REVISITING PRETRAINING OBJECTIVES FOR TABULAR DEEP LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent deep learning models for tabular data currently compete with the traditional ML models based on decision trees (GBDT). Unlike GBDT, deep models can additionally benefit from pretraining, which is a workhorse of DL for vision and NLP. For tabular problems, several pretraining methods were proposed, but it is not entirely clear if pretraining provides consistent noticeable improvements and what method should be used, since the methods are often not compared to each other or comparison is limited to the simplest MLP architectures.

In this work, we aim to identify the best practices to pretrain tabular DL models that can be universally applied to different datasets and architectures. Among our findings, we show that using the object target labels during the pretraining stage is beneficial for the downstream performance and advocate several target-aware pretraining objectives. Overall, our experiments demonstrate that properly performed pretraining significantly increases the performance of tabular DL models, which often leads to their superiority over GBDTs.

1 INTRODUCTION

Tabular problems are ubiquitous in industrial ML applications, which include data described by a set of heterogeneous features, such as learning-to-rank, click-through rate prediction, credit scoring, and many others. Despite the current dominance of deep learning models in the ML literature, for tabular problems, the "old-school" decision tree ensembles (e.g., GBDT) are often the top choice for practitioners. Only recently, several works have proposed the deep models that challenge the supremacy of GBDT in the tabular domain Arik & Pfister (2020); Gorishniy et al. (2021); Somepalli et al. (2021); Gorishniy et al. (2022) and suggest that the question "tabular DL or GBDT" is yet to be answered.

An important advantage of deep models over GBDT is that they can potentially achieve higher performance via pretraining their parameters with a properly designed objective. These pretrained parameters, then, serve as a better than random initialization for subsequent finetuning for downstream tasks. For computer vision and NLP domains, pretraining is a de facto standard and is shown to be necessary for the state-of-the-art performance He et al. (2021); Devlin et al. (2019). For tabular problems such a consensus is yet to be achieved as well as the best practices of tabular pretraining are to be established. Pretraining for tabular problems is typically performed directly on the downstream target datasets, unlike pretraining in vision or NLP problems, for which huge "extra" data is available on the Internet. While a large number of prior works address the pretraining of tabular DL models Yoon et al. (2020); Bahri et al. (2022); Ucar et al. (2021); Darabi et al. (2021), it is challenging to make reliable conclusions about pretraining efficacy in tabular DL from the literature since experimental setups vary significantly. Some evaluation protocols assume the unlabeled data is abundant but use a small subset of labels from each dataset during finetuning for evaluation – demonstrating pretraining efficacy, but somewhat limiting the performance of supervised baselines.

In contrast, in our work, we focus on the setup with fully labeled tabular datasets to understand if pretraining helps tabular DL in a fully supervised setting and compare pretraining methods to the strong supervised baselines. To this end, we perform a systematic experimental evaluation of several pretraining objectives, identify the superior ones, and describe the practical details of how to perform tabular pretraining optimally. Our main findings, which are important for practitioners, are summarized below:

- Pretraining provides substantial gains over well-tuned supervised baselines in the fully supervised setup.
- Simple self-prediction based pretraining objectives are comparable to the objective based on contrastive learning. To the best of our knowledge, this behavior was not reported before in tabular DL.
- The object labels can be exploited for more effective pretraining. In particular, we describe several "target-aware" objectives and demonstrate that they often outperform their "unsupervised" counterparts.
- The pretraining provides the most noticeable improvements for the vanilla MLP architecture. In particular, their performance after pretraining becomes comparable to the state-of-the-art models trained from scratch, which is important for practitioners, who are interested in simple and efficient solutions.
- The ensembling of pretrained models is beneficial. It indicates that the pretraining stage does not significantly decrease the diversity of the models, despite the fact that all the models are initialized by the same set of parameters.

Overall, our work provides a set of recipes for practitioners interested in tabular pretraining, which results in higher performance for most of the tasks. The code of our experiments is available online.[1]

2 Related Work

Here we briefly review the lines of research that are relevant to our study.

Status quo in tabular deep learning. A plethora of recent works have proposed a large number of deep models for tabular data Klambauer et al. (2017); Popov et al. (2020); Arik & Pfister (2020); Song et al. (2019); Wang et al. (2017); Badirli et al. (2020); Hazimeh et al. (2020); Huang et al. (2020); Somepalli et al. (2021); Gorishniy et al. (2021); Kossen et al. (2021). Several systematic studies, however, reveal that these models typically do not consistently outperform the decision tree ensembles, such as GBDT (Gradient Boosting Decision Tree) Chen & Guestrin (2016); Prokhorenkova et al. (2018); Ke et al. (2017), which are typically the top-choice in various ML competitions Gorishniy et al. (2021); Shwartz-Ziv & Armon (2021). Additionally, several works have shown that the existing sophisticated architectures are not consistently superior to properly tuned simple models, such as MLP and ResNet Gorishniy et al. (2021); Kadra et al. (2021). Finally, the recent work Gorishniy et al. (2022) has highlighted that the appropriate embeddings of numerical features in the high-dimensional space are universally beneficial for different architectures. In our work, we experiment with pretraining of both traditional MLP-like models and advanced embedding-based models proposed in Gorishniy et al. (2022).

Pretraining in deep learning. For domains with structured data, like natural images or texts, pretraining is currently an established stage in the typical pipelines, which leads to higher general performance and better model robustness He et al. (2021); Devlin et al. (2019). Pretraining with the auto-encoding objective was also previously studied as a regularization strategy helping in the optimization process Erhan et al. (2010); El-Nouby et al. (2021) without large scale pretraining datasets. During the last years, several families of successful pretraining methods have been developed. An impactful line of research on pretraining is based on the paradigm of contrastive learning, which effectively enforces the invariance of the learned representations to the human-specified augmentations Chen et al. (2020); He et al. (2020). Another line of methods exploits the idea of self-prediction, i.e., these methods require the model to predict certain parts of the input given the remaining parts He et al. (2021); Devlin et al. (2019). In the vision community, the self-prediction based methods are shown to be superior to the methods that use contrastive learning objectives He et al. (2021). In our experiments, we demonstrate that self-prediction based objectives are comparable to the contrastive learning ones on tabular data, while being much simpler.

Pretraining for the tabular domain. Numerous pretraining methods were recently proposed in several recent works on tabular DL Arik & Pfister (2020); Yoon et al. (2020); Darabi et al. (2021); Ucar et al. (2021); Somepalli et al. (2021); Kossen et al. (2021). However, most of these works do

[1]Code: `https://anonymous.4open.science/r/pretrains-DD02`

not focus on the pretraining objective per se and typically introduce it as a component of their tabular DL pipeline. Moreover, the experimental setup varies significantly between methods. Therefore, it is difficult to extract conclusive evidence about pretraining effectiveness from the literature. To the best of our knowledge, there is only one systematic study on the tabular pretraining Bahri et al. (2022), but its experimental evaluation is performed only with the simplest MLP models, and we found that the superiority of the contrastive pretraining, reported in Bahri et al. (2022), does not hold for tuned models in our setup, where contrastive objective is comparable to the simpler self-prediction objectives.

3 REVISITING PRETRAINING OBJECTIVES

In this section, we evaluate the typical pretraining objectives under the unified experimental setup on the number of datasets from the literature on tabular DL. Our goal is to answer whether pretraining generally provides significant benefits over tuned models trained from scratch and to identify the pretraining objectives that lead to the best downstream task performance.

3.1 EXPERIMENTAL SETUP

We mostly follow the experimental setup from Gorishniy et al. (2021) and describe its main details here for completeness.

Notation. Each tabular dataset $(\mathbb{X}, \mathbb{Y})$ is represented by a set of pairs $\{(x_i, y_i)\}_{i=1}^{n}$, where $x_i = (x_i^1, \ldots, x_i^m) \in \mathbb{X}$ are the object's features (both numerical and categorical) and $y_i \in \mathbb{Y}$ is the target variable. The downstream task is either regression $\mathbb{Y} = \mathbb{R}$ or classification $\mathbb{Y} = \{1, \ldots, k\}$. Each model has the backbone $f(x|\theta)$ that is followed by two separate heads: a pretraining head $h(z|\mu)$ and a downstream task head $g(z|\lambda)$, with learnable parameters θ, μ, λ respectively, and $z = f(x|\theta)$ denotes the output of the backbone for an input object x.

Datasets. We evaluate the pretraining methods on a curated set of eleven middle to large scale datasets used in prior literature on tabular deep learning. The benchmark is biased towards tasks, where tuned MLP models were shown to be inferior to GBDT Gorishniy et al. (2021) since we aim to understand if pretraining can help the deep models to beat the "shallow" ones. The datasets represent a diverse set of tabular data problems with classification and regression targets. The main dataset properties are summarized in Table 1. We report ROC-AUC for all binary classification datasets,

	GE	CH	CA	HO	OT	HI	FB	AD	WE	CO	MI
#objects	9873	10000	20640	22784	61878	98049	197080	48842	397099	581012	1200192
#num.features	32	10	8	16	93	28	50	6	123	54	136
#cat.features	0	1	0	0	0	0	1	8	0	0	0
metric	Acc.	AUC	RMSE	RMSE	Log. Loss	AUC	RMSE	AUC	RMSE	Acc.	RMSE
#classes	5	2	–	–	9	2	–	2	–	7	–

Table 1: Datasets used for the experiments

accuracy for multi-class classification datasets and RMSE for regression datasets, with OT being the one exception, where we report log-loss, as it was used as a default metric in the corresponding Kaggle competition. We use the quantile-transform from the Scikit-learn library Pedregosa et al. (2011) to preprocess the numerical features for all datasets except OT, where the absence of such transformation was shown to be superior Gorishniy et al. (2022). Additional information about the datasets is provided in Appendix B.

Models. We use MLP as a simple deep baseline to compare and ablate the methods. Our implementation of MLP exactly follows Gorishniy et al. (2021), the model is regularized by dropout and weight decay. As more advanced deep models, we evaluate MLP equipped with numerical feature embeddings, specifically, target-aware piecewise linear encoding (MLP-T-LR) and embeddings with periodic activations (MLP-PLR) from Gorishniy et al. (2022). These models represent the current state-of-the-art solution for tabular DL Gorishniy et al. (2022), and are of interest as most prior work on pretraining in tabular DL focus on pretraining with the simplest MLP models in evaluation.

The implementation of models with numerical embeddings follows Gorishniy et al. (2022). We use AdamW Loshchilov & Hutter (2019) optimizer, do not use learning rate schedules and fix batch sizes for each dataset based on the dataset size (see Appendix B).

Pretraining. Pretraining is always performed directly on the target dataset and does not exploit additional data. The learning process thus comprises two stages. On the first stage, the model parameters are optimized w.r.t. the pretraining objective. On the second stage, the model is initialized with the pretrained weights and finetuned on the downstream classification or regression task. We focus on the fully-supervised setup, i.e., assume that target labels are provided for all dataset objects. Typically, pretraining stage involves the input corruption: for instance, to generate positive pairs in contrastive-like objectives or to corrupt the input for reconstruction in self-prediction based objectives. We use random feature resampling as a proven simple baseline for input corruption in tabular data Bahri et al. (2022); Yoon et al. (2020). Learning rate and weight decay are shared between the two stages (see Table 15 for the ablation). We fix the maximum number of pretraining iterations for each dataset at $100k$. On every $10k$-th iteration, we compute the value of the pretraining objective using the hold-out validation objects for early-stopping on large-scale WE, CO and MI datasets. On other datasets we directly finetune the current model every $10k$-th iteration and perform early-stopping based on the target metric after finetuning (we do not observe much difference between early stopping by loss or by downstream metric, see Table 16).

Hyperparameters & Evaluation. Hyperparameter tuning is crucial for a fair comparison, therefore, we use Optuna Akiba et al. (2019) to optimize the model and pretraining hyperparameters for each method on each dataset. We use the validation subset of each dataset for hyperparameter tuning. The exact search spaces for the hyperparameters of each method are provided in Appendix C.

We run the tuned configuration of each pretraining method with 15 random seeds and report the average metric on the test splits. When comparing to GBDT, we obtain three ensembles by splitting the fifteen single model predictions into three disjoint subsets of five models and averaging predictions within each subset. Then, we report the average metric over the three ensembles.

3.2 Comparing pretraining objectives

Here we compare the contrastive learning and self-prediction objectives from prior work in the described setup. For contrastive learning, we follow the method described in Bahri et al. (2022): use InfoNCE loss, consider corrupted inputs $\hat{x}$ as positives for x and the rest of the batch as negatives. For self-prediction methods, we evaluate two objectives: the first one is the reconstruction of the original x, given the corrupted input $\hat{x}$ (the reconstruction loss is computed for all columns), the second one is the binary mask prediction, where the objective is to predict the mask vector m indicating the corrupted columns from the corrupted input $\hat{x}$. The results of the comparison are in Table 2. We summarize our key findings below.

Contrastive is not superior. Both the reconstruction and the mask prediction objectives are preferable to the contrastive objective. The two self-prediction objectives have the advantage of being conceptually simpler, and easier to implement, while also being less resource-intensive (no need for the second view of augmented examples in each batch, simpler loss function). We thus recommend the self-prediction based objectives as a practical solution for pretraining in tabular DL.

Pretraining is beneficial for the state-of-the-art models. Models with the numerical feature embeddings also benefit from pretraining with either reconstruction or mask prediction demonstrating the top performance on the downstream task. However, the improvement is typically less noticeable compared to the vanilla MLPs.

There is no universal solution between self-prediction objectives. We observe that for some datasets the reconstruction objective outperforms the mask prediction (OT, WE, CO, MI), while on others the mask prediction is better (GE, CH, HI, AD). We also note that the mask prediction objective sometimes leads to unexpected performance drops for models with numerical embeddings (WE, MI), we do not observe significant performance drops for the reconstruction objective.

The main takeaway: simple pretraining strategies based on self-prediction lead to significant improvements in the downstream accuracy compared to the tuned supervised baselines learned from scratch across different tabular DL models and datasets. In practice, we recommend trying both re-

construction and mask prediction as tabular pretraining baselines, as either one might show superior performance depending on the dataset being used.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
MLP											
no pretraining	0.635	0.849	0.506	3.156	0.479	0.801	5.737	0.908	1.909	0.963	0.749
contrastive	0.672	**0.855**	0.455	**3.056**	0.469	**0.813**	5.697	0.910	1.881	0.960	0.748
rec	0.662	0.853	**0.445**	**3.044**	**0.466**	0.805	**5.641**	**0.910**	**1.875**	**0.965**	**0.746**
mask	**0.691**	**0.857**	0.454	3.113	0.472	**0.814**	**5.681**	**0.912**	1.883	0.964	0.748
MLP-PLR											
no pretraining	0.668	**0.858**	0.469	**3.008**	0.483	0.809	**5.608**	0.926	1.890	0.969	0.746
rec	0.667	0.852	**0.439**	**3.031**	**0.472**	0.808	**5.571**	**0.926**	**1.877**	**0.971**	**0.745**
mask	**0.685**	**0.863**	**0.434**	**3.007**	0.477	**0.818**	**5.586**	**0.927**	1.911	**0.970**	0.748
MLP-T-LR											
no pretraining	0.634	**0.866**	0.444	3.113	0.482	0.805	5.520	0.925	1.897	0.968	0.749
rec	**0.652**	0.857	**0.424**	3.109	**0.472**	0.808	**5.363**	0.924	**1.861**	**0.969**	**0.746**
mask	**0.654**	**0.868**	**0.424**	**3.045**	**0.472**	**0.818**	5.544	**0.926**	1.916	**0.969**	0.748

Table 2: Results for pretraining deep models with different objectives. We report metrics averaged over 15 seeds, bold entries correspond to results that are statistically significantly better (we use Tukey HSD test). The comparisons are separate for different models. ↑ corresponds to accuracy and ROC-AUC metrics, ↓ corresponds to RMSE and log-loss for OT. "scratch" stands for the supervised baseline, initialized with random weights

4 TARGET-AWARE PRETRAINING OBJECTIVES

In this section, we show that exploiting the target variables during the pretraining stage can further increase the downstream performance. Specifically, we evaluate several strategies to leverage information about targets during pretraining, identify the best ones and compare them to GBDT. Below we describe a list of target-aware pretraining objectives that we investigate.

Supervised loss with augmentations. A straightforward way to incorporate the target variable into the pretraining is by using the input corruption as an augmentation for the standard supervised learning objective. An important difference of this baseline in our setup to the one in Bahri et al. (2022) is that we treat learning on corrupted samples as a pretraining stage and finetune the entire model on the full uncorrupted dataset afterwards (we ablate this in subsection 5.3).

Supervised loss with augmentations + self-prediction. We evaluate a natural extension to the above baseline: a combination of the supervised objective with the unsupervised self-prediction. Note, that during the pretraining stage both losses are calculated on corrupted inputs, while the finetuning is performed on the non-corrupted dataset. For the self-prediction objectives we evaluate both the reconstruction and the mask prediction. We use different prediction heads for supervised and self-prediction objectives. We sum supervised and self-prediction losses with equal weights.

Target-aware pretraining. An alternative to the approaches described above is the modification of the pretraining task itself. An example of this approach is supervised contrastive learning Khosla et al. (2020), where the target variable is used to sample positive and negative examples. We introduce the target variable into the self-prediction based objectives with two modifications.

First, we condition the mask prediction or the reconstruction head on the original input's target by concatenating the hidden representation from the backbone network $z = f(\hat{x})$ with the target variable representation before passing it to the pretraining head to obtain predictions $p = h(\texttt{concat}[z, y])$. For classification datasets we encode y with one-hot-encoding, for regression targets we use the standard scaling.

Second, we change the input corruption scheme by sampling the replacement from the feature target conditional distribution where a target is different to the original. Intuitively, corrupting the input object x in the direction of the target different to the original makes the pretraining task more correlated with the downstream target prediction. Concretely, given an object-target pair (x_i, y_i), we

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓	Avg. Rank
						MLP						
no pretraining	0.635	0.849	0.506	3.156	0.479	0.801	5.737	0.908	1.909	0.963	0.749	5.5 ± 1.4
mask	0.691	0.857	0.454	3.113	0.472	0.814	5.681	0.912	1.883	0.964	0.748	3.8 ± 1.4
rec	0.662	0.853	0.445	**3.044**	0.466	0.805	5.641	0.910	1.875	0.965	0.746	3.6 ± 1.5
sup	0.693	0.856	0.441	3.077	**0.459**	0.814	5.689	0.914	1.883	0.968	0.748	3.0 ± 1.0
mask + target	0.683	0.857	0.434	**3.056**	0.468	**0.819**	5.633	0.914	1.876	0.965	0.748	2.9 ± 1.3
rec + target	0.659	0.853	0.454	**3.044**	0.463	0.806	5.636	0.909	1.884	0.965	**0.745**	3.7 ± 1.9
mask + sup	0.693	0.857	0.436	3.099	**0.458**	0.817	5.685	0.915	1.873	0.967	0.748	2.7 ± 1.2
rec + sup	0.684	0.854	0.436	**3.012**	**0.456**	0.815	5.672	0.911	**1.862**	0.967	0.747	2.6 ± 1.5
						MLP-PLR						
no pretraining	0.668	0.858	0.469	**3.008**	0.483	0.809	5.608	0.926	1.890	0.969	0.746	3.5 ± 1.7
mask	0.685	**0.863**	0.434	**3.007**	0.477	0.818	5.586	0.927	1.911	0.970	0.748	2.8 ± 1.7
rec	0.667	0.852	0.439	**3.031**	0.472	0.808	5.571	0.926	1.877	**0.971**	**0.745**	2.6 ± 1.2
sup	**0.710**	0.859	0.433	3.136	0.479	0.811	5.521	0.924	1.873	**0.971**	0.748	2.5 ± 1.2
mask + target	0.694	0.862	**0.425**	**3.023**	0.474	**0.821**	5.537	**0.929**	1.911	0.969	0.749	2.5 ± 1.9
rec + target	0.688	0.860	0.445	**3.064**	0.475	0.812	5.507	0.927	1.887	**0.971**	0.748	2.7 ± 1.3
mask + sup	**0.711**	**0.866**	0.441	3.129	0.480	0.813	5.480	0.925	1.875	0.969	**0.745**	2.5 ± 1.4
rec + sup	**0.709**	0.858	0.433	**3.059**	0.465	0.807	5.571	0.927	**1.865**	**0.971**	**0.745**	1.9 ± 1.2
						MLP-T-LR						
no pretraining	0.634	**0.866**	0.444	3.113	0.482	0.805	5.520	0.925	1.897	0.968	0.749	3.9 ± 1.7
mask	0.654	**0.868**	**0.424**	**3.045**	0.472	0.818	5.544	0.926	1.916	0.969	0.748	2.8 ± 1.7
rec	0.652	0.857	**0.424**	3.109	0.472	0.808	**5.363**	0.924	**1.861**	0.969	**0.746**	2.5 ± 1.4
sup	0.682	0.860	0.430	3.135	0.471	0.807	5.525	0.927	1.893	**0.971**	0.747	2.8 ± 1.5
mask + target	0.649	**0.865**	**0.421**	**3.058**	0.474	**0.820**	5.644	**0.929**	1.924	0.969	0.749	2.8 ± 2.1
rec + target	0.668	**0.864**	0.440	3.113	0.473	0.806	5.493	0.927	**1.862**	0.969	**0.746**	2.5 ± 1.4
mask + sup	0.676	0.858	0.429	3.199	0.468	0.814	5.510	0.926	1.869	**0.971**	0.748	2.5 ± 0.8
rec + sup	0.678	**0.865**	0.437	3.112	0.462	0.807	5.516	0.927	**1.862**	**0.970**	0.748	2.4 ± 1.2

Table 3: Variations of the target-aware pretraining schemes. Notation follows Table 2. Bold results indicate statistically significant winners across all models and methods. "+ target" denotes target-conditioned pretraining, "+ sup" denotes auxiliary supervised head.

sample a new target $\hat{y}_i$ from a uniform distribution over the set $\{y \mid y \neq y_i\}$ [2], then each feature x^j is replaced with a sample from the $p(x^j|\hat{y}_i)$ distribution, instead of $p(x^j)$.

4.1 Comparing target-aware objectives

Here we compare the strategies of incorporating the target variable into pretraining. The results of the comparison are in Table 3. Our key findings are formulated below.

Supervised loss with augmentations is another strong baseline for MLP. Pretraining with the supervised loss on corrupted data consistently improves over supervised training from scratch for the MLP. This objective is a strong baseline along with the self-prediction based objectives. However, for models with numerical embeddings the supervised objective with corruptions is less consistent and sometimes is inferior to training from scratch, thus for these models we recommend the self-prediction objectives alone as baselines.

Target-aware objectives demonstrate the best performance. Both the supervised loss with self-prediction and modified self-prediction objectives improve over the unsupervised pretraining baselines across datasets and model architectures.

For the objective with the combination of supervised and self-prediction losses the variation with the reconstruction loss is the most consistent across models and dataset with no performance drops below the pretraining-free baseline. The variant with mask prediction shows similar performance and stability for the MLP, but is not as good as reconstruction for models with numerical embeddings.

[2]For regression problems, when the target variable is continuous, we preliminarily discretize it into n uniform bins, where n is chosen according to the Freedman–Diaconis rule Freedman & Diaconis (1981).

For the target-aware self-prediction objectives, the modified mask prediction delivers significant improvements over its unsupervised counterpart. Modified reconstruction objective, however, does not improve over unsupervised reconstruction objective.

Main takeaways: Target-aware objectives help further increase the downstream performance, improving upon their unsupervised counterparts. For the reconstruction based self-prediction baseline the addition of the supervised loss is most beneficial ("rec + sup" from Table 4), for the mask prediction objective it's target-aware modification provides more improvements ("mask + target" from Table 4). A simple MLP model pretrained with those "target-aware" objectives often reaches or surpasses complex models with numerical embeddings trained from scratch. In practice, we recommend first trying the baseline pretraining objectives ("rec", "mask", "sup" from Table 4), choosing the suitable baseline for the dataset and improving it accordingly: supervised loss for the reconstruction and target-aware modification for the mask prediction.

4.2 Comparison to GBDT

Here we compare MLPs and MLPs with numerical feature embeddings pretrained with the supervised loss with reconstruction and the target-aware mask prediction objectives to the GBDTs. Table 4 shows the results of the comparison.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓	Avg. Rank
CatBoost	0.692	**0.864**	0.430	3.093	0.450	0.807	5.226	0.928	1.801	0.967	**0.741**	2.6 ± 1.4
XGBoost	0.683	0.860	0.434	3.152	0.454	0.805	5.338	0.927	**1.782**	0.969	**0.742**	3.0 ± 1.5
MLP												
no pretraining	0.656	0.852	0.482	3.055	0.467	0.805	5.666	0.910	1.850	0.968	0.747	4.8 ± 1.1
mask + target	0.709	0.860	0.414	2.949	0.457	**0.828**	5.551	0.916	1.809	0.969	0.746	2.8 ± 1.2
rec + sup	0.709	0.859	0.419	2.951	**0.442**	0.817	5.531	0.913	1.801	0.973	0.745	2.5 ± 1.2
MLP-P-LR												
no pretraining	0.695	**0.864**	0.454	2.953	0.470	0.814	5.324	0.928	1.835	**0.974**	0.744	2.6 ± 1.2
mask + target	**0.719**	**0.866**	**0.407**	2.952	0.458	**0.828**	5.373	**0.930**	1.849	0.973	0.745	2.1 ± 1.2
rec + sup	**0.737**	0.862	0.424	2.964	0.449	0.811	**5.124**	**0.929**	1.813	**0.974**	0.744	2.0 ± 1.0
MLP-T-LR												
no pretraining	0.662	**0.868**	0.437	3.028	0.472	0.808	5.424	0.927	1.850	0.972	0.747	3.7 ± 1.1
mask + target	0.673	**0.868**	**0.410**	**2.894**	0.460	**0.827**	5.458	**0.930**	1.849	0.972	0.746	2.4 ± 1.4
rec + sup	0.705	**0.866**	0.425	3.057	**0.444**	0.814	5.422	0.927	1.811	**0.974**	0.746	2.5 ± 1.1

Table 4: Comparison of pretrained models to GBDT. Notation follows Table 2. Results represent ensembles of models. Bold entries correspond to the overall statistically significant best entries.

We observe that both pretraining with target aware objectives and using numerical feature embeddings consistently improve the performance of the simple MLP backbone. In particular, MLP coupled with target aware pretraining starts to outperform GBDT on 4 datasets (GE, CA, OT, HI). Combined with numerical feature embeddings, pretraining improves MLP performance further, making it superior to GBDT on the majority of the datasets, with two exceptions in WE and MI.

5 Analysis

5.1 Investigating the properties of pretrained models

In this section, we provide a possible explanation of why the incorporation of the target variable into pretraining can lead to better downstream task performance. We do this through the experiments on the controllable synthetic data. Here we describe the properties and the generation process of the data and our observations on the differences of the pretraining schemes.

We follow the synthetics generation protocol described in Gorishniy et al. (2021) with a modification that allows for the manual control of the feature importance for the particular prediction task. Concretely, we generate the objects features $\{x_i\}_{i=1}^{n}$ as samples from the multivariate Gaus-

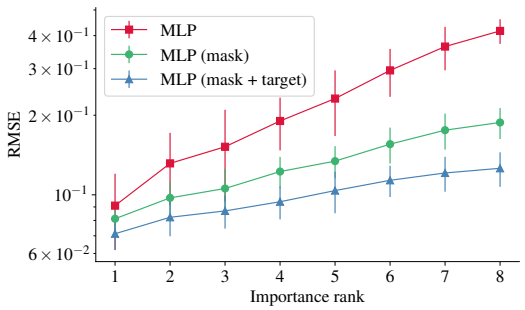

Figure 1: The decodability of object feature from the intermediate representations computed by the pretrained models and the models trained from scratch. The pretrained models decently capture the information about all the features, while the randomly initialized models capture the most informative features and suppress the others.

sian distribution with zero mean and covariance matrix Σ with identical diagonal and a constant $c = 0.5$ everywhere else. To generate the corresponding objects targets $\{y_i\}_{i=1}^n$ we sample a vector $p \in \mathbb{R}^m$ from the Dirichlet distribution $p \sim \text{Dir}(1_m)$ and let p define the influence of the objects features on the target. Then, we build an ensemble of 10 random oblivious decision trees $\{T_i(x)\}_{i=1}^{10}$ of depth 10, where on each tree level we sample a feature $j \sim \text{Cat}(p)$ and a threshold $t \sim \text{Unif}(\min(x^j), \max(x^j))$ for a decision rule. For each tree leaf, we sample a scalar $l \sim \mathcal{N}(0, 1)$, representing a logit in binary classification. We define the binary targets as follows: $y(x) = \text{I}\left\{\frac{1}{10}\sum_{i=1}^{10} T_i(x) > 0\right\}$. Intuitively if a particular feature is often used for splitting in the nodes of a decision tree, it would have more influence on the target variable. Indeed, we find the feature importances[3] correlate well with the predefined vector p. We set the size of the dataset $n = 50.000$ and the number of features $m = 8$ and generate 50 datasets with different feature importance vectors p for the analysis.

For each generated dataset, we then check whether the finetuned models capture the information about object features in their intermediate representations. Specifically, we train an MLP to predict the value of the i-th object feature given the frozen embeddings produced by a finetuned network initialized from (a) random initialization, (b) mask prediction pretraining, (c) target-aware mask prediction pretraining. The separate MLP is used for each feature and the RMSE learning objective is used. Then we report the RMSE on the test set for all features $i \in [0, m]$ along with their importance rank in the dataset on Figure 1. Here, the lower ranks correspond to the more important features. Figure 1 reveals that the target-aware pretraining enables the model to capture more information about the informative features compared to the "unsupervised" pretraining and, especially, to the learning from scratch. The latter one successfully captures the most informative feature from the training data, while suppressing the less important, but still significant features. We conjecture that this is the source of superiority of the target-aware pretraining.

5.2 Efficient ensembling

In this section we show, that it is possible to construct ensembles from one pretraining checkpoint (pretrained with the target conditioned mask prediction objective). To this end, we run finetuning with 15 different random seeds starting from the one pretrained checkpoint. Table 5 shows the results.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
single	0.683	0.857	0.434	3.056	0.468	0.819	5.633	0.914	1.876	0.965	0.748
standard ensemble	0.709	0.860	0.414	2.949	0.457	0.828	5.551	0.916	1.809	0.969	0.746
efficient ensemble	0.702	0.861	0.411	2.967	0.461	0.825	5.590	0.917	1.820	0.969	0.746

Table 5: Efficient ensembling for MLP mask + target

Both ensembling the models from a shared pretrain checkpoint and ensembling multiple independent pretraining runs produces strong ensembles, which shows that it is sufficient to pretrain once and create ensembles by several independent finetuning processes. This is important in practice since

[3]Computed with the CatBoost method "`get_feature_importance()`"

finetuninig is typically cheaper (i.e. requires fewer iterations), and still is able to produce diverse models from the one pretraining checkpoint for ensembles of comparable quality.

5.3 On importance of finetuning on clean data

Here we show, that the second stage of finetuning the model on the entire dataset without input corruption is often necessary for the best downstream performance. To this end we compare finetuning the models on clean data with using models right after pretraining for two objectives: supervised loss and supervised loss with the reconstruction objective. Across all datasets for both methods

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
no pretraining	0.635	0.849	0.506	3.156	0.479	0.801	5.737	0.908	1.909	0.963	0.749
aug	0.693	0.856	0.441	3.077	0.459	0.814	5.689	0.914	1.883	0.968	0.748
aug | no finetune	0.674	0.853	0.464	3.251	0.461	0.808	6.167	0.910	1.938	0.958	0.752
rec + sup	0.684	0.854	0.436	3.012	0.456	0.815	5.672	0.911	1.862	0.967	0.747
rec + sup | no finetune	0.683	0.853	0.467	3.232	0.474	0.811	6.044	0.907	1.901	0.956	0.752

Table 6: Finetuning MLP on clean data versus using the model trained on corrupted inputs only

finetuning on uncorrupted data with the supervised loss proves to be essential for the best performance. Sometimes excluding the second finetuning stage degrades the performance below the tuned supervised baseline of training from scratch.

5.4 Does pretraining require more compute?

In this section we investigate how much more time is spent on pretraining, compared to training the models from scratch. We run pretraining with "rec + sup" objective with 50k, 100k and 150k pretraining iterations thresholds (early-stopping, in theory, could make 100k and more iterations equivalent to the 50k, but in practice it was not the case). We report downstream performance along with average time spent to pretrain and finetune a model in Table 7. Pretraining often requires by

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
no pretraining	0.635	0.849	0.506	3.156	0.479	0.801	5.737	0.908	1.909	0.963	0.749
	29s	10s	25s	26s	38s	29s	159s	12s	57s	352s	211s
rec + sup | 50k	0.679	0.857	0.441	3.064	0.462	0.813	5.650	0.910	1.879	0.966	0.747
	327s	227s	280s	277s	355s	256s	624s	403s	242s	593s	292s
rec + sup | 100k	0.684	0.854	0.436	3.012	0.456	0.815	5.672	0.911	1.862	0.967	0.747
	661s	312s	533s	455s	570s	526s	740s	759s	439s	649s	472s
rec + sup | 150k	0.692	0.859	0.435	3.012	0.456	0.816	5.629	0.910	1.866	0.968	0.746
	891s	338s	758s	509s	712s	862s	916s	1069s	663s	927s	665s

Table 7: Comparison of time spent for MLP training from scratch and "rec + sup" pretraining with 50k, 100k and 150k max-iterations threshold. Second row in each group reports average time spent training one model in seconds on an A100 GPU.

an order of magnitude more compute, it is especially apparent on smaller scale datasets like GE, CH, CA, HO, OT, HI, AD. However, the absolute time spent on pretraining is still acceptable, as the original training from scratch takes seconds on small datasets. Generally the more iterations you use for pretraining, the better downstream quality you get.

6 Conclusion

In this work, we have systematically evaluated typical pretraining objectives for tabular deep learning. We have revealed several important recipes for optimal pretraining performance that can be universally beneficial across various problems and models. Our findings confirm that pretraining can significantly improve the performance of tabular deep models and provide additional evidence that tabular DL can become a strong alternative to GBDT.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *KDD*, 2019.

Sercan O. Arik and Tomas Pfister. Tabnet: Attentive interpretable tabular learning. *arXiv*, 1908.07442v5, 2020.

Sarkhan Badirli, Xuanqing Liu, Zhengming Xing, Avradeep Bhowmik, Khoa Doan, and Sathiya S. Keerthi. Gradient boosting neural networks: Grownet. *arXiv*, 2002.07971v2, 2020.

Dara Bahri, Heinrich Jiang, Yi Tay, and Donald Metzler. Scarf: Self-supervised contrastive learning using random feature corruption. In *International Conference on Learning Representations*, 2022. URL `https://openreview.net/forum?id=CuV_qYkmKb3`.

P. Baldi, P. Sadowski, and D. Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature Communications*, 5, 2014.

Jock A. Blackard and Denis J. Dean. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture*, 24(3):131–151, 2000.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *SIGKDD*, 2016.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Sajad Darabi, Shayan Fazeli, Ali Pazoki, Sriram Sankararaman, and Majid Sarrafzadeh. Contrastive mixup: Self-and semi-supervised learning for tabular domain. *arXiv preprint arXiv:2108.12296*, 2021.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv*, 1810.04805v2, 2019.

Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*, 2021.

Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. Why does unsupervised pre-training help deep learning? *J. Mach. Learn. Res.*, 11: 625–660, mar 2010. ISSN 1532-4435.

David Freedman and Persi Diaconis. On the histogram as a density estimator:l 2 theory. *Z. Wahrscheinlichkeitstheorie verw Gebiete*, 57(4):453–476, December 1981.

Yura Gorishniy, Ivan Rubachev, and Artem Babenko. On embeddings for numerical features in tabular deep learning. *arXiv preprint arXiv:2203.05556*, 2022.

Yury Gorishniy, Ivan Rubachev, Valentin Khrulkov, and Artem Babenko. Revisiting deep learning models for tabular data. In *NeurIPS*, 2021.

Hussein Hazimeh, Natalia Ponomareva, Petros Mol, Zhenyu Tan, and Rahul Mazumder. The tree ensemble layer: Differentiability meets conditional computation. In *ICML*, 2020.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.

Xin Huang, Ashish Khetan, Milan Cvitkovic, and Zohar Karnin. Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv*, 2012.06678v1, 2020.

Arlind Kadra, Marius Lindauer, Frank Hutter, and Josif Grabocka. Well-tuned simple nets excel on tabular datasets. In *NeurIPS*, 2021.

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. *Advances in neural information processing systems*, 30:3146–3154, 2017.

R. Kelley Pace and Ronald Barry. Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297, 1997.

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.

Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *NIPS*, 2017.

Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: a decision-tree hybrid. In *KDD*, 1996.

Jannik Kossen, Neil Band, Clare Lyle, Aidan N. Gomez, Tom Rainforth, and Yarin Gal. Self-attention between datapoints: Going beyond individual input-output pairs in deep learning. In *NeurIPS*, 2021.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

Renata C. B. Madeo, Clodoaldo Ap. M. Lima, and Sarajane Marques Peres. Gesture unit segmentation using support vector machines: segmenting gestures from rest positions. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing, SAC*, 2013.

Andrey Malinin, Neil Band, German Chesnokov, Yarin Gal, Mark John Francis Gales, Alexey Noskov, Andrey Ploskonosov, Liudmila Prokhorenkova, Ivan Provilkov, Vatsal Raina, Vyas Raina, Mariya Shmatova, Panos Tigas, and Boris Yangel. Shifts: A dataset of real distributional shift across multiple large-scale tasks. *ArXiv*, abs/2107.07455, 2021.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Sergei Popov, Stanislav Morozov, and Artem Babenko. Neural oblivious decision ensembles for deep learning on tabular data. In *ICLR*, 2020.

Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. Catboost: unbiased boosting with categorical features. In *NeurIPS*, 2018.

Tao Qin and Tie-Yan Liu. Introducing LETOR 4.0 datasets. *arXiv*, 1306.2597v1, 2013.

Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *arXiv*, 2106.03253v1, 2021.

Kamaljot Singh, Ranjeet Kaur Sandhu, and Dinesh Kumar. Comment volume prediction using neural networks and decision trees. In *IEEE UKSim-AMSS 17th International Conference on Computer Modelling and Simulation, UKSim*, 2015.

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. *arXiv*, 2106.01342v1, 2021.

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 *CIKM*, 2019.

Talip Ucar, Ehsan Hajiramezanali, and Lindsay Edwards. Subtab: Subsetting features of tabular data for self-supervised representation learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luís Torgo. Openml: networked science in machine learning. *arXiv*, 1407.7722v1, 2014.

Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. Deep & cross network for ad click predictions. In *ADKDD*, 2017.

Jinsung Yoon, Yao Zhang, James Jordon, and Mihaela van der Schaar. Vime: Extending the success of self- and semi-supervised learning to tabular domain. In *NeurIPS*, 2020.

A APPENDIX

B DATASETS

We used the following datasets:

- Gesture Phase Prediction (Madeo et al. (2013))
- Churn Modeling[4]
- California Housing (real estate data, Kelley Pace & Barry (1997))
- House 16H[5]
- Adult (income estimation, Kohavi (1996))
- Otto Group Product Classification[6]
- Higgs (simulated physical particles, Baldi et al. (2014); we use the version with 98K samples available at the OpenML repository Vanschoren et al. (2014))
- Facebook Comments (Singh et al. (2015))
- Covertype (forest characteristics, Blackard & Dean. (2000))
- Microsoft (search queries, Qin & Liu (2013)). We follow the pointwise approach to learning-to-rank and treat this ranking problem as a regression problem.
- Weather (temperature, Malinin et al. (2021)). We take 10% of the dataset for our experiments due to the its large size.

C HYPERPARAMETERS

C.1 CATBOOST

We fix and do not tune the following hyperparameters:

- `early-stopping-rounds` $= 50$
- `od-pval` $= 0.001$
- `iterations` $= 2000$

For tuning on the MI and CO datasets, we set the `task_type` parameter to "GPU". In all other cases (including the evaluation on these two datasets), we set this parameter to "CPU".

[4]https://www.kaggle.com/shrutimechlearn/churn-modelling
[5]https://www.openml.org/d/574
[6]https://www.kaggle.com/c/otto-group-product-classification-challenge/data

Parameter	Distribution
Max depth	$\text{UniformInt}[1, 10]$
Learning rate	$\text{LogUniform}[0.001, 1]$
Bagging temperature	$\text{Uniform}[0, 1]$
L2 leaf reg	$\text{LogUniform}[1, 10]$
Leaf estimation iterations	$\text{UniformInt}[1, 10]$
# Iterations	100

Table 8: CatBoost hyperparameter space

C.2 XGBoost

We fix and do not tune the following hyperparameters:

- `booster` = "gbtree"
- `early-stopping-rounds` = 50
- `n-estimators` = 2000

Parameter	Distribution
Max depth	$\text{UniformInt}[3, 10]$
Min child weight	$\text{LogUniform}[0.0001, 100]$
Subsample	$\text{Uniform}[0.5, 1]$
Learning rate	$\text{LogUniform}[0.001, 1]$
Col sample by tree	$\text{Uniform}[0.5, 1]$
Gamma	$\{0, \text{LogUniform}[0.001, 100]\}$
Lambda	$\{0, \text{LogUniform}[0.1, 10]\}$
# Iterations	100

Table 9: XGBoost hyperparameter space.

C.3 MLP

We fix and do not tune the following hyperparameters:

- `Layer size` = 512
- `Head hidden size` = 512

Parameter	Distribution
# Layers	$\text{UniformInt}[1, 8]$
Dropout	$\{0, \text{Uniform}[0, 0.5]\}$
Learning rate	$\text{LogUniform}[5e\text{-}5, 0.005]$
Weight decay	$\{0, \text{LogUniform}[1e\text{-}6, 1e\text{-}3]\}$
Corrupt Probability	$\{0, \text{Uniform}[0.2, 0.8]\}$
# Iterations	100

Table 10: MLP hyperparameter space.

C.4 EMBEDDING HYPERPARAMETERS

We fix and do not tune the following hyperparameters:

- `Layer size` $= 512$
- `Head hidden size` $= 512$

The distribution for the output dimensions of linear layers is $\mathrm{UniformInt}[1, 128]$.

`PLR`, `T-LR`. We share the same hyperparameter space for models with embeddings across all datasets.

For the target-aware embeddings (tree-based) `T-LR`, the distribution for the number of leaves is $\mathrm{UniformInt}[2, 256]$, the distribution for the minimum number of items per leaf is $\mathrm{UniformInt}[1, 128]$ and the distribution for the minimum information gain required for making a split is $\mathrm{LogUniform}[1e\text{-}9, 0.01]$.

For the periodic embeddings `PLR`. The distribution for k is $\mathrm{UniformInt}[1, 128]$, the distribution for the σ parameter is $\mathrm{LogUniform}[0.01, 100]$

D PRETRAINING THE FT-TRANSFORMER MODEL

Here we evaluate the recommended pretraining objectives for pretraining the FT-Transformer model from Gorishniy et al. (2021). We evaluate pretraining for the FT-Transformer on a subset of datasets for the reconstruction, mask prediction, and their supervised variations (mask+target and rec+sup). The results are in the Table 11.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	AD ↑	Avg. Rank
MLP-PLR								
no pretraining	0.668	0.858	0.469	**3.008**	0.483	0.809	0.926	2.7 ± 1.2
mask	0.685	**0.863**	0.434	**3.007**	0.477	0.818	0.927	2.1 ± 1.1
rec	0.667	0.852	0.439	**3.031**	0.472	0.808	0.926	2.1 ± 1.1
rec + supervised	**0.709**	0.858	0.433	**3.059**	**0.465**	0.807	0.927	1.7 ± 1.2
mask + target	0.694	**0.862**	**0.425**	**3.023**	0.474	0.821	**0.929**	1.9 ± 1.1
FT-Transformer								
no pretraining	0.631	**0.860**	0.470	3.240	0.479	0.809	0.912	3.7 ± 1.5
mask	0.696	**0.864**	**0.429**	3.113	**0.466**	**0.845**	0.916	1.9 ± 1.4
rec	0.687	**0.860**	0.438	3.120	0.478	0.840	0.917	2.6 ± 1.3
rec + supervised	**0.718**	**0.862**	**0.430**	3.078	**0.463**	0.825	0.917	1.9 ± 1.1
mask + target	0.694	**0.861**	0.436	3.083	0.467	**0.842**	0.918	2.0 ± 0.9

Table 11: Results for pretraining the FT-Transformer model, along with the results for pretraining the MLP-PLR. Bold accounts for the best-performing model between both MLP-PLR and FT-Transformer.

We can see that pretraining provides performance improvements for the FT-Transformer model comparable to those for MLP-PLR, with MLP-PLR remaining state-of-the-art with or without pretraining. Including transformers does not reveal new patterns or behavior, but provides additional evidence that pretraining does help neural networks on tabular data.

E COMPARISON TO THE LEARNING RATE SCHEDULERS

Since in our setup pretraining is a technique which helps in the optimization process, we compare it to various commonly used learning rate schedules which are also used as an aid in the optimization. The results for various learning rate schedulers along with two self-training based pretraining baselines are reported in Table 12.

We can see that the learning rate schedules provide some improvement and help the optimization process, however, the effect is almost negligible compared to the one from pretraining. We conjec-

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓	Avg. Rank
scratch	0.635	0.849	0.506	3.156	0.479	0.801	5.737	0.908	1.909	0.963	0.749	2.7 ± 0.6
cyclic	0.616	0.850	0.498	3.148	0.479	0.801	5.835	0.908	1.920	0.962	0.749	3.2 ± 0.9
warmup	0.655	**0.854**	0.491	3.123	0.478	0.799	**5.662**	0.908	1.912	0.964	0.749	2.5 ± 1.1
warmup + lin	0.669	0.851	0.489	3.123	0.480	0.801	**5.669**	0.908	1.912	**0.968**	0.748	2.2 ± 0.8
warmup + cos	0.639	0.851	0.494	3.107	0.478	0.801	5.823	0.908	1.893	0.962	0.748	2.6 ± 0.6
mask	**0.691**	**0.857**	**0.454**	3.113	0.472	**0.814**	**5.681**	**0.912**	1.883	0.964	0.748	1.5 ± 0.7
rec	0.662	0.853	**0.445**	**3.044**	**0.466**	0.805	**5.641**	**0.910**	**1.875**	0.965	**0.746**	1.5 ± 0.5

Table 12: Comparison of self-prediction based pretrains to learning rate schedulers. All the results are for the MLP model

ture that exhaustive hyperparameter tuning diminishes the benefits of using the advanced learning rate schedules.

F DATASETS WITH MORE CATEGORICAL FEATURES

Additionally, we experiment with datasets that contain more categorical features. We add 4 datasets, the information about the datasets is presented in Table 13.

Abbr	Name	# Train	# Validation	# Test	# Num	# Cat	Task type	Batch size
DI	Diamond	34521	8631	10788	6	3	Regression	128
BL	Black Friday	106764	26692	33365	5	4	Regression	128
BR	Brazilian houses	6842	1710	2136	8	4	Regression	128
BA	Bank	28934	7234	9043	7	9	Binclass	128

Table 13: Datasets with categorical features

The results for the MLP and MLP-PLR with self-prediction-based pretrains along with their supervised variations are presented in Table 14.

	DI ↓	BL ↓	BR ↓	BA ↑
CatBoost	0.134	0.682	0.407	**0.932**
XGBoost	0.136	**0.680**	0.407	0.930
MLP				
no pretraining	0.140	0.697	0.406	0.927
rec	0.140	0.695	0.400	0.926
mask	0.138	0.691	**0.385**	0.929
rec + supervised	0.135	0.689	0.394	0.930
mask + target	0.135	0.691	**0.376**	**0.933**
MLP-PLR				
no pretraining	0.134	0.689	0.399	0.929
rec	0.133	0.684	0.390	0.931
mask	**0.131**	0.685	0.392	0.931
rec + supervised	0.133	0.681	0.393	0.928
mask + target	**0.132**	0.685	**0.385**	**0.933**

Table 14: The results for single models with various pretraining objectives on datasets with categorical features.

We can see that results on the datasets with more categorical features align with the results in the main text: pretraining strategies based on self-prediction improve upon the supervised baselines and target-aware objectives help further increase the downstream performance.

G SHARE OR SPLIT LEARNING RATE AND WEIGHT DECAY BETWEEN PRETRAINING AND FINETUNING?

Here we demonstrate that tuning and using the same learning rate and weight decay for both pretraining and finetuning results in similar performance to tuning these parameters separately for the two stages. We opt for sharing the learning rate and weight decay for pretraining and finetuning in all the experiments in the paper.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
shared wd / shared lr	$0.683_{\pm 1e\text{-}2}$	$0.857_{\pm 2e\text{-}3}$	$0.434_{\pm 7e\text{-}3}$	$3.056_{\pm 4e\text{-}2}$	$0.468_{\pm 2e\text{-}3}$	$0.819_{\pm 2e\text{-}3}$	$5.633_{\pm 4e\text{-}2}$	$0.914_{\pm 1e\text{-}3}$	$1.876_{\pm 5e\text{-}3}$	$0.965_{\pm 7e\text{-}4}$	$0.748_{\pm 4e\text{-}4}$
shared wd / split lr	$0.697_{\pm 9e\text{-}3}$	$0.857_{\pm 3e\text{-}3}$	$0.431_{\pm 7e\text{-}3}$	$3.032_{\pm 3e\text{-}2}$	$0.469_{\pm 2e\text{-}3}$	$0.819_{\pm 2e\text{-}3}$	$5.647_{\pm 4e\text{-}2}$	$0.915_{\pm 9e\text{-}4}$	$1.934_{\pm 8e\text{-}3}$	$0.964_{\pm 9e\text{-}4}$	$0.748_{\pm 4e\text{-}4}$
split wd / split lr	$0.688_{\pm 9e\text{-}3}$	$0.856_{\pm 3e\text{-}3}$	$0.430_{\pm 4e\text{-}3}$	$3.046_{\pm 4e\text{-}2}$	$0.471_{\pm 3e\text{-}3}$	$0.821_{\pm 7e\text{-}4}$	$5.734_{\pm 5e\text{-}2}$	$0.914_{\pm 7e\text{-}4}$	$1.891_{\pm 6e\text{-}3}$	$0.964_{\pm 1e\text{-}3}$	$0.748_{\pm 3e\text{-}4}$
split wd / split lr	$0.694_{\pm 1e\text{-}2}$	$0.858_{\pm 2e\text{-}3}$	$0.431_{\pm 6e\text{-}3}$	$3.066_{\pm 3e\text{-}2}$	$0.468_{\pm 2e\text{-}3}$	$0.821_{\pm 2e\text{-}3}$	$5.632_{\pm 4e\text{-}2}$	$0.914_{\pm 1e\text{-}3}$	$1.878_{\pm 4e\text{-}3}$	$0.966_{\pm 1e\text{-}3}$	$0.748_{\pm 3e\text{-}4}$

Table 15: Results for single models with MLP mask + target pretraining

H EARLY-STOPPING CRITERIONS

Here we demonstrate that early stopping the pretraining by the value of the pretraining objective on the hold-out validation set is comparable to the early stopping by the downstream metric on the hold-out validation set after finetuning. See Table 16 for the results. This is an important practical observation, as computing pretraining objective is much faster than the full finetuning of the model, especially on large scale datasets.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑
finetune early stop	0.683	0.857	0.434	3.056	0.468	0.819	5.633	0.914
pretrain early stop	0.674	0.855	0.434	3.031	0.469	0.818	5.738	0.914

Table 16: Results for single models with MLP mask + target pretraining

I EXTENDED TABLES WITH EXPERIMENTAL RESULTS

The scores with standard deviations for single models and ensembles are provided in 17 and 18 respectively.

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
CatBoost	$0.683_{\pm 4.7e-3}$	$0.864_{\pm 8.1e-4}$	$0.433_{\pm 1.7e-3}$	$3.115_{\pm 1.8e-2}$	$0.457_{\pm 1.3e-3}$	$0.806_{\pm 3.4e-4}$	$5.324_{\pm 4.0e-2}$	$0.927_{\pm 3.1e-4}$	$1.837_{\pm 2.1e-3}$	$0.966_{\pm 3.2e-4}$	$0.743_{\pm 3.0e-4}$
XGBoost	$0.678_{\pm 4.8e-3}$	$0.858_{\pm 2.3e-3}$	$0.436_{\pm 2.5e-3}$	$3.160_{\pm 6.9e-3}$	$0.457_{\pm 6.0e-3}$	$0.804_{\pm 1.5e-3}$	$5.383_{\pm 2.8e-2}$	$0.927_{\pm 7.0e-4}$	$1.802_{\pm 2.0e-3}$	$0.969_{\pm 6.1e-4}$	$0.742_{\pm 1.5e-4}$
MLP											
no pretraining	$0.635_{\pm 1.3e-2}$	$0.849_{\pm 1.6e-3}$	$0.506_{\pm 8.6e-3}$	$3.156_{\pm 2.1e-2}$	$0.479_{\pm 1.4e-3}$	$0.801_{\pm 9.5e-4}$	$5.737_{\pm 6.1e-2}$	$0.908_{\pm 1.0e-3}$	$1.909_{\pm 4.6e-3}$	$0.963_{\pm 8.7e-4}$	$0.749_{\pm 3.6e-4}$
mask	$0.691_{\pm 1.0e-2}$	$0.857_{\pm 2.5e-3}$	$0.454_{\pm 5.0e-3}$	$3.113_{\pm 4.3e-2}$	$0.472_{\pm 3.0e-3}$	$0.814_{\pm 1.7e-3}$	$5.681_{\pm 3.1e-2}$	$0.912_{\pm 8.0e-4}$	$1.883_{\pm 2.9e-3}$	$0.964_{\pm 9.3e-4}$	$0.748_{\pm 3.1e-4}$
rec	$0.662_{\pm 1.0e-2}$	$0.853_{\pm 2.2e-3}$	$0.445_{\pm 4.0e-3}$	$3.044_{\pm 3.3e-2}$	$0.466_{\pm 2.1e-3}$	$0.805_{\pm 1.3e-3}$	$5.641_{\pm 3.2e-2}$	$0.910_{\pm 1.2e-3}$	$1.875_{\pm 3.4e-3}$	$0.965_{\pm 5.4e-4}$	$0.746_{\pm 2.3e-4}$
contrastive	$0.672_{\pm 1.4e-2}$	$0.855_{\pm 2.0e-3}$	$0.455_{\pm 4.5e-3}$	$3.056_{\pm 5.2e-2}$	$0.469_{\pm 2.6e-3}$	$0.813_{\pm 1.3e-3}$	$5.697_{\pm 2.9e-2}$	$0.910_{\pm 1.3e-3}$	$1.881_{\pm 4.5e-3}$	$0.960_{\pm 1.2e-3}$	$0.748_{\pm 3.6e-4}$
sup	$0.693_{\pm 1.1e-2}$	$0.856_{\pm 1.8e-3}$	$0.441_{\pm 5.3e-3}$	$3.077_{\pm 2.9e-2}$	$0.459_{\pm 2.1e-3}$	$0.814_{\pm 7.4e-4}$	$5.689_{\pm 2.1e-2}$	$0.914_{\pm 7.8e-4}$	$1.883_{\pm 4.6e-3}$	$0.968_{\pm 5.2e-4}$	$0.748_{\pm 3.0e-4}$
supcon	$0.666_{\pm 1.4e-2}$	$0.850_{\pm 2.2e-3}$	$0.454_{\pm 4.2e-3}$	$3.108_{\pm 2.5e-2}$	$0.480_{\pm 1.9e-3}$	$0.806_{\pm 7.3e-4}$	$5.680_{\pm 2.4e-2}$	$0.911_{\pm 6.0e-4}$	$1.873_{\pm 2.4e-3}$	$0.966_{\pm 5.8e-4}$	$0.747_{\pm 3.2e-4}$
mask + sup	$0.693_{\pm 8.2e-3}$	$0.857_{\pm 2.3e-3}$	$0.436_{\pm 6.9e-3}$	$3.099_{\pm 2.4e-2}$	$0.458_{\pm 1.7e-3}$	$0.817_{\pm 5.6e-4}$	$5.685_{\pm 3.6e-2}$	$0.915_{\pm 5.3e-4}$	$1.873_{\pm 5.1e-3}$	$0.967_{\pm 4.0e-4}$	$0.748_{\pm 2.8e-4}$
rec + sup	$0.684_{\pm 7.7e-3}$	$0.854_{\pm 4.5e-3}$	$0.436_{\pm 4.4e-3}$	$3.012_{\pm 4.0e-2}$	$0.456_{\pm 1.9e-3}$	$0.815_{\pm 5.8e-4}$	$5.672_{\pm 3.6e-2}$	$0.911_{\pm 1.4e-3}$	$1.862_{\pm 2.8e-3}$	$0.967_{\pm 6.6e-4}$	$0.747_{\pm 4.9e-4}$
mask + target	$0.683_{\pm 1.0e-2}$	$0.857_{\pm 2.1e-3}$	$0.434_{\pm 7.2e-3}$	$3.056_{\pm 4.0e-2}$	$0.468_{\pm 1.9e-3}$	$0.819_{\pm 1.6e-3}$	$5.633_{\pm 3.7e-2}$	$0.914_{\pm 1.1e-3}$	$1.876_{\pm 4.8e-3}$	$0.965_{\pm 6.6e-4}$	$0.748_{\pm 4.5e-4}$
- target sampling	$0.680_{\pm 9.7e-3}$	$0.857_{\pm 3.1e-3}$	$0.432_{\pm 4.9e-3}$	$3.019_{\pm 3.5e-2}$	$0.468_{\pm 1.9e-3}$	$0.815_{\pm 1.7e-3}$	$5.697_{\pm 3.1e-2}$	$0.912_{\pm 6.7e-4}$	$1.887_{\pm 3.1e-3}$	$0.964_{\pm 1.1e-3}$	$0.748_{\pm 3.1e-4}$
rec + target	$0.659_{\pm 8.6e-3}$	$0.853_{\pm 3.2e-3}$	$0.454_{\pm 6.7e-3}$	$3.044_{\pm 4.9e-2}$	$0.463_{\pm 1.6e-3}$	$0.806_{\pm 1.5e-3}$	$5.636_{\pm 3.1e-2}$	$0.909_{\pm 9.0e-4}$	$1.884_{\pm 2.3e-3}$	$0.965_{\pm 8.7e-4}$	$0.745_{\pm 3.9e-4}$
- target sampling	$0.641_{\pm 5.6e-3}$	$0.853_{\pm 3.4e-3}$	$0.455_{\pm 4.6e-3}$	$3.046_{\pm 2.4e-2}$	$0.463_{\pm 2.0e-3}$	$0.806_{\pm 1.3e-3}$	$5.640_{\pm 1.9e-2}$	$0.910_{\pm 1.1e-3}$	$1.877_{\pm 3.3e-3}$	$0.966_{\pm 4.4e-4}$	$0.746_{\pm 3.9e-4}$
MLP-PLR											
no pretraining	$0.668_{\pm 1.4e-2}$	$0.858_{\pm 4.7e-3}$	$0.469_{\pm 5.2e-3}$	$3.008_{\pm 2.3e-2}$	$0.483_{\pm 1.6e-3}$	$0.809_{\pm 2.3e-3}$	$5.608_{\pm 5.6e-2}$	$0.926_{\pm 6.1e-4}$	$1.890_{\pm 5.0e-3}$	$0.969_{\pm 1.0e-3}$	$0.746_{\pm 3.7e-4}$
mask	$0.685_{\pm 5.6e-3}$	$0.863_{\pm 1.8e-3}$	$0.434_{\pm 4.3e-3}$	$3.007_{\pm 4.6e-2}$	$0.477_{\pm 2.5e-3}$	$0.818_{\pm 8.4e-4}$	$5.586_{\pm 2.4e-2}$	$0.927_{\pm 5.1e-4}$	$1.911_{\pm 5.4e-3}$	$0.970_{\pm 5.1e-4}$	$0.748_{\pm 3.9e-4}$
rec	$0.667_{\pm 7.2e-3}$	$0.852_{\pm 7.2e-3}$	$0.439_{\pm 5.2e-3}$	$3.031_{\pm 3.8e-2}$	$0.472_{\pm 3.0e-3}$	$0.808_{\pm 1.2e-3}$	$5.571_{\pm 1.2e-1}$	$0.926_{\pm 6.4e-4}$	$1.877_{\pm 4.0e-3}$	$0.971_{\pm 4.8e-4}$	$0.745_{\pm 3.6e-4}$
sup	$0.710_{\pm 4.6e-3}$	$0.859_{\pm 4.1e-3}$	$0.433_{\pm 3.6e-3}$	$3.136_{\pm 8.1e-2}$	$0.479_{\pm 1.9e-3}$	$0.811_{\pm 1.0e-3}$	$5.521_{\pm 4.6e-2}$	$0.924_{\pm 1.5e-3}$	$1.873_{\pm 2.0e-3}$	$0.971_{\pm 4.5e-4}$	$0.748_{\pm 8.0e-4}$
mask + sup	$0.711_{\pm 7.1e-3}$	$0.866_{\pm 2.0e-3}$	$0.441_{\pm 4.9e-3}$	$3.129_{\pm 4.1e-2}$	$0.480_{\pm 1.8e-3}$	$0.813_{\pm 8.2e-4}$	$5.480_{\pm 4.6e-2}$	$0.925_{\pm 1.0e-3}$	$1.875_{\pm 2.2e-3}$	$0.969_{\pm 6.0e-4}$	$0.745_{\pm 2.4e-4}$
rec + sup	$0.709_{\pm 5.1e-3}$	$0.858_{\pm 1.9e-3}$	$0.433_{\pm 2.7e-3}$	$3.059_{\pm 3.6e-2}$	$0.465_{\pm 2.2e-3}$	$0.807_{\pm 6.2e-4}$	$5.571_{\pm 1.2e-1}$	$0.927_{\pm 5.8e-4}$	$1.865_{\pm 3.1e-3}$	$0.971_{\pm 4.4e-4}$	$0.745_{\pm 2.4e-4}$
mask + target	$0.694_{\pm 9.1e-3}$	$0.862_{\pm 1.7e-3}$	$0.425_{\pm 4.2e-3}$	$3.023_{\pm 4.3e-2}$	$0.474_{\pm 2.0e-3}$	$0.821_{\pm 1.1e-3}$	$5.537_{\pm 3.4e-2}$	$0.929_{\pm 3.3e-4}$	$1.911_{\pm 6.2e-3}$	$0.969_{\pm 5.8e-4}$	$0.749_{\pm 1.2e-3}$
- target sampling	$0.690_{\pm 9.6e-3}$	$0.864_{\pm 2.8e-3}$	$0.421_{\pm 4.5e-3}$	$2.971_{\pm 4.1e-2}$	$0.479_{\pm 2.0e-3}$	$0.821_{\pm 1.0e-3}$	$5.440_{\pm 8.1e-2}$	$0.928_{\pm 5.6e-4}$	$1.906_{\pm 5.5e-3}$	$0.970_{\pm 3.9e-4}$	$0.748_{\pm 5.3e-4}$
rec + target	$0.688_{\pm 8.2e-3}$	$0.860_{\pm 1.4e-3}$	$0.445_{\pm 2.7e-3}$	$3.064_{\pm 3.4e-2}$	$0.475_{\pm 1.9e-3}$	$0.812_{\pm 1.1e-3}$	$5.507_{\pm 1.0e-1}$	$0.927_{\pm 3.9e-4}$	$1.887_{\pm 2.4e-3}$	$0.971_{\pm 3.9e-4}$	$0.748_{\pm 7.2e-4}$
- target sampling	$0.687_{\pm 7.9e-3}$	$0.855_{\pm 4.8e-3}$	$0.453_{\pm 6.4e-3}$	$3.008_{\pm 2.7e-2}$	$0.471_{\pm 2.1e-3}$	$0.812_{\pm 5.5e-4}$	$5.592_{\pm 9.8e-2}$	$0.927_{\pm 3.4e-4}$	$1.876_{\pm 3.1e-3}$	$0.970_{\pm 4.4e-4}$	$0.745_{\pm 3.5e-4}$
MLP-T-LR											
no pretraining	$0.634_{\pm 6.9e-3}$	$0.866_{\pm 1.3e-3}$	$0.444_{\pm 1.8e-3}$	$3.113_{\pm 4.5e-2}$	$0.482_{\pm 1.7e-3}$	$0.805_{\pm 9.3e-4}$	$5.520_{\pm 3.6e-2}$	$0.925_{\pm 6.8e-4}$	$1.897_{\pm 4.5e-3}$	$0.968_{\pm 5.0e-4}$	$0.749_{\pm 5.2e-4}$
mask	$0.654_{\pm 6.4e-3}$	$0.868_{\pm 1.0e-3}$	$0.424_{\pm 2.4e-3}$	$3.045_{\pm 3.7e-2}$	$0.472_{\pm 2.5e-3}$	$0.818_{\pm 1.8e-3}$	$5.544_{\pm 3.5e-2}$	$0.926_{\pm 7.1e-4}$	$1.916_{\pm 3.1e-3}$	$0.969_{\pm 4.6e-4}$	$0.748_{\pm 3.5e-4}$
rec	$0.652_{\pm 7.4e-3}$	$0.857_{\pm 4.4e-3}$	$0.424_{\pm 3.1e-3}$	$3.109_{\pm 3.7e-2}$	$0.472_{\pm 2.0e-3}$	$0.808_{\pm 1.0e-3}$	$5.363_{\pm 6.6e-2}$	$0.924_{\pm 1.8e-4}$	$1.861_{\pm 3.9e-3}$	$0.969_{\pm 7.3e-4}$	$0.746_{\pm 4.8e-4}$
sup	$0.682_{\pm 5.1e-3}$	$0.860_{\pm 4.1e-3}$	$0.430_{\pm 1.9e-3}$	$3.135_{\pm 1.7e-2}$	$0.471_{\pm 1.2e-3}$	$0.807_{\pm 6.2e-4}$	$5.525_{\pm 2.3e-2}$	$0.927_{\pm 3.8e-4}$	$1.893_{\pm 2.4e-3}$	$0.971_{\pm 5.8e-4}$	$0.747_{\pm 3.1e-4}$
mask + sup	$0.676_{\pm 7.7e-3}$	$0.858_{\pm 7.9e-3}$	$0.429_{\pm 1.8e-3}$	$3.199_{\pm 2.4e-2}$	$0.468_{\pm 9.9e-4}$	$0.814_{\pm 8.2e-4}$	$5.510_{\pm 3.6e-2}$	$0.926_{\pm 8.7e-4}$	$1.869_{\pm 2.9e-3}$	$0.971_{\pm 3.9e-4}$	$0.748_{\pm 2.5e-4}$
rec + sup	$0.678_{\pm 6.7e-3}$	$0.865_{\pm 1.1e-3}$	$0.437_{\pm 2.0e-3}$	$3.112_{\pm 4.1e-2}$	$0.462_{\pm 2.2e-3}$	$0.807_{\pm 2.6e-3}$	$5.516_{\pm 2.2e-2}$	$0.927_{\pm 3.7e-4}$	$1.862_{\pm 3.4e-3}$	$0.970_{\pm 5.0e-4}$	$0.748_{\pm 4.2e-4}$
mask + target	$0.649_{\pm 7.3e-3}$	$0.865_{\pm 1.7e-3}$	$0.421_{\pm 3.9e-3}$	$3.058_{\pm 4.3e-2}$	$0.474_{\pm 1.9e-3}$	$0.820_{\pm 1.2e-3}$	$5.644_{\pm 4.6e-2}$	$0.929_{\pm 3.5e-4}$	$1.924_{\pm 7.6e-3}$	$0.969_{\pm 4.8e-4}$	$0.749_{\pm 5.1e-4}$
- target sampling	$0.649_{\pm 8.3e-3}$	$0.861_{\pm 3.2e-3}$	$0.417_{\pm 4.9e-3}$	$3.050_{\pm 3.3e-2}$	$0.476_{\pm 1.7e-3}$	$0.819_{\pm 1.6e-3}$	$5.492_{\pm 3.2e-2}$	$0.928_{\pm 5.6e-4}$	$1.874_{\pm 3.7e-3}$	$0.969_{\pm 4.0e-4}$	$0.749_{\pm 1.4e-3}$
rec + target	$0.668_{\pm 7.0e-3}$	$0.864_{\pm 1.7e-3}$	$0.440_{\pm 3.5e-3}$	$3.113_{\pm 3.5e-2}$	$0.473_{\pm 4.3e-3}$	$0.806_{\pm 1.0e-3}$	$5.493_{\pm 4.4e-2}$	$0.927_{\pm 5.6e-4}$	$1.862_{\pm 3.5e-3}$	$0.969_{\pm 6.5e-4}$	$0.746_{\pm 3.4e-4}$
- target sampling	$0.667_{\pm 1.0e-2}$	$0.859_{\pm 2.7e-3}$	$0.432_{\pm 3.3e-3}$	$3.104_{\pm 2.9e-2}$	$0.470_{\pm 2.2e-3}$	$0.806_{\pm 1.5e-3}$	$5.391_{\pm 3.9e-2}$	$0.927_{\pm 4.1e-4}$	$1.867_{\pm 3.8e-3}$	$0.968_{\pm 7.5e-4}$	$0.746_{\pm 3.1e-4}$

Table 17: Extended results for single models

	GE ↑	CH ↑	CA ↓	HO ↓	OT ↓	HI ↑	FB ↓	AD ↑	WE ↓	CO ↑	MI ↓
CatBoost	$0.692_{\pm 1.8e-3}$	$0.864_{\pm 6.8e-5}$	$0.430_{\pm 1.1e-3}$	$3.093_{\pm 5.1e-3}$	$0.450_{\pm 3.5e-4}$	$0.807_{\pm 7.5e-5}$	$5.226_{\pm 1.2e-2}$	$0.928_{\pm 1.3e-4}$	$1.801_{\pm 1.2e-3}$	$0.967_{\pm 1.3e-4}$	$0.741_{\pm 1.4e-4}$
XGBoost	$0.683_{\pm 1.3e-3}$	$0.860_{\pm 4.3e-4}$	$0.434_{\pm 7.0e-4}$	$3.152_{\pm 1.2e-3}$	$0.454_{\pm 2.5e-3}$	$0.805_{\pm 8.3e-4}$	$5.338_{\pm 1.8e-2}$	$0.927_{\pm 2.1e-4}$	$1.782_{\pm 4.9e-4}$	$0.969_{\pm 8.8e-5}$	$0.742_{\pm 5.3e-5}$
MLP											
no pretraining	$0.656_{\pm 5.9e-3}$	$0.852_{\pm 5.2e-4}$	$0.482_{\pm 2.9e-3}$	$3.055_{\pm 8.4e-3}$	$0.467_{\pm 2.0e-3}$	$0.805_{\pm 2.9e-4}$	$5.666_{\pm 2.6e-3}$	$0.910_{\pm 2.7e-4}$	$1.850_{\pm 1.0e-3}$	$0.968_{\pm 2.5e-4}$	$0.747_{\pm 8.6e-5}$
mask	$0.722_{\pm 1.6e-3}$	$0.859_{\pm 6.4e-4}$	$0.437_{\pm 8.1e-4}$	$3.026_{\pm 6.3e-3}$	$0.451_{\pm 1.6e-3}$	$0.824_{\pm 6.5e-4}$	$5.578_{\pm 6.8e-3}$	$0.913_{\pm 3.0e-4}$	$1.828_{\pm 1.0e-3}$	$0.967_{\pm 1.1e-4}$	$0.746_{\pm 7.7e-5}$
rec	$0.679_{\pm 2.0e-3}$	$0.856_{\pm 2.8e-4}$	$0.424_{\pm 1.0e-4}$	$2.967_{\pm 6.9e-3}$	$0.453_{\pm 1.2e-3}$	$0.812_{\pm 2.5e-4}$	$5.571_{\pm 1.4e-2}$	$0.912_{\pm 7.5e-5}$	$1.811_{\pm 1.4e-3}$	$0.972_{\pm 2.0e-4}$	$0.744_{\pm 7.8e-5}$
contrastive	$0.708_{\pm 4.4e-3}$	$0.857_{\pm 8.8e-4}$	$0.434_{\pm 4.0e-3}$	$2.952_{\pm 4.4e-3}$	$0.451_{\pm 6.1e-4}$	$0.820_{\pm 4.5e-5}$	$5.634_{\pm 1.4e-2}$	$0.912_{\pm 1.3e-4}$	$1.804_{\pm 2.8e-3}$	$0.964_{\pm 2.1e-4}$	$0.746_{\pm 1.7e-4}$
sup	$0.717_{\pm 2.2e-3}$	$0.857_{\pm 7.4e-4}$	$0.424_{\pm 9.4e-4}$	$3.022_{\pm 1.9e-2}$	$0.443_{\pm 1.9e-3}$	$0.816_{\pm 1.4e-4}$	$5.602_{\pm 6.5e-3}$	$0.916_{\pm 7.0e-5}$	$1.828_{\pm 4.7e-3}$	$0.973_{\pm 3.6e-4}$	$0.746_{\pm 2.1e-4}$
supcon	$0.686_{\pm 5.2e-3}$	$0.851_{\pm 8.4e-4}$	$0.434_{\pm 3.0e-3}$	$3.014_{\pm 6.0e-3}$	$0.465_{\pm 1.3e-3}$	$0.809_{\pm 1.3e-4}$	$5.579_{\pm 3.4e-3}$	$0.912_{\pm 3.2e-4}$	$1.827_{\pm 9.4e-4}$	$0.970_{\pm 2.1e-4}$	$0.745_{\pm 5.3e-5}$
mask + sup	$0.716_{\pm 5.7e-3}$	$0.859_{\pm 1.1e-3}$	$0.418_{\pm 2.2e-3}$	$3.066_{\pm 1.2e-2}$	$0.443_{\pm 1.5e-3}$	$0.819_{\pm 1.3e-4}$	$5.601_{\pm 4.7e-3}$	$0.916_{\pm 2.0e-4}$	$1.810_{\pm 2.9e-4}$	$0.973_{\pm 3.9e-5}$	$0.747_{\pm 1.3e-4}$
rec + sup	$0.709_{\pm 3.7e-3}$	$0.859_{\pm 1.8e-3}$	$0.419_{\pm 2.1e-3}$	$2.951_{\pm 1.9e-2}$	$0.442_{\pm 8.6e-4}$	$0.817_{\pm 1.0e-4}$	$5.531_{\pm 3.0e-3}$	$0.913_{\pm 5.2e-4}$	$1.801_{\pm 4.0e-3}$	$0.973_{\pm 2.6e-5}$	$0.745_{\pm 1.1e-4}$
mask + target	$0.709_{\pm 7.3e-3}$	$0.860_{\pm 1.6e-3}$	$0.414_{\pm 1.1e-3}$	$2.949_{\pm 1.9e-2}$	$0.457_{\pm 5.9e-4}$	$0.828_{\pm 6.3e-4}$	$5.551_{\pm 7.2e-3}$	$0.916_{\pm 4.6e-4}$	$1.809_{\pm 5.3e-4}$	$0.969_{\pm 5.2e-5}$	$0.746_{\pm 1.9e-4}$
- target sampling	$0.706_{\pm 4.9e-3}$	$0.860_{\pm 1.1e-3}$	$0.410_{\pm 1.5e-3}$	$2.955_{\pm 2.1e-2}$	$0.456_{\pm 3.8e-4}$	$0.822_{\pm 1.1e-3}$	$5.601_{\pm 2.0e-2}$	$0.914_{\pm 2.5e-4}$	$1.837_{\pm 1.0e-3}$	$0.968_{\pm 2.8e-4}$	$0.746_{\pm 2.1e-4}$
rec + target	$0.677_{\pm 4.6e-3}$	$0.857_{\pm 3.9e-4}$	$0.433_{\pm 1.1e-3}$	$2.926_{\pm 2.3e-2}$	$0.448_{\pm 9.4e-4}$	$0.816_{\pm 5.3e-4}$	$5.555_{\pm 5.8e-3}$	$0.910_{\pm 2.2e-4}$	$1.825_{\pm 2.5e-3}$	$0.972_{\pm 6.1e-5}$	$0.743_{\pm 1.2e-4}$
- target sampling	$0.669_{\pm 3.7e-3}$	$0.858_{\pm 1.4e-3}$	$0.435_{\pm 8.2e-4}$	$3.003_{\pm 1.1e-2}$	$0.451_{\pm 1.2e-3}$	$0.815_{\pm 7.0e-4}$	$5.577_{\pm 1.1e-2}$	$0.913_{\pm 2.4e-4}$	$1.822_{\pm 6.6e-4}$	$0.972_{\pm 2.7e-4}$	$0.744_{\pm 7.6e-5}$
MLP-PLR											
no pretraining	$0.695_{\pm 3.7e-3}$	$0.864_{\pm 7.6e-4}$	$0.454_{\pm 1.2e-3}$	$2.953_{\pm 7.4e-3}$	$0.470_{\pm 7.5e-4}$	$0.814_{\pm 7.9e-4}$	$5.324_{\pm 3.2e-2}$	$0.928_{\pm 7.7e-5}$	$1.835_{\pm 1.5e-3}$	$0.974_{\pm 2.2e-4}$	$0.744_{\pm 1.2e-4}$
mask	$0.725_{\pm 4.9e-3}$	$0.865_{\pm 7.0e-4}$	$0.421_{\pm 1.7e-3}$	$2.921_{\pm 1.0e-2}$	$0.457_{\pm 8.4e-4}$	$0.827_{\pm 1.1e-4}$	$5.444_{\pm 5.4e-3}$	$0.928_{\pm 1.0e-4}$	$1.850_{\pm 3.4e-3}$	$0.974_{\pm 2.3e-4}$	$0.745_{\pm 6.5e-5}$
rec	$0.698_{\pm 1.5e-3}$	$0.857_{\pm 1.5e-3}$	$0.418_{\pm 1.2e-3}$	$2.954_{\pm 8.1e-3}$	$0.454_{\pm 1.9e-3}$	$0.813_{\pm 7.8e-4}$	$5.124_{\pm 2.4e-2}$	$0.928_{\pm 2.2e-4}$	$1.818_{\pm 2.5e-3}$	$0.975_{\pm 2.9e-4}$	$0.743_{\pm 2.1e-4}$
sup	$0.733_{\pm 2.2e-3}$	$0.867_{\pm 9.6e-4}$	$0.421_{\pm 1.0e-3}$	$3.054_{\pm 4.5e-2}$	$0.465_{\pm 1.3e-3}$	$0.816_{\pm 1.4e-4}$	$5.407_{\pm 3.8e-2}$	$0.926_{\pm 4.3e-4}$	$1.834_{\pm 3.3e-4}$	$0.975_{\pm 1.6e-4}$	$0.746_{\pm 2.1e-4}$
mask + sup	$0.732_{\pm 2.0e-3}$	$0.869_{\pm 3.5e-4}$	$0.424_{\pm 8.9e-4}$	$3.055_{\pm 1.6e-2}$	$0.468_{\pm 7.8e-4}$	$0.817_{\pm 3.4e-4}$	$5.366_{\pm 1.1e-2}$	$0.927_{\pm 2.1e-4}$	$1.848_{\pm 1.7e-3}$	$0.974_{\pm 2.1e-4}$	$0.744_{\pm 7.9e-5}$
rec + sup	$0.737_{\pm 2.0e-3}$	$0.862_{\pm 1.3e-3}$	$0.424_{\pm 9.9e-4}$	$2.964_{\pm 2.3e-3}$	$0.449_{\pm 2.3e-4}$	$0.811_{\pm 5.3e-4}$	$5.124_{\pm 2.4e-2}$	$0.929_{\pm 1.7e-4}$	$1.813_{\pm 1.9e-3}$	$0.974_{\pm 2.2e-4}$	$0.744_{\pm 7.2e-5}$
mask + target	$0.719_{\pm 3.5e-3}$	$0.866_{\pm 4.3e-4}$	$0.407_{\pm 8.2e-4}$	$2.952_{\pm 3.5e-3}$	$0.458_{\pm 4.2e-4}$	$0.828_{\pm 6.0e-4}$	$5.373_{\pm 1.8e-2}$	$0.930_{\pm 1.0e-4}$	$1.849_{\pm 2.1e-3}$	$0.973_{\pm 1.4e-4}$	$0.745_{\pm 2.5e-4}$
- target sampling	$0.724_{\pm 7.0e-3}$	$0.867_{\pm 1.0e-3}$	$0.403_{\pm 1.1e-3}$	$2.877_{\pm 1.0e-2}$	$0.466_{\pm 9.4e-4}$	$0.828_{\pm 2.5e-4}$	$5.175_{\pm 1.5e-2}$	$0.930_{\pm 2.0e-4}$	$1.833_{\pm 3.6e-3}$	$0.974_{\pm 3.3e-4}$	$0.744_{\pm 2.0e-4}$
rec + target	$0.705_{\pm 2.3e-3}$	$0.862_{\pm 2.0e-4}$	$0.431_{\pm 1.4e-3}$	$2.983_{\pm 1.2e-2}$	$0.465_{\pm 9.2e-4}$	$0.816_{\pm 1.7e-4}$	$5.096_{\pm 1.8e-2}$	$0.928_{\pm 1.9e-4}$	$1.860_{\pm 2.0e-3}$	$0.974_{\pm 3.4e-5}$	$0.745_{\pm 2.7e-4}$
- target sampling	$0.712_{\pm 2.3e-3}$	$0.860_{\pm 9.3e-4}$	$0.437_{\pm 3.3e-3}$	$2.933_{\pm 2.0e-2}$	$0.450_{\pm 1.7e-3}$	$0.815_{\pm 3.5e-4}$	$5.173_{\pm 2.7e-2}$	$0.928_{\pm 5.6e-5}$	$1.811_{\pm 1.5e-3}$	$0.974_{\pm 3.6e-4}$	$0.744_{\pm 6.3e-5}$
MLP-T-LR											
no pretraining	$0.662_{\pm 7.6e-3}$	$0.868_{\pm 5.0e-4}$	$0.437_{\pm 8.2e-4}$	$3.028_{\pm 1.8e-2}$	$0.472_{\pm 4.7e-4}$	$0.808_{\pm 1.5e-4}$	$5.424_{\pm 2.2e-2}$	$0.927_{\pm 2.8e-4}$	$1.850_{\pm 9.0e-4}$	$0.972_{\pm 1.5e-4}$	$0.747_{\pm 7.6e-5}$
mask	$0.679_{\pm 4.7e-3}$	$0.868_{\pm 2.3e-4}$	$0.413_{\pm 1.0e-3}$	$2.930_{\pm 1.2e-2}$	$0.450_{\pm 7.3e-4}$	$0.826_{\pm 1.3e-3}$	$5.370_{\pm 8.7e-3}$	$0.927_{\pm 3.1e-4}$	$1.836_{\pm 2.7e-3}$	$0.973_{\pm 8.8e-5}$	$0.745_{\pm 1.1e-4}$
rec	$0.694_{\pm 3.7e-3}$	$0.861_{\pm 1.6e-4}$	$0.414_{\pm 1.5e-3}$	$3.035_{\pm 1.9e-2}$	$0.459_{\pm 3.4e-4}$	$0.812_{\pm 2.7e-4}$	$5.039_{\pm 1.8e-2}$	$0.925_{\pm 1.2e-4}$	$1.803_{\pm 2.3e-3}$	$0.973_{\pm 4.9e-5}$	$0.744_{\pm 4.5e-4}$
sup	$0.698_{\pm 3.7e-3}$	$0.865_{\pm 7.0e-4}$	$0.424_{\pm 1.1e-3}$	$3.107_{\pm 6.8e-3}$	$0.463_{\pm 3.5e-4}$	$0.809_{\pm 2.5e-4}$	$5.442_{\pm 1.6e-2}$	$0.928_{\pm 1.4e-4}$	$1.849_{\pm 5.4e-4}$	$0.975_{\pm 2.8e-4}$	$0.746_{\pm 1.4e-5}$
mask + sup	$0.698_{\pm 4.0e-3}$	$0.866_{\pm 1.1e-3}$	$0.421_{\pm 8.6e-4}$	$3.088_{\pm 5.1e-3}$	$0.460_{\pm 4.6e-4}$	$0.818_{\pm 2.8e-4}$	$5.407_{\pm 3.3e-3}$	$0.927_{\pm 5.1e-4}$	$1.824_{\pm 3.0e-3}$	$0.975_{\pm 1.4e-4}$	$0.747_{\pm 8.2e-5}$
rec + sup	$0.705_{\pm 2.4e-3}$	$0.866_{\pm 4.6e-4}$	$0.425_{\pm 5.1e-4}$	$3.057_{\pm 1.0e-2}$	$0.444_{\pm 1.4e-3}$	$0.814_{\pm 6.8e-4}$	$5.422_{\pm 1.8e-3}$	$0.927_{\pm 6.6e-5}$	$1.811_{\pm 1.0e-3}$	$0.974_{\pm 1.1e-4}$	$0.746_{\pm 3.6e-4}$
mask + target	$0.673_{\pm 1.0e-3}$	$0.868_{\pm 4.6e-4}$	$0.410_{\pm 7.8e-4}$	$2.894_{\pm 1.8e-2}$	$0.460_{\pm 1.4e-3}$	$0.827_{\pm 3.4e-4}$	$5.458_{\pm 2.8e-2}$	$0.930_{\pm 1.7e-5}$	$1.849_{\pm 4.2e-3}$	$0.972_{\pm 2.5e-4}$	$0.746_{\pm 2.3e-4}$
- target sampling	$0.677_{\pm 5.1e-3}$	$0.866_{\pm 1.1e-3}$	$0.397_{\pm 3.7e-4}$	$2.938_{\pm 1.8e-2}$	$0.462_{\pm 4.4e-4}$	$0.826_{\pm 1.2e-4}$	$5.384_{\pm 1.5e-2}$	$0.929_{\pm 1.6e-4}$	$1.840_{\pm 2.2e-3}$	$0.973_{\pm 1.1e-4}$	$0.747_{\pm 5.0e-4}$
rec + target	$0.693_{\pm 2.5e-3}$	$0.866_{\pm 3.1e-4}$	$0.432_{\pm 6.1e-4}$	$3.045_{\pm 1.1e-2}$	$0.456_{\pm 2.4e-3}$	$0.812_{\pm 6.1e-4}$	$5.344_{\pm 8.1e-3}$	$0.928_{\pm 1.4e-4}$	$1.830_{\pm 2.4e-3}$	$0.972_{\pm 1.1e-4}$	$0.744_{\pm 2.1e-4}$
- target sampling	$0.700_{\pm 2.5e-3}$	$0.863_{\pm 1.3e-3}$	$0.423_{\pm 9.8e-4}$	$3.029_{\pm 1.4e-2}$	$0.454_{\pm 1.8e-3}$	$0.811_{\pm 5.1e-4}$	$5.083_{\pm 4.7e-3}$	$0.928_{\pm 1.7e-4}$	$1.809_{\pm 1.0e-3}$	$0.972_{\pm 5.5e-5}$	$0.744_{\pm 3.0e-5}$

Table 18: Extended results for ensemble models