2016-2019	1,408	43		✓
Insurance-Premium-Data	1,338	6		
Moneyball	1,232	14	✓	
socmob	1,156	5	✓	
MIP-2016-regression	1,090	116	✓	
geographical_origin_of_music	1,059	116	✓	
concrete_compressive_strength	1,030	8	✓	
Household-monthly-electricity-bill	1,000	9		
stock	950	9		
QSAR_fish_toxicity	908	6	✓	
cars	804	17	✓	
energy_efficiency	768	8	✓	
kdd_el_nino-small	709	8		
student_performance_por	649	30	✓	
strikes	625	6		
sensory	576	11	✓	
meta	528	21		
forest_fires	517	12	✓	
rmftsa_ladata	508	10		
boston	506	13	✓	
no2	500	7		
Diabetes(scikit-learn)	442	10		
NBA-2k20-player-dataset	439	14		✓
baseball-hitter	263	22		
bodyfat	252	14		
Lisbon-House-Prices	246	13		
tecator	240	124	✓	

D Benchmark Datasets

This appendix elaborates on the benchmark presented in §4. We consider all datasets proposed by *AutoML Multimodal Benchmark* (SHI) [63], *Vectorizing* (VEC) [32], and *CARTE-Benchmark* (CRT) [47], resulting in a final set of 50 datasets. We deduplicate datasets that appear as-is in more than one benchmark. In addition, since CARTE explores the concept of multi-table learning, they introduce highly-overlapping datasets for which we remove one variant (see 4.3 and B.2 in their paper).

Table 11 presents the classification datasets and Table 12 the regression ones. Each table includes an internal *ID* used for reference, the *Dataset* name, the number of examples n and of features m , and the number of classes C for classification.²⁶ Finally, we also indicate the benchmark sources where each dataset appears. In addition, Table 13 presents the full benchmark with a short description per dataset, and Table 14 details the datasets removed during the deduplication process. Most of the excluded datasets are regression datasets from the *CARTE-Benchmark*, because of its high-overlapping nature.

Table 11: The 14 classification datasets of the benchmark, with their n examples, m features, C classes, and presence in the SHI, VEC and CRT benchmarks.

ID	Dataset	n	m	C	SHI	VEC	CRT
C01	women_clothing_review	18,788	10	5	✓		
C02	us-accidents	7,728,394	42	4		✓	✓
C03	data_scientist_salary	15,841	6	6	✓		
C04	imdb_genre_prediction	800	11	2	✓		
C05	product_sentiment_machine_hack	5,091	2	4	✓		
C06	google_qa_question_type_reason	4,863	39	5	✓		
C07	michelin-guide-restaurants-2021	17,735	11	5			✓
C08	fake_job_postings2	12,725	5	2	✓		
C09	jigsaw_unintended_bias100K	100,000	40	2	✓		
C10	yelp-reviews-dataset	10,000	5	5			✓
C11	news_channel	20,284	17	6	✓		
C12	wine_reviews	84,123	5	30	✓	✓	✓
C13	kick_starter_funding	86,502	9	2	✓		
C14	melbourne_airbnb	18,316	89	10	✓		

E Baselines

In this appendix we first discuss models excluded from the evaluation due to data leakage concerns, and then cover implementation details for the baselines used in our main experiments §4.

E.1 Excluded Baselines

As opposed to GDBTs or single-dataset deep learning methods, evaluating pretrained tabular models introduces additional complexity. Indeed, leakage can come in multiple forms. When LLMs are involved, there is a risk of memorization [9], and models trained on synthetic datasets [38] which try to mimic real-world distributions, can be unintentionally biased towards popular benchmarks.

While these two forms of leakage are subtle and hard to detect, a more direct form must be strictly avoided: When the same dataset (or a variant of it) is used during pretraining, and then it is evaluated as a downstream task. In such scenario there is inevitable severe data leakage, especially when running with multiple random test splits. The rest of the section explains how both *TP-BERTa* [84] and *CM2* [86] suffer from such contamination with respect to our benchmark. As we briefly mention in §7, we advocate for improving TFM research by encouraging models that are practical to evaluate, by releasing several versions of each model, and providing default hyperparameters.

TP-BERTa We exclude TP-BERTa from our evaluation for two key reasons. First, their implementation assumes that every example is treated as a serialized single sequence, which allows for

²⁶We treat ranking problems with up to 10 discrete values as multiclass problems.

Table 12: The 36 regression datasets of the benchmark, with their n examples, m features, and presence in the SHI, VEC and CRT benchmarks.

ID	Dataset	n	m	SHI	VEC	CRT
R01	used-cars-dataset-cardekho	37,814	112		✓	
R02	second-hand-mercedes-benz	16,392	7		✓	
R03	animeplanet-recommendation	14,391	14		✓	
R04	ML/DS-Salaries	119,628	9		✓	
R05	Babies-R-Us	5,085	12		✓	
R06	employee_salaries	9,228	11	✓	✓	
R07	spotify-tracks-dataset	114,000	18	✓		
R08	california_house_price	37,951	39	✓		
R09	fifa	19,178	28		✓	
R10	coffee-scrap-coffeereview	2,440	17		✓	
R11	BikeWale	9,003	6	✓	✓	
R12	used-car-prices-in-pakistan	72,655	9		✓	
R13	bookprice_prediction	4,989	8	✓		
R14	ae_price_prediction	22,662	12	✓		
R15	Employee-remuneration	44,574	5	✓	✓	
R16	filmtv-movies-dataset	41,399	17		✓	
R17	free-7-million-company-dataset	7,173,426	7	✓	✓	
R18	museums	22,290	21		✓	
R19	vivino-wine-data	8,650	6		✓	
R20	wikiliq-dataset	12,569	12		✓	
R21	beer-profile-and-ratings	3,197	24		✓	
R22	korean-drama	1,647	9		✓	
R23	videogamesales	16,598	5		✓	
R24	zomato-bangalore-restaurants	41,665	15	✓	✓	
R25	the-movies-dataset	45,460	20		✓	
R26	nba-draft-basketball	1,669	22		✓	
R27	Goodreads	3,967	14	✓		
R28	Rotten-Tomatoes	7,158	15		✓	
R29	saudi-arabia-used-cars-dataset	8,035	12		✓	
R30	top-ramen-ratings-2022	4,105	4	✓	✓	
R31	Journal-Score-SJR	31,136	21		✓	
R32	chocolate-bar-ratings	1,795	8		✓	
R33	mercari_price_suggestion100K	100,000	9	✓		
R34	wine-price-on-polish-market	2,247	18		✓	
R35	clear-corpus	4,724	30	✓	✓	
R36	jc_penney_products	10,860	5	✓		

a maximum length of 512 tokens, as elaborated in Appendix A.1. While this decision is efficient for datasets with a low amount of features and no free-text presence, around half of the datasets in our benchmark are too long for that limitation, as they either contain too many features or long free-texts. Second, TP-BERTa’s pretraining uses datasets that appear in our evaluation set, as listed in Table 6 of their paper [84]. It is evident that several datasets overlap directly with datasets in our benchmark §4 (e.g., *1510_fifa*, *1368_IMDb-Ratings*, *1639_Melbourne*), disqualifying them for our purposes. Furthermore, we observe a concerning overlap between their pretraining and downstream task datasets (e.g., *airlines*, *sf police* and *diabetes*). We believe that this questions the validity of their evaluation, and that such contamination poses a serious challenge for the TFM community which could be substantially addressed by better tabular data repositories [76].

CM2 CM2 was pretrained over *OpenTabs*, a compilation of more than 2,000 datasets drawn from public tabular data repositories, including OpenML and Kaggle. While this collection is valuable, pretraining a model over these datasets compromises further evaluation of any of them. Naturally, the overlap with our benchmark here is extremely high, making it infeasible to use as a

Table 13: Benchmark Datasets Description

ID	Description
C01	Women Clothing E-Commerce Reviews
C02	US Accidents between 2016 and 2023
C03	Indian Data Scientist Salary Prediction
C04	IMDB Movies Genre Prediction
C05	Product Sentiment Analysis
C06	Google QA Question Type Reason Explanation
C07	Michelin Guide Restaurants Awards
C08	Fake Job Posting Detection
C09	Online Social Media Comments Toxicity
C10	YELP Dataset Reviews
C11	News Channel Prediction
C12	Wine Reviews for Variety Prediction
C13	Kickstarter Funding Prediction
C14	Melbourne AirBnB Listings
R01	User cars and listing price in the website Cardekho
R02	Second-hand cars Mercedes Benz price Italy
R03	Anime-Planet Recommendation Database 2020
R04	Salaries of ML/DS Professionals Worldwide
R05	Prices Prediction for baby product from Babies R Us website
R06	Employee Salary in Montgomery County, MD
R07	Spotify Tracks Popularity
R08	California Houses 2020 Prices
R09	FIFA 2022 Players Wages
R10	Coffee Review Rating
R11	Bike and scooters from bikewale website in India
R12	Used car prices in Pakistan 2021
R13	Book Price Prediction
R14	American Eagle Retailer Price Prediction
R15	Employee Remuneration and Expenses - Vancouver
R16	FilmTV movies ataset rating
R17	Company size prediction
R18	General information on the US museums
R19	Vivino Spanish Wine Data
R20	WikiliQ - Alcohol dataset (May, 2022)
R21	Tasting profiles and consumer reviews for beers
R22	Korean Dramas
R23	Video Games Sales
R24	Zomato Restaurants in Bengaluru
R25	Metadata of movies released until 2017 for box-office revenues
R26	NBA Draft Basketball Player Data 1989-2021
R27	Books ratings
R28	Rotten Tomatoes Movie Ratings
R29	Saudi Arabia Used Cars Price from Syarah Website
R30	Ramen Ratings
R31	Academic impact for Scientific Journals
R32	Chocolate Bar expert ratings
R33	Mercari Online Marketplace Product Prices
R34	Information about wines on the polish market
R35	Readability scores for text passages spanning various genres and time periods
R36	JC Penney Product Prices in Retailer Website

Table 14: Excluded datasets, with their benchmark origin and reason for removal: (1) *Duplicate* dataset, (2) *Unavailable*, for datasets with inconvenient or unavailable hosting outside tabular repositories, and (3) *Pretraining*, for two (regression) datasets mistakenly used for the pretraining.

Dataset	Benchmark	Reason	Duplicate
google_qa_answer_type	SHI	Duplicate	google_qa_question_type
news_popularity2	SHI	Pretraining	
US Presidential	VEC	Unavailable	
Journal Influence	VEC	Duplicate	Journal-Score-SJR
Buy Buy Baby	CRT	Duplicate	Babies-R-Us
Bikedekho	CRT	Duplicate	BikeWale
Journal Score JCR	CRT	Duplicate	Journal-Score-SJR
Japanese Anime	CRT	Duplicate	animeplanet-recommendation
Mydramalist	CRT	Duplicate	korean-drama
Prescription Drugs	CRT	Unavailable	
Roger Ebert	CRT	Unavailable	
US Presidential	CRT	Unavailable	
Used Cars 24	CRT	Duplicate	used-car-prices-in-pakistan
Whisky	CRT	Pretraining	
Wine.com	CRT	Duplicate	wine_reviews
WineEnthusiasts	CRT	Duplicate	wine_reviews

baseline. Interestingly, their repository²⁷ lists TP-BERTa as a method trained on a subset of OpenTabs, reinforcing that the leakage is shared between the models.

E.2 Baselines Implementation and Hyperparameters

This section outlines the implementation and hyperparameter tuning strategy used for the baselines reported in §4. While each baseline has its own model-specific preprocessing pipeline, we apply two shared preprocessing steps to both TabSTAR (as detailed in Appendix A.1) and all the baselines: (1) We perform date preprocessing by using skrub’s *DatetimeEncoder*, and (2) Apply a clipped z-score transformation for target variables in regression datasets.

Textual Feature Handling CARTE natively supports textual inputs, and the TabPFN-v2 API client²⁸ does as well, although its implementation details remain undisclosed. As all other baselines do not natively support free-text features,²⁹ we preprocess these features into fixed-size embeddings using skrub’s *TextEncoder*, which internally applies a frozen *e5-small-v2* encoder to each semantic column. This aligns with the encoder used in TabSTAR, enabling a fair comparison across models. There are, however, two key differences in how TabSTAR handles these embeddings: First, the embeddings are specifically finetuned for the task, contributing significantly to its strong performance as shown in §5 and further analyzed in §G.2. The second detail is that skrub applies dimensionality reduction to 30 dimensions, as proposed by [32]. This compressed representation performs comparably to the full embedding space, while offering improved inference efficiency.

Hyperparameter Tuning For the tuned models we use the *Optuna* package³⁰ with random search, with a budget of 4 hours for every run and parallelizing trials on 8 CPU cores. We use 5-fold cross-validation and take the best configuration selected based on this mean score. We use it then to retrain the model on the full training data.

²⁷<https://github.com/Chao-Ye/CM2>

²⁸<https://github.com/PriorLabs/tabPFN-client>

²⁹CatBoost includes a built-in text module, but it underperforms compared to dense text embeddings.

³⁰<https://pypi.org/project/optuna/>

E.2.1 TabPFN-v2

We run TabPFN-v2 using their API client which supports text features.³¹ While the intrinsic details of their textual handling remain undocumented, it's reasonable to assume that it resembles the processing we apply to GBDTs, as their model leverages ICL and their architecture has no textual encoder.

E.2.2 TabICL

We run TabICL using its package³² with the default configuration. Since TabICL can only be evaluated on classification tasks, we exclude it from the regression analysis.

E.2.3 TabDPT

We run TabDPT using its package³³ with the default configuration. While it achieves state-of-the-art performance on some regression tasks, we encounter severe performance issues on others. Due to this inconsistency, we exclude TabDPT from the reported results and are in contact with the authors to investigate potential implementation issues.

E.2.4 CARTE

We run CARTE using its package,³⁴ which inherently performs k -fold cross-validation. After consulting with the authors, we set $k = 5$ for efficiency instead of 10. We do grid search over their recommended learning rates,³⁵ and we take the best-performing variant per dataset split.

E.2.5 RealMLP

We run RealMLP using its official implementation in the pytabkit package.³⁶ Following the advice of its authors, we disable label smoothing and optimize for *cross_entropy* for binary classification and *1-*auc_ovr** for multiclass classification, and keep all other default hyperparameters. For the tuned version, we follow the hyperparameter search space suggested by [39] as detailed in Table 15.

Table 15: RealMLP-Tuned Hyperparameters Search Space

Hyperparameter	Search Space
num_emb_type	Choice([None, PBLD, PL, PLR])
add_front_scale	Choice([True, False], p=[0.6, 0.4])
lr	$\log \mathcal{U}(2e-2, 3e-1)$
p_drop	Choice([0.0, 0.15, 0.3], p=[0.3, 0.5, 0.2])
act	Choice([ReLU, SELU, Mish])
hidden_sizes	Choice([[256, 256, 256], [64, 64, 64, 64], [512]], p=[0.6, 0.2, 0.2])
wd	Choice([0.0, 2e-2])
plr_sigma	$\log \mathcal{U}(0.05, 0.5)$ if classification else 0.1

E.2.6 CatBoost

We run CatBoost using the *catboost* package³⁷ and run the default configuration suggested by [26] by setting *early_stopping_rounds* = 50, *od_pval* = 0.001, *iterations* = 2000. For the tuned version, we follow the hyperparameter search suggested by [38] as detailed in Table 16.

³¹We use v2.0.8, the latest version available at the time of running the experiments.

³²<https://pypi.org/project/tabicl/>

³³<https://pypi.org/project/tabdpt/>

³⁴<https://github.com/soda-inria/carte>

³⁵ $\{2.5 \times 10^{-4}, 5 \times 10^{-4}, 7.5 \times 10^{-4}, 2.5 \times 10^{-3}, 5 \times 10^{-3}, 7.5 \times 10^{-3}\}$

³⁶<https://pypi.org/project/pytabkit/>

³⁷<https://pypi.org/project/catboost/>

Table 16: CatBoost-Tuned Hyperparameters Search Space

Hyperparameter	Search Space
learning_rate	$\log \mathcal{U}(e^{-5}, 1)$
random_strength	$\mathcal{U}\{1, 2, \dots, 20\}$
l2_leaf_reg	$\log \mathcal{U}(1, 10)$
bagging_temperature	$\mathcal{U}(0.0, 1.0)$
leaf_estimation_iterations	$\mathcal{U}\{1, 2, \dots, 20\}$
iterations	$\mathcal{U}\{100, 101, \dots, 4000\}$

E.2.7 XGBoost

We run XGBoost using the *xgboost* package.³⁸ For the default configuration, we follow the suggestion of [26] and use: *booster* = "gbtree", *early_stopping_rounds* = 50, *n_estimators* = 2000. For the tuned variant, we follow the hyperparameter search space suggested by [38], as shown in Table 17.

Table 17: XGBoost-Tuned Hyperparameters Search Space

Hyperparameter	Search Space
learning_rate	$\log \mathcal{U}(e^{-7}, 1)$
max_depth	$\mathcal{U}\{1, 2, \dots, 10\}$
subsample	$\mathcal{U}(0.2, 1)$
colsample_bytree	$\mathcal{U}(0.2, 1)$
colsample_bylevel	$\mathcal{U}(0.2, 1)$
min_child_weight	$\log \mathcal{U}(e^{-16}, e^5)$
alpha	$\log \mathcal{U}(e^{-16}, e^2)$
reg_lambda	$\log \mathcal{U}(e^{-16}, e^2)$
gamma	$\log \mathcal{U}(e^{-16}, e^2)$
n_estimators	$\mathcal{U}\{100, 101, \dots, 4000\}$

E.2.8 LightGBM

We run LightGBM using the *lightgbm* package.³⁹ and its default implementation. For the tuned variant, we follow the hyperparameter search space suggested by [38], as shown in Table 18.

Table 18: LightGBM-Tuned Hyperparameters Search Space

Hyperparameter	Search Space
num_leaves	$\mathcal{U}\{5, 6, \dots, 50\}$
max_depth	$\mathcal{U}\{3, 4, \dots, 20\}$
learning_rate	$\log \mathcal{U}(e^{-3}, 1)$
n_estimators	$\mathcal{U}\{50, 51, \dots, 2000\}$
min_child_weight	$\{10^{-5}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 10^2, 10^3, 10^4\}$
subsample	$\mathcal{U}(0.2, 0.8)$
colsample_bytree	$\mathcal{U}(0.2, 0.8)$
reg_alpha	$\{0, 10^{-1}, 1, 2, 5, 7, 10, 50, 100\}$
reg_lambda	$\{0, 10^{-1}, 1, 5, 10, 20, 50, 100\}$

E.2.9 Random Forest

We treat Random Forest as a weak baseline to establish a lower-bound reference for each dataset split. We run it with the *sklearn* package⁴⁰ and use its default configuration with *n_estimators* = 100.

³⁸<https://pypi.org/project/xgboost/>

³⁹<https://pypi.org/project/lightgbm/>

⁴⁰<https://scikit-learn.org/>

Table 19: Classification performance per dataset (up to 10K). Best score before rounding is bolded. We report average AUROC with 95% CIs. Models: CARTE (CRT), CatBoost (CTB), TabICL (ICL), LightGBM (LGB), RealMLP (MLP), TabDPT (DPT), TabM (TBM), TabPFNv2 (PFN), RandomForest (RF), TabSTAR (STR), XGBoost (XGB). Tuned models are marked with a '+'.

ID	CRT	CTB	CTB+	DPT	ICL	LGB	LGB+	MLP	MLP+	PFN	RF	STR	TBM	TBM+	XGB	XGB+
C01	88.4 ±0.3	90.2 ±0.3	90.3 ±0.4	90.2 ±0.3	90.7 ±0.4	89.5 ±0.4	89.6 ±0.4	90.2 ±0.5	90.2 ±0.3	90.3 ±0.3	88.8 ±0.3	90.8 ±0.3	90.5 ±0.3	90.5 ±0.3	89.3 ±0.4	90.2 ±0.4
C02	—	97.3 ±0.5	97.4 ±0.4	92.3 ±0.6	95.2 ±0.4	97.2 ±0.4	97.3 ±0.5	96.4 ±0.5	96.6 ±0.5	—	96.3 ±0.5	97.9 ±0.5	96.8 ±0.5	97.2 ±0.4	97.2 ±0.3	97.6 ±0.3
C03	82.5 ±0.3	82.0 ±0.3	81.2 ±0.6	85.4 ±0.3	81.4 ±0.3	79.8 ±0.3	79.9 ±0.3	82.6 ±0.4	82.7 ±0.3	82.4 ±0.3	77.3 ±0.3	83.0 ±0.3	83.3 ±0.3	83.8 ±0.3	80.5 ±0.3	82.1 ±0.3
C04	—	84.5 ±1.9	85.4 ±1.8	83.3 ±2.8	86.9 ±2.3	83.0 ±2.3	82.8 ±2.3	82.4 ±2.1	82.8 ±2.2	88.3 ±1.5	82.4 ±3.0	83.7 ±2.3	85.0 ±3.1	85.1 ±2.6	80.7 ±2.9	85.4 ±2.0
C05	88.5 ±0.9	90.6 ±1.0	90.9 ±0.5	92.2 ±0.6	92.9 ±0.5	88.2 ±1.3	88.4 ±0.6	91.0 ±0.8	91.4 ±0.6	91.2 ±0.7	88.0 ±1.1	91.3 ±0.8	91.1 ±0.7	90.7 ±0.6	88.1 ±1.1	90.3 ±0.6
C06	80.9 ±1.5	81.3 ±1.3	82.6 ±1.1	73.5 ±2.2	74.3 ±1.8	81.9 ±1.1	82.8 ±0.9	75.6 ±1.4	77.5 ±1.4	87.7 ±0.5	73.6 ±1.2	87.0 ±0.6	79.0 ±1.7	79.2 ±1.5	81.5 ±1.0	83.7 ±0.9
C07	90.1 ±0.4	90.3 ±0.3	90.6 ±0.3	89.9 ±0.2	86.7 ±0.3	88.9 ±0.4	89.3 ±0.4	90.0 ±0.3	90.2 ±0.2	89.8 ±0.4	85.1 ±0.7	91.5 ±0.3	91.1 ±0.3	91.1 ±0.3	87.2 ±0.6	89.3 ±0.5
C08	90.8 ±1.6	93.0 ±0.9	93.2 ±1.0	93.2 ±1.4	95.1 ±0.7	93.1 ±0.7	93.2 ±0.9	92.1 ±0.9	91.9 ±1.3	91.3 ±1.1	90.2 ±1.1	93.5 ±1.3	92.1 ±1.3	92.1 ±1.5	91.9 ±1.1	94.4 ±0.7
C09	—	82.4 ±1.3	82.6 ±1.4	81.3 ±1.2	84.3 ±1.0	81.1 ±1.2	82.9 ±1.3	82.6 ±1.4	81.3 ±2.5	82.5 ±1.1	76.4 ±1.4	94.0 ±0.8	81.8 ±4.1	83.6 ±4.1	79.4 ±1.0	83.2 ±1.2
C10	—	86.7 ±0.2	87.0 ±0.2	87.2 ±0.2	87.7 ±0.2	85.7 ±0.2	85.9 ±0.3	87.7 ±0.2	87.6 ±0.2	87.6 ±0.4	84.6 ±0.2	89.1 ±0.2	87.5 ±0.2	87.5 ±0.3	85.4 ±0.3	86.8 ±0.2
C11	78.4 ±0.5	79.7 ±0.5	80.7 ±0.4	79.6 ±0.4	81.5 ±0.5	78.9 ±0.4	79.2 ±0.4	80.4 ±0.4	79.6 ±0.4	81.4 ±0.4	76.2 ±0.4	79.2 ±0.6	81.2 ±0.3	81.5 ±0.3	78.2 ±0.5	80.2 ±0.5
C12	96.0 ±0.2	97.5 ±0.1	97.7 ±0.1	93.8 ±0.3	97.8 ±0.1	96.7 ±0.1	96.7 ±0.2	97.6 ±0.1	97.6 ±0.1	—	95.3 ±0.2	98.3 ±0.1	97.8 ±0.1	97.9 ±0.1	96.6 ±0.2	97.3 ±0.1
C13	70.9 ±0.8	73.7 ±1.1	74.1 ±1.0	70.7 ±0.7	74.0 ±0.8	72.6 ±1.0	72.7 ±1.0	73.0 ±0.7	72.4 ±1.2	72.3 ±0.9	71.1 ±1.1	75.0 ±0.7	74.3 ±0.6	74.4 ±0.7	70.2 ±0.8	74.1 ±1.1
C14	—	83.5 ±0.3	84.0 ±0.5	82.4 ±0.3	80.4 ±0.4	81.9 ±0.4	82.3 ±0.4	81.9 ±0.5	83.2 ±0.4	—	80.1 ±0.3	84.0 ±0.3	84.8 ±0.2	85.4 ±0.3	81.4 ±0.4	83.2 ±0.4

F Extended Main Results

In this appendix we provide the main results for the experiment as reported in §5. As elaborated in §4, each model is evaluated on each dataset across 10 splits. Since performance scales vary between datasets, we follow the normalization approach proposed by [38], rescaling all scores to the $[0, 1]$ range. The final reported performance for each model is the average over these normalized runs, and we compute 95% confidence intervals using the standard normal approximation: $\hat{\mu} \pm 1.96 \frac{\hat{\sigma}}{\sqrt{n}}$.

F.1 Technical Limitations

As discussed in §5, TabPFN-v2 is unable to run on 4 datasets: C12, because it is a multiclass problem with more than 10 classes, and C02, C14 and R01 because they support inference for up to 500,000 cells. Attempts to run the model over a subset of the examples led to a significantly worse performance, and thus we decide not to report them to allow a fair comparison.

Additionally, CARTE is unable to run over 15 of the datasets in the benchmark due to a known bug⁴¹ in their implementation for the *PowerTransformation*, which struggles in the presence of features with too less unique values. Furthermore, TabDPT’s performance over a few regression datasets is especially low. Since we focus on classification tasks, we exclude this model’s performance on this task. Finally, RealMLP, fails to run on R01 in the unlimited setting due to memory constraints.

F.2 Dataset Level Performance

We report AUROC for classification and R^2 for regression, with 95% CIs computed over the 10 runs for each dataset. Tables 19 and 20 summarize classification performance on datasets with up to 10K and over 10K examples, respectively. Tables 21 and 22 to the same for regression tasks. For conciseness, datasets are referred by their ID from Appendix D.

⁴¹<https://github.com/soda-inria/carte/issues/23>

Table 20: Classification performance per dataset (above 10K). Best score before rounding is bolded. We report average AUROC with 95% CIs. Models: CARTE (CRT), CatBoost (CTB), TabICL (ICL), LightGBM (LGB), RealMLP (MLP), TabDPT (DPT), TabM (TBM), TabPFNv2 (PFN), RandomForest (RF), TabSTAR (STR), XGBoost (XGB). Unlimited models are marked with a '!'.

ID	CRT	CTB	CTB!	DPT	ICL	ICL!	LGB	LGB!	MLP	MLP!	PFN	STR	STR!	TBM	TBM!	XGB	XGB!
C01	88.4 ±0.3	90.3 ±0.4	90.5 ±0.4	90.2 ±0.3	90.7 ±0.4	90.9 ±0.4	89.6 ±0.4	89.8 ±0.4	90.2 ±0.3	90.6 ±0.3	90.3 ±0.3	90.8 ±0.3	91.2 ±0.3	90.5 ±0.3	90.8 ±0.4	90.2 ±0.4	90.4 ±0.3
C02	— ±0.4	97.4 ±0.3	98.3 ±0.6	92.3 ±0.4	95.2 ±0.4	96.9 ±0.5	97.3 ±0.3	98.3 ±0.3	96.6 ±0.3	98.5 ±0.2	— ±0.5	97.9 ±0.2	98.4 ±0.2	97.2 ±0.4	98.5 ±0.3	97.6 ±0.3	98.2 ±0.3
C03	82.5 ±0.3	81.2 ±0.6	81.5 ±0.8	85.4 ±0.3	81.4 ±0.3	81.2 ±0.3	79.9 ±0.3	80.4 ±0.3	82.7 ±0.3	83.4 ±0.3	82.4 ±0.3	83.0 ±0.3	83.8 ±0.3	83.8 ±0.5	84.1 ±0.4	82.1 ±0.4	82.3 ±0.3
C07	90.1 ±0.4	90.6 ±0.3	91.0 ±0.4	89.9 ±0.2	86.7 ±0.3	87.1 ±0.3	89.3 ±0.3	90.0 ±0.4	90.2 ±0.3	90.5 ±0.2	89.8 ±0.4	91.5 ±0.4	91.9 ±0.3	91.1 ±0.3	91.6 ±0.5	89.3 ±0.5	89.9 ±0.5
C08	90.8 ±1.6	93.2 ±1.0	93.1 ±1.2	93.2 ±1.4	95.1 ±0.7	95.3 ±0.7	93.2 ±0.7	93.4 ±0.9	91.9 ±0.9	91.9 ±1.3	91.3 ±1.3	93.5 ±1.1	95.1 ±1.5	92.1 ±0.9	93.9 ±1.5	94.4 ±1.1	94.4 ±0.9
C09	— ±1.4	82.6 ±1.2	85.9 ±1.2	81.3 ±1.2	84.3 ±1.0	86.3 ±1.1	82.9 ±1.3	85.6 ±1.2	81.3 ±2.5	84.8 ±0.9	82.5 ±1.1	94.0 ±0.8	96.2 ±0.3	83.6 ±1.0	86.1 ±1.3	83.2 ±1.2	85.5 ±1.2
C11	78.4 ±0.5	80.7 ±0.4	81.8 ±0.4	79.6 ±0.5	81.5 ±0.5	82.3 ±0.4	79.2 ±0.4	80.5 ±0.5	79.6 ±0.5	81.0 ±0.5	81.4 ±0.4	79.2 ±0.4	81.0 ±0.5	82.8 ±0.4	80.2 ±0.5	81.4 ±0.4	81.4 ±0.4
C12	96.0 ±0.2	97.7 ±0.1	98.5 ±0.1	93.8 ±0.3	97.8 ±0.1	98.6 ±0.1	96.7 ±0.1	79.7 ±0.1	97.6 ±0.1	98.5 ±0.1	— ±0.1	98.3 ±0.1	99.1 ±0.1	97.9 ±0.1	98.7 ±0.1	97.3 ±0.1	98.4 ±0.1
C13	70.9 ±0.8	74.1 ±1.0	76.7 ±0.8	70.7 ±0.7	74.0 ±0.8	76.7 ±0.8	72.7 ±1.0	75.8 ±0.9	72.4 ±1.2	76.4 ±0.6	72.3 ±0.9	75.0 ±0.7	77.9 ±1.0	74.4 ±0.7	77.6 ±1.0	74.1 ±1.1	76.9 ±0.9
C14	— ±0.5	84.0 ±0.4	84.8 ±0.3	82.4 ±0.3	80.4 ±0.3	81.3 ±0.3	82.3 ±0.4	83.5 ±0.3	83.2 ±0.4	84.3 ±0.3	— ±0.3	84.0 ±0.3	85.1 ±0.3	85.4 ±0.3	86.2 ±0.3	83.2 ±0.4	84.2 ±0.4

Table 21: Regression performance per dataset (up to 10K). Best score before rounding is bolded. We report average R^2 with 95% CIs. Models: CARTE (CRT), CatBoost (CTB), LightGBM (LGB), RealMLP (MLP), TabM (TBM), TabPFNv2 (PFN), RandomForest (RF), TabSTAR (STR), XGBoost (XGB). Tuned models are marked with a '+'.

ID	CRT	CTB	CTB+	LGB	LGB+	MLP	MLP+	PFN	RF	STR	TBM	TBM+	XGB	XGB+	
R01	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	99.9 ±0.1	100.0 ±0.0	— ±0.0	100.0 ±0.0	100.0 ±0.0	99.8 ±0.1	99.8 ±0.1	99.9 ±0.1	100.0 ±0.0	100.0 ±0.0
R02	— ±1.0	98.3 ±0.9	98.3 ±1.0	98.1 ±1.0	98.1 ±1.0	98.2 ±1.0	98.2 ±1.0	98.4 ±0.9	98.0 ±1.0	98.0 ±1.1	99.0 ±0.8	99.0 ±0.8	97.9 ±1.1	97.9 ±1.1	98.3 ±0.9
R03	71.8 ±0.5	73.8 ±0.4	74.1 ±0.3	72.2 ±0.4	72.4 ±0.4	69.7 ±0.7	70.1 ±0.8	74.9 ±0.8	68.4 ±0.6	70.9 ±0.8	75.3 ±0.5	75.4 ±0.5	69.1 ±0.5	73.7 ±0.5	73.7 ±0.5
R04	— ±7.4	86.2 ±7.4	85.1 ±7.1	85.7 ±6.5	84.6 ±6.2	88.7 ±4.4	90.3 ±3.2	86.1 ±5.2	85.1 ±6.8	81.5 ±8.7	84.0 ±3.1	84.1 ±3.1	85.8 ±7.5	87.5 ±7.5	87.5 ±7.5
R05	93.2 ±0.8	89.5 ±1.3	90.0 ±1.2	87.4 ±1.5	88.2 ±1.4	90.9 ±1.2	91.2 ±1.5	92.7 ±0.9	86.7 ±1.3	93.2 ±1.2	92.1 ±1.1	92.3 ±1.1	87.1 ±1.4	90.1 ±1.3	90.1 ±1.3
R06	97.7 ±0.8	98.1 ±0.5	98.2 ±0.5	97.7 ±0.7	97.8 ±0.7	97.5 ±0.9	97.8 ±0.9	98.5 ±0.6	97.2 ±1.0	98.1 ±0.3	97.8 ±0.8	97.9 ±0.8	97.1 ±0.7	97.9 ±0.7	97.9 ±0.7
R07	71.6 ±1.3	66.9 ±0.8	67.8 ±0.8	62.4 ±1.0	63.4 ±1.0	68.9 ±0.8	67.4 ±1.3	62.2 ±1.3	61.5 ±1.2	71.1 ±1.8	67.9 ±0.7	69.6 ±0.9	61.9 ±1.5	68.8 ±0.8	68.8 ±0.8
R08	93.2 ±0.8	93.0 ±0.7	93.2 ±0.7	92.9 ±0.7	93.0 ±0.7	92.9 ±0.8	92.7 ±0.8	93.9 ±0.7	92.2 ±0.7	92.8 ±0.6	93.4 ±0.7	93.3 ±0.7	92.3 ±0.7	93.1 ±0.7	93.1 ±0.7
R09	89.2 ±0.5	89.2 ±0.6	89.3 ±0.4	88.7 ±0.5	88.9 ±0.5	88.6 ±0.5	89.2 ±0.5	89.8 ±0.5	88.6 ±0.7	88.8 ±0.5	89.5 ±0.7	89.7 ±0.7	88.2 ±0.7	89.3 ±0.5	89.3 ±0.5
R10	99.5 ±0.2	99.0 ±0.6	98.9 ±0.6	98.5 ±0.3	98.5 ±0.3	99.3 ±0.2	99.4 ±0.2	99.1 ±0.2	98.1 ±0.3	99.5 ±0.1	98.9 ±0.1	99.1 ±0.1	98.6 ±0.2	99.4 ±0.2	99.4 ±0.2
R11	— ±0.9	94.2 ±0.9	94.0 ±0.9	93.6 ±0.9	93.6 ±1.0	94.2 ±1.0	94.4 ±1.0	94.5 ±0.9	92.5 ±1.3	93.9 ±1.1	94.5 ±0.8	94.6 ±0.8	93.3 ±1.1	94.2 ±0.8	94.2 ±0.8
R12	— ±0.3	98.5 ±0.2	98.5 ±0.3	98.3 ±0.2	98.4 ±0.3	98.4 ±0.2	98.4 ±0.2	98.5 ±0.2	98.0 ±0.3	98.3 ±0.2	98.6 ±0.2	98.7 ±0.2	98.2 ±0.3	98.5 ±0.2	98.5 ±0.2
R13	52.8 ±2.1	57.6 ±2.7	58.2 ±2.6	54.7 ±2.8	55.0 ±2.6	55.4 ±2.5	55.1 ±2.5	53.8 ±2.7	51.8 ±3.5	53.4 ±3.0	58.7 ±1.9	58.4 ±2.1	49.0 ±3.3	57.5 ±2.7	57.5 ±2.7
R14	96.6 ±0.2	97.3 ±0.1	97.3 ±0.1	97.0 ±0.1	97.0 ±0.1	96.8 ±0.1	96.7 ±0.2	96.2 ±0.2	96.8 ±0.1	96.4 ±0.1	96.3 ±0.1	96.3 ±0.1	97.1 ±0.1	97.1 ±0.1	97.1 ±0.1
R15	— ±0.9	79.5 ±1.1	79.9 ±0.9	77.8 ±0.9	77.7 ±1.4	76.2 ±1.5	77.4 ±1.5	79.4 ±1.0	77.3 ±1.2	78.9 ±1.5	80.6 ±0.7	79.9 ±0.7	76.8 ±1.2	80.3 ±1.0	80.3 ±1.0

Continued on next page

Table 21: Regression performance per dataset (up to 10K). Best score before rounding is bolded. We report average R^2 with 95% CIs. Models: CARTE (CRT), CatBoost (CTB), LightGBM (LGB), RealMLP (MLP), TabM (TBM), TabPFNv2 (PFN), RandomForest (RF), TabSTAR (STR), XGBoost (XGB). Tuned models are marked with a '+'.

ID	CRT	CTB	CTB+	LGB	LGB+	MLP	MLP+	PFN	RF	STR	TBM	TBM+	XGB	XGB+
R16	98.7 ±0.0	98.7 ±0.0	98.7 ±0.0	98.7 ±0.0	98.7 ±0.0	97.6 ±0.1	97.8 ±0.1	98.8 ±0.0	97.7 ±0.1	98.7 ±0.0	97.9 ±0.1	97.9 ±0.1	97.7 ±0.1	97.8 ±0.1
R17	95.3 ±2.0	94.8 ±2.4	95.2 ±2.0	94.6 ±2.3	95.2 ±2.0	94.4 ±2.4	95.4 ±1.9	95.1 ±2.1	94.2 ±2.5	94.7 ±2.3	96.7 ±1.1	96.7 ±1.1	94.5 ±2.6	95.4 ±1.9
R18	98.1 ±0.4	98.2 ±0.4	98.2 ±0.4	98.1 ±0.4	98.1 ±0.4	98.0 ±0.5	97.7 ±0.5	93.1 ±1.3	97.7 ±0.4	97.4 ±0.6	98.2 ±0.3	98.1 ±0.4	97.5 ±0.8	98.3 ±0.4
R19	— ±0.8	85.5 ±0.9	85.8 ±0.9	85.2 ±0.9	85.2 ±0.9	82.7 ±1.3	83.3 ±1.1	84.0 ±1.1	85.3 ±0.9	83.8 ±1.1	85.0 ±1.0	85.1 ±0.9	83.9 ±1.0	82.4 ±2.0
R20	96.4 ±0.5	96.0 ±0.6	96.2 ±0.5	95.7 ±0.5	95.8 ±0.5	95.3 ±0.8	96.0 ±0.5	93.9 ±1.0	95.7 ±0.6	95.6 ±0.9	95.9 ±0.7	95.8 ±0.7	95.3 ±0.9	96.2 ±0.5
R21	92.3 ±1.2	92.5 ±1.1	92.6 ±1.1	92.3 ±1.2	92.3 ±1.2	92.0 ±1.1	91.7 ±1.3	93.3 ±1.1	91.8 ±1.1	92.3 ±1.0	91.6 ±1.4	92.1 ±1.2	91.6 ±1.0	92.6 ±1.0
R22	— ±5.5	45.2 ±5.9	45.2 ±5.6	42.4 ±5.5	44.7 ±5.5	43.6 ±6.4	44.4 ±6.3	43.8 ±5.5	39.1 ±6.4	39.7 ±5.5	47.9 ±5.5	48.1 ±5.4	36.1 ±7.1	46.8 ±5.8
R23	85.0 ±1.8	85.4 ±1.2	85.5 ±1.1	84.3 ±1.5	85.5 ±1.2	83.3 ±1.8	85.6 ±1.5	84.8 ±1.6	81.8 ±1.5	83.0 ±1.6	86.4 ±0.9	86.6 ±0.8	83.3 ±1.2	85.1 ±1.2
R24	86.0 ±0.6	84.0 ±0.8	85.8 ±0.7	78.8 ±0.9	79.5 ±0.9	84.6 ±0.8	82.8 ±0.9	70.3 ±1.6	79.7 ±1.0	81.8 ±1.8	84.9 ±0.9	85.5 ±0.8	82.9 ±0.7	85.9 ±0.9
R25	94.3 ±0.6	95.4 ±0.5	95.3 ±0.5	95.2 ±0.6	95.2 ±0.5	93.2 ±0.5	94.1 ±0.6	86.7 ±0.7	94.8 ±1.0	94.8 ±1.5	94.3 ±0.5	94.5 ±0.4	94.6 ±0.6	95.4 ±0.5
R26	99.8 ±0.1	99.4 ±0.1	99.4 ±0.1	99.3 ±0.2	99.3 ±0.2	99.9 ±0.1	99.9 ±0.0	99.8 ±0.1	99.0 ±0.2	99.6 ±0.1	98.1 ±0.7	98.7 ±0.8	99.1 ±0.2	99.5 ±0.1
R27	81.9 ±1.6	85.2 ±1.3	85.3 ±1.4	84.4 ±1.4	84.8 ±1.3	83.3 ±1.5	76.7 ±11.2	82.2 ±1.6	84.8 ±1.3	82.0 ±1.6	83.8 ±1.5	84.0 ±1.7	83.2 ±1.2	85.5 ±1.2
R28	52.5 ±2.6	53.4 ±2.7	53.3 ±2.8	50.0 ±3.0	51.0 ±2.9	52.3 ±3.2	52.6 ±3.0	61.7 ±2.1	46.0 ±2.7	51.5 ±2.9	54.7 ±2.7	55.6 ±2.6	45.3 ±2.3	53.5 ±2.7
R29	— ±0.7	94.4 ±0.7	94.4 ±0.7	94.8 ±0.7	94.8 ±0.8	94.3 ±0.7	94.4 ±0.8	95.7 ±0.6	93.0 ±1.0	94.3 ±0.9	94.9 ±0.8	95.0 ±0.7	93.4 ±0.8	94.5 ±0.8
R30	23.8 ±4.2	22.7 ±4.1	24.6 ±3.9	22.3 ±3.7	23.9 ±4.2	23.9 ±3.9	15.7 ±5.4	18.8 ±4.8	20.9 ±3.5	23.5 ±5.0	15.7 ±5.0	22.9 ±4.8	23.4 ±4.2	17.5 ±3.7
R31	92.1 ±0.3	92.1 ±0.4	92.0 ±0.4	91.6 ±0.4	91.6 ±0.4	91.5 ±0.3	91.5 ±0.4	93.2 ±0.3	89.9 ±0.5	91.7 ±0.4	92.4 ±0.5	92.6 ±0.5	90.3 ±0.5	92.0 ±0.4
R32	— ±6.3	28.8 ±5.8	28.2 ±6.4	25.0 ±5.7	27.3 ±5.7	19.8 ±8.8	19.8 ±7.1	26.2 ±5.4	28.6 ±5.9	19.4 ±5.6	25.8 ±5.3	27.3 ±6.0	22.7 ±7.3	31.6 ±6.2
R33	46.6 ±1.5	47.4 ±1.7	47.8 ±1.7	44.8 ±1.8	45.2 ±1.6	40.4 ±2.1	44.6 ±1.6	44.9 ±1.6	40.2 ±1.6	46.0 ±1.8	48.4 ±1.1	48.4 ±1.2	40.3 ±2.1	47.9 ±1.6
R34	— ±2.5	90.9 ±2.1	91.4 ±2.9	88.9 ±2.4	89.9 ±3.0	88.6 ±1.4	90.7 ±1.8	92.3 ±3.1	88.7 ±2.9	89.1 ±1.7	91.1 ±1.7	90.9 ±1.8	88.9 ±3.6	91.6 ±2.1
R35	85.9 ±0.6	84.4 ±0.5	84.5 ±0.5	83.7 ±0.6	83.7 ±0.6	83.1 ±0.9	84.2 ±0.8	85.2 ±0.7	81.2 ±0.4	85.7 ±0.8	84.5 ±0.6	84.5 ±0.5	81.6 ±0.8	84.4 ±0.3
R36	96.3 ±0.9	95.2 ±0.9	95.5 ±0.9	94.7 ±1.0	94.9 ±1.0	95.4 ±1.1	95.7 ±0.9	91.2 ±1.1	94.3 ±1.0	95.7 ±1.0	96.2 ±0.8	96.4 ±0.6	94.4 ±0.9	95.5 ±0.9

E.3 Head-to-head comparisons

We compare the performance of TabSTAR against each of the models in head-to-head comparisons. We report win rate, which can be seen as a private case of the normalized metric with only two models. We exclude failed runs when comparing to CARTE and TabPFN-v2. Table 23 shows the performance of TabSTAR against all models competing up to 10K examples, for both regression and classification, with 95% CIs over the win rate. Table 24 does the same for TabSTAR-Unlimited.

F.4 Running Times and Compute Information

Hardware All baselines are evaluated using a single *NVIDIA A100-SXM4* GPU with 40GB memory, and 8 CPU cores of type *AMD EPYC 7742 64-Core Processor*. The only exclusions are TabPFN-v2,

Table 22: Regression performance per dataset (above 10K). Best score before rounding is bolded. We report average R^2 with 95% CIs. Models: CARTE (CRT), CatBoost (CTB), LightGBM (LGB), RealMLP (MLP), TabM (TBM), TabPFNv2 (PFN), RandomForest (RF), TabSTAR (STR), XGBoost (XGB). Unlimited models are marked with a '!'.

ID	CRT	CTB	CTB!	LGB	LGB!	MLP	MLP!	PFN	STR	STR!	TBM	TBM!	XGB	XGB!
R01	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	— —	100.0 ±0.0	100.0 ±0.0	99.8 ±0.1	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0
R02	— ±0.9	98.3 ±0.9	98.8 ±0.9	98.1 ±1.0	98.8 ±1.0	98.2 ±1.0	98.8 ±0.9	98.4 ±1.1	98.0 ±1.0	98.4 ±0.8	99.0 ±0.9	98.5 ±0.9	98.3 ±0.9	98.8 ±0.9
R03	71.8 ±0.5	74.1 ±0.3	75.1 ±0.4	72.4 ±0.4	72.7 ±0.5	70.1 ±0.8	71.0 ±0.6	74.9 ±0.4	70.9 ±0.8	72.3 ±0.6	75.4 ±0.5	76.4 ±0.5	73.7 ±0.5	74.6 ±0.6
R04	— ±7.1	85.1 ±7.1	91.3 ±0.8	84.6 ±6.2	91.0 ±1.3	91.4 ±3.2	86.1 ±0.8	81.5 ±5.2	91.0 ±8.7	84.1 ±0.7	91.0 ±3.1	87.5 ±0.7	91.3 ±3.7	91.3 ±0.6
R07	71.6 ±1.3	67.8 ±0.8	85.2 ±0.3	63.4 ±1.0	70.1 ±0.5	67.4 ±1.3	85.0 ±0.3	62.2 ±1.3	71.1 ±1.8	79.9 ±1.8	69.6 ±0.9	84.0 ±1.3	68.8 ±0.8	86.1 ±0.7
R08	93.2 ±0.8	93.2 ±0.7	93.8 ±0.6	93.0 ±0.7	93.6 ±0.5	92.7 ±0.8	93.5 ±0.7	93.9 ±0.7	92.8 ±0.6	93.3 ±0.6	93.3 ±0.7	93.9 ±0.6	93.1 ±0.7	93.9 ±0.5
R09	89.2 ±0.5	89.3 ±0.4	89.5 ±0.4	88.9 ±0.5	89.2 ±0.5	89.2 ±0.5	89.3 ±0.5	89.8 ±0.5	88.8 ±0.5	89.2 ±0.6	89.7 ±0.7	89.3 ±0.4	89.3 ±0.5	89.5 ±0.4
R12	— ±0.2	98.5 ±0.2	98.9 ±0.1	98.4 ±0.2	98.9 ±0.2	98.4 ±0.1	99.0 ±0.1	98.5 ±0.2	98.3 ±0.2	98.4 ±0.2	98.7 ±0.2	99.0 ±0.2	98.5 ±0.2	98.9 ±0.2
R14	96.6 ±0.2	97.3 ±0.1	97.8 ±0.1	97.0 ±0.1	97.5 ±0.1	96.7 ±0.2	97.4 ±0.1	96.2 ±0.2	96.4 ±0.1	96.8 ±0.2	96.3 ±0.1	97.2 ±0.1	97.1 ±0.1	97.5 ±0.1
R15	— ±1.1	79.9 ±0.7	86.5 ±0.9	77.7 ±0.6	82.0 ±1.5	77.4 ±0.6	85.7 ±1.0	79.4 ±1.5	78.9 ±1.4	83.9 ±1.4	79.9 ±0.7	87.2 ±0.8	80.3 ±1.0	87.0 ±0.7
R16	98.7 ±0.0	98.7 ±0.0	98.8 ±0.0	98.7 ±0.0	98.8 ±0.0	97.8 ±0.1	98.0 ±0.1	98.8 ±0.0	98.7 ±0.0	98.8 ±0.1	97.9 ±0.1	98.1 ±0.1	97.8 ±0.1	98.0 ±0.1
R17	95.3 ±2.0	95.2 ±2.0	97.6 ±0.6	95.2 ±2.0	97.6 ±0.7	95.4 ±1.9	97.5 ±0.6	95.1 ±2.1	94.7 ±2.3	97.6 ±0.6	96.7 ±1.1	97.2 ±0.7	95.4 ±1.9	97.7 ±0.6
R18	98.1 ±0.4	98.2 ±0.4	99.0 ±0.2	98.1 ±0.4	98.8 ±0.3	97.7 ±0.5	98.4 ±0.4	93.1 ±1.3	97.4 ±0.6	97.7 ±0.6	98.1 ±0.4	98.7 ±0.3	98.3 ±0.4	98.9 ±0.3
R20	96.4 ±0.5	96.2 ±0.5	96.3 ±0.4	95.8 ±0.5	95.8 ±0.5	96.0 ±0.5	95.6 ±0.5	93.9 ±1.0	95.6 ±0.9	95.3 ±1.0	95.8 ±0.7	95.9 ±0.6	96.2 ±0.5	96.3 ±0.5
R23	85.0 ±1.8	85.5 ±1.1	86.3 ±0.7	85.5 ±1.2	85.8 ±0.7	85.6 ±1.5	86.6 ±0.8	84.8 ±1.6	83.0 ±1.6	85.2 ±1.7	86.6 ±0.8	88.3 ±0.4	85.1 ±1.2	86.1 ±0.8
R24	86.0 ±0.6	85.8 ±0.7	97.1 ±0.3	79.5 ±0.9	85.4 ±0.6	82.8 ±0.9	96.5 ±0.4	70.3 ±1.6	81.8 ±1.6	95.5 ±1.8	85.5 ±0.7	97.8 ±0.3	85.9 ±0.9	97.3 ±0.3
R25	94.3 ±0.6	95.3 ±0.5	95.9 ±0.4	95.2 ±0.5	95.7 ±0.5	94.1 ±0.7	95.0 ±0.5	86.7 ±1.0	94.8 ±0.5	95.1 ±0.5	94.5 ±0.4	95.9 ±0.4	95.4 ±0.5	95.8 ±0.5
R31	92.1 ±0.3	92.0 ±0.4	93.0 ±0.3	91.6 ±0.4	92.5 ±0.3	91.5 ±0.4	92.5 ±0.4	93.2 ±0.3	91.7 ±0.4	92.7 ±0.4	92.6 ±0.5	93.6 ±0.3	92.0 ±0.4	92.9 ±0.3
R33	46.6 ±1.5	47.8 ±1.7	55.9 ±1.1	45.2 ±1.6	49.8 ±1.3	44.6 ±1.6	54.5 ±1.3	44.9 ±1.6	46.0 ±1.8	57.2 ±1.3	48.4 ±1.2	57.8 ±1.2	47.9 ±1.6	55.7 ±1.3
R36	96.3 ±0.9	95.5 ±0.9	95.5 ±0.9	94.9 ±1.0	94.9 ±1.0	95.7 ±0.9	95.8 ±0.9	91.2 ±1.1	95.7 ±0.8	95.5 ±0.5	96.4 ±0.6	96.1 ±0.9	95.5 ±0.9	95.5 ±0.9

Table 23: Win rates of TabSTAR (up to 10K) against baselines. Win rate with 95% CI.

Model	Classification	Regression
CARTE	93.3 ± 5.2	35.7 ± 5.9
CatBoost	80.7 ± 6.6	33.6 ± 4.9
CatBoost-Tuned	75.0 ± 7.2	32.2 ± 4.8
LightGBM	92.1 ± 4.5	49.7 ± 5.2
LightGBM-Tuned	90.7 ± 4.8	45.3 ± 5.1
RandomForest	95.7 ± 3.4	66.9 ± 4.9
RealMLP	82.1 ± 6.4	51.4 ± 5.2
RealMLP-Tuned	83.6 ± 6.2	49.4 ± 5.2
TabDPT	73.6 ± 7.3	66.7 ± 4.9
TabICL	70.0 ± 7.6	—
TabM	64.3 ± 8.0	39.2 ± 5.0
TabM-Tuned	63.3 ± 8.0	37.4 ± 5.0
TabPFN-v2	66.4 ± 8.9	40.9 ± 5.2
XGBoost	96.4 ± 3.1	69.4 ± 4.8
XGBoost-Tuned	78.6 ± 6.8	30.3 ± 4.8

Table 24: Win rates of TabSTAR-Unlimited (above 10K) against baselines. Win rate with 95% CI.

Model	Classification	Regression
CARTE	98.6 ± 2.8	63.5 ± 7.5
CatBoost-Tuned	96.0 ± 3.9	53.0 ± 6.9
CatBoost-Tuned-Unlimited	77.0 ± 8.3	22.0 ± 5.8
LightGBM-Tuned	98.0 ± 2.8	69.5 ± 6.4
LightGBM-Tuned-Unlimited	92.0 ± 5.3	44.0 ± 6.9
RealMLP-Tuned	97.0 ± 3.4	72.0 ± 6.2
RealMLP-Tuned-Unlimited	86.0 ± 6.8	38.9 ± 7.0
TabDPT	85.0 ± 7.0	—
TabICL	85.0 ± 7.0	—
TabICL-Unlimited	82.0 ± 7.6	—
TabM-Tuned	76.8 ± 8.4	65.2 ± 6.7
TabM-Tuned-Unlimited	53.0 ± 9.8	22.5 ± 5.8
TabPFN-v2	90.0 ± 7.1	61.1 ± 7.0
TabSTAR	94.0 ± 4.7	82.0 ± 5.3
XGBoost-Tuned	95.0 ± 4.3	59.5 ± 6.8
XGBoost-Tuned-Unlimited	81.0 ± 7.7	28.5 ± 6.3

which runs on its API version, and CARTE, for which we employ several *NVIDIA GeForce GTX 1080 Ti* GPUs with 11GB memory, as their high latency and lack of default hyperparameters forced us to rely on more accessible hardware.

Running Times Table 25 presents the average running times of the different models per dataset, with up to 10,000 examples, excluding the tuned variants, TabPFN-v2 and CARTE. TabSTAR’s average running time for a downstream task ranges from 59s (C04) to 9,332s (R24). Although higher than a single run of other baselines, these times remain far below the 14,400s (4 hours) typically required for their tuned counterparts. TabICL achieves considerably faster running times; however, their ICL-based approach limits them to smaller datasets, and inference is as costly as training (see §5). TabPFN-v2, though unreported, is subject to the same limitation.

E.5 TabSTAR’s Cost Analysis by Number of Features

While §5 reports average computational costs, these values vary considerably across datasets, depending on factors such as the number of examples, the feature count, and the amount of textual fields. In this subsection, we examine how TabSTAR’s computational cost scales with the number of features.

Table 25: Average model training running time in seconds, per dataset, for up to 10K examples.

ID	CatBoost	LightGBM	RandomForest	RealMLP	TabICL	TabSTAR	XGBoost
C01	116	35	48	148	37	397	36
C02	251	57	75	179	58	1,135	58
C03	106	51	61	147	49	553	48
C04	29	17	19	25	20	59	18
C05	26	13	16	63	13	132	13
C06	170	79	83	126	77	938	81
C07	163	78	98	186	80	1,044	78
C08	89	72	86	180	74	1,057	69
C09	33	25	33	128	29	414	25
C10	49	35	43	126	36	572	35
C11	43	22	29	121	20	521	20
C12	614	39	37	145	28	729	36
C13	70	55	68	163	55	551	50
C14	1,938	307	333	418	311	4,068	329
R01	488	161	1,204	368	—	2,908	161
R02	28	10	45	128	—	283	10
R03	111	62	282	171	—	592	61
R04	21	9	17	138	—	240	7
R05	65	20	68	90	—	366	23
R06	60	24	119	129	—	501	23
R07	66	39	285	145	—	673	44
R08	213	105	951	222	—	1,448	112
R09	12	8	56	115	—	604	7
R10	102	45	104	72	—	255	47
R11	26	10	45	98	—	197	10
R12	34	10	42	133	—	335	11
R13	78	44	145	94	—	484	46
R14	67	24	91	159	—	1,296	27
R15	38	18	102	132	—	390	17
R16	108	76	316	193	—	898	80
R17	62	54	1,110	160	—	261	53
R18	124	74	1,705	201	—	553	68
R19	22	13	38	107	—	198	12
R20	70	52	401	176	—	705	49
R21	53	29	86	62	—	222	30
R22	57	31	62	52	—	171	33
R23	41	19	137	140	—	281	18
R24	461	396	650	520	—	9,332	392
R25	191	104	2,583	223	—	750	110
R26	29	12	22	30	—	205	11
R27	103	49	164	92	—	334	49
R28	175	92	460	158	—	647	97
R29	34	8	21	100	—	308	7
R30	26	14	47	56	—	114	14
R31	148	78	505	200	—	797	77
R32	20	11	23	31	—	82	11
R33	89	55	271	161	—	608	47
R34	66	22	66	50	—	122	24
R35	99	50	146	101	—	411	50
R36	58	38	223	146	—	823	39

Table 26: Cost Analysis for TabSTAR per number of features. Median training and inference times and peak memory usage on GPU and CPU, aggregated across 45 datasets.

Feature Group	Train			Inference		
	Time (s)	GPU (GB)	CPU (GB)	Time (s)	GPU (GB)	CPU (GB)
Up to 10	85.6	1.9	4.4	0.5	1.6	4.4
30-50	228.8	3.1	4.5	1.0	1.7	4.4
100+	417.9	7.8	4.5	2.5	2.3	4.4

Table 26 summarizes both latency and memory consumption of TabSTAR, aggregated over a group of up to 10 features, a group of 30-50 features and a group of more than 100 features. These datasets were randomly sampled from the pretraining corpus; consequently, the reported times may differ from those observed on the benchmark datasets. We observe that latency can increase by a factor of 5, and GPU memory consumption by a factor of 4. This behavior stems from the transformer’s architecture, whose computational complexity is quadratic in the sequence length. This raises concerns for training TabSTAR over datasets with hundreds of features, and calls for improvements to expand to datasets of thousands of features.

Additionally, TabSTAR’s memory consumption in classification tasks is affected by the number of classes due to its target-aware tokens. We assume that for TabSTAR, the effective number of features equals the actual number of features plus the number of classes. For highly multiclass datasets, this can become a limiting factor.

F.6 Number of Trials for Tuned models

The tuned models are optimized with a budget of 4 hours using 8 CPU cores. Each trial is executed on a single core, ensuring that at least one trial could be completed within the allocated time. Table 27 presents the number of hyperparameter trials per dataset.

G Extended Analysis

In this section we expand on the analysis results discussed in §6.

G.1 Evaluation Datasets for Analysis

All experiments are conducted on 20 datasets from the benchmark described in Appendix D. Each experiment reports performance using AUROC for classification and R^2 for regression. Each table reports both regression and classification tasks, distinguishable by their ID. Furthermore, each experiment compares its variants to other models, excluding TabPFN-v2 and CARTE, which could not be executed on all datasets. Although this comparison covers only a subset of the benchmark and should not be interpreted as conclusive, the results nonetheless offer valuable reference points.

G.2 The Role of the Encoder Unfreezing (Q1)

Table 28 shows the results for each variant of the experiment presented in §6. It is evident that unfreezing the textual encoder yields a significant performance boost across datasets, as seen in Figure 6: The frozen TabSTAR variant is much worse than any other baseline! Furthermore, while finetuning a single layer gives a significant boost, it underperforms compared to 6 unfrozen layers.

G.3 The Effect of Pretraining (Q2)

We pretrain three TabSTAR variants on nested dataset subsets of size 16, 64, and 256. The 64-dataset variant contains the original 16 plus 48 new datasets, and the 256-dataset variant builds on those 64 by adding another 192. This cumulative design minimizes variance between variants so that performance differences reflect only the effect of increasing data volume.

While LoRA [42] is a powerful technique, it can’t be applied to a randomly initialized model. Therefore, we perform full finetuning of the non-pretrained model, as explained in Appendix B.2.

Table 27: Average number of hyperparameter trials per dataset. Tuned models: CatBoost (CTB), LightGBM (LGB), RealMLP (MLP), XGBoost (XGB). Unlimited models are marked with a '!'.

Dataset	CTB	CTB!	LGB	LGB!	MLP	MLP!	XGB	XGB!
C01	27.1	24.4	770.2	812.0	56.0	26.7	97.7	90.0
C02	18.1	9.4	417.5	51.3	31.9	8.0	47.2	19.6
C03	18.0	15.8	312.0	338.1	37.4	25.2	42.4	39.5
C04	105.8	—	11283.8	—	642.1	—	458.5	—
C05	122.1	—	3373.3	—	98.3	—	228.3	—
C06	18.1	—	406.9	—	67.7	—	44.6	—
C07	17.3	18.5	335.6	376.6	31.8	23.7	53.4	37.7
C08	101.1	67.8	2911.4	2394.6	43.1	37.3	508.3	566.1
C09	169.0	81.3	3969.5	475.7	50.5	9.2	883.9	207.2
C10	73.7	—	1598.1	—	55.4	—	165.2	—
C11	58.3	48.5	812.4	863.4	44.8	26.0	124.1	96.3
C12	9.9	8.5	207.2	34.4	35.3	8.2	41.8	17.3
C13	85.8	27.4	1848.3	329.6	40.8	8.7	234.6	105.7
C14	8.6	8.4	50.9	36.3	14.9	9.7	13.7	11.6
R01	14.2	12.2	278.5	169.6	11.3	—	32.8	15.4
R02	138.9	118.6	4463.0	5214.6	48.0	30.7	742.4	1054.5
R03	55.5	64.2	1033.6	1265.9	35.8	30.0	94.3	97.8
R04	123.8	31.6	4309.9	714.9	44.1	8.7	1324.0	364.8
R05	78.8	—	4428.5	—	87.9	—	307.0	—
R06	88.8	—	2305.3	—	47.5	—	232.2	—
R07	109.9	58.8	1370.0	344.1	42.0	8.5	145.6	72.5
R08	30.5	30.1	568.6	122.9	21.4	10.4	65.7	47.1
R09	194.6	208.7	6391.0	5975.1	50.4	31.2	589.0	680.6
R10	53.7	—	2806.9	—	135.9	—	241.8	—
R11	142.9	—	4669.1	—	58.7	—	662.6	—
R12	126.8	45.6	2990.9	814.7	46.9	8.6	385.8	178.0
R13	80.9	—	2814.6	—	91.2	—	204.8	—
R14	59.9	67.6	1725.3	1542.6	39.5	15.2	200.6	177.0
R15	107.4	69.2	2896.1	1177.4	52.5	14.0	337.5	242.7
R16	65.8	46.3	1446.3	663.7	34.9	10.9	160.9	104.9
R17	79.7	44.6	1517.3	316.3	45.3	9.1	254.0	104.4
R18	39.1	43.8	1572.4	1041.4	29.4	16.6	258.8	144.5
R19	117.5	—	4046.2	—	63.1	—	1063.2	—
R20	62.6	89.1	1611.5	2176.4	36.4	30.1	250.5	258.8
R21	101.6	—	3284.9	—	122.3	—	159.4	—
R22	67.8	—	3584.3	—	259.9	—	127.2	—
R23	91.8	105.4	3179.2	2773.0	53.4	27.6	420.1	314.5
R24	52.4	38.1	867.2	417.7	35.2	10.4	104.7	60.3
R25	40.4	33.5	806.9	357.6	23.6	10.8	122.4	72.0
R26	166.5	—	12567.7	—	279.4	—	625.0	—
R27	53.9	—	2034.2	—	85.3	—	155.5	—
R28	40.9	—	670.8	—	37.5	—	59.3	—
R29	135.3	—	7615.2	—	63.8	—	1024.3	—
R30	194.3	—	5275.3	—	114.8	—	431.7	—
R31	36.6	39.5	774.5	504.3	26.4	12.6	90.3	54.1
R32	143.2	—	6605.7	—	191.7	—	337.3	—
R33	69.3	30.8	1053.9	212.0	36.7	8.3	119.8	51.1
R34	69.1	—	3962.2	—	148.9	—	314.9	—
R35	65.8	—	3351.6	—	76.3	—	139.9	—
R36	108.0	120.2	2633.9	4634.0	49.9	48.0	289.9	228.8

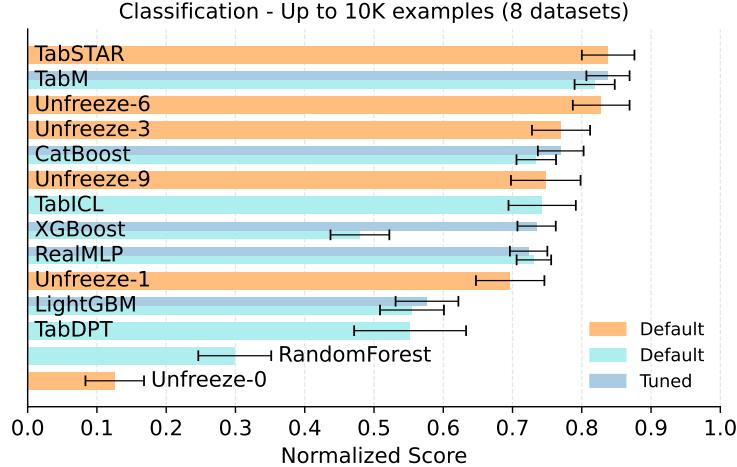


Figure 6: Comparison of normalized scores with 95% CIs for the encoder unfreezing experiment, comparing TabSTAR variants against baseline models in classification tasks.

Table 28: Downstream performance for Q1: The Role of the Encoder Unfreezing. Results for 20 datasets with 95% CI, for varying number of unfrozen layers. The top performance score is bolded first, and then all scores are rounded. We report AUROC for classification and R^2 for regression.

ID	0	1	3	6	9
C01	87.8±0.3	90.4±0.4	90.8±0.4	91.0±0.4	90.8±0.4
C02	94.9±1.2	97.8±0.3	97.8±0.3	98.1±0.3	97.7±0.3
C03	77.6±0.9	82.4±0.2	82.8±0.3	83.0±0.5	83.0±0.4
C05	87.0±0.7	90.2±0.9	90.8±0.6	91.5±0.9	89.4±0.8
C07	82.1±1.4	89.1±0.9	90.6±0.5	91.3±0.4	91.0±0.4
C11	77.4±0.5	78.4±0.6	78.7±0.7	78.8±0.5	78.3±0.6
C12	94.4±0.2	98.3±0.1	98.3±0.1	98.3±0.1	98.2±0.1
C13	66.5±0.9	71.7±1.1	72.9±0.7	74.1±0.7	73.0±0.8
R02	97.0±1.7	97.9±1.2	98.0±1.1	97.9±1.2	98.0±1.1
R03	67.2±0.8	69.8±0.8	70.9±0.9	71.2±0.8	70.8±0.6
R05	80.8±2.6	91.8±1.5	93.2±0.5	93.1±0.7	92.9±0.8
R09	88.1±0.7	88.7±0.7	89.0±0.5	89.0±0.6	88.9±0.8
R12	97.0±0.6	98.2±0.3	98.2±0.2	98.3±0.3	98.2±0.3
R13	37.6±3.6	45.7±2.1	51.6±3.2	51.4±2.9	52.1±1.8
R18	96.3±1.1	97.1±0.6	97.4±0.6	97.3±0.6	97.0±0.6
R23	79.3±2.4	81.6±1.6	82.8±1.6	83.2±2.2	82.1±1.9
R27	81.3±1.6	81.3±1.6	81.5±1.6	81.3±1.7	81.4±1.6
R30	14.2±4.6	15.6±5.6	19.2±3.1	19.1±4.0	17.3±4.5
R33	36.0±2.0	43.6±1.5	44.3±1.4	46.0±1.8	43.5±2.2
R34	84.3±2.8	87.4±2.9	88.4±3.4	87.9±3.5	88.7±2.7

Table 29 show the normalized results for classification and regression, and Table 30 shows the dataset level results. It is evident that for most of the datasets, improvement is observed when scaling, with the 256 datasets variant winning for almost all datasets. Figure 7 highlights this effect, showing that adding more datasets enables TabSTAR to outperform the baselines.

G.4 Numerical Verbalization (Q3)

We show the full results for the experiment in §6, with Table 31 illustrating the verbalizations in each variant. Note that we do not include an exact value verbalization, since it would increase the number of unique text inputs and place extra memory demands. The two variants which integrate numerical information into the verbalization dominate the experiment, although the improvement seems to be

Table 29: Normalized score with 95% CIs by the number of datasets used during TabSTAR pretraining.

Pretraining Datasets	0	16	64	256
Classification	0.352 ± 0.086	0.450 ± 0.084	0.558 ± 0.086	0.786 ± 0.076
Regression	0.338 ± 0.073	0.395 ± 0.068	0.642 ± 0.066	0.811 ± 0.055

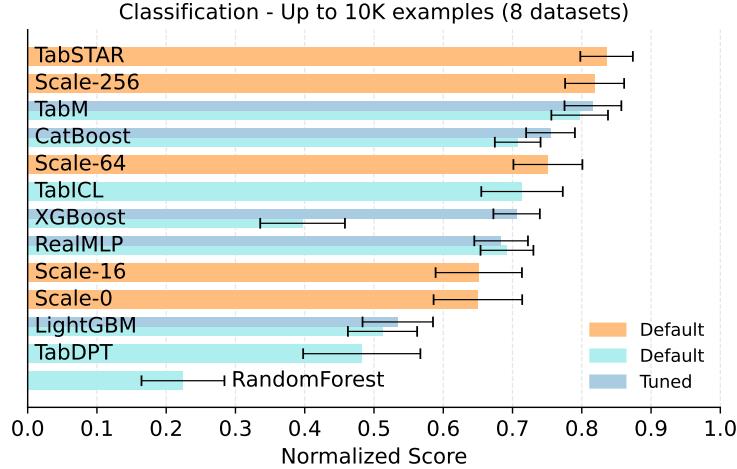


Figure 7: Comparison of normalized scores with 95% CIs for the pretraining effect experiment, comparing TabSTAR variants against baseline models in classification tasks.

marginal for some datasets. Interestingly, some datasets significantly underperform, with the *R27* dataset completely failing the task. The addition of the quantile information on top of the bin seems to have limited impact, although marginally winning on the average performance. Figure 8 presents a comparison with baseline models, highlighting the importance of incorporating richer numerical verbalizations beyond relying solely on the column name.

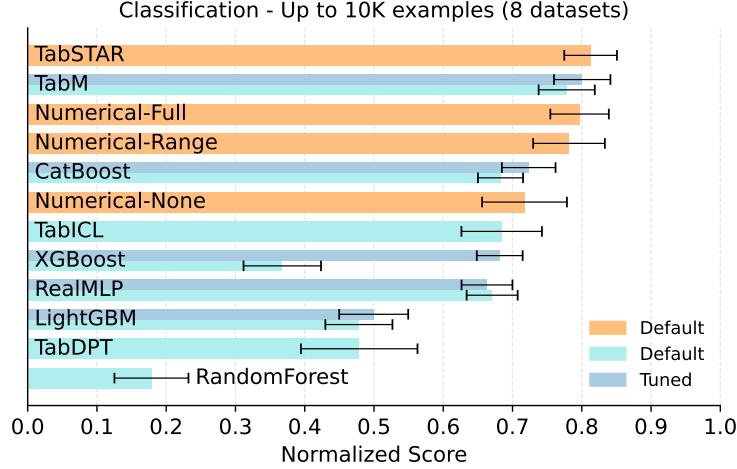


Figure 8: Comparison of normalized scores with 95% CIs for the numerical verbalization experiment, comparing TabSTAR variants against baseline models in classification tasks.

Table 30: Downstream performance for Q2: The Effect of Pretraining. Results for 20 datasets with 95% CI, for varying number of pretraining datasets. The top performance score is bolded first, and then all scores are rounded. We report AUROC for classification and R^2 for regression.

ID	0	16	64	256
C01	90.8±0.4	90.7±0.4	90.7±0.3	91.0±0.4
C02	98.0±0.4	97.4±0.8	97.8±0.3	98.1±0.3
C03	69.2±8.9	83.1±0.4	83.2±0.3	83.0±0.5
C05	90.7±0.8	90.3±1.1	90.6±0.5	91.5±0.9
C07	87.9±0.5	87.6±0.8	90.9±0.3	91.3±0.4
C11	78.2±0.6	77.4±0.9	78.0±0.4	78.8±0.5
C12	98.3±0.1	98.2±0.1	98.2±0.1	98.3±0.1
C13	74.0±0.5	73.5±0.8	73.7±1.0	74.1±0.7
R02	97.2±1.6	97.6±1.3	97.9±1.2	97.9±1.2
R03	67.3±1.9	68.5±1.0	71.2±0.7	71.2±0.8
R05	88.0±2.8	90.6±2.2	92.2±1.0	93.1±0.7
R09	88.1±1.0	88.4±0.7	88.8±0.5	89.0±0.6
R12	97.7±0.3	97.9±0.3	98.1±0.2	98.3±0.3
R13	49.4±2.7	49.1±3.6	48.0±3.0	51.4±2.9
R18	95.0±2.2	94.9±1.2	96.7±0.6	97.3±0.6
R23	82.2±1.6	81.2±2.3	81.8±2.3	83.2±2.2
R27	80.9±1.7	81.3±1.6	81.5±1.6	81.3±1.7
R30	13.1±4.1	18.5±4.9	18.5±3.9	19.1±4.0
R33	45.4±2.2	42.8±3.3	45.0±1.5	46.0±1.8
R34	83.8±4.3	86.0±3.3	88.0±3.4	87.9±3.5

Table 31: Illustrative verbalization of a numerical feature (*Age*) for the Q3: Numerical Verbalization experiment.

Value	Name	Name + Bin	TabSTAR
17	Age: Numeric	Age: Lower than 18	Age: Lower than 18 (Quantile 0%)
20	Age: Numeric	Age: 18–23	Age: 18–23 (Quantile 0–10%)
25	Age: Numeric	Age: 23–27	Age: 23–27 (Quantile 10–20%)
29	Age: Numeric	Age: 27–31	Age: 27–31 (Quantile 20–30%)
33	Age: Numeric	Age: 31–35	Age: 31–35 (Quantile 30–40%)
38	Age: Numeric	Age: 35–40	Age: 35–40 (Quantile 40–50%)
42	Age: Numeric	Age: 40–45	Age: 40–45 (Quantile 50–60%)
48	Age: Numeric	Age: 45–51	Age: 45–51 (Quantile 60–70%)
55	Age: Numeric	Age: 51–58	Age: 51–58 (Quantile 70–80%)
63	Age: Numeric	Age: 58–67	Age: 58–67 (Quantile 80–90%)
83	Age: Numeric	Age: 67–87	Age: 67–87 (Quantile 90–100%)
93	Age: Numeric	Age: Higher than 87	Age: Higher than 87 (Quantile 100%)
–	Age: Unknown Value	Age: Unknown Value	Age: Unknown Value

Table 32: Downstream performance for Q3: Numerical Verbalization. Results for 20 datasets with 95% CI, for different verbalizations. The top performance score is bolded first, and then all scores are rounded. We report AUROC for classification and R^2 for regression.

ID	Name	Name + Bin	TabSTAR
C01	91.0±0.4	91.2±0.4	91.0±0.4
C02	98.1±0.2	97.9±0.3	98.1±0.3
C03	83.3±0.4	83.3±0.3	83.0±0.5
C05	90.8±0.9	91.1±1.2	91.5±0.9
C07	91.2±0.2	91.4±0.4	91.3±0.4
C11	78.2±0.5	78.2±0.6	78.8±0.5
C12	98.2±0.1	98.4±0.1	98.3±0.1
C13	71.6±0.7	73.8±0.9	74.1±0.7
R02	98.0±1.1	98.0±1.1	97.9±1.2
R03	67.4±0.7	71.3±0.6	71.2±0.8
R05	93.1±1.1	93.2±0.8	93.1±0.7
R09	88.9±0.5	88.9±0.6	89.0±0.6
R12	98.2±0.2	98.3±0.2	98.3±0.3
R13	50.1±3.3	49.7±3.4	51.4±2.9
R18	97.1±0.6	97.1±0.7	97.3±0.6
R23	81.9±2.4	82.7±1.9	83.2±2.2
R27	16.7±6.6	82.0±1.5	81.3±1.7
R30	16.4±4.4	17.7±5.3	19.1±4.0
R33	45.6±1.1	46.2±2.2	46.0±1.8
R34	88.0±3.6	88.6±2.5	87.9±3.5