



HAL
open science

Sparse tree-based initialization for neural networks

Patrick Lutz, Ludovic Arnould, Claire Boyer, Erwan Scornet

► **To cite this version:**

Patrick Lutz, Ludovic Arnould, Claire Boyer, Erwan Scornet. Sparse tree-based initialization for neural networks. 2022. hal-03792008

HAL Id: hal-03792008

<https://hal.science/hal-03792008>

Preprint submitted on 29 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SPARSE TREE-BASED INITIALIZATION FOR NEURAL NETWORKS

Patrick Lutz¹, Ludovic Arnould², Claire Boyer², and Erwan Scornet³

¹Department of Computer Science, Boston University

²Laboratoire de Probabilités, Statistique et Modélisation, Sorbonne Université

³CMAP, Ecole Polytechnique

September 29, 2022

Abstract

Dedicated neural network (NN) architectures have been designed to handle specific data types (such as CNN for images or RNN for text), which ranks them among state-of-the-art methods for dealing with these data. Unfortunately, no architecture has been found for dealing with tabular data yet, for which tree ensemble methods (tree boosting, random forests) usually show the best predictive performances. In this work, we propose a new sparse initialization technique for (potentially deep) multilayer perceptrons (MLP): we first train a tree-based procedure to detect feature interactions and use the resulting information to initialize the network, which is subsequently trained via standard stochastic gradient strategies. Numerical experiments on several tabular data sets show that this new, simple and easy-to-use method is a solid concurrent, both in terms of generalization capacity and computation time, to default MLP initialization and even to existing complex deep learning solutions. In fact, this wise MLP initialization raises the resulting NN methods to the level of a valid competitor to gradient boosting when dealing with tabular data. Besides, such initializations are able to preserve the sparsity of weights introduced in the first layers of the network through training. This fact suggests that this new initializer operates an implicit regularization during the NN training, and emphasizes that the first layers act as a sparse feature extractor (as for convolutional layers in CNN).

1 Introduction

Neural networks have become hegemonic in machine learning in particular when dealing with very structured data. They indeed provide state-of-the-art performances for applications with images or text. However, they still perform

poorly on tabular inputs, for which tree ensemble methods remain the gold standards (Grinsztajn et al., 2022).

Tree ensemble methods Tree-based methods are widely used in the ML community, especially for processing tabular data. Two main approaches exist depending on whether the tree building process is parallel (e.g. Random Forest, RF, see Breiman, 2001a) or sequential (e.g. Gradient Boosting Decision Trees, GBDT, see Friedman, 2001). In these tree ensemble procedures, the final prediction relies on averaging predictions of randomized decision trees, coding for particular partitions of the input space. The two most successful and most widely used implementations of these methods are XGBoost and LightGBM (see Chen and Guestrin, 2016; Ke et al., 2017) which both rely on the sequential GBDT approach.

Neural networks Neural Networks (NN) are efficient methods to unveil the patterns of spatial or temporal data, such as images (Krizhevsky et al., 2012) or texts (Liu et al., 2016). Their performance results notably from the fact that several architectures directly encode relevant structures in the input: convolutional neural networks (CNN, LeCun et al., 1995) use convolutions to detect spatially-invariant patterns in images, and recurrent neural networks (RNN, Rumelhart et al., 1985) use a hidden temporal state to leverage the natural order of a text. However, a dedicated *natural* architecture has yet to be introduced to deal with tabular data. Indeed, designing such an architecture would require to detect and leverage the structure of the relations between variables, which is much easier for images or text (spatial or temporal correlation) than for tabular data (unconstrained covariance structure).

NN initialization and training In the absence of a suitable architecture for handling tabular data, the Multi-Layer Perceptron (MLP) architecture (Rumelhart et al., 1986) remains the obvious choice due to its generalist nature. Apart from the large number of parameters, one difficulty of MLP training arises from the non-convexity of the loss function (see, e.g., Sun, 2020). In such situations, the initialization of the network parameters (weights and biases) are of the utmost importance, since it can influence both the optimization stability and the quality of the minimum found. Typically, such initializations are drawn according to independent uniform distributions with a variance decreasing w.r.t. the size of the layer (He et al., 2015). Therefore, one may wonder how to capitalize on methods that are inherently capable of recognizing patterns in tabular data (e.g., tree-based methods) to propose a new NN architecture suitable for tabular data and an initialization procedure that leads to faster convergence and better generalization performance.

1.1 Related works

How MLP can be used to handle tabular data remains unclear, especially since a corresponding prior in the MLP architecture adapted to the correlations of the input is not obvious, to say the least. Indeed, none of the existing NN architectures can consistently match the performance of state-of-the-art tree-based predictors on tabular data (Shwartz-Ziv and Armon, 2022; Gorishniy et al., 2021; and in particular Table 2 in Borisov et al., 2021).

Self-attention architectures Specific NN architectures have been proposed to deal with tabular data. For example, TabNet (Arik and Pfister, 2021) uses a sequential self-attention structure to detect relevant features and then applies several networks for prediction. SAINT (Somepalli et al., 2021), on the other hand, uses a two-dimensional attention structure (on both features and samples) organized in several layers to extract relevant information which is then fed to a classical MLP. These methods typically require a large amount of data, since the self-attention layers and the output network involve numerous MLP.

Trees and neural networks Several solutions have been proposed to leverage information from pre-trained tree-based methods to develop NN that are capable of efficiently processing tabular data. For example, TabNN (Ke et al., 2018) first trains a GBDT on the available data, then extracts a group of features per individual tree, compresses the resulting groups, and uses a tailored Recursive Encoder based on the structure of these groups (with an initialization based on the tree leaves). Another approach consists in using pre-trained tree-based methods to initialize MLP. Translating decision trees into MLP was first achieved by Sethi (1990) and Brent (1991), and later exploited by Welbl (2014), Richmond et al. (2015) and Biau et al. (2019). Such procedures can be seen as a way to relax and generalize the partition geometry produced by trees and their aggregation. However, these approaches are restricted to shallow NN, as a decision tree is unfolded into a 3-layer neural network.

1.2 Contributions

In this work, we propose a new method to initialize a potentially deep MLP for learning tasks with tabular data. Our initialization consists in first training a tree-based method (RF, GBDT or Deep Forest, see Section 2.1) and then using its translation into a MLP as initialization. This procedure is shown to outperform the widely used uniform initialization of MLP (default initialization in Pytorch Paszke et al., 2019) in the following manner.

1. **Improved performances.** For tabular data, the predictive performances of the MLP after training are improved compared to MLP that use a random initialization. Our procedure also outperforms more complex deep learning procedures based on self-attention and is on par with classical tree-based methods (such as XGBoost).

2. **Faster optimization.** The optimization following a tree-based initialization is boosted in the sense that it enjoys a faster convergence towards a (better) empirical minimum: a tree-based initialization results in faster training of the MLP.

Regarding deep MLP, we show that initializing the first layers with tree-based methods is sufficient, other layers being randomly initialized in a conventional way. This supports the idea that in our method, the (first) tree-based initialized layers act as relevant embeddings that allow this standard NN architecture to detect correlations in the inputs.

Outline In Section 2, we introduce the predictors in play and describe how tree-based methods can be translated into MLP. The core of our analysis is contained in Section 3, where we describe in detail the MLP initialization process and provide extensive numerical evaluations showing the benefits of this method.

2 Equivalence between trees and MLP

Consider the classical setting of supervised learning in which we are given a set of input/output samples $\{(X_i, Y_i)\}_{i=1}^n$ drawn i.i.d. from some unknown joint distribution. Our goal is to construct a (MLP) function to predict the output from the input. To do so, we leverage the translation of tree-based methods into MLP.

2.1 Presentation of the predictors in play

Tree-based methods We consider three different tree ensemble methods: Random Forests (RF), Gradient Boosting Decision Trees (GBDT) and Deep Forests (DF). They all share the same base component: the Decision Tree (DT, for details see Breiman et al., 1984). We call its terminal nodes *leaf nodes* and the other nodes the *inner nodes*. RF (Breiman, 2001b) is a predictor consisting of a collection of M independently trained and randomized trees. Its final prediction is made by averaging the predictions of all its DT in regression or by a majority vote in classification. GBDT (Friedman, 2001) aims at minimizing a prediction loss function by successively aggregating DT that approximate the opposite gradient of that loss function (see Chen and Guestrin, 2016, for details on XGBoost). DF (Zhou and Feng, 2019) is a hybrid learning procedure in which random forests are used as elementary components (neurons) of a neural-network-like architecture (see Figure 5 and Appendix A for details).

Multilayer Perceptron (MLP) The multilayer perceptron is a predictor consisting of a composition of multiple affine functions, with (potentially different) nonlinear activation functions between them. Standard activation functions include, for instance, the rectified linear unit or the hyperbolic tangent. Deep MLP




Figure 1: Illustration of a decision tree, its induced feature space partition and its corresponding MLP translation on a problem with two explanatory variables. Figure taken from Biau et al. (2019).

are a much richer class of predictors than tree-based methods which build simple partitions of the space and output piecewise constant predictions. Therefore, any of the tree-based models presented above can be approximated and in fact exactly rewritten as an MLP as follows.

2.2 An exact translation of tree-based methods into MLP

From decision tree to 3-layer MLP Recall that a decision tree codes for a partition of the input space in as many parts as there are leaf nodes in the tree. Given an input x , we can identify the leaf where x falls by examining for each hyperplane of the partition whether x falls on the right or left side of the hyperplane. The prediction is then made by averaging the outputs of all the points falling into the leaf of x . A DT can be thus translated into a two-hidden-layer MLP:

1. the first layer contains a number of neurons equal to the number of hyperplanes in the partition, each neuron encoding by ± 1 whether x falls on the left or right side of the hyperplane,
2. the second layer contains a number of neurons equal to the number of leaves in the DT. Based on the first layer, it identifies in which leaf x falls and outputs a vector with a single 1 at the leaf position and -1 everywhere else,
3. the last layer contains a single output neuron that returns the tree prediction. Its weights encode the value stored in each leaf of the tree.

The translation procedure is explained in details (and formally written mathematically) in Biau et al. (2019) and in Appendix B for completeness.

From RF/GBDT to 3-layer MLP Although RF and GBDT are constructed in different ways, they both average multiple DT predictions to give the final result. Thus, to translate a RF or a GBDT into an MLP, we simply turn each tree into a 3-layer MLP as described above, and concatenate all the obtained networks to form a wider 3-layer MLP. When concatenating, we set all weights between the MLP translations of the different trees to 0, since the trees do not

interact with each other in predicting the target value for a new feature vector. Note that no additional layer is required in the MLP translation for averaging the results of multiple decision trees, as this operation can directly be carried out by the third layer. To this end, the last layer of the MLP translation uses an identity activation function and can therefore include the averaging step across trees, which is itself a linear transformation of the tree outputs.

From Deep Forests to deeper MLP As a Deep Forest is a cascade of Random Forests, it can be therefore translated into an MLP containing the MLP translations of the different RF in cascade, resulting in a deeper and wider MLP (note that the obtained MLP counts a number of hidden layers as a multiple of 3). Furthermore, in the Deep Forest architecture, the input vector is concatenated to the output of each intermediate layer, as shown in Figure 5. To mimic these skip connections in the MLP, we add additional neurons to each layer, except for the last three, which encode an identity mapping. Figure 6 illustrates this concept. Note that recasting a deep forest into a deep MLP using this method may suffer from numerical instabilities altering the predictive behaviour, see Appendix C. This is due to a phenomenon of catastrophic cancellation, more likely to occur with deep MLP translations. This does not impact the sequel of the study, relying on the MLP approximations considered hereafter.

2.3 Relaxing tree-based translation to allow gradient descent training

As shown in the previous section, one can construct an MLP that exactly reproduces a tree-based predictor. However this translation involves (i) piecewise constant activation functions (sign) and (ii) different activation functions in a same layer (sign and identity when translating DF). These constraints can hinder the MLP training, relying on stochastic gradient strategies (requiring differentiability), as well as efficient implementation tricks, given that automatic differentiation libraries only support one activation function per layer. Therefore, given a pre-trained tree-based predictor (RF, GBDT or DF), we aim at relaxing its translation into a MLP, mimicking its behavior as closely as possible but in a compatible way with standard NN training.

From tree-based methods to differentiable MLP To do so, Welbl (2014); Biau et al. (2019) consider the differentiable tanh activation, well suited for approximating both the sign and identity functions. Indeed, this can be achieved by multiplying or dividing the output of a neuron by a large constant before applying the function tanh and rescaling the result accordingly if necessary, i.e. for $a, c > 0$ large enough, $\text{sign}(x) \approx \tanh(ax)$ and $x \approx c \tanh\left(\frac{x}{c}\right)$.

We cannot choose a arbitrarily large as this would make gradients vanish during the network optimization (the function being flat on most of the space), and hinder the training. We therefore introduce 4 hyper-parameters for the MLP encoding of any tree-based method that regulate the degree of approximation for the activation functions after the first, second and third layers of a

decision tree translation, as well as for the identity mapping, respectively denoted by `strength01`, `strength12`, `strength23` and `strength_id`.

Hyperparameter choice The use of the tanh activation function involves extra hyper-parameters. We study the influence of each one, by making them vary in some range (keeping the others fixed to 10^{10} , resulting in an almost perfect approximation of the sign and identity functions), see Appendix C.1 for details. Our analysis shows that increasing the hyperparameters beyond some limit value is no longer beneficial (as the activation functions are already perfectly approximated) and, across multiple data sets, these limit values are similar. We also exhibit relevant search spaces that will allow us to find optimal HP values for each application.

3 A new initialization method for MLP training

In this section, we study the impact of tree-based initialization methods for MLP training when dealing with tabular data. The latter empirically proves to be always preferable to standard random initialization and makes MLP a competitive predictor for tabular data.

3.1 Our proposal

Random initialization is the most common technique for initializing MLP before stochastic gradient training. It consists in setting all layer parameters to random values of small magnitude centered at 0. More precisely, all parameter values of the j -th layer are uniformly drawn in $[-1/\sqrt{d_j}, 1/\sqrt{d_j}]$ where d_j is the layer input dimension; this is the default behaviour of most MLP implementations such as `nn.Linear` in PyTorch (Paszke et al., 2019).

We introduce new ways of initializing an MLP for learning with tabular data, by leveraging the recasting of tree-based methods in a neural network fashion:

- **RF/GBDT initialization.** First, a RF/GBDT is fitted to the training data and transformed into a 3-layer neural network, following the procedure described in Section 2. The first two layers of this network are used to initialize the first two layers of the network of interest. Thus, upon initialization, these first two layers encode the RF partition. The parameters of the third and all subsequent layers are randomly initialized as described above.
- **DF initialization.** Similarly as above, a Deep Forest (DF) using ℓ forest layers is first fitted to the training data. The first $3\ell - 1$ layers of the MLP are then initialized using the first $3\ell - 1$ layers of the MLP encoding of this pre-trained DF. The parameters of the 3ℓ -th and all subsequent layers are randomly initialized as explained above.

The tree-based initialization techniques may seem far-fetched at first glance, but they are actually consistent with recent approaches for adapting Deep Learning models for tabular data. The key to interpreting them is to think of the first (tree-based initialized) layers of the MLP as a feature extractor that produces an abstract representation of the input data (in fact, this is a vector encoding the tree-based predictor space partition in which the observation lies). The subsequent randomly initialized layers, once trained, then perform the prediction task based on this abstract representation (see MLP sparsity in Section 3.5 for details).

3.2 Experimental setup

Datasets & learning tasks We compare prediction performances on a total of 10 datasets: 3 regression datasets (Airbnb, Diamonds and Housing), 5 binary classification datasets (Adult, Bank, Blastchar, Heloc, Higgs) and 2 multi-class classification datasets (Covertypes and Volkert). We mostly chose data sets that are used for benchmarking in relevant literature: Adult, Heloc, Housing, Higgs and Covertypes are used by Borisov et al. (2021) and Bank, Blastchar and Volkert are used by Somepalli et al. (2021). Moreover, we add Airbnb and Diamonds to balance the different types of prediction tasks. The considered datasets are all medium-sized (10–60k observations) except for Covertypes and Higgs (approx. 500k observations). Details about the datasets are given in Appendix D.1.

Predictors We consider the following tree-based predictors: Random Forest (RF), Deep Forest (DF, Zhou and Feng, 2017) and XGBoost (denoted by GBDT, Chen and Guestrin, 2016). The latter usually achieves state-of-the-art performances on tabular data sets (see, e.g., Shwartz-Ziv and Armon, 2022; Gorishniy et al., 2021; Borisov et al., 2021). We also consider deep learning approaches: MLP with default uniform initialization (MLP rand. init.) or tree-based initialization (resp. MLP RF init., MLP GBDT init. and MLP DF init.); and a transformer architecture SAINT Somepalli et al. (2021). This complex architecture is specifically designed for applications on tabular data and includes self-attention and inter-sample attention layers that extract feature correlations that are then passed on to an MLP. For regression and classification tasks, we use the mean-squared error (MSE) and cross-entropy loss for NN training, respectively. We choose SAINT as a baseline model as it is reported to outperform all other NN predictors on most of our data sets (all except Airbnb and Diamonds, see Borisov et al., 2021; Somepalli et al., 2021).

Parameter optimization All NN are trained using the Adam optimizer (see Kingma and Ba, 2014, and details in Appendix D.5.2). The MLP hyper-parameters (HP) are determined empirically using the optuna library (Akiba et al., 2019) for Bayesian optimization.

For most HP, we use the default search spaces of Borisov et al. (2021), and for all HP tuning the translation between tree-base methods and MLP, we have identified relevant search spaces (see Appendix C.1). An overview of all search spaces used for each method can be found in Appendix D.5. The quantity minimized during HP tuning is the model’s validation loss, and the smallest validation loss that occurred during training for all MLP-based models.

3.3 A better MLP initialization for a better optimization

In this subsection, the optimization of standard MLP is shown to benefit from the proposed initialization technique. Experiments have been conducted on 6 out of the 10 data sets.

Experimental protocol 1 (P1) To obtain comparable optimization processes, we ensure that all purely MLP-related hyper-parameters, i.e. the MLP width, depth and learning rate, are identical for all MLP in play regardless of the initialization technique. These HP are chosen to maximize the predictive performance of the standard randomly initialized MLP. All HP related to the initialization technique (HP of the tree-based predictor, HP tuning the translation between trees and MLP) are optimized independently for each tree-based initialization method.

Results Figure 2 shows that for most data sets, the use of tree-based initialization methods for MLP training provides a *faster convergence* towards a *better minimum* (in terms of generalization) than random initialization. This is all the more remarkable since Protocol P1 has been calibrated in favor of random initialization. Among tree-based initializers, GBDT initializations outperform or are on par with RF initializations in terms of the optimization behavior on all regression and binary classification problems. However, for multi-class classification problems, the advantages of tree-based initialization seem to be limited. This is probably due to the fact that the MLP architecture at play is tailored for random initialization, being thus too restrictive for tree-based initializers.

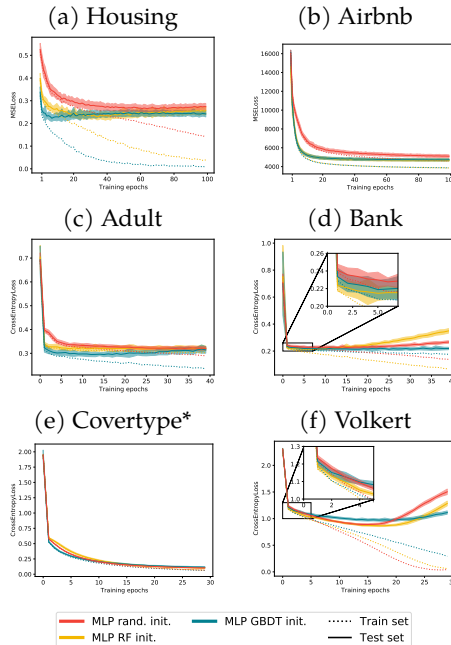


Figure 2: Optimization behaviour of randomly, RF and GBDT initialized MLP evaluated over a 5 times repeated (satisfied) 5-fold of each data set, according to Protocol P1. The lines and shaded areas report the mean and standard deviation. *evaluation on a single 5-fold cross validation.

Model \ Data set	Housing	Airbnb	Diamonds	Adult	Bank	Blastchar	Heloc	Higgs	Covertype	Volkert
	MSE ↓	MSE ↓ (x10 ³)	MSE ↓ (x10 ³)	AUC ↑ (in %)	AUC ↑ (in %)	AUC ↑ (in %)	AUC ↑ (in %)	AUC ↑ (in %)	Acc. ↑ (in %)	Acc. ↑ (in %)
Random Forest	0.263±0.009	5.39±0.13	9.80±0.35	91.6±0.3	92.8±0.3	84.5±1.2	91.3±0.6	80.4±0.1	83.6±0.1	64.2±0.3
GBDT	0.208±0.010	4.71±0.15	7.38±0.28	92.7±0.3	93.3±0.3	84.7±1.0	92.1±0.4	82.8±0.1	97.0±0.0	71.3±0.4
Deep Forest	0.225±0.008	4.68±0.16	8.23±0.29	91.8±0.3	92.9±0.2	83.7±1.2	90.3±0.5	81.2±0.0*	92.4±0.1*	66.3±0.4
MLP rand. init.	0.258±0.011	5.07±0.16	15.5±12.5	90.5±0.4	91.0±0.3	81.4±1.2	80.1±0.1	83.2±0.3	<u>96.7±0.0</u>	72.2±0.4
MLP RF init.	0.222±0.009	4.66±0.16	<u>7.93±0.22</u>	<u>92.1±0.3</u>	<u>92.4±0.4</u>	84.4±1.2	91.7±0.4	83.6±0.1	<u>96.7±0.0</u>	74.1±0.4
MLP GBDT init.	0.206±0.007	4.70±0.09	8.15±0.35	<u>92.2±0.3</u>	<u>92.5±0.3</u>	84.6±1.2	91.5±0.6	83.0±0.0	96.2±0.0	73.5±0.5
MLP DF init.	0.234±0.016	4.81±0.13	8.28±0.24	91.9±0.4	92.2±0.3	84.2±1.0	91.4±0.6	83.3±0.1*	94.5±0.3*	71.3±0.5
SAINT	0.258±0.011	4.81±0.15	17.7±3.83	91.6±0.3	92.2±0.4	84.0±0.8	90.2±0.7	83.7±0.1*	96.6±0.1*	70.1±0.4

Table 1: Best scores for Protocol 2. For each data set, predictors performing at least as well as the best (resp. DL) score up to its standard deviation are highlighted in **bold** (resp. underlined). The scores are based on 5 times repeated (stratified) 5-fold cross validation. For each model, HP have been chosen via the “optuna” library with 100 iterations. *score based on a simple 5-fold cross validation. See Appendix D.4.1 for a comparison with literature results.

Experiments presented in Appendix D.3 with fixed arbitrary widths corroborate this idea: in this case, the RF initialization is beneficial for the optimization process. For the Adult, Bank, and Volkert data sets, Figure 2 also shows the performance of each method at initialization. None of these procedures leads to a better MLP performance at initialization (due to both the non-exact translation from trees to MLP and to the additional randomly initialized layers), but rather help to guide the MLP in its learning process.

3.4 A better MLP initialization for a better generalization

In this subsection, tree-based initialization methods are shown to systematically improve the predictive power of neural networks compared to random initialization.

Experimental protocol 2 (P2) Each MLP is trained on 100 epochs, but with HP tuned depending on the initialization technique. For maximum comparability, the optimization budget is strictly the same for all methods (100 “optuna” iterations each, where one optuna iteration includes a hold-out validation). In particular, when using a tree-based initializer, we use 25 HP optimization iterations to find optimal HP for the tree-based predictor, fix these HPs, and then use the remaining 75 iterations to determine optimal HP for the MLP. For all NN approaches, the model with the best performance on the validation set during training is kept (using the classical *early-stopping* procedure). Performances are measured via the AUROC score for binary classification, accuracy for multi-class classification and MSE for regression (using 5 runs of 5-fold cross-validation).

Results Table 1 shows that RF or GBDT initialization strictly outperform random initialization, in terms of final generalization performance, for all data sets

except Covertypes (for which performances are similar). Additionally, MLP using both RF and GBDT initialization techniques outperform SAINT on all medium-sized data sets and fall short on large data sets (Higgs and Covertypes). Despite its simplicity and the additional random/noisy layers, the proposed methods (based on RF or GBDT) are on par with GBDT half of the data sets, ranking MLP as relevant predictors for tabular data. Among the tree-based initializers, GBDT initialization performs better than the RF one on regression and binary classification problems but not for multi-class classification, which may result from how GBDT deals with multiclass labels (see Appendix D.3 for explanation). DF initialization, for its part, cannot compete with RF and GBDT initialization, despite showing some improvement over the random one (except for Covertypes and Volkert). This underlines that injecting prior information via tree-based methods into the first layers of a MLP is sufficient to improve its performance.

The interested reader may find a comparison of the optimization procedures of all MLP methods and SAINT (Figure 11) and a table summarizing all HP (Table 9) in Appendix D.5.2. We remark that tree-based initializers generally bring into play wider networks with similar depths (fixed width of 2048 and an optimized depth between 4 and 7) compared to MLP with default initialization. But, for most data sets, the overall procedure is computationally more efficient than state-of-the-art deep learning architectures such as SAINT. This hold for the number of parameters, the training time and the test time (see Table 5-7 in Appendix D.4).

3.5 Analyzing key elements of the new initialization methods

Influence of the MLP width We mainly use standard search spaces from (Borisov et al., 2021) to determine the optimal hyper-parameters for each model. However, the MLP width is an exception to this. The standard search spaces used in the literature usually involve MLP with a few hundred neurons per layer (e.g. up to 100 neurons in Borisov et al., 2021); yet, in this work, we consider MLP with a width up to 2048 neurons. Large MLP are actually very beneficial for tree-based initialization methods as they allow the use of more expressive tree-based models in the initialization step.

Figure 3 compares the performance of an MLP with random/GBDT initializations and different widths. There is no gain in prediction by using wider (and therefore more complex) NN, when randomly initialized. This is corroborated by the results of Table 4: for all regression and binary classification data sets, the performance of our (potentially much wider) MLP with random initialization is consistently very close to the literature values, and only increases for multi-class classifica-




Figure 3: Influence of width on the final generalization performance for random and GBDT initializations. Mean values over 5 times repeated 5-fold cross-validation on the housing data set.

tion problems. However, an MLP initialized with GBDT significantly benefits from enlarging the NN width (justifying a width of 2048 for tree-based initialized MLP in Table D.4.1). This confirms the idea that tree-based initialization helps revealing relevant features to the MLP, all the more as the NN width increases, and by doing so, boost the MLP performance after training.

Performance of the initializer For the regression and binary classification problems considered here, the GBDT model trained before the NN initialization always performs significantly better than the RF model; for the multi-class classification problems, the differences are tenuous (see Figure 12 in Appendix D.6). Note, however, that a better performance of the tree-based predictor used for initialization does not always lead to a better performance after training (see the examples of Airbnb and Volkert data sets in Figure 12). This observation suggests that the predictive performance of the initializer (GBDT or RF) is not sufficient in itself to explain the predictive power of the trained network: other aspects, such as the expressiveness of the feature interactions captured by the initializer or the structure it induces on the MLP, must also play a significant role.

MLP sparsity Figure 4 shows the histograms of the weights of the three first and last layers before and after MLP training, for random, RF and GBDT initializations on the Housing data set (see Appendix D.7 for the Adult data set).

The weight distribution on the first two layers changes significantly during training when the MLP parameters are randomly initialized: the weights are uniformly distributed at epoch 0 but appears to be Gaussian (closely centered around zero) after training. When RF or GBDT initializers are used instead, the weights of the first two layers at epoch 0 are sparsely distributed by construction, and their distribution is preserved during training (notice the logarithmic y-axis for these plots in Figure 4). Note that the (uniform) distribution of the weights in other layers is also preserved through training (third and last lines of Figure 4). This means that our initialization technique, in combination with SGD optimization strategies, introduces an *implicit regularization* of the NN optimization: the obtained MLP is structured and consists of sparse feature extraction layers (with a structure that is particularly suited for capturing features in tab-




Figure 4: Histograms of the first three and the last layers’ weights before and after the MLP training on the Housing data set. Comparison between random, RF and GBDT initializations.

ular data), followed by dense prediction layers. This is very similar to the architecture of CNN (constrained by design), a very successful class of NN designed specifically for image processing.

It has been argued (see, e.g., Section 2.2 in Neal, 2012) that “with Gaussian priors the contributions of individual hidden [neurons] are all negligible, and consequently, these [neurons] do not represent ‘hidden features’ that capture important aspects of the data.” The argument further states that a layer can only capture patterns in the data if it has a sufficient number of individual neurons with non-negligible weights. Indeed, empirical evaluations show that MLP whose weights are tightly centered around zero have worse predictive abilities than those that use a wider range of weights (Blundell et al., 2015). In this line of reasoning, the sparse weight distribution of the first two layers we observe before and after training of tree-based-initialized MLP shows how well these layers are adapted to detect features in the data.

4 Conclusion and Future work

This work builds upon the permeability that can exist between tree methods and neural networks, in particular how the former can help the latter during training, especially with tabular inputs. More precisely, we have proposed new methods for smartly initializing standard MLP using pre-trained tree-based methods. The sparsity of this initialization is preserved during training which shows that it encodes relevant correlations between the data features. Among deep learning methods, such initializations (via GBDT or RF) of MLP always improve the performance compared to the widely used random initialization, and provide an easy-to-use and more efficient alternative to SAINT, the state-of-the-art attention-based deep learning method for tabular data. Beyond deep learning models, the performance of this wisely-initialized MLP is remarkably approaching that of XGBoost, which so far reigns supreme for learning tasks on tabular data.

There is still undoubtedly room for improvement to democratize such a network initialization. For example, one could be interested in extending this method to MLP with arbitrary activation functions (other than tanh). A first step would be to code sign and identity functions, for instance using the ReLU activation.

Our approach also yields a default method for choosing the MLP width, which results from the number of trees and their depth used in the initialization technique. Optimizing DF naturally leads to a default choice of MLP depth. However DF performances are not sufficient to legitimate such procedures for now. Other ways of translating tree structures into deep networks should be considered to build upon our work and propose new data-driven choices of MLP depth.

References

- Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on tabular data? *arXiv preprint arXiv:2207.08815*, 2022.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001a.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- 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, 2017.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. Recurrent neural network for text classification with multi-task learning. *arXiv preprint arXiv:1605.05101*, 2016.
- Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10): 1995, 1995.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- David E Rumelhart, Geoffrey E Hinton, James L McClelland, et al. A general framework for parallel distributed processing. *Parallel distributed processing: Explorations in the microstructure of cognition*, 1(45-76):26, 1986.
- Ruo-Yu Sun. Optimization for deep learning: An overview. *Journal of the Operations Research Society of China*, 8(2):249–294, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90, 2022. ISSN 1566-2535. doi: <https://doi.org/10.1016/j.inffus.2021.11.011>. URL <https://www.sciencedirect.com/science/article/pii/S1566253521002360>.

- Yury Gorishniy, Ivan Rubachev, Valentin Khrukov, and Artem Babenko. Revisiting deep learning models for tabular data, 2021. URL <https://arxiv.org/abs/2106.11959>.
- Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. Deep neural networks and tabular data: A survey, 2021. URL <https://arxiv.org/abs/2110.01889>.
- Sercan Ö Arik and Tomas Pfister. Tabnet: Attentive interpretable tabular learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6679–6687, 2021.
- 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, 2021. URL <https://arxiv.org/abs/2106.01342>.
- Guolin Ke, Jia Zhang, Zhenhui Xu, Jiang Bian, and Tie-Yan Liu. Tabnn: A universal neural network solution for tabular data. 2018.
- Ishwar Krishnan Sethi. Entropy nets: from decision trees to neural networks. *Proceedings of the IEEE*, 78(10):1605–1613, 1990.
- Richard P Brent. Fast training algorithms for multilayer neural nets. *IEEE Transactions on Neural Networks*, 2(3):346–354, 1991.
- Johannes Welbl. Casting random forests as artificial neural networks (and profiting from it). In *German Conference on Pattern Recognition*, pages 765–771. Springer, 2014.
- David L Richmond, Dagmar Kainmueller, Michael Y Yang, Eugene W Myers, and Carsten Rother. Relating cascaded random forests to deep convolutional neural networks for semantic segmentation. *arXiv preprint arXiv:1507.07583*, 2015.
- G erard Biau, Erwan Scornet, and Johannes Welbl. Neural random forests. *Sankhya A*, 81(2):347–386, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- L Breiman, JH Friedman, R Olshen, and CJ Stone. Classification and regression trees. 1984.

- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, Oct 2001b. ISSN 1573-0565. doi: 10.1023/A:1010933404324. URL <https://doi.org/10.1023/A:1010933404324>.
- Z. Zhou and J. Feng. Deep forest. *National Science Review*, 6(1):74–86, 2019.
- Zhi-Hua Zhou and Ji Feng. Deep forest: Towards an alternative to deep neural networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3553–3559, 2017. doi: 10.24963/ijcai.2017/497. URL <https://doi.org/10.24963/ijcai.2017/497>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.
- Wei Fan, Haixun Wang, Philip S Yu, and Sheng Ma. Is random model better? on its accuracy and efficiency. In *Third IEEE International Conference on Data Mining*, pages 51–58. IEEE, 2003.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- R. Kelley Pace and Ronald Barry. Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297, 1997. ISSN 0167-7152. doi: [https://doi.org/10.1016/S0167-7152\(96\)00140-X](https://doi.org/10.1016/S0167-7152(96)00140-X). URL <https://www.sciencedirect.com/science/article/pii/S016771529600140X>.
- Pierre Baldi, Peter Sadowski, and Daniel Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5(1): 1–9, 2014.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf>.

Ryan Turner, David Eriksson, Michael McCourt, Juha Kiili, Eero Laaksonen, Zhen Xu, and Isabelle Guyon. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020, 2021. URL <https://arxiv.org/abs/2104.10201>.

C.M. Bishop, P.N.C.C.M. Bishop, G. Hinton, and Oxford University Press. *Neural Networks for Pattern Recognition*. Advanced Texts in Econometrics. Clarendon Press, 1995. ISBN 9780198538646. URL <https://books.google.fr/books?id=T0S0BgAAQBAJ>.

A Details on Deep Forest (DF) and its translation

The layers of DF are composed of an assortment of Breiman’s Random Forests and Completely-Random Forests (CRF, Fan et al. (2003)) and are trained one after another in a cascade manner. At a given layer, the outputs of all forests are concatenated, together with the raw input data. This new vector serves as input for the next DF layer. This process is repeated for each layer and the final output is obtained by averaging the forest outputs of the best layer (without raw data).




Figure 5: Illustration of the Deep Forest cascade structure for a classification problem with 3 classes. Each level of the cascade consists of two Breiman RFs (black) and two completely random forests (blue). The original input feature vector is concatenated to the output of each intermediate layer. Figure taken from (Zhou and Feng, 2017).

B Details of the translation of a decision tree into an MLP

Recall that a decision tree codes for a partition of the input space in as many parts as there are leaf nodes in the tree. To know in which partition cell an input feature vector $x \in \mathbb{R}^d$ falls into, we move in the tree from the root to the




Figure 6: Illustration of the MLP translation of a Deep Forest. Yellow nodes use the $x \mapsto 2 \cdot \mathbb{1}_{\{x > 0\}} - 1$ activation function and green nodes use the identity activation function.

corresponding leaf using simple rules: at each m -th inner node, x is passed onto the left child node if its i_m -th coordinate is less than or equal to some threshold t_m , and to the right child node otherwise. The decision rule at each inner node of the tree introduces a split of the feature space into two subsets $\mathcal{H}_m^- = \{x \in \mathbb{R}^d \mid x^{(i_m)} \leq t_m\}$ and $\mathcal{H}_m^+ = \{x \in \mathbb{R}^d \mid x^{(i_m)} > t_m\}$. Consistent with how the MLP translation works, we intentionally define \mathcal{H}_m^- and \mathcal{H}_m^+ such that at each inner node m , $\mathcal{H}_m^- \cup \mathcal{H}_m^+ = \mathbb{R}^d$. Let N be the number of inner nodes of the decision tree; note that the decision tree has exactly $N + 1$ leaf nodes, since it is by definition a complete binary tree, see Figure 1 for an illustration. For a leaf node $\ell \in \{1, \dots, N + 1\}$ of the tree, let $\mathcal{P}_\ell^- \subset \{1, \dots, N\}$ (respectively \mathcal{P}_ℓ^+) be the set of all inner nodes whose left (respectively right) subtree contains ℓ , that is, $\mathcal{P}_\ell^+ \cup \mathcal{P}_\ell^-$ is the set of all parent nodes of ℓ . Then, the decision tree sorts an observation $x \in \mathbb{R}^d$ into its leaf \mathcal{R}_ℓ if and only if

$$x \in \mathcal{R}_\ell = \left(\bigcap_{m \in \mathcal{P}_\ell^-} \mathcal{H}_m^- \right) \cap \left(\bigcap_{m \in \mathcal{P}_\ell^+} \mathcal{H}_m^+ \right). \quad (1)$$

In fact, $\{\mathcal{R}_\ell\}_{\ell \in \mathcal{L}}$ is the feature space partition coded by the tree, see Figure 1 for an example. Finally, the tree returns the average response of all training samples that fall into the same leaf as the input data; let us call a_ℓ the average response of all training samples in \mathcal{R}_ℓ . The final prediction of the decision tree g can therefore be expressed as

$$g(x) = \sum_{\ell=1}^{N+1} a_\ell \mathbb{1}_{\{x \in \mathcal{R}_\ell\}}.$$

Let us now explore how an MLP can be designed to reproduce the prediction of a decision tree. Consider an MLP of depth 3 with N neurons on the first layer. For each inner node $m \in \{1, \dots, N\}$, the m -th neuron of the first layer indicates on which side of the split introduced by this inner node a given feature vector

lies: it equals -1 if the feature vector lies in \mathcal{H}_m^- and $+1$ if it lies in \mathcal{H}_m^+ . This can be achieved applying the following affine transformation and a sign activation function to the feature vector,

$$A_1 : x \in \mathbb{R}^d \mapsto x^{(i_m)} - t_m \quad \text{and} \quad \varphi_1 : x \mapsto \begin{cases} -1 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0. \end{cases}$$

The second layer of the 3-layer MLP has $N + 1$ neurons. For each leaf node $\ell \in \{1, \dots, N + 1\}$, the ℓ -th neuron of the second layer indicates whether a given feature vector $x \in \mathbb{R}^d$ lies in \mathcal{R}_ℓ or not: it equals $+1$ if $x \in \mathcal{R}_\ell$ and -1 if $x \notin \mathcal{R}_\ell$. Using equation (1), this can be achieved by applying the following affine transformation and a sign activation function to the output of the first layer,

$$A_2 : x \in \mathbb{R}^N \mapsto \sum_{m \in \mathcal{P}_\ell^+} x^{(m)} - \sum_{m \in \mathcal{P}_\ell^-} x^{(m)} - |\mathcal{P}_\ell^+ \cup \mathcal{P}_\ell^-| + \frac{1}{2} \quad \text{and} \quad \varphi_2 : x \mapsto \begin{cases} -1 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0. \end{cases}$$

The last layer of the MLP contains a single output neuron that returns the tree prediction. Using the output of the second layer, this can be achieved by applying the following affine transformation and an identity activation function,

$$A_3 : x \in \mathbb{R}^{N+1} \mapsto \frac{1}{2} \left(\sum_{\ell=1}^{N+1} x^{(\ell)} a_\ell + \sum_{\ell=1}^{N+1} a_\ell \right) \quad \text{and} \quad \varphi_3 : x \mapsto x \quad (2)$$

where a_ℓ is the average response of all training samples in \mathcal{R}_ℓ . Note that $\{a_\ell\}_{\ell=1}^{N+1}$ is a set of real numbers in regression problems and a set of probability vectors representing class distributions in classification problems. An illustration of the MLP translation of a decision tree is shown in Figure 1. This translation procedure is explained, for example, in Biau et al. (2019) with more details.

C Detail on the MLP translation accuracy

C.1 On the choice of hyper-parameters

In Section 2.3, four hyper-parameters were introduced to approximate the sign and identity functions through the layers of an elementary MLP. We address here the choice of the HPs and propose an optimal range for these parameters in the sense that they are as small as possible while guaranteeing a faithful MLP translation.

We focus on the analysis of deep forest translation, as the structure of all other tree-based methods can be seen as a truncated variant of a deep forest. The deep forest is trained and translated into an MLP on each data set (see Section 2) for different values of the HPs. To identify the influence of each HP, we make them vary in some range while the other three HPs are fixed to 10^{10} ,

resulting in an almost perfect approximation of the respective sign and identity functions. Figure 7 shows the predictive performance of a deep forest and its MLP translation playing with different HPs.




Figure 7: Comparison of the performance of a trained deep forest and its neural network encoding. Deep forest architecture: maximal depth of 8 per tree, 8 trees per forest, 1 forests per layer, 3 layers.

Figure 7 shows in particular that


- (i) increasing the HPs beyond some limit value is no longer beneficial as the activation functions are already perfectly approximated;
- (ii) across multiple data sets, these limit values are similar.

One could note that the coefficients in the first layer of a decision tree translation should be of a larger order of magnitude than those corresponding to the other activation functions to achieve an accurate translation. To give some insight into why this is the case, recall that the m -th neuron of the first layer determines whether the input vector belongs to \mathcal{H}_m^- or \mathcal{H}_m^+ , and note that its outputs can be of arbitrarily small size because the vector can be arbitrarily close to the decision boundaries. Note also that an MLP translation would better compromise on translation accuracy to ensure sufficient gradient flow. Based on these observations, we remark that choosing the HP of the following orders allows for maximum gradient flow while still providing an accurate translation: $\text{strength01} \in [1, 10^4]$, $\text{strength12} \in [10^{-2}, 10^2]$, $\text{strength23} \in [10^{-2}, 10^2]$ and $\text{strength_id} \in [10^{-2}, 10^2]$. This will actually help us later on to calibrate the search spaces when empirically tuning these HPs for each data set.

C.2 A fundamental numerical instability of the neural network encoding

The encoding of a decision tree by a neural network proposed in Section 2.3 is numerically unstable, i.e., it does not necessarily yield the same result as the tree itself, even when using the original, non-approximated activation functions. This is the result of a catastrophic cancellation that occurs within the MLP translation. The term catastrophic cancellation describes the remarkable loss of precision that occurs when two nearly equal numbers are numerically subtracted. For example, take the numbers $a = 1$ and $b = 10^{-10}$, and perform the computation $(a + b) - a$ on a machine with limited precision, say to 8 significant digits. The machine will return $(a + b) - a = 1 - 1 = 0$, although this result is clearly not correct. This phenomenon occurs in the third layer of the MLP encoding, see equation (2). The two sums calculated in this layer are almost equal in magnitude but have opposite signs, resulting in a catastrophic cancellation that has a greater impact the more partitions of the input space the decision tree uses, i.e. the deeper it is.


Figure 8 illustrates the effect of this phenomenon, comparing the mean approximation error between a simple decision tree and its neural network encoding on the airbnb data set. In Figure 8a, the result at the output layer of the tree was replaced by the exact training mean of the corresponding decision tree partition, compensating for the catastrophic cancellation. No such compensation was done for Figure 8b. This shows the grave implications of this instability: the mean error grows exponentially with the depth of an individual tree.




(a) replacing the output layer's result with the exact training mean of the corresp. tree partition (b) using the output layer's result with catastrophic cancellation

Figure 8: Illustration of the fundamental numerical instability of the decision tree encoding.

Although the errors introduced by this phenomenon may not be large for a given decision tree, they might accumulate when several such trees are composed, for example in Random or Deep Forests. Figure 9 compares the mean approximation error between Random/Deep Forests of different complexities and their corresponding neural net encoding on the Airbnb data set. It shows



(a) Random Forest.



(b) Deep Forest

Figure 9: Effects of numerical instabilities on more complex tree-based predictors. Airbnb data set. Random Forests are composed of trees of depth 7. Deep forest architecture: tree depth of 7, 5 trees per forest, 1 forest per layer and a variable number of layers.

that the composition of several trees in a cascade manner, as performed by the Deep Forest, leads to a stronger amplification of their individual inaccuracies than the parallel composition of trees, as performed by the Random Forest. This result is to be expected because decision trees composed in parallel do not influence each other’s predictions, whereas in a cascade architecture the results of the first layer of decision trees affect the input of the subsequent layers and inaccuracies can thus develop stronger effects.

We note that this catastrophic cancellation can be easily circumvented by introducing an additional layer. If this maps the output of the second layer from $\{-1, 1\}$ to $\{0, 1\}$, the last layer could then simply multiply each of these outputs by the average response of a partition set. However, Figure 9 also shows that the error introduced by the catastrophic cancellation remains relatively small, except for deep forests with many layers. Therefore, we did not immediately address this issue and planned to fall back on this analysis if the MLP coding did not produce the expected results later in our analysis. However, this somewhat imprecise MLP coding worked well for all our purposes.

D Supplements to numerical evaluations

D.1 Data sets

Data sets description In the sequel, we run numerical experiments on 10 real-world, heterogeneous, tabular data sets, all but two of which have already been used to benchmark deep learning methods, see Borisov et al. (2021); Somepalli et al. (2021). The chosen data sets represent a variety of different learning tasks and sample sizes. Tables 2 & 3 respectively give links to the platforms storing the data sets (four of them are available on the UCI Machine Learning Repository, Dua and Graff, 2017) and an overview of their main properties.

Data set	Link
Housing	Scikit-learn
Airbnb	Inside Airbnb
Diamond	OpenML
Adult	UCI Machine Learning Repository
Bank	UCI Machine Learning Repository
Blastchar	Kaggle
Heloc	FICO
Higgs	UCI Machine Learning Repository
Coverttype	UCI Machine Learning Repository
Volkert	AutoML

Table 2: Links to data sets.

	Housing	Airbnb	Diamonds	Adult	Bank	Blastchar	Heloc	Higgs	Coverttype	Volkert
Dataset size	20 640	119 268	53 940	32 561	45 211	7 043	9 871	550 000	581 012	58 310
# Num. features	8	10	6	6	7	3	21	27	44	147
# Cat. features	0	3	3	8	9	17	2	1	10	0
Task	Regr.	Regr.	Regr.	Classif.	Classif.	Classif.	Classif.	Classif.	Classif.	Classif.
# Classes	-	-	-	2	2	2	2	2	7	10

Table 3: Main properties of the data sets.

The Housing data set contains U.S. Census household attributes and the associated learning task is to predict the median house value for California districts (Kelley Pace and Barry, 1997). The Airbnb data set is provided by the company itself and holds attributes on different Airbnb listings in Berlin, such as the location of the apartment, the number of reviews, etc. The goal is to predict the price of each listing. Similarly, the diamond data set contains characteristics of different diamonds (e.g., carat weight or cut quality), and the goal is to predict the price of a diamond. The Adult data set contains Census information on adults (over 16-year olds) and its prediction task is to determine whether a person earns over \$50k a year. The Bank data set is related with direct marketing campaigns (phone calls) of a Portuguese banking institution, the classification goal is to predict whether the client will subscribe a term deposit. The Blastchar data set features information on customers of a fictional company that provides phone and internet services. The classification goal is to predict whether a customer cancels their contract in the upcoming month. The Heloc data set contains personal and credit record information on people that recently took on a line of credit, the classification task being to predict whether they will repay this credit within 2 years. On the Higgs data set (Baldi et al., 2014), the classification problem is to distinguish between signal processes that produce Higgs bosons and background processes that do not. For this purpose, it contains kinematic properties measured by the particle detectors in the accelerator that have been produced using Monte Carlo simulations. The Coverttype data set contains cartographic variables on forest cells and it’s task is to predict the

forest cover type. Finally, for the Volkert data set, different patches of the same size have been cut from images that belong to 10 different landscape scenes (coast, forest, mountain, plain, etc.). Each observation contains visual descriptors of one patch, the goal of this classification problem is to find the landscape type of the original picture.

D.2 Implementation details

RFs are implemented using `sklearn`'s `RandomForestRegressor` and `RandomForestClassifier` classes with default configuration for all parameters that are not mentioned explicitly. DFs are implemented using the `ForestLayer` library (Zhou and Feng, 2017) and GBDTs are implemented using the `XGBoost` library (Chen and Guestrin, 2016). MLPs are implemented and trained with `pytorch`, using the mean-squared error and the cross entropy as objective function for regression and classification problems respectively. The SAINT model is implemented using the library provided by Somepalli et al. (2021).

All methods are trained on a 32 GB RAM machine using 12 Intel Core i7-8700K CPUs, and one NVIDIA GeForce RTX 2080 GPU when possible (only the GDBT and MLP implementations including SAINT use the GPU). Hyper-parameter searches are parallelized on up to 4 of these machines.

Hyper-parameter optimization We tune all hyper-parameters using the `optuna` library (Akiba et al., 2019) with a fixed number of iterations for all models. In this context, an *iteration* corresponds to a set of hyper-parameters whose performance is evaluated with respect to a given method. The `optuna` library uses Bayesian optimization and, in particular, the tree-structured Parzen estimator model (Bergstra et al., 2011) to determine the parameters to be explored at each iteration of hyper-parameter optimization. This approach has been reported to outperform random search for hyper-parameter optimization (Turner et al., 2021).

Data pre-processing Machine learning pipelines often include pre-processing transformations of the input data before the training phase, a necessary step, especially when using neural networks (Bishop et al., 1995). We follow the pre-processing that is used in Borisov et al. (2021) and Somepalli et al. (2021). Hence, we normalize all continuous input features to zero mean and unit variance. This corresponds to linearly transform the input features as follows

$$\tilde{\mathbf{x}}_{:j} = \frac{\mathbf{x}_{:j} - \mu}{\sigma}$$

where $\mathbf{x}_{:j}$ is the j -th continuous feature of either train, validation or test observations, μ and σ are the mean and standard deviation calculated over the train set only. This way we assure that no information from the validation or test sets is used in the normalization step. Moreover, all categorical features are label encoded, i.e. each level of a categorical variable is replaced with an integer in $\{1, \dots, \# \text{ levels}\}$.

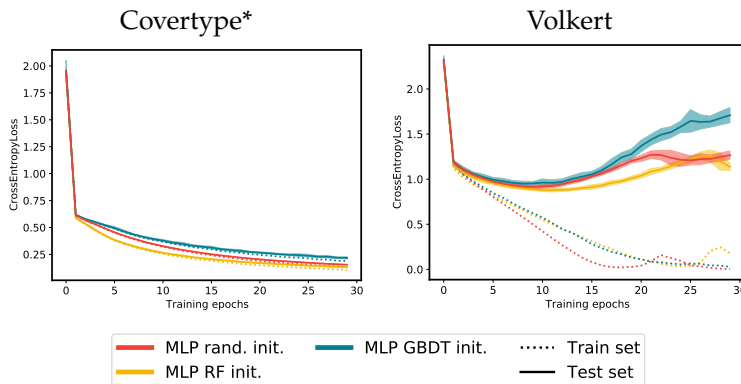


Figure 10: Optimization behaviour of randomly, RF and GBDT initialized MLP and SAINT evaluated over a 5 times repeated (stratified) 5-fold of each data set, according to Protocol P1, but where the MLP width is fixed to 2048 for all methods. The lines and shaded areas report the mean and standard deviation. *evaluation on a single 5-fold cross validation.

D.3 Working with an arbitrary width in P1 (optimization behaviour)

Figure 10 shows the optimization behaviour of the randomly, RF and GBDT initialized MLP on the multi-class classification problems. Note that in contrast to Figure 2 in this setting, which is less restrictive for RF initialization, this method does indeed lead to a *faster convergence* and a *better minimum* (in terms of generalization).

However, for these multi-class classification problems, the GBDT initialization tends to deteriorate the optimization compared to RF or random initialization methods. Indeed, RF are genuinely multiclassification predictors whose splits are built using all output classes simultaneously whereas splits in GBDT are only built following a one-vs-all strategy. This implies that, with a fixed budget of splits (and therefore of neurons), RF are likely to be more versatile than GBDT.

D.4 Additional material for Protocol P2 (generalization behaviour)

D.4.1 Extension of Table 1 (best performances)

Table 4 provides a comparison of the performances obtained by ourselves and the literature (where available) for each model. Notice that our results are broadly consistent with those in the literature, with two exceptions. First, our random initialized MLP tends to perform better than in the literature, which can be explained by the fact that we use a much larger search space than usual for the MLP width (see Section 3.5 for a discussion on this). Second, our performance on Higgs is significantly lower than in the literature. This can be ex-

Data set Model	Housing (†)		Airbnb	Diamonds	Covertypes (†)		Volkert (§)	
	MSE ↓		MSE ↓ × 10 ³	MSE ↓ × 10 ⁻³	Accuracy ↑ in %		Accuracy ↑ in %	
	perf. in literature	our results	our results	our results	perf. in literature	our results	perf. in literature	our results
Random Forest	0.272±0.006	0.263±0.009	5.39±0.13	9.80±0.35	78.1±0.1	83.6±0.1	66.3±1.3	64.2±0.3
GBDT	0.206±0.005	0.208±0.010	4.71±0.15	7.38±0.28	97.3±0.0	97.0±0.0	69.0±0.5	71.3±0.4
Deep Forest	-	0.225±0.008	4.68±0.16	8.23±0.29	-	92.4±0.1*	-	66.3±0.4
MLP rand. init.	0.263±0.008	0.258±0.011	5.07±0.16	15.5±12.5	91.0±0.4	<u>96.7±0.0</u>	63.0±1.56	72.2±0.4
MLP RF init.	-	0.222±0.009	4.66±0.16	7.93±0.22	-	96.7±0.0	-	74.1±0.4
MLP GBDT init.	-	0.206±0.007	4.70±0.09	8.15±0.35	-	96.2±0.0	-	73.5±0.5
MLP DF init.	-	0.234±0.016	4.81±0.13	8.28±0.24	-	94.5±0.3*	-	71.3±0.5
SAINT	0.226±0.004	0.258±0.011	4.81±0.15	17.7±3.83	96.3±0.1	96.6±0.1*	70.1±0.6	70.1±0.4

Data set Model	Adult (†)		Bank (§)		Blastchar (§)		Heloc (†)		Higgs (†)	
	AUC ↑ in %		AUC ↑ in %		AUC ↑ in %		AUC ↑ in %		AUC ↑ in %	
	perf. in literature	our results	perf. in literature	our results	perf. in literature	our results	perf. in literature	our results	perf. in literature	our results
Random Forest	91.7±0.2	91.6±0.3	89.1±0.3	92.8±0.3	80.6±0.7	84.5±1.2	90.0±0.2	91.3±0.6	79.7±0.0	80.4±0.1
GBDT	92.8±0.1	92.7±0.3	93.0±0.2	93.3±0.3	81.8±0.3	84.7±1.0	92.2±0.0	92.1±0.4	85.9±0.0	82.8±0.1
Deep Forest	-	91.8±0.3	-	92.9±0.2	-	83.7±1.2	-	90.3±0.5	-	81.2±0.0*
MLP rand. init.	90.3±0.2	90.5±0.4	91.5±0.2	91.0±0.3	59.6±0.3	81.4±1.2	80.3±0.1	80.1±0.1	85.6±0.0	83.2±0.3
MLP RF init.	-	<u>92.1±0.3</u>	-	<u>92.4±0.4</u>	-	84.4±1.2	-	91.7±0.4	-	83.6±0.1
MLP GBDT init.	-	<u>92.2±0.3</u>	-	<u>92.5±0.3</u>	-	84.6±1.2	-	91.5±0.6	-	83.0±0.0
MLP DF init.	-	91.9±0.4	-	92.2±0.3	-	84.2±1.0	-	91.4±0.6	-	83.3±0.1*
SAINT	91.6±0.4	91.6±0.3	93.3±0.1	92.2±0.4	84.7±0.3	84.0±0.8	90.7±0.2	90.2±0.7	88.3±0.0	83.7±0.1*

Table 4: Best scores for Protocol P2. For each data set, our best overall score is highlighted in **bold** and our best Deep Learning score is underlined. Our scores are based on 5 times repeated (stratified) 5-fold cross validation. For each of our models, HP were selected via the `optuna` library (100 iterations). Sources for literature values: Borisov et al. (2021) (†) and Somepalli et al. (2021) (§). *score based on a single 5-fold cross validation.

plained by the fact that we only include 5% of the original data set’s observations in our analysis due to hardware limitations that do not allow us to train large MLP on 11M samples.

D.4.2 Number of parameters of best neural networks

In Table 5, we compare the number of parameters of each NN method. Although the tree-based initialised MLP contain more parameters than the randomly initialized ones, the former are mostly sparse and the execution times are close (see Table 6). Finally note that the number of parameters of the RF/GBDT init. MLP is globally on par with that of SAINT (sometimes more, sometimes less) but for a smaller execution times (Table 6) and mostly better performances (Table 4).

D.4.3 Comparison of the execution times of the best neural networks

Table 6 presents a comparison of the execution times of the training of different NN methods using the hyper-parameters determined by the protocol P2. For each model, the total time during training is given up to the point where the best validation loss is reached (“early stopping”). It shows that RF/GBDT initialized MLP train faster than SAINT and a bit slower than randomly initialized MLP.

Model \ Data set	Housing	Airbnb	Diamonds	Adult	Bank	Blastchar	Heloc	Higgs	Covertypes	Volkert
MLP rand. init.	2.47M	1.86M	363k	1.09M	52.4K	13.1M	11.3M	11.6M	1.14M	9.03M
MLP RF init.	33.6M	12.6M	8.42M	29.4M	8.43M	25.2M	16.8M	4.26M	21.1M	17.1M
MLP GBDT init.	8.41M	12.6M	12.6M	33.6M	8.43M	16.8M	25.2M	8.46M	4.32M	21.3M
MLP DF init.	88.1M	34.0M	59.3M	42.0M	46.2M	34.36M	25.8M	43.2M	57.6M	34.1M
SAINT	56.8M	27.0M	53.1M	7.20M	6.12M	322M	98.2M	43.2M	6.44M	169M

Table 5: Comparison of number of parameters for each model.

Table 7 presents a comparison of the executions times of one model forward pass on the whole data set.

Model \ Data set	Housing	Airbnb	Diamonds	Adult	Bank	Blastchar	Heloc	Higgs	Covertypes	Volkert
MLP rand. init.	5.96 (32)	91.9 (98)	13.3 (31)	11.8 (37)	21.6 (62)	6.78 (34)	14.3 (60)	467 (32)	312 (69)	12.3 (31)
MLP RF init.	35.6 (29)	129 (44)	24.6 (25)	16.0 (19)	19.6 (23)	5.77 (18)	4.83 (15)	247 (39)	2040 (91)	24.6 (25)
MLP GBDT init.	17.5 (49)	276 (95)	48.9 (37)	32.2 (31)	3.20 (3)	1.69 (7)	3.89 (8)	58.8 (5)	435 (66)	48.9 (37)
MLP DF init.	218 (72)	355 (48)	175 (31)	91 (54)	96 (26)	22.3 (52)	9.04 (19)	3260 (76)	5570 (95)	175 (31)
SAINT	81.9 (37)	640 (83)	394 (84)	15.6 (11)	52.7 (32)	7.23 (2)	51.0 (31)	2310 (19)	6580 (97)	394 (84)

Table 6: Comparison of the execution time in seconds for model training until the best validation lost is reached. The number of training epochs is indicated in parentheses.

Model \ Data set	Housing	Airbnb	Diamonds	Adult	Bank	Blastchar	Heloc	Higgs	Covertypes	Volkert
MLP rand. init.	13.4	103	12.2	98.7	141	10.4	36.9	1060	1070	160
MLP RF init.	45.0	97.7	42.8	125	180	14.5	12.6	1020	1830	191
MLP GBDT init.	19.6	167	22.1	147	182	11.7	12.5	378	1170	218
MLP DF init.	102	227	105	200	415	28.5	13.4	4680	4060	330
SAINT	129	508	130	304	429	77.9	114	2870	3670	645

Table 7: Comparison of the execution time in milliseconds of one model forward pass on the whole data set.

D.4.4 Optimization behaviour

For completeness, Figure 11 shows the optimization behaviour of the randomly, RF and GBDT initialized MLP as well as SAINT.

D.5 Hyper-parameter setting

D.5.1 Search spaces

Table 8 shows the HP search spaces that were used to determine an optimal HP setting. The same search spaces were used for the experimental protocols P1 and P2. Note that, in Table 8, n.classes corresponds to the number of classes for classification problems and is 1 for regression problems. Furthermore, the different search spaces given for SAINT were used for smaller/larger data sets, where a data set qualifies as smaller if it has less than 50 explanatory variables.

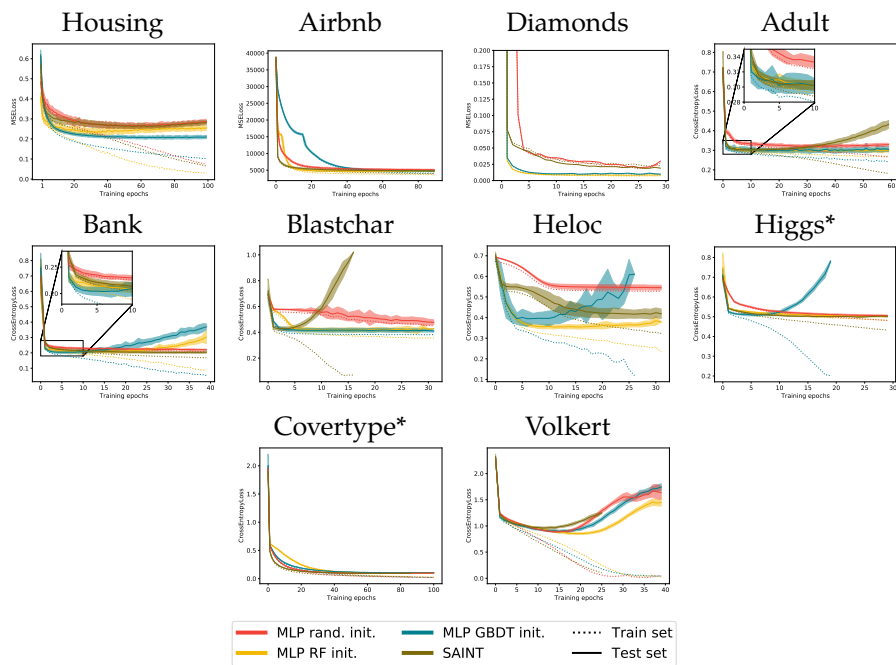


Figure 11: Optimization behaviour of randomly, RF and GBDT initialized MLP and SAINT evaluated over a 5 times repeated (statisfied) 5-fold of each data set, according to Protocol P2. The lines and shaded areas report the mean and standard deviation. *evaluation on a single 5-fold cross validation.

Method	Parameter	Search space	Function
Random Forests	max_depth	$\{1, \dots, 12\}$	see here
	n_estimators	$\{1000\}$	
	max_features	$[0, 1]$	
GBDT	max_depth	$\{1, \dots, 12\}$	see here
	n_estimators	$\{1000\}$	
	reg_alpha	$[10^{-8}, 1]$	
	reg_lambda	$[10^{-8}, 1]$	
	learning_rate	$[0.01, 0.3]$	
Deep Forest	forest_depth	$\{1, 2, 3\}$	Number of Deep Forest layers
	n_forests	$\{1\}$	Number of forests per Deep Forest layer
	n_estimators	$\{1000\}$	RF parameters, see here
	max_depth	$\{1, \dots, 12\}$	
	max_features	$[0, 1]$	
MLP random init.	learning_rate	$[10^{-6}, 10^{-1}]$	learning rate of SGD training
	depth	$\{1, \dots, 10\}$	number of layer
	width	$\{1, \dots, 2048\}$	number of neurons per layer
	epochs	$\{100\}$	number of SGD training epochs
	batch_size	$\{256\}$	batch size of SGD training
MLP RF init.	max_depth	$\{1, \dots, 11\}$	Parameters of the RF initializer, see here
	n_estimators	$2048/2^{\max_depth}$	
	max_features	$[0, 1]$	learning rate of SGD training number of layer number of neurons per layer number of SGD training epochs batch size of SGD training
	learning_rate	$[10^{-6}, 10^{-1}]$	
	depth	$\{3, \dots, 10\}$	
	width	$\{2048\}$	
	epochs	$\{100\}$	
	batch_size	$\{256\}$	MLP translation parameters, see Section 2.3
	strength01	$[1, 10^4]$	
	strength12	$[0.01, 100]$	
MLP GBDT init.	max_depth	$\{1, \dots, 11\}$	Parameters of the GBDT initializer, see here
	n_estimators	$2048/(n_classes \cdot 2^{\max_depth})$	
	reg_alpha	$[10^{-8}, 1]$	learning rate of SGD training number of layer number of neurons per layer number of SGD training epochs batch size of SGD training
	reg_lambda	$[10^{-8}, 1]$	
	learning_rate_GBDT	$[0.01, 0.3]$	
	learning_rate	$[10^{-6}, 10^{-1}]$	
	depth	$\{3, \dots, 10\}$	
	width	$\{2048\}$	MLP translation parameters, see Section 2.3
	epochs	$\{100\}$	
	batch_size	$\{256\}$	
strength01	$[1, 10^4]$		
strength12	$[0.01, 100]$		
MLP DF init.	forest_depth	$\{1, 2, 3\}$	Number of Deep Forest layers
	n_forests	$\{1\}$	Number of forests per Deep Forest layer
	n_estimators	$2048/2^{\max_depth}$	RF parameters, see here
	max_depth	$\{1, \dots, 12\}$	
	max_features	$[0, 1]$	
	learning_rate	$[10^{-6}, 10^{-1}]$	learning rate of SGD training
	depth	$\{3, \dots, 10\}$	number of layer
	width	$\{2048\}$	number of neurons per layer
	epochs	$\{100\}$	number of SGD training epochs
	batch_size	$\{256\}$	batch size of SGD training
	strength01	$[1, 10^4]$	MLP translation parameters, see Section 2.3
	strength12	$[0.01, 100]$	
	strength23	$[0.01, 100]$	
strength_id	$[0.01, 100]$		
epochs	$\{100\}$	number of SGD training epochs	
batch_size	$\{256\}/\{64\}$	batch size of SGD training	
SAINT	dim	$[32, 64, 128]/[8, 16]$	number of neurons per layer in attention block
	depth	$\{1, 2, 3\}$	number of layers in each attention block
	heads	$\{2, 4, 8\}$	number of head in each attention layer
	dropout	$\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$	dropout used during SGD training

Table 8: Hyper-parameter search spaces used for numerical evaluations.

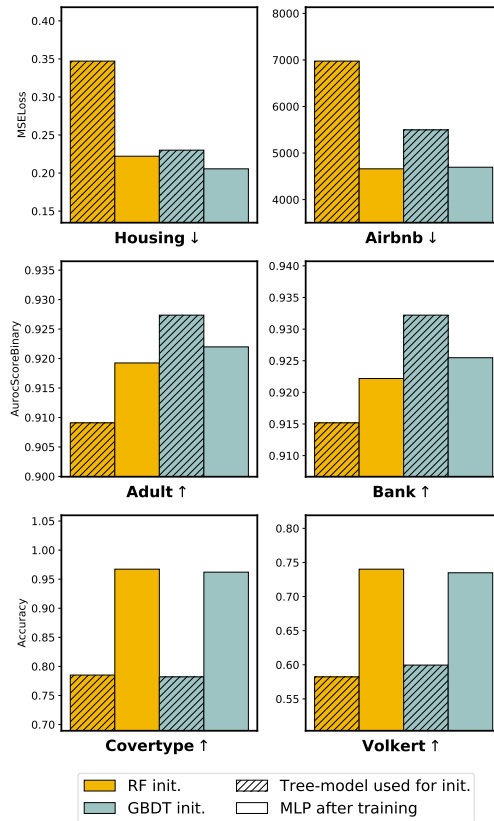


Figure 12: Comparison of the performance of the RF and GBDT models used for initialization and the final performance of the corresponding MLPs.

D.5.2 Experimental protocol P2

Tables 9 and 10 show the HP setting used for the experimental protocol P2. For the search spaces and descriptions of the function of each HP see Table 8.

D.6 Performances of tree-based methods used for initialization of MLP

Figure 12 compares the performance of RF and GBDT models and the performance of optimized MLP, initialized with RF and GBDT respectively. We can notice that the difference in performance between GBDT and RF does not systematically turn into the same difference in performance for the corresponding trained networks. This suggests that beyond their respective performances, the very structures of RF and GBDT predictors play an important role in the final MLP performances.

Method	Parameter	Housing	Airbnb	Adult	Bank	Covertime	Volkert
Random Forests	max_depth	12	12	11	12	12	12
	n_estimators	1000	1000	1000	1000	1000	1000
	max_features	0.437	0.623	0.596	0.943	0.811	0.688
GBDT	max_depth	12	9	6	7	11	10
	n_estimators	1000	1000	1000	1000	1000	1000
	reg_alpha	0.305	4.60×10^{-6}	2.39×10^{-5}	1.52×10^{-4}	0.728	4.47×10^{-6}
	reg_lambda	1.13×10^{-2}	1.75×10^{-8}	1.35×10^{-6}	1.07×10^{-3}	6.51×10^{-4}	1.71×10^{-6}
	learning_rate	3.82×10^{-2}	0.238	1.08×10^{-2}	1.34×10^{-2}	0.181	0.107
Deep Forest	forest_depth	4	9	2	2	9	3
	n_forests	1	1	1	1	1	1
	n_estimators	1000	1000	1000	1000	1000	1000
	max_depth	5	12	11	9	12	12
	max_features	0.361	0.410	0.166	0.206	0.218	0.134
MLP random init.	learning_rate	9.01×10^{-4}	4.21×10^{-4}	2.07×10^{-4}	1.1×10^{-4}	1.15×10^{-4}	2.29×10^{-4}
	depth	4	4	4	4	4	6
	width	1100	959	1175	856	738	1482
	epochs	100	100	100	100	100	100
	batch_size	256	256	256	256	256	256
MLP RF init.	max_depth	8	10	8	8	10	8
	n_estimators	8	2	8	8	2	8
	max_features	0.442	0.321	0.613	0.650	0.897	0.825
	learning_rate	1.04×10^{-4}	1.72×10^{-4}	1.55×10^{-5}	1.01×10^{-4}	1.04×10^{-5}	1.45×10^{-4}
	depth	10	5	5	4	7	6
	width	2048	2048	2048	2048	2048	2048
	epochs	100	100	100	100	100	100
	batch_size	256	256	256	256	256	256
	strength01	1090	668	537	71.4	13.7	1.02
	strength12	0.0749	1.09	62.7	34.5	1.05×10^{-2}	5.53×10^{-2}
MLP GBDT init.	max_depth	3	4	4	4	8	4
	n_estimators	256	128	128	128	1	12
	reg_alpha	1.30×10^{-7}	1.10×10^{-2}	1.26×10^{-8}	0.413	1.33×10^{-2}	6.76×10^{-6}
	reg_lambda	1.57×10^{-7}	9.52×10^{-4}	7.85×10^{-4}	7.48×10^{-3}	0.643	1.99×10^{-7}
	learning_rate_GBDT	0.211	0.297	0.202	0.285	0.112	0.272
	learning_rate	1.11×10^{-5}	1.97×10^{-5}	4.77×10^{-5}	6.22×10^{-4}	6.19×10^{-5}	1.63×10^{-4}
	depth	4	5	6	4	3	7
	width	2048	2048	2048	2048	2048	2048
	epochs	100	100	100	100	100	100
	batch_size	256	256	256	256	256	256
	strength01	575	7830	132	20.5	7280	4.08
strength12	5.60	0.461	66.0	5.52	93.4	7.11×10^{-2}	
MLP DF init.	forest_depth	6	3	3	2	3	2
	n_forests	1	1	1	1	1	1
	n_estimators	16	2	64	32	2	8
	max_depth	7	10	5	6	10	8
	max_features	0.350	0.598	0.992	0.322	0.633	0.342
	learning_rate	1.04×10^{-5}	6.67×10^{-5}	1.54×10^{-5}	3.08×10^{-5}	1.58×10^{-5}	2.31×10^{-4}
	depth	23	10	9	13	15	9
	width	2048	2048	2048	2048	2048	2048
	epochs	100	100	100	100	100	100
	batch_size	256	256	256	256	256	256
	strength01	515	36.6	41.0	15.5	51.6	1.41
	strength12	0.162	0.242	10.6	0.213	0.124	0.154
	strength23	1.94	10.4	47.8	1.94	4.26×10^{-2}	0.149
strength_id	3.63×10^{-2}	6.34×10^{-2}	7.44	2.75×10^{-2}	5.09×10^{-2}	3.69	
SAINT	epochs	100	100	100	100	100	100
	batch_size	256	256	256	256	64	256
	dim	128	64	32	32	8	16
	depth	3	2	2	1	2	2
	heads	2	8	2	8	4	8
	dropout	0.2	0	0.4	0.8	0.5	0.8

Table 9: Hyper-parameters used for the experimental protocol P2.

Method	Parameter	Diamonds	Blastchar	Heloc	Higgs
Random Forests	max_depth	12	6	9	12
	n_estimators	1000	1000	1000	1000
	max_features	0.967	0.547	0.607	0.577
GBDT	max_depth	7	1	1	11
	n_estimators	1000	1000	1000	1000
	reg_alpha	0.341	7.15×10^{-7}	0.123	2.29×10^{-8}
	reg_lambda	5.15×10^{-4}	1.59×10^{-7}	1.44×10^{-2}	0.391
	learning_rate	9.17×10^{-2}	1.48×10^{-2}	0.282	2.46×10^{-2}
Deep Forest	forest_depth	4	7	10	3
	n_forests	1	1	1	1
	n_estimators	1000	1000	1000	1000
	max_depth	12	2	4	12
	max_features	0.454	0.641	0.196	0.163
MLP random init.	learning_rate	2.35×10^{-4}	1.05×10^{-4}	1.14×10^{-6}	2.26×10^{-5}
	depth	9	8	8	9
	width	1011	1475	1369	1284
	epochs	100	100	100	100
	batch_size	256	256	256	256
MLP RF init.	max_depth	10	5	7	9
	n_estimators	2	64	16	4
	max_features	0.904	0.425	0.728	0.670
	learning_rate	6.67×10^{-5}	5.07×10^{-6}	7.33×10^{-6}	2.17×10^{-5}
	depth	4	8	6	3
	width	2048	2048	2048	2048
	epochs	100	100	100	100
	batch_size	256	256	256	256
	strength01	19.8	4500	331	1.43
strength12	0.420	42.9	1.06	0.329	
MLP GBDT init.	max_depth	3	1	3	5
	n_estimators	256	1024	256	64
	reg_alpha	4.56×10^{-2}	1.63×10^{-5}	6.21×10^{-7}	2.58×10^{-6}
	reg_lambda	6.17×10^{-4}	2.19×10^{-4}	3.03×10^{-4}	3.20×10^{-6}
	learning_rate_GBDT	0.214	4.72×10^{-2}	8.42×10^{-2}	0.290
	learning_rate	8.94×10^{-5}	5.60×10^{-6}	4.54×10^{-4}	1.36×10^{-4}
	depth	5	6	8	4
	width	2048	2048	2048	2048
	epochs	100	100	100	100
	batch_size	256	256	256	256
strength01	3870	4690	6550	4780	
strength12	56.6	21.0	31.8	0.423	
MLP DF init.	forest_depth	3	2	2	3
	n_forests	1	1	1	1
	n_estimators	4	128	64	8
	max_depth	9	4	5	8
	max_features	0.695	0.516	0.280	0.572
	learning_rate	2.04×10^{-5}	2.00×10^{-6}	1.91×10^{-5}	9.33×10^{-6}
	depth	16	10	8	12
	width	2048	2048	2048	2048
	epochs	100	100	100	100
	batch_size	256	256	256	256
	strength01	21.0	93.0	97.8	1.12
	strength12	0.119	20.0	0.987	9.22×10^{-2}
	strength23	5.34×10^{-2}	0.283	27.1	0.207
strength_id	0.358	0.475	9.70	0.152	
SAINT	epochs	32100	100	100	100
	batch_size	256	256	256	64
	dim	64	128	64	16
	depth	3	3	3	2
	heads	4	8	2	8
dropout	0.2	0.5	0.8	0.8	

Table 10: Hyper-parameters used for the experimental protocol P2.

D.7 Additional Figures to Section 3.5 (Analyzing key elements of the new initialization methods)

Figures 13 and 14 show the same histograms as Figure 4 evaluated on the other data set considered in protocol P2. Note the logarithmic y-axis for the first two RF and GBDT initialized layers.




Figure 13: Histograms of the first three first and the last layers' weights before and after the MLP training on the Airbnb, Diamonds, Adult, Bank and Blastchar data sets. Comparison between random, RF and GBDT initializations.




Figure 14: Histograms of the first three first and the last layers' weights before and after the MLP training on the Heloc, Higgs, Coverttype and Volkert data sets. Comparison between random, RF and GBDT initializations.