Revisiting Deep Learning Models for Tabular Data

Anonymous Author(s)
Affiliation
Address
email

Abstract

The necessity of deep learning for tabular, structured data is still an unanswered question addressed by a large number of research efforts. The recent literature on tabular DL proposes several deep architectures reported to be superior to traditional "shallow" models like GBDT. However, since existing works often use different benchmarks and tuning protocols, it is unclear if the proposed models universally outperform GBDT. Moreover, the models are often not compared to each other, therefore, it is challenging to identify the best deep model for practitioners.

In this work, we start from a thorough review of the main families of DL models recently developed for tabular data. We carefully tune and evaluate them on a wide range of datasets and reveal two significant findings. First, we show that the choice between GBDT and DL models highly depends on data and there is still no universally superior solution. Second, we demonstrate that a simple ResNet-like architecture is a surprisingly effective baseline, which outperforms most of the sophisticated models from the DL literature. Finally, as an additional contribution, we introduce a new attention-based architecture that outperforms the simple ResNet baseline, leading to a new state-of-the-art DL solution for most of the benchmarks. The source code is available at https://github.com/AnonymouzAuthorz/RevisitingTabularDL.

1 Introduction

Due to the tremendous success of deep learning for unstructured data like raw images, audio and texts (Goodfellow et al., 2016), there has been a lot of research interest to extend this success to problems with structured data stored in tabular form. In these problems, data points are represented as vectors of heterogeneous features, which is typical for industrial applications and ML competitions, where non-deep models, e.g., GBDT (Chen and Guestrin, 2016; Ke et al., 2017; Prokhorenkova et al., 2018) are currently the top-choice solutions. Along with potentially higher performance, using deep learning for tabular data is appealing as it would allow to construct multi-modal pipelines for problems, where only one part of the input is structured, and other parts include images, audio and other DL-friendly data. Such pipelines can then be trained end-to-end by gradient optimization for both structured and unstructured modalities. For these reasons, a large number of DL solutions were recently proposed, and new models continue to emerge (Arik and Pfister, 2020; Badirli et al., 2020; Hazimeh et al., 2020; Huang et al., 2020; Klambauer et al., 2017; Popov et al., 2020; Song et al., 2019; Wang et al., 2017, 2020).

Unfortunately, due to the lack of established benchmarks (such as ImageNet (Deng et al., 2009) for computer vision or GLUE (Wang et al., 2019a) for NLP) and the fast development of the field, existing papers use different datasets for evaluation, and proposed DL models are often not adequately compared to each other. Therefore, from the current literature, it is unclear whether GBDT is surpassed by DL models or what DL model generally performs better than others. These obscurities impede the research process and make the observations from the papers not conclusive enough.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Given the increasing number of works on tabular DL, we believe it is timely to review the recent developments from the field to identify more justified conclusions that can serve as a basis for future studies. To this end, we perform a large-scale evaluation of many recent open-sourced DL models on a wide range of datasets with the same protocol of hyperparameter tuning. Our large-scale study has revealed several curious findings. First, we do not observe a definitive superiority of GBDT and DL models w.r.t. each other, and their relative performance largely depends on the particular task. For instance, GBDT is superior for data with heterogeneous features, but lags behind DL solutions on non-heterogeneous data and poorly scales to classification problems when the number of classes is large (1K+). Second, we investigate the relative performance of DL models to each other. Surprisingly, we discover that most recent models are inferior to a simple ResNet-like architecture. Given its high performance and simplicity, it can serve as a strong baseline for future works on tabular DL. Based on these observations, we design a new attention-based model that outperforms the ResNet baseline on tasks with heterogeneous features, providing a new state-of-the-art DL solution for most of the problems. Overall, we hope that our evaluation protocol and the described baselines will lay the foundation for more decisive studies on tabular DL. We encourage other researchers to add their datasets and models to our repository.

We summarize the contributions of our paper as follows:

1. We thoroughly evaluate the main models for tabular DL on a wide range of problems to investigate their relative performance and for the appropriate comparison to GBDT.
2. We reveal that GBDT and deep models are superior for different types of problems.
3. We demonstrate that a simple ResNet-like architecture is a very strong baseline for tabular DL, which was overlooked by existing literature. Given its simplicity, this baseline is necessary for comparison in future tabular DL works.
4. We design a new attention-based architecture that outperforms the state-of-the-art DL models on tasks with heterogeneous data.

2 Related work

The "shallow" state-of-the-art for problems with tabular data is currently ensembles of decision trees, such as GBDT (Gradient Boosting Decision Tree) (Friedman, 2001), which are typically the top-choice in various ML competitions. At the moment, there are several established GBDT libraries, such as XGBoost (Chen and Guestrin, 2016), LightGBM (Ke et al., 2017), CatBoost (Prokhorenkova et al., 2018), which are widely used by both ML researchers and practitioners. While these implementations vary in detail, on most of the tasks, their performances do not differ much (Prokhorenkova et al., 2018).

During several recent years, a large number of deep learning models for tabular data have been developed (Arik and Pfister, 2020; Badirli et al., 2020; Hazimeh et al., 2020; Huang et al., 2020; Klambauer et al., 2017; Popov et al., 2020; Song et al., 2019; Wang et al., 2017). Most of these models can be roughly categorized into three groups, which we briefly describe below.

Differentiable trees. The first group of models is motivated by the strong performance of decision tree ensembles for tabular data. Since decision trees are not differentiable and do not allow gradient optimization, they cannot be used as a component for pipelines trained in the end-to-end fashion. To address this issue, several works (Hazimeh et al., 2020; Kontschieder et al., 2015; Popov et al., 2020; Yang et al., 2018) propose to "smooth" decision functions in the internal tree nodes to make the overall tree function and tree routing differentiable. While the methods of this family can outperform GBDT on some tasks (Popov et al., 2020), in our experiments, they did not outperform ResNet.

Attention-based models. Due to the ubiquitous success of attention-based architectures for different domains (Dosovitskiy et al., 2021; Vaswani et al., 2017), several authors propose to employ attention-like modules for tabular DL as well (Arik and Pfister, 2020; Huang et al., 2020; Song et al., 2019). In our experiments, we show that the properly tuned ResNet outperforms the existing attention-based models. Nevertheless, we show that careful incorporation of attention modules can improve the performance and develop a new DL architecture that outperforms ResNet.

Explicit modeling of multiplicative interactions. In the literature on recommender systems and click-through-rate prediction, several works criticize MLP since it is unsuitable for modeling mul-

tiplicative interactions between features (Beutel et al., 2018; Qin et al., 2021; Wang et al., 2017). Inspired by this motivation, some works (Beutel et al., 2018; Wang et al., 2017, 2020) have proposed different ways to incorporate feature products into MLP. In our experiments, however, we did not find such methods to be superior to properly tuned baselines.

The literature also proposes some other architectural designs (Badirli et al., 2020; Klambauer et al., 2017) that cannot be explicitly assigned to any of the groups above. Overall, the community has developed a variety of models that are evaluated on different benchmarks and are rarely compared to each other. Our work aims to establish a fair comparison of them and identify the solutions that consistently provide high performance.

3 Comparing the models for tabular data

3.1 Preliminaries

Notation. In this work, we consider supervised learning problems. $D{=}\{(x_i,\ y_i)\}_{i=1}^{n}$ denotes a dataset, where $x_i{=}(x_i^{(num)},\ x_i^{(cat)}) \in \mathbb{X}$ represents numerical $x_{ij}^{(num)}$ and categorical $x_{ij}^{(cat)}$ features of an object and $y_i \in \mathbb{Y}$ denotes the corresponding object label. The total number of features is denoted as k. The dataset is split into three disjoint subsets: $D = D_{train} \cup D_{val} \cup D_{test}$, where D_{train} is used for training, D_{val} is used for early stopping and hyperparameter tuning and D_{test} is used for the final evaluation. We consider three types of tasks: binary classification $\mathbb{Y} = \{0,\ 1\}$, multiclass classification $\mathbb{Y} = \{1,\ \ldots,\ C\}$ and regression $\mathbb{Y} = \mathbb{R}$.

Scope of the comparison. In our work, we focus on the relative performance of different architectures and do not employ various model-agnostic DL practices, such as pretraining, data augmentation, distillation, learning rate warmup, learning rate decay and many others. While these practices can potentially improve the performance, our goal is to evaluate the impact of inductive biases imposed by the different model architectures.

3.2 Baselines

First, we include two established "shallow" models that are currently trusted by practitioners as well as the simplest DL models that should serve as a natural sanity check for testing novel deep architectures against them:

- **XGBoost** (Chen and Guestrin, 2016). One of the most popular GBDT implementations.
- **CatBoost** (Prokhorenkova et al., 2018). GBDT implementation that uses oblivious decision trees (Lou and Obukhov, 2017) as week learners.
- **MLP**. The simplest feed-forward model consisting of `Linear-ReLU-Dropout` blocks (details are provided in supplementary).
- Additionally, we employ **ResNet** — an MLP-like architecture with skip connections (He et al., 2015) and batch normalizations (Ioffe and Szegedy, 2015) (details are provided supplementary). Surprisingly, this simple baseline is typically overlooked in existing papers.

Second, we cover the following deep architectures that were specifically designed for tabular data:

- **SNN** (Klambauer et al., 2017). An MLP-like architecture with the SELU activation that enables training deeper models.
- **NODE** (Popov et al., 2020). A differentiable ensemble of oblivious decision trees.
- **TabNet** (Arik and Pfister, 2020). A recurrent architecture that alternates dynamical reweighing of features and conventional feed-forward modules.
- **GrowNet** (Badirli et al., 2020). Gradient boosted weak MLPs. The official implementation supports only classification and regression problems.
- **DCN V2** (Wang et al., 2020). Consists of an MLP-like module and the feature crossing module (a combination of linear layers and multiplications).
- **AutoInt** (Song et al., 2019). Transforms features to tokens and applies self-attention on the token level.

All technical details are provided in supplementary.

3.3 Datasets and metrics.

Datasets. We use a diverse set of eleven public datasets (see supplementary for the detailed description). For each dataset, there is exactly one train-val-test split, so all algorithms use the same splits. We informally call the dataset's features "heterogeneous" if they describe different physical properties of the object and are of diverse units of measurements, which is somewhat opposite to homogeneous features, such as pixels for images or words for text data. We apply this term to California Housing (CA, real estate data, Kelley Pace and Barry (1997)), Covertype (CO, forest characteristics, Blackard and Dean. (2000)), Yahoo (YA, search queries, Chapelle and Chang (2011)) and Microsoft (MI, search queries, Qin and Liu (2013)). Non-heterogeneous datasets include Higgs (HI, physical particles, Baldi et al. (2014)), ALOI (AL, images, Geusebroek et al. (2005)), Epsilon (EP, simulated physics experiments) and Year (YE, audio features, Bertin-Mahieux et al. (2011)). We do not assign anonymous datasets (Helena (HE) and Jannis (JA), Guyon et al. (2019)) and datasets with few numerical features (Adult (AD), Kohavi (1996)) to either heterogeneous or non-heterogeneous. We follow the pointwise approach to learning-to-rank and treat ranking problems (Microsoft, Yahoo) as regression problems.
Datasets scope. In our work, dataset sizes vary from 20K to 1M+, and the number of features is significantly less than the number of samples.
Metrics. For binary classification problems, we report Area Under the ROC Curve (AUC). For multiclass problems, we report accuracy. For regression problems, we report Root-Mean-Square Error.

3.4 Implementation details

Data preprocessing. Data preprocessing is known to be vital for DNN models. For each dataset, the same preprocessing was used for all models for fair comparison. By default, we used the quantile transformation from the Scikit-learn library (Pedregosa et al., 2011). We applied standardization (mean subtraction and scaling) to Helena and ALOI. The latter one represents image data and standardization is a normal practice in computer vision. We do not preprocess the Epsilon dataset, since we observed it to be highly detrimental to deep models' performance. We apply standardization (mean subtraction and scaling) to regression targets for all algorithms.

Tuning. For every dataset, we carefully tune each model's hyperparameters. The best hyperparameters are the ones that perform best on the validation set, so the test set is never used for tuning. For most algorithms, we use the Optuna library (Akiba et al., 2019) to run Bayesian optimization (the Tree-Structured Parzen Estimator algorithm), which is reported to be superior to random search (Turner et al., 2021). For the rest, we iterate over a predefined set of configurations recommended by corresponding papers. For details, see supplementary.

Evaluation. For each tuned configuration, we run 15 experiments with different random seeds and report the performance on the test set. For some algorithms, we also report the performance of default configurations without hyperparameter tuning.

Neural networks. We minimize cross-entropy for classification problems and mean squared error for regression problems. We use the AdamW optimizer (Loshchilov and Hutter, 2019) and do not apply learning rate schedules. For each dataset, we use a predefined batch size for all algorithms unless special instructions on batch sizes are given in the corresponding papers (see supplementary). We continue training until there are $\texttt{patience} + 1$ consecutive epochs without improvements on the validation set; we set $\texttt{patience} = 16$ for all algorithms.

See supplementary for information on hardware and training time.

3.5 The comparison results

Table 1 reports the comparison of the deep models w.r.t. each other and GBDT. Since GBDT is essentially an ensembling technique, we also compare all the competitors in the ensembling mode. We build three ensemble predictions by splitting the 15 single model predictions into three disjoint groups of size five each. As shown in Table 2, ensembling significantly improves the performance of deep architectures, which is consistent with existing literature on this phenomenon (Fort et al., 2020).

Table 1: Results for single models. See supplementary for confidence intervals. For each dataset, top results for baseline neural networks are in **bold**, top results for baseline neural networks and FT-Transformer are in **blue**, the overall top results are in **red**. "Top" means "the gap between this result and the result with the best mean score is not statistically significant". Datasets with heterogeneous features are underlined. FT-Transformer is introduced in section 3.6. Notation: "d" ~ "default", ↓ ~ lower is better, ↑ ~ higher is better.

	CA ↓	AD ↑	HE ↑	JA ↑	HI ↑	AL ↑	EP ↑	YE ↓	CO ↑	YA ↓	MI ↓
Baseline Neural Networks											
SNN	0.507	**0.816**	0.3728	0.718	0.721	0.954	0.8970	8.881	0.9465	0.769	0.7521
TabNet	0.513	0.796	0.3782	0.724	0.717	0.954	0.8902	9.032	0.9335	0.819	0.7565
GrowNet	0.500	0.793	–	–	0.724	–	**0.8977**	8.866	–	0.775	0.7549
DCN2	0.486	0.784	0.3853	0.714	0.720	0.955	**0.8975**	8.939	0.9491	0.766	0.7500
AutoInt	0.479	0.801	0.3722	0.716	**0.726**	0.945	0.8948	8.875	0.9312	0.795	0.7517
MLP	0.494	0.796	0.3832	0.719	0.721	0.954	0.8968	8.861	0.9499	0.776	0.7521
NODE	**0.464**	0.791	0.3593	**0.726**	0.724	0.918	0.8958	**8.774**	0.9436	**0.762**	**0.7474**
ResNet	0.487	**0.816**	**0.3960**	**0.727**	**0.727**	**0.963**	0.8971	8.845	**0.9560**	0.766	0.7493
FT-Transformer											
FT-Transformer$_d$	0.470	0.799	0.3812	0.725	0.723	0.953	0.8959	8.869	0.9617	**0.758**	**0.7475**
FT-Transformer	**0.464**	0.807	0.3913	**0.731**	**0.728**	0.960	**0.8982**	8.820	**0.9641**	**0.758**	**0.7469**
GBDT											
CatBoost$_d$	**0.430**	0.797	0.3814	0.721	0.724	0.946	0.8882	8.913	0.9076	0.751	0.7454
CatBoost	**0.431**	0.791	0.3853	0.723	0.725	–	0.8880	8.877	0.9658	0.743	0.7429
XGBoost$_d$	0.463	0.775	0.3502	0.721	0.705	0.925	0.8803	9.446	0.9640	0.773	0.7719
XGBoost	0.433	0.796	0.3755	0.724	0.725	–	0.8857	8.947	**0.9695**	**0.736**	**0.7424**

Table 2: Results for ensembles. Notation mostly follows Table 1, but top results are now defined by the mean score. Standard deviations are reported in supplementary.

	CA ↓	AD ↑	HE ↑	JA ↑	HI ↑	AL ↑	EP ↑	YE ↓	CO ↑	YA ↓	MI ↓
Baseline Neural Networks											
SNN	0.485	**0.822**	0.3804	0.722	0.725	0.962	0.8971	8.747	0.9542	0.762	0.7487
TabNet	0.492	0.809	0.3908	**0.734**	0.724	0.961	0.8952	8.773	0.9497	0.814	0.7505
GrowNet	0.483	0.794	–	–	0.731	–	**0.8986**	8.718	–	0.766	0.7504
DCN2	0.478	0.786	0.3884	0.721	0.721	0.960	0.8977	8.764	0.9575	0.762	0.7488
AutoInt	**0.461**	0.805	0.3822	0.728	**0.732**	0.959	0.8967	8.730	0.9519	0.782	0.7477
MLP	0.489	0.806	0.3902	0.722	0.726	0.960	0.8970	8.716	0.9547	0.768	0.7510
NODE	0.461	0.791	0.3609	0.728	0.725	0.921	0.8968	**8.710**	0.9579	**0.758**	**0.7464**
ResNet	0.478	0.818	**0.3981**	0.733	0.731	**0.966**	0.8978	8.733	**0.9607**	0.760	0.7467
FT-Transformer											
FT-Transformer$_d$	0.455	0.801	0.3948	0.735	0.730	0.966	0.8969	8.719	**0.9695**	**0.748**	**0.7429**
FT-Transformer	**0.450**	0.810	**0.3983**	**0.737**	0.731	**0.967**	0.8984	8.722	0.9692	**0.748**	0.7434
GBDT											
CatBoost$_d$	0.428	0.798	0.3863	0.724	0.726	0.948	0.8894	8.885	0.9096	0.7490	0.7440
CatBoost	**0.423**	0.794	0.3885	0.727	0.726	–	0.8899	8.837	0.9685	0.7401	**0.7413**
XGBoost$_d$	0.463	0.775	0.3502	0.721	0.705	0.925	0.8803	9.446	0.9640	0.7732	0.7719
XGBoost	0.431	0.796	0.3767	0.725	0.725	–	0.8880	8.819	**0.9696**	**0.7320**	0.7421

3.5.1 Which DNN architecture is the best?

Table 1 indicates that ResNet yields the most consistent performance among single models, as it is the best model for six datasets out of eleven and it is very close to the leader on Epsilon. On the remaining datasets, NODE is the only model that shows non-negligible improvement over ResNet. However, NODE has an ensemble-like structure (and requires significantly more resources for training), so

it is not a truly "single" model, and the gap is mostly gone in the ensembling mode, see Table 2. Surprisingly, MLP is often on par or even better than some of the recently proposed models (SNN, TabNet, GrowNet, DCN V2).
The main take away: ResNet is a strong baseline that none of the recently proposed architectures can outperform and MLP is still a good sanity check.

3.5.2 How GBDT and DNN are compared?

Both Table 1 and Table 2 demonstrate that GBDT is strictly superior to DNN on heterogeneous data (see section 3.3), while on non-heterogeneous data situation is rather the opposite. We also faced slow training of GBDT on multiclass problems with a large number of classes (ALOI ~ 1000 classes) and failed to tune CatBoost and XGBoost due to this phenomenon. This is a well-known issue and sophisticated modifications (such as Si et al. (2017)) are required to alleviate it, but none of them are available in popular GBDT implementations, so DNN solutions should be preferred for this type of tasks.
The main take away: DL research efforts aimed at surpassing GBDT should cover a diverse set of heterogeneous datasets in addition to non-heterogeneous datasets, which should serve only as a sanity check.

3.6 FT-Transformer

In this section, we introduce FT-Transformer (**F**eature **T**okenizer + **Transformer**) — a simple adaptation of the Transformer architecture (Vaswani et al., 2017) for the tabular domain. Figure 1 demonstrates the main parts of FT-Transformer. In a nutshell, our model transforms all features (categorical and numerical) to tokens and runs a stack of Transformer layers over the tokens, so every Transformer layer operates on the *feature* level of *one* object. We compare FT-Transformer to conceptually similar AutoInt in section 4.2.

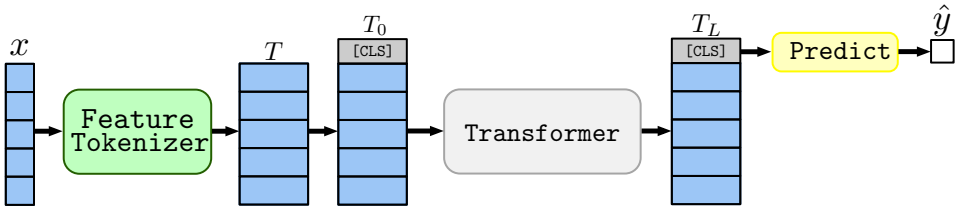

Figure 1: The FT-Transformer architecture. Firstly, Feature Tokenizer transforms features to tokens. The tokens are then processed by the Transformer module and the [CLS] token is used for prediction.

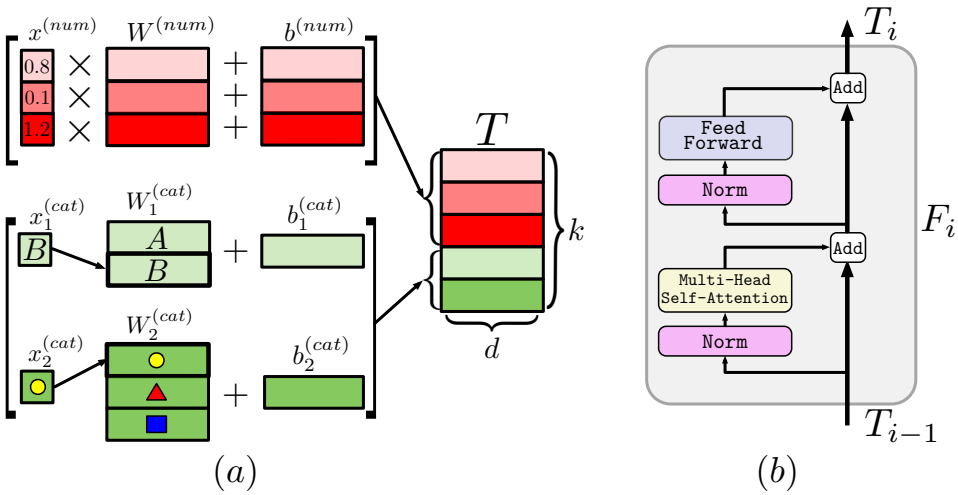

Figure 2: (a) Feature Tokenizer; in the example, there are three numerical and two categorical features; (b) One Transformer layer.

Feature Tokenizer. The Feature Tokenizer module (see Figure 2) transforms the input features x to tokens $T \in \mathbb{R}^{k \times d}$. The token for a given feature x_j is computed as follows:

$$T_j = b_j + f_j(x_j) \in \mathbb{R}^d \qquad f_j : \mathbb{X}_j \to \mathbb{R}^d.$$

where b_j is the j-th *feature bias*, $f_j^{(num)}$ is implemented as the element-wise multiplication with the *direction* $W_j^{(num)} \in \mathbb{R}^d$ and $f_j^{(cat)}$ is implemented as the lookup table $W_j^{(cat)} \in \mathbb{R}^{S_j \times d}$ for categorical features. Overall:

$$\begin{aligned}
T_j^{(num)} &= b_j^{(num)} + x_j^{(num)} \cdot W_j^{(num)} && \in \mathbb{R}^d, \\
T_j^{(cat)} &= b_j^{(cat)} + e_j^T W_j^{(cat)} && \in \mathbb{R}^d, \\
T &= \texttt{stack}\left[T_1^{(num)}, \ \ldots, \ T_{k^{(num)}}^{(num)}, \ T_1^{(cat)}, \ \ldots, \ T_{k^{(cat)}}^{(cat)}\right] && \in \mathbb{R}^{k \times d}.
\end{aligned}$$

where e_j^T is a one-hot vector for the corresponding categorical feature.

Transformer. At this stage, the `[CLS]` token (or "classification token", or "output token", see Devlin et al. (2019)) is appended to T and L Transformer layers $F_1, \ \ldots, \ F_L$ are applied:

$$T_0 = \texttt{stack}\,[\,\texttt{[CLS]}, \ T] \qquad T_i = F_i(T_{i-1}).$$

We use the PreNorm variant for easier optimization (Wang et al., 2019b), see Figure 2. See the original paper (Vaswani et al., 2017) for the background on Multi-Head Self-Attention (MHSA) and the Feed Forward module. See supplementary for details such as activations, placement of normalizations and dropout modules (Srivastava et al., 2014).

Prediction. The final representation of the `[CLS]` token is used for prediction:

$$\hat{y} = \texttt{Linear}(\texttt{ReLU}(\texttt{LayerNorm}(T_L^{\texttt{[CLS]}}))).$$

Limitations. FT-Transformer requires more resources (both hardware and time) for training than simple models such as ResNet and may not be easily scaled to datasets when the number of features is "too large" (it is determined by the available hardware and time budget). The main reason lies in the quadratic complexity of the vanilla MHSA with respect to the number of features. However, the issue can be alleviated by using efficient approximations of MHSA (Tay et al., 2020). Additionally, it is still possible to distill FT-Transformer into simpler architectures for better inference performance. We report training times and the used hardware in supplementary.

3.6.1 FT-Transformer vs baseline neural networks

Table 1 and Table 2 indicate the advantage of FT-Transformer

- FT-Transformer is superior to all recently proposed DL architectures.
- FT-Transformer is superior to ResNet on heterogeneous data and performs on par with ResNet on non-heterogeneous data (see section 3.3). We develop this intuition further in section 4.1.
- FT-Transformer strongly benefits from ensembling: in this mode, FT-Transformer closes the gap with ResNet on two datasets (Helena and ALOI).

Interestingly, FT-Transformer appears to be not very sensitive to hyperparameters:

- FT-Transformer allows for building powerful ensembles without tuning: ensembles of default FT-Transformers demonstrate similar performance to ensembles of tuned ones.
- on heterogeneous data, even the default configuration of FT-Transformer already performs well and beats most competitors.

3.6.2 FT-Transformer vs GBDT

In this section, we compare FT-Transformer with CatBoost and XGBoost.

We start with observations on the overall performance using Table 1 and Table 2.

- **Default hyperparameters.** Ensembled FT-Transformer is superior to GBDT, with California as the only exception. As for single models, FT-Transformer consistently demonstrates adequate off-the-shelf performance. The default XGBoost turns out to be uncompetitive (we observed extreme overfitting). CatBoost is more reliable in that regard, with only one exception (Covertype).
- **Tuned hyperparameters.** Once hyperparameters are properly tuned, GBDTs start dominating on the heterogeneous datasets, while FT-Transformer keeps leadership on DNN-friendly problems. Nevertheless, a small "success story" happens on one heterogeneous dataset (Covertype), where FT-Transformer becomes the only DNN that performs on par with GBDT (see Table 2).
- **Ensembling.** The benefit from ensembling for GBDT is clearly limited compared to FT-Transformer. It makes sense, since a single GBDT is already an ensemble-like model, which is not the case for FT-Transformer.

4 Analysis

4.1 When FT-Transformer is better than ResNet?

In this section, we make the first step towards understanding the difference in behavior between FT-Transformer and ResNet. To achieve that, we design a sequence of synthetic tasks where the difference between performance of the two models gradually changes from negligible to dramatic. Namely, we generate and *fix* objects $\{x_i\}_{i=1}^n$, perform the train-val-test split *once* and interpolate between two regression targets: f_{GBDT}, which is supposed to be "GBDT friendly" and f_{DNN}, which is expected to be easier for DNN. Formally, for one object:

$$x \sim \mathcal{N}(0, I_k),$$
$$y = \alpha \cdot f_{GBDT}(x) + (1 - \alpha) \cdot f_{DNN}(x).$$

where $f_{GBDT}(x)$ is an average prediction of 30 randomly constructed decision trees, and $f_{DNN}(x)$ is an MLP with three randomly initialized hidden layers. Both f_{GBDT} and f_{DNN} are generated once, i.e. the same functions are applied to all objects (see supplementary for details). The resulting targets are standardized before training. The results are visualized in Figure 3. ResNet and FT-Transformer perform similarly well on the DNN-friendly tasks and outperform CatBoost on those tasks. However, the ResNet's relative performance drops significantly when the target becomes more GBDT friendly. By contrast, FT-Transformer yields competitive performance across the whole range of tasks.

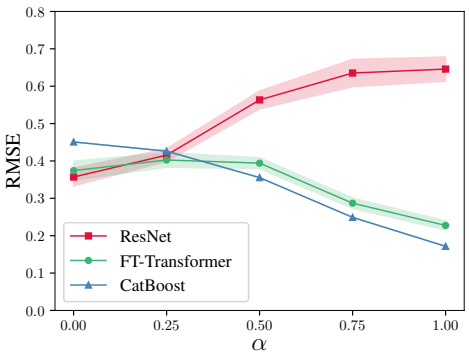

Figure 3: RMSE for the same test set with five different regression targets. The results are averaged over five seeds, shadows represent standard deviations.

The conducted experiment reveals a type of functions that are better approximated by FT-Transformer than by ResNet. Additionally, the fact that these functions are based on decision trees correlates with the results in Table 1 and Table 2, where FT-Transformer shows the most convincing improvements over other deep models exactly on those datasets where GBDT dominates.

4.2 Ablation study

In this section, we test some design choices of FT-Transformer.

First, we compare FT-Transformer with AutoInt (Song et al., 2019), since it is the closest competitor in its spirit. AutoInt also converts all features to tokens and applies self-attention on top of them. However, in its details, AutoInt significantly differs from FT-Transformer: its "tokenizer" does not include feature biases, its backbone significantly differs from the vanilla Transformer (Vaswani et al., 2017), and the inference mechanism does not use the `[CLS]` token.

Second, we check whether feature biases in Feature Tokenizer are essential for good performance.

Table 3: The results of the comparison between FT-Transformer and two attention-based alternatives: AutoInt and FT-Transformer without feature biases. Notation follows Table 1

	CA ↓	HE ↑	JA ↑	HI ↑	AL ↑	YE ↓	CO ↑	MI ↓
AutoInt	0.479	0.372	0.716	**0.726**	0.945	8.875	0.931	0.752
FT-Transformer (w/o feat. enc.)	0.472	0.381	0.724	0.723	0.958	**8.821**	0.962	0.751
FT-Transformer	**0.464**	**0.391**	**0.731**	**0.728**	**0.960**	**8.820**	**0.964**	**0.747**

Table 4: Rank correlation (takes values in $[-1,\ 1]$) between permutation test's feature importances ranking and two alternative rankings: Attention Maps (AM) and Integrated Gradients (IG). Means and standard deviations over five runs are reported.

	CA	HE	JA	HI	AL	YE	CO	MI
AM	0.81 (0.05)	0.77 (0.03)	0.78 (0.05)	0.91 (0.03)	0.84 (0.01)	0.92 (0.01)	0.84 (0.04)	0.86 (0.02)
IG	0.84 (0.08)	0.74 (0.03)	0.75 (0.04)	0.72 (0.03)	0.89 (0.01)	0.50 (0.03)	0.90 (0.02)	0.56 (0.02)

We tune and evaluate FT-Transformer without feature biases following the same protocol as in section 3.4 and reuse the remaining numbers from section 3.5. The results averaged over 15 runs are reported in Table 3 and demonstrate both the superiority of the Transformer's backbone to that of Autoint and the necessity of feature biases.

4.3 Obtaining feature importances from attention maps

In this section, we evaluate attention maps as a source of information on feature importances for FT-Transformer for a given set of samples. For the i-th sample, we calculate the average attention map p_i for the `[CLS]` token from Transformer's forward pass. Then, the obtained individual distributions are averaged into one distribution p that represents the feature importances:

$$p = \frac{1}{n_{samples}} \sum_i p_i \qquad p_i = \frac{1}{n_{heads} \times L} \sum_{h,l} p_{ihl}.$$

where p_{ihl} is the h-th head's attention map for the `[CLS]` token from the forward pass of the l-th layer on the i-th sample. The main advantage of the described heuristic technique is its efficiency: it requires a single forward for one sample.

In order to evaluate our approach, we compare it with Integrated Gradients (IG, Sundararajan et al. (2017)), a general technique applicable to any differentiable model. We use permutation test (PT, Breiman (2001)) as a reasonable interpretable method which allows to establish a constructive metric, namely, rank correlation. We run all the methods on the train set and summarize results in Table 4. Interestingly, the proposed method yields reasonable feature importances and performs similarly to IG (note that this does not imply similarity to IG's feature importances). Given that IG can be orders of magnitude slower and the "baseline" in the form of PT requires $(n_{features} + 1)$ forward passes (versus one for the proposed method), we conclude that the simple averaging of attention maps can be a good choice in terms of cost-effectiveness.

5 Conclusion

In this work, we have investigated the status quo in the field of deep learning for tabular data. By the extensive experimental comparison of the recent DL models on various datasets, we have identified that most of them do not consistently outperform the simplest ResNet-like baseline as well as established GBDT implementations. Given these observations, we have proposed a new attention-based architecture that outperforms ResNet on many tasks. The code and all the details of the study are open-sourced [1], and we hope that our evaluation and two simple baselines will serve as a basis for further developments on tabular DL.

[1] https://github.com/AnonymouzAuthorz/RevisitingTabularDL

References

T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. Optuna: A next-generation hyperparameter optimization framework. In *KDD*, 2019.

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

S. Badirli, X. Liu, Z. Xing, A. Bhowmik, K. Doan, and S. S. Keerthi. Gradient boosting neural networks: Grownet. *arXiv*, 2002.07971v2, 2020.

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

T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere. The million song dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)*, 2011.

A. Beutel, P. Covington, S. Jain, C. Xu, J. Li, V. Gatto, and E. H. Chi. Latent cross: Making use of context in recurrent recommender systems. In *WSDM 2018: The Eleventh ACM International Conference on Web Search and Data Mining*, 2018.

J. A. Blackard and D. 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.

L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

O. Chapelle and Y. Chang. Yahoo! learning to rank challenge overview. In *Proceedings of the Learning to Rank Challenge*, volume 14, 2011.

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

J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv*, 1810.04805v2, 2019.

A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

S. Fort, H. Hu, and B. Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv*, 1912.02757v2, 2020.

J. H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.

J. M. Geusebroek, G. J. Burghouts, , and A. W. M. Smeulders. The amsterdam library of object images. *Int. J. Comput. Vision*, 61(1):103–112, 2005.

I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

I. Guyon, L. Sun-Hosoya, M. Boullé, H. J. Escalante, S. Escalera, Z. Liu, D. Jajetic, B. Ray, M. Saeed, M. Sebag, A. Statnikov, W. Tu, and E. Viegas. Analysis of the automl challenge series 2015-2018. In *AutoML*, Springer series on Challenges in Machine Learning, 2019.

H. Hazimeh, N. Ponomareva, P. Mol, Z. Tan, and R. Mazumder. The tree ensemble layer: Differentiability meets conditional computation. In *ICML*, 2020.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv*, 1512.03385v1, 2015.

X. Huang, A. Khetan, M. Cvitkovic, and Z. Karnin. Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv*, 2012.06678v1, 2020.

S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30: 3146–3154, 2017.

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

G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter. Self-normalizing neural networks. In *NIPS*, 2017.

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

P. Kontschieder, M. Fiterau, A. Criminisi, and S. Rota Bulo. Deep neural decision forests. In *Proceedings of the IEEE international conference on computer vision*, 2015.

I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

Y. Lou and M. Obukhov. Bdt: Gradient boosted decision tables for high accuracy and scoring efficiency. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017.

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.

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

L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin. Catboost: unbiased boosting with categorical features. In *NeurIPS*, 2018.

T. Qin and T. Liu. Introducing LETOR 4.0 datasets. *arXiv*, 1306.2597v1, 2013.

Z. Qin, L. Yan, H. Zhuang, Y. Tay, R. K. Pasumarthi, X. Wang, M. Bendersky, and M. Najork. Are neural rankers still outperformed by gradient boosted decision trees? In *ICLR*, 2021.

S. Si, H. Zhang, S. S. Keerthi, D. Mahajan, I. S. Dhillon, and C. Hsieh. Gradient boosted decision trees for high dimensional sparse output. In *ICML*, 2017.

W. Song, C. Shi, Z. Xiao, Z. Duan, Y. Xu, M. Zhang, and J. Tang. Autoint: Automatic feature interaction learning via self-attentive neural networks. In *CIKM*, 2019.

N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.

M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *ICML*, 2017.

Y. Tay, M. Dehghani, D. Bahri, and D. Metzler. Efficient transformers: A survey. *arXiv*, 2009.06732v1, 2020.

R. Turner, D. Eriksson, M. McCourt, J. Kiili, E. Laaksonen, Z. Xu, and I. Guyon. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. *arXiv*, https://arxiv.org/abs/2104.10201v1, 2021.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NIPS*, 2017.

A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*, 2019a.

Q. Wang, B. Li, T. Xiao, J. Zhu, C. Li, D. F. Wong, and L. S. Chao. Learning deep transformer models for machine translation. In *ACL*, 2019b.

R. Wang, B. Fu, G. Fu, and M. Wang. Deep & cross network for ad click predictions. In *ADKDD*, 2017.

R. Wang, R. Shivanna, D. Z. Cheng, S. Jain, D. Lin, L. Hong, and E. H. Chi. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. *arXiv*, 2008.13535v2, 2020.

Y. Yang, I. G. Morillo, and T. M. Hospedales. Deep neural decision trees. *arXiv*, 1806.06988v1, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See the "Limitations" paragraph of section 3.6
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] The increase in resource requirements and possible mitigations are mentioned in the "Limitations" paragraph of section 3.6. We also report hardware and training times in supplementary.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A] We do not include theoretical results.
   (b) Did you include complete proofs of all theoretical results? [N/A] We do not include theoretical results.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See https://github.com/AnonymouzAuthorz/RevisitingTabularDL
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See https://github.com/AnonymouzAuthorz/RevisitingTabularDL and section 3.4
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We report them in supplementary due to the lack of space in the main paper.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See supplementary. Similar details are also available for many experiments at https://github.com/AnonymouzAuthorz/RevisitingTabularDL
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See section 3.3
   (b) Did you mention the license of the assets? [Yes] We did at https://github.com/AnonymouzAuthorz/RevisitingTabularDL
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We packed our splits of public datasets as a single pack, the URL can be obtained at https://github.com/AnonymouzAuthorz/RevisitingTabularDL
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] The used datasets are publicly available and well known in the community.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] The used datasets do not contain such information. We briefly described what the used datasets are about in section 3.3.

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We did not used crowdsourcing. We did not conduct research with human subjects.

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] We did not used crowdsourcing. We did not conduct research with human subjects.

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] We did not used crowdsourcing. We did not conduct research with human subjects.