

Geodesic Flow Kernels for Semi-Supervised Learning on Mixed-Variable Tabular Dataset

Yoontae Hwang¹, Yongjae Lee²,

¹University of Oxford

²Ulsan National Institute of Science and Technology (UNIST)

yoontae.hwang@eng.ox.ac.uk, yongjaelee@unist.ac.kr

Abstract

Tabular data poses unique challenges due to its heterogeneous nature, combining both continuous and categorical variables. Existing approaches often struggle to effectively capture the underlying structure and relationships within such data. We propose GFTab (Geodesic Flow Kernels for Semi-Supervised Learning on Mixed-Variable Tabular Dataset), a semi-supervised framework specifically designed for tabular datasets. GFTab incorporates three key innovations: 1) Variable-specific corruption methods tailored to the distinct properties of continuous and categorical variables, 2) A Geodesic flow kernel based similarity measure to capture geometric changes between corrupted inputs, and 3) Tree-based embedding to leverage hierarchical relationships from available labeled data. To rigorously evaluate GFTab, we curate a comprehensive set of 21 tabular datasets spanning various domains, sizes, and variable compositions. Our experimental results show that GFTab outperforms existing ML/DL models across many of these datasets, particularly in settings with limited labeled data.

Code — <https://github.com/Yoontae6719/Geodesic-Flow-Kernels-for-Semi-Supervised-Learning-on-Mixed-Variable-Tabular-Dataset>

Extended manuscript — <https://arxiv.org/abs/2412.12864>

Introduction

One of the long-standing goals in ML has been to discover structure and pattern in data. While these efforts have been very successful in handling images, audio, and text, there is still room for improvement in handling tabular data. There are *three main reasons* why handling tabular data is difficult. One possible reason for this is that, unlike images and text, which share certain structural aspects within Euclidean space, tabular data do not appear to have such solid structural similarities. For instance, in an image, adjacent pixels are often observed to have similar color values and belong to the same object. However, in a tabular dataset, adjacent rows or columns may not necessarily have any inherent relationship or similarity. In most cases, the nature of tabular data would not change at all even if the order of rows and columns are shuffled.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Second, tabular datasets may present certain challenges due to the presence of both continuous and categorical variables, each with distinct characteristics. Continuous variables are characterized by an infinite number of potential values within a specified range, whereas categorical variables are typically represented by values from a finite set. Moreover, while continuous variables can be ordered, categorical variables cannot be ordered in general. It would be beneficial to consider the inherent differences between these variable types (Gorishniy et al. 2021), rather than treating them identically as is often done in previous studies. Additionally, in real world, the proportion of continuous and categorical variables can vary considerably across different datasets, which can make it challenging to develop approaches that are widely applicable. This heterogeneity contributes to the domain-specific nature of tabular research and hinders the creation of universally applicable methods. It is possible that studies on tabular datasets have overlooked these complexities, evaluating performance only on benchmark data with predominantly continuous variables, which may not fully represent real-world scenarios.

Finally, real-world tabular datasets often lack labels, which are crucial for supervised learning tasks. In the absence of labels, discovering meaningful patterns and structures within the data becomes an even more daunting task. Without labels to guide the learning process, the model must rely on the intrinsic structure and relationships within the data to uncover relevant patterns, which can be particularly difficult given the lack of inherent structural similarities and the heterogeneous nature of tabular datasets.

Some studies (Bahri et al. 2022; Ucar, Hajiramezanali, and Edwards 2021; Yoon et al. 2020) have demonstrated encouraging outcomes in addressing tabular datasets with limited labeled data. These methods employ corruption techniques to maximize the similarity between the original input and corrupted versions, thereby enabling the models to learn meaningful representations. However, these methods disregard variable types and have primarily been evaluated on datasets containing only continuous variables. Consequently, their effectiveness may be limited when dealing with real-world tabular datasets that often contain a mix of continuous and categorical variables (Borisov et al. 2022).

In this paper, we introduce GFTab: **Geodesic Flow Kernel** on a **Mixed-variable Tabular** dataset, an approach de-

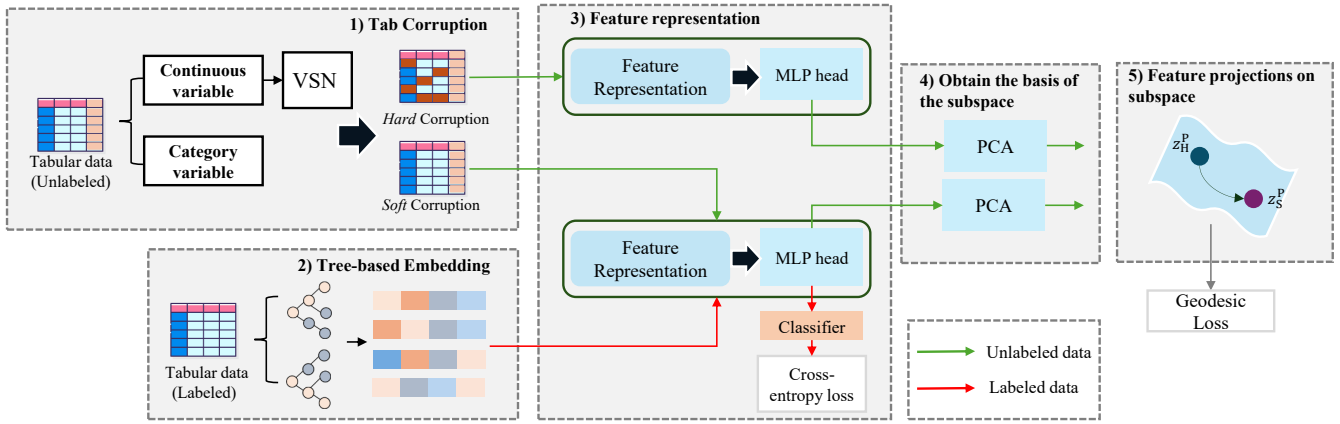


Figure 1: The proposed model (GFTab) is a semi-supervised framework specifically designed for handling tabular data.

signed to address the unique challenges inherent in tabular data, particularly in semi-supervised learning settings. GFTab consists of three key components:

- Variable-specific corruption methods that account for the distinct characteristics of continuous and categorical variables.
- Geodesic flow kernel that captures the geometric relationships between corrupted representations on the Grassmann manifold.
- Tree-based embedding technique that leverages labeled data to learn hierarchical feature relationships.

By integrating these components, GFTab is able to effectively learn robust representations of tabular data, even in scenarios with limited labeled samples or noisy labels. Our framework is designed to be applicable across a wide range of tabular datasets and domains. In the following sections, we describe each component of GFTab in detail, present our proposed benchmark datasets for evaluating tabular learning methods, and demonstrate GFTab’s effectiveness through extensive experiments.

Related Works

Tabular DL Models. Due to the importance of tabular data in real-world industries, ML/DL models for tabular data have been actively studied recently. The TabTransformer (Huang et al. 2020) generated embeddings on categorical variables to learn the context of tabular data in a robust way, and it performed well in terms of prediction Accuracy. However, they only used a simple MLP layer for continuous variables. AutoInt (Song et al. 2019) is designed to specifically learn the interactions between features in high-dimensional tabular data sets. The FT-transformer (Gorishniy et al. 2021) used Transformer layers, but it treated all features in the same way regardless of their variable types (categorical and continuous). Furthermore, it becomes computationally inefficient when there are a large number of features. NODE (Popov, Morozov, and Babenko 2020) proposed a method for training deep ensembles of oblivious differentiable decision trees on tabular data and demonstrated strong per-

formance on only a few datasets. The TabGNN (Guo et al. 2021) modeled the relationships within the tabular data samples using graph neural networks in an explicit and systematic manner. However, its numerical experiments are limited to large-sized tabular datasets only. NPTs (Kossen et al. 2021) used self-attention to explicitly infer the relationships between data points like TabGNN. However, it is computationally expensive and has only been demonstrated on a small number of datasets. The GRANDE (Marton et al. 2024) model proposed an ensemble learning method to integrate multiple neural networks. TabPFN (Hollmann et al. 2023) introduced a probabilistic approach for tabular data, leveraging Bayesian principles to provide robust uncertainty estimates and improve prediction accuracy. However, to the best of our knowledge, most of the models discussed above have been tested using datasets with 100% labels, which is not representative of real world applications.

Tabular Self- and Semi-Supervised Learning Self-supervised learning techniques have been actively studied in computer vision (Chen et al. 2020; Jing and Tian 2020; Zhai et al. 2019) and natural language processing (Lee, Chang, and Toutanova 2019; Qiu et al. 2020; Ruder and Plank 2018; Song et al. 2020), but not much research has been conducted for tabular data. The TABBIE (Iida et al. 2021), TabBERT (Yin et al. 2020) and TaPas (Herzig et al. 2020) were developed to learn cell-wise, column-wise, and row-wise representations, respectively, in order to understand complex table semantics or numerical trends from PDF files or tables. The SAINT (Somepalli et al. 2021), TabNet (Arik and Pfister 2021) and VIME (Yoon et al. 2020) were designed to be able to restore the original values of data samples after being injected with various types of noise. However, these works did not make a distinction between continuous and categorical variables. SCARF (Bahri et al. 2022) was inspired by popular techniques used in computer vision, such as color distortion, random cropping, and blurring. It used these methods to corrupt a random subset of features and form a view for tabular data. SubTab (Ucar, Hajiramezani, and Edwards 2021) introduced a new framework that divides input features into multiple subsets and reconstructs

the data from these subsets in an autoencoder setting. Similar to the other models, SCARF and SubTab also applied the same corruption methods regardless of variable type (continuous or categorical). In contrast to existing models, our model uses different corruption methods for continuous and categorical variables that are specifically designed to reflect the characteristics of each variable type. The VIME was the first to propose a semi-supervised learning framework for tabular data based on consistency regularization. Furthermore, to the best of our knowledge, there are no other semi-supervised learning studies for tabular data. Although it used different loss functions for continuous and categorical variables in the reconstruction and mask vector estimation, it still applied the same corruption method to both variable types as discussed before.

Methodology

When learning from mixed-variable tabular dataset in a semi-supervised manner, a fundamental question arises:

How should we design inductive biases for tabular data?

We propose GFTab (Geodesic Flow Kernel on Mixed-variable Tabular data) to exploit the unique characteristics of mixed-variable tabular datasets. Our framework includes: 1) Variable-specific corruption methods for continuous and categorical variables, 2) Geodesic flow kernel based similarity measure to capture geometric changes between corrupted inputs, and 3) Tree-based embedding to learn hierarchical relationships from labeled data.

The GFTab Framework I : Corruption Process

Semi-Supervised Setting Suppose that there are N_l labeled samples $D_l = \{x_i, y_i\}_{i=1}^{N_l} \subseteq \mathbb{R}^M \times \mathbb{R}$ and N_u unlabeled samples $D_u = \{x_i\}_{i=1}^{N_u} \subseteq \mathbb{R}^M$, with N_u significantly greater than N_l . In this scenario, $x_i \in \mathbb{R}^M$ represents an example, and $y_i \in \mathbb{R}$ is its corresponding label. We simplify by considering y_i as a scalar, representing a class $k \in \mathcal{K}$. For analytical reasons, we introduce y_i^k as a binary indicator, which is set to 1 if $y_i = k$ and 0 otherwise. In this work, x_i is broken down into continuous components $x_i^{\text{cont}} \in \mathbb{R}^{M_{\text{cont}}}$ and categorical components $x_i^{\text{cat}} \in \mathbb{R}^{M_{\text{cat}}}$. For convenience, we will omit the subscript i in the following discussions.

Continuous Variable Corruption For continuous variables, we firstly employ a variable selection network (VSN) to extract an effective feature representation. The VSN consists of two Gated Residual Networks (GRNs) (Lim et al. 2021): one for each element $x^{m,\text{cont}}$ of x^{cont} (parameterized by θ_1), and the other for the entire x^{cont} (parameterized by θ_2). The output of the first GRN is denoted as $\xi^m = \text{GRN}_{(\theta_1)}(x^{m,\text{cont}})$, while the output of the second GRN is passed through a softmax function to obtain importance weights $v = \text{Softmax}(\text{GRN}_{(\theta_2)}(x^{\text{cont}}))$. The m -th element of v , denoted as v^m , serves as an importance weight for the corresponding continuous variable $x^{m,\text{cont}}$. See Appendix C for empirical validation of this feature selection capability. The final representation of continuous variables, Ξ^{cont} , is ob-

tained by combining the weighted outputs of the first GRN:

$$\Xi^{\text{cont}} = \sum_{m=1}^{M_{\text{cont}}} v^m \xi^m \quad (1)$$

To create corrupted views of continuous variables, we generate soft and hard *views*, $x_{\text{soft}}^{\text{cont}}$ and $x_{\text{hard}}^{\text{cont}}$ by randomly shuffling the row and columns within the Ξ^{cont} . Here, we define the permutation matrix \mathbf{P} of size $M_{\text{cont}} \times M_{\text{cont}}$.

$$\begin{aligned} x_{\text{soft}}^{\text{cont}} &= \lambda \Xi^{\text{cont}} + (1 - \lambda) \Xi^{\text{cont}} \mathbf{P} \\ x_{\text{hard}}^{\text{cont}} &= (1 - \lambda) \Xi^{\text{cont}} + \lambda \Xi^{\text{cont}} \mathbf{P} \end{aligned} \quad (2)$$

In this case, the equation (2) generate corrupted *views* of the continuous features Ξ^{cont} . The soft view $x_{\text{soft}}^{\text{cont}}$ mixes the original features with a randomly permuted version, while the hard view $x_{\text{hard}}^{\text{cont}}$ puts more emphasis on the permuted features. Thus, the soft views retains more of the original structure, while the hard views is more randomized.

Categorical Variable Corruption Unlike continuous variables, Because ordered categories allow for the concept of relative closeness, unordered categories lack this inherent structure. categorical variables present unique challenges "how to corrupt categorical variable?" due to their discrete nature and the distinction between ordered and unordered types. The motivation behind this and related analysis results can be found in Figure ??.

Also, existing methods often overlook this distinction, leading to suboptimal results. For example, simply injecting noise (Bahri et al. 2022; Ucar, Hajiramezanali, and Edwards 2021) or estimating masking (Yoon et al. 2020) may not sufficiently capture their complexity, especially when dealing with imbalanced distributions or variables with identical discrete values across different items. To address this, we propose a corruption technique tailored specifically to categorical variables. For *categorical* variables, we determine the corrupted ratio of each variable. Let $k = [k^1, \dots, k^{M_{\text{cat}}}]$ be the mask vector where k^m 's are independently sampled from over the set $[-s^m(\gamma), s^m(\gamma)] \setminus \{0\}$. Here, $s^m(\gamma)$ is size of the neighborhood for k^m with corruption rate γ . Sampling from the set $[-1, 1] \setminus \{0\}$ would result in the most corruption. We then create soft and hard *views* for categorical variables as follows:

$$\begin{aligned} x_{\text{soft}}^{\text{cat}} &= x^{\text{cat}} + \text{nh}_{s(\gamma)}(k) \\ x_{\text{hard}}^{\text{cat}} &= x^{\text{cat}} + \text{nh}_{s=1}(k) \end{aligned} \quad (3)$$

Where the neighborhood function $\text{nh}_{s(\gamma)}$ is defined as a random function that shifts each category c to c' within a neighborhood of size $s(\gamma)$. The size $s(\gamma)$ is determined such that the probability of any category shift from c to c' matches the corruption rate γ . The following remark outlines how the neighborhood size is determined when the corruption rate is set arbitrarily. If the shifted value c' falls outside the range $[0, n]$, it is reset to c .

Remark: Consider a categorical variable with $n \geq 2$ categories. To achieve a corruption rate of at least γ , the minimum size s of the neighborhood $\text{nh}_{s(\gamma)}$ must satisfy the inequality $s \geq \lceil 2n(1 - \gamma) - 1 \rceil$.

Feature Representation Module. After creating separate views for continuous and categorical variables, the feature

representation module aims to capture the relationships between all variables to learn a more comprehensive representation of the data. To achieve this, we combine the information from the soft and hard views of the data samples, denoted as $\mathbf{x}_{\text{soft}} = (\mathbf{x}_{\text{soft}}^{\text{cont}}, \mathbf{x}_{\text{soft}}^{\text{cat}})$ and $\mathbf{x}_{\text{hard}} = (\mathbf{x}_{\text{hard}}^{\text{cont}}, \mathbf{x}_{\text{hard}}^{\text{cat}})$, respectively. We apply the Transformer layers to the column embeddings $\mathbf{e}(x_{\text{soft}}) \in \mathbb{R}^{d_{\text{emb}}}$ and $\mathbf{e}(x_{\text{hard}}) \in \mathbb{R}^{d_{\text{emb}}}$. First, a self-attention (Vaswani et al. 2017) aggregates column embeddings with normalized importance:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{d_{\text{attn}}}\right)\mathbf{V} \quad (4)$$

Here, $\mathbf{Q} = \mathbf{W}_Q \mathbf{e} \in \mathbb{R}_{\text{attn}}^{d_{\text{emb}}}$, $\mathbf{K} = \mathbf{W}_K \mathbf{e} \in \mathbb{R}_{\text{attn}}^{d_{\text{emb}}}$ and $\mathbf{V} = \mathbf{W}_V \mathbf{e} \in \mathbb{R}_V^{d_{\text{emb}}}$ are queries, keys, and values, respectively. \mathbf{e} is the collection of column embeddings, \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are learnable weight matrices. In actual implementation, multi-head attention is used to further improve the feature extraction. The output is then transformed back into an embedding of dimension d_{lin} through a fully connected layer. Then, the representation embedding $z_{\text{hard}} \in \mathbb{R}^{d_{\text{lin}}}$ and $z_{\text{soft}} \in \mathbb{R}^{d_{\text{lin}}}$ are obtained after passing through feed-forward layers.

The GFTab Framework II: Geodesic Flow Kernel

The motivation for using a geodesic flow kernel in mixed-variable tabular data lies in the inherent challenges of measuring distances. In the context of tabular data, the distance between two points cannot always be easily or meaningfully quantified using conventional distance metrics such as InfoNCE (Oord, Li, and Vinyals 2018), Barlow (Zbontar et al. 2021) and so on.

For example, a small change in a continuous variable might not significantly affect the overall relationship between data points. However, a single change in a categorical variable could result in a completely different context or classification. This variability necessitates a more sophisticated approach to capturing the relationships between data points. The related analysis results can be found in Figure ??.

In this section, we address this by modeling geodesic flow kernel. For this, we leverage the assumption that data lie in a low-dimensional linear subspace to derive features from both z_{soft} and z_{hard} , modeling them with a basis $\mathbf{P} \in \mathbb{R}^{d_{\text{lin}} \times D}$.

DEF 1 (Grassmannian) Grassmannian $\mathbf{G}(D, d_{\text{lin}})$, which is the collection of all D -dimensional linear subspaces $\mathbb{R}^{d_{\text{lin}} \times D}$, is a smooth Riemannian manifold. Also, an element \mathbf{P} of $\mathbf{G}(D, d_{\text{lin}})$ can be specified by a basis, i.e., $d_{\text{lin}} \times D$ matrix with orthogonal columns.

In the context of the GFTab, the concept of Grassmannians facilitates a robust approach to feature representation from a tabular dataset by allowing us to model the space of subspaces formed by features. By representing different corrupted *views*, such as z_{soft} and z_{hard} , as points on the Grassmannian, we gain the ability to learn their geometric relationship. This approach aids in identifying invariant features that are crucial for learning under corruption intensity.

To capture this relationship, we introduce the concept of geodesic flow (Gallivan et al. 2003), denoted as $\mathbf{GF}(\pi)$. Let \mathbf{P}_{hard} and \mathbf{P}_{soft} denote the sets of basis vector for the subspaces corresponding to the hard and soft representations,

respectively. The geodesic flow between \mathbf{P}_{hard} and \mathbf{P}_{soft} , denoted by \mathbf{GF} , is parameterized as $\mathbf{GF} : \pi \in [0, 1] \rightarrow \mathbf{GF}(\pi) \in \mathbf{G}(D, d_{\text{lin}})$ such that $\mathbf{GF}(0) = \mathbf{P}_{\text{hard}}$ and $\mathbf{GF}(1) = \mathbf{P}_{\text{soft}}$. For $\pi \in (0, 1)$,

$$\begin{aligned} \mathbf{GF}(\pi) &= [\mathbf{P}_{\text{hard}} \quad \mathbf{R}_{\text{hard}}] \begin{bmatrix} \mathbf{U}_1 \Gamma(\pi) & 0 \\ 0 & -\mathbf{U}_2 \Sigma(\pi) \end{bmatrix} \\ &= \mathbf{P}_{\text{hard}} \mathbf{U}_1 \Gamma(\pi) - \mathbf{R}_{\text{hard}} \mathbf{U}_2 \Sigma(\pi) \end{aligned} \quad (5)$$

Here, the $\mathbf{R}_{\text{hard}} \in \mathbb{R}^{d_{\text{lin}} \times (d_{\text{lin}} - D)}$ is the orthogonal complement of \mathbf{P}_{hard} .

According to the gSVD (Van Loan 1976), the matrices $\mathbf{P}_{\text{hard}}^\top \mathbf{P}_{\text{soft}}$ and $\mathbf{R}_{\text{hard}}^\top \mathbf{P}_{\text{soft}}$ share identical right singular vectors, \mathbf{V} . Using gSVD, we can decompose as here:

$$\begin{aligned} \mathbf{P}_{\text{hard}}^\top \mathbf{P}_{\text{soft}} &= \mathbf{U}_1 \Gamma \mathbf{V}^\top \\ \mathbf{R}_{\text{hard}}^\top \mathbf{P}_{\text{soft}} &= \mathbf{U}_2 \Sigma \mathbf{V}^\top \end{aligned} \quad (6)$$

where Γ and Σ have diagonal elements $\cos(\theta_i)$ and $\sin(\theta_i)$ for $i = 1, \dots, D$. The θ_i are the principal angle (Mandolesi 2019) between \mathbf{P}_{hard} and \mathbf{P}_{soft} .

By varying the parameter π from 0 to 1, we generate intermediate subspaces that capture the gradual transition from the z_{hard} to the z_{soft} . This transition is facilitated by the geodesic flow kernel, which smoothly varies z_{hard} along the geodesic path to z_{soft} . The purpose of this smooth variation is to learn the underlying relationships in the tabular data, allowing us to obtain a representation that is robust to real-world noise. To quantify this transition and capture these noise-resistant relationships, we project the encoded representations onto the subspace $\mathbf{GF}(\pi)$ at each point along the geodesic. To formalize this process and provide a measure of similarity between the encoded representations, we introduce the geodesic flow kernel.

DEF 2. (Geodesic Flow Kernel) Let z_{hard} and z_{soft} are the encoded representation, respectively. The, the geodesic flow kernel is defined as:

$$z_{\text{hard}}^\top \mathbf{A} z_{\text{hard}} = \int_0^1 (\mathbf{GF}(\pi)^\top z_{\text{hard}})^\top (\mathbf{GF}(\pi)^\top z_{\text{soft}}) d\pi \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{D \times D}$ is positive semi-definite matrix.

Fortunately, we do not necessary to compute all projection into the subspace. By (Gong et al. 2012; Simon, Koniusz, and Harandi 2021), we can get matrix \mathbf{A} in a closed form:

$$\mathbf{A} = [\mathbf{P}_{\text{hard}} \mathbf{U}_1 \quad \mathbf{R}_{\text{soft}} \mathbf{U}_2] \begin{bmatrix} \mathbf{D} & \mathbf{E} \\ \mathbf{E} & \mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{U}_1^\top \mathbf{P}_{\text{hard}}^\top \\ \mathbf{U}_2^\top \mathbf{R}_{\text{soft}}^\top \end{bmatrix} \quad (8)$$

where $d_{1i} = 1 + \frac{\sin(2\theta_i)}{2\theta_i}$, $e_{2i} = \frac{\cos(2\theta_i) - 1}{2\theta_i}$ and $g_{3i} = 1 - \frac{\sin(2\theta_i)}{2\theta_i}$, which are the i -th elements of the diagonal matrices \mathbf{D} , \mathbf{E} , and \mathbf{G} respectively.

The matrix $\mathbf{A} \in \mathbb{R}^{D \times D}$ captures the geometric relationship between z_{soft} and z_{hard} subspaces on $\mathbf{G}(D, d_{\text{lin}})$. In other words, by learn the variation between z_{hard} and z_{soft} , our model extracts meaningful patterns, even with noise values common in real-world tabular datasets. The geodesic similarity loss function \mathcal{L}_{sim} in Equation (9) encourages the model to minimize the dissimilarity between the z_{hard} and

Panel A. Datasets with more categorical variables								
model	Diabetes	Insurance	adult	bank	cmc	credit-approval	credit-g	
GFTab	0.3826 ± 0.0139	0.4500 ± 0.0022	0.8023 ± 0.0070	0.7336 ± 0.0040	0.4625 ± 0.0215	0.6738 ± 0.0019	0.7433 ± 0.0424	
GRANDE	0.3635 ± 0.0012	0.4225 ± 0.0001	0.7596 ± 0.0033	0.7017 ± 0.0042	0.4306 ± 0.0229	0.8521 ± 0.0008	0.6570 ± 0.0224	
TabPFN	0.2559 ± 0.0089	0.4312 ± 0.0001	0.7530 ± 0.0024	0.6102 ± 0.0047	0.4566 ± 0.0027	0.8975 ± 0.0091	0.4802 ± 0.0409	
SCARF	0.2335 ± 0.0001	0.4312 ± 0.0002	0.4157 ± 0.0244	0.4699 ± 0.0009	0.3479 ± 0.0425	0.6993 ± 0.0518	0.4451 ± 0.0377	
SubTab	0.2569 ± 0.0187	0.4312 ± 0.0001	0.7506 ± 0.0035	0.6626 ± 0.0049	0.5012 ± 0.0141	0.6728 ± 0.0202	0.6474 ± 0.0388	
VIME	0.2611 ± 0.0313	0.4314 ± 0.0001	0.7369 ± 0.0057	0.6932 ± 0.0179	0.4942 ± 0.0336	0.7177 ± 0.1637	0.5397 ± 0.0891	
XGBoost	0.3534 ± 0.0002	0.4312 ± 0.0003	0.7344 ± 0.0006	0.5246 ± 0.0024	0.4540 ± 0.0213	0.8276 ± 0.0096	0.5968 ± 0.0171	
CatBoost	0.3552 ± 0.0003	0.4312 ± 0.0001	0.7428 ± 0.0006	0.4844 ± 0.0103	0.3694 ± 0.0430	0.8469 ± 0.0082	0.5590 ± 0.0310	
model	dresses-sales	fars	jasmine	kick	okcupid-stem	online-shoppers	shruntime	
GFTab	0.4165 ± 0.0057	0.6199 ± 0.0104	0.7775 ± 0.0041	0.4928 ± 0.0044	0.4219 ± 0.0195	0.7871 ± 0.0106	0.7088 ± 0.0025	
GRANDE	0.4611 ± 0.0256	0.5600 ± 0.0122	0.7630 ± 0.0044	0.4654 ± 0.0012	0.4412 ± 0.0117	0.7405 ± 0.0042	0.7355 ± 0.0034	
TabPFN	0.3671 ± 0.0010	0.5247 ± 0.0104	0.7491 ± 0.0036	0.4749 ± 0.0010	0.3889 ± 0.0019	0.7954 ± 0.0012	0.7130 ± 0.0069	
SCARF	0.4694 ± 0.0838	0.1314 ± 0.0095	0.5839 ± 0.0138	0.4753 ± 0.0010	0.2787 ± 0.0000	0.4579 ± 0.0003	0.4459 ± 0.0044	
SubTab	0.5300 ± 0.0348	0.5020 ± 0.0035	0.7078 ± 0.0063	0.4840 ± 0.0030	0.4021 ± 0.0118	0.6624 ± 0.0096	0.6808 ± 0.0026	
VIME	0.5069 ± 0.0346	0.5900 ± 0.0041	0.7548 ± 0.0050	0.4758 ± 0.0006	0.4758 ± 0.0006	0.4583 ± 0.0015	0.5323 ± 0.0421	
XGBoost	0.5041 ± 0.0553	0.4196 ± 0.0002	0.7851 ± 0.0061	0.4749 ± 0.0002	0.3546 ± 0.0043	0.7964 ± 0.0046	0.6863 ± 0.0002	
CatBoost	0.4056 ± 0.0804	0.3571 ± 0.0002	0.7940 ± 0.0029	0.4749 ± 0.0001	0.3062 ± 0.0059	0.7847 ± 0.0065	0.6215 ± 0.0107	

Panel B. Datasets with more continuous variables								
model	KDD	Shipping	churn	eye-movements	nomao	qsar	road-safety	
GFTab	0.7998 ± 0.0070	0.6493 ± 0.0020	0.7865 ± 0.0365	0.5534 ± 0.0205	0.9411 ± 0.0030	0.7867 ± 0.0140	0.7553 ± 0.0042	
GRANDE	0.7848 ± 0.0069	0.6186 ± 0.0045	0.7783 ± 0.0247	0.5676 ± 0.0102	0.9113 ± 0.0067	0.7729 ± 0.0110	0.7566 ± 0.0017	
TabPFN	0.7722 ± 0.0045	0.6432 ± 0.0024	0.7657 ± 0.0174	0.5854 ± 0.0020	0.8893 ± 0.0032	0.8254 ± 0.0000	0.7389 ± 0.0015	
SCARF	0.5603 ± 0.0183	0.6213 ± 0.0141	0.4624 ± 0.0005	0.4878 ± 0.0167	0.5065 ± 0.0217	0.6126 ± 0.0528	0.4976 ± 0.0023	
SubTab	0.6634 ± 0.0269	0.5514 ± 0.0091	0.7539 ± 0.014	0.5711 ± 0.0060	0.9290 ± 0.0038	0.8404 ± 0.0139	0.6750 ± 0.0006	
VIME	0.7042 ± 0.0164	0.6358 ± 0.0148	0.7051 ± 0.0304	0.5524 ± 0.0277	0.9361 ± 0.0028	0.8538 ± 0.0158	0.7528 ± 0.0025	
XGBoost	0.8001 ± 0.0081	0.6247 ± 0.0081	0.5231 ± 0.0001	0.5533 ± 0.0187	0.9078 ± 0.0002	0.8001 ± 0.0035	0.7452 ± 0.0250	
CatBoost	0.8138 ± 0.0073	0.6416 ± 0.0038	0.5163 ± 0.0413	0.5685 ± 0.0161	0.8936 ± 0.0040	0.7824 ± 0.0252	0.7545 ± 0.0153	

Table 1: Comparison of F1 score between GFTab and baseline models on 21 tabular benchmark datasets in 20% labeled training setting. The best performing method is highlighted in red and the second best in blue, while the third best is bold.

z_{soft} representations, taking into account their geometric relationship on $\mathbf{G}(D, d_{\text{lin}})$.

$$\mathcal{L}_{\text{sim}} = 1 - \frac{z_{\text{soft}}^{\top} \mathbf{A} z_{\text{hard}}}{\left\| \sqrt{\mathbf{A} z_{\text{soft}}} \right\| \left\| \sqrt{\mathbf{A} z_{\text{hard}}} \right\|} \quad (9)$$

The GFTab Framework III: Tree-Based Embedding

For tabular data, the relationships between columns are crucial, as they can significantly impact the overall data structure and interpretation. However, most existing tabular DL models fail to capture comprehensive relationships between many different columns. To address this limitation, we employ a tree-based embedding approach inspired by DATE (Kim et al. 2020) to effectively leverage labeled data. Note that tree-based models are known to be effective in capturing various relationships among different columns. A GBDT $\mathcal{T} = \{\mathcal{V}, \mathcal{E}\}$ is used, where $\mathcal{V} = \{v_R\} \cup \mathcal{V}_I \cup \mathcal{V}_L$ represents nodes (root, internal, and leaf) and \mathcal{E} is set of edges. Each leaf node $v_L \in \mathcal{V}_L$ is assigned a learnable embedding vector $\mathbf{s}_L \in \mathbb{R}^{\text{d}_{\text{emb}}}$. For data point x_i , activated leaves are represented by binary vector \mathbf{p}_i . The tree-based embedding $\mathbf{S}_i \in \mathbb{R}^{L \times \text{d}_{\text{emb}}}$ is:

$$\mathbf{S}_i = \phi(\mathbf{p}_{i,1} \mathbf{s}_1, \mathbf{p}_{i,2} \mathbf{s}_2, \dots, \mathbf{p}_{i,|\mathcal{V}_L|} \mathbf{s}_{|\mathcal{V}_L|}), \quad (10)$$

where $\phi(\cdot)$ removes all-zero rows. The tree embeddings are then processed through the feature representation module and a feed-forward layer. See Appendix D for empirical validation of this embeddings capability. So, GFTab’s overall loss combines geodesic similarity loss \mathcal{L}_{sim} and supervised

cross-entropy loss \mathcal{L}_{ce} :

$$\mathcal{L}_{\text{GFTab}} = \mathcal{L}_{\text{sim}} + \beta \mathcal{L}_{\text{ce}}, \quad (11)$$

where \mathcal{L}_{ce} is cross-entropy loss for labeled data and β is a balance parameter that controls the relative importance of the supervised loss.

Tabular Benchmarks Dataset

We configured benchmark datasets to properly evaluate the performance of our proposed method. We selected 21 datasets after carefully reviewing more than 4,000 datasets including OpenML (3,953 datasets), AMLB (Gijsbers et al. 2022) (71 datasets), and (Grinsztajn, Oyallon, and Varoquaux 2022)(22 datasets). Detailed criteria are given below.

- **Preprocessing.** Datasets with more than 30% missing values were excluded. For the remaining datasets, columns with more than 30% missing values were removed. Also, redundant categorical variables, which have only one category, were removed.
- **Variable types.** In order to evaluate tabular models in a more real-world like environment, we selected datasets with both continuous and categorical variables. Surprisingly, around 60% of the entire datasets did not satisfy this condition.
- **Data distribution.** We assume that data samples are i.i.d. Hence, datasets with certain distributional structure (sequential or temporal) were excluded. Also, we eliminated datasets with too simple distributions, which can be easily predicted with high Accuracy by naive models. Artificially generated datasets were excluded as well. Lastly,

as this study focuses on classification tasks, datasets for regression tasks were not considered.

- **Dataset size.** Most previous studies did not evaluate their models with datasets of different sizes. For more comprehensive evaluation, we selected datasets with different sizes: small-sized ($\sim 10,000$ samples), medium-sized ($10,000 \sim 50,000$), and large-sized ($50,000 \sim$).

More details on dataset can be found in Appendix A.

Experiment

We evaluate GFTab on a diverse set of tabular datasets to assess its performance and robustness across different domains and data characteristics.

Implementation details

We present the implementation details of evaluation metrics and baseline models. For other settings such as hyperparameters and experiment settings, see Appendix B.

Baseline models & hyperparameter. We use GRANDE (Marton et al. 2024), TabPFN (Hollmann et al. 2023) SCARF (Bahri et al. 2022), VIME (Yoon et al. 2020), SubTab (Ucar, Hajiramezani, and Edwards 2021), XGBoost (Chen and Guestrin 2016) and CatBoost (Prokhorenkova et al. 2018) as baseline models. For XGBoost and CatBoost, we performed hyperparameter optimization using Optuna (Akiba et al. 2019), conducting 250 experiments for each model. The search space for hyperparameters was kept consistent with the one used in GRANDE (Marton et al. 2024) to ensure a fair comparison. Regarding the deep learning models, such as GRANDE, TabPFN, SCARF, VIME, and SubTab, we provide detailed information about the hyperparameter tuning process in Appendix B.

Evaluation metrics. For efficient presentation in our ablation study, we define a *win matrix*. A win matrix is a k by k matrix, where k is the number of method being compared in this study. The entry $W_{i,j}$ is calculated as the proportion of the number of times method i outperformed method j over m datasets. It is represented as follows:

$$W_{i,j} = \sum_{d=1}^m \mathbb{I}[\text{method } i \text{ beats method } j \text{ on dataset } d] / m \quad (12)$$

Is GFTab really effective for tabular datasets?

We evaluate GFTab’s performance on 21 diverse benchmark datasets, comparing it against seven baseline models under two primary conditions: 20% labeled training data (Table 1) and 20% labeled training data with 20% label noise (Table 2). The datasets are divided into two panels: Panel A for datasets with more categorical variables and Panel B for those with more continuous variables.

In Table 1, which presents results for the 20% labeled data settings, GFTab demonstrates superior performance across diverse datasets. It achieves the highest F1 score in nearly half of the cases (10 out of 21) and ranks among the top three in an additional seven. GFTab particularly excels with categorical-dominant datasets (Panel A), leading in 7 out of 14 cases and placing in the top three for 5 more. For continuous-dominant datasets (Panel B), GFTab maintains

strong performance, topping 3 out of 7 cases and ranking in the top three for 2 others. Table 2 shows GFTab’s robustness when 20% label noise is introduced. It retains competitive performance overall, securing the highest F1 score in 8 datasets and ranking among the top three in 10 others. Notably, GFTab’s resilience is particularly evident in datasets with a higher proportion of categorical variables, even under these noisy conditions.

While certain methods perform consistently well on some datasets (SubTab on dresses-sales and qsar, CatBoost on jasmine and KDD, GRANDE on okcupid-stem and shrutime, XGboost on online-shoppers), we emphasize that there is no one-size-fits-all solution for all tabular datasets (Marton et al. 2024). Our proposed GFTab, while not universally superior on all datasets, clearly performs robustly regardless of label noise when compared to the top three performers. Also, we found that a label percentage of 10% yielded similar results to the 20% setting. Detailed results for this additional experiment can be found in Appendix E.

How to corrupt categorical variables effectively?

When designing corruption methods for categorical variables, we must consider their ordered or unordered nature. For ordered variables, while precise quantification of distances between values is impossible, we can still determine relative closeness, allowing corruption to neighboring values (GFTab). Unordered variables lack this concept of similarity, requiring random replacement with another category (Permute). However, challenges arise with imbalanced distributions, as categorical variables can exhibit identical discrete values across different items, complicating the creation of *views*. To address this, we can employ a row-wise random selection strategy (Random). Unfortunately, it is challenging to choose one of the aforementioned methods separately for each categorical column based on its characteristics. Therefore, an alternative solution is to project the categorical variable into the embedding space and corrupt it in this space, regardless of the type of the categorical variable (Embed).

To evaluate the effectiveness of these approaches, we tested five methods on our benchmark datasets: GFTab, Embed, Random, Permute, and None (where None indicates no categorical feature corruptions). The results, as shown in Figure E.1(Appendix), demonstrate the efficacy of our corruption methods for categorical variables, consistently outperforming other methods. Notably, GFTab shows marked performance enhancement compared to Permute and Random, particularly in settings with noisy labels. This provides compelling evidence that our selected corruption method is beneficial in handling categorical variables within tabular datasets.

Is geodesic flow useful for tabular datasets?

In this section, we explore the effect of different loss functions on the performance of GFTab. The *win-matrix* in Figure Figure E.2(Appendix) shows that GFTab loss function had the most number of wins compared to the other loss functions (Uniform-Alignment (Wang and Isola 2020) Barlow Twins (Zbontar et al. 2021), and InfoNCE (Oord, Li, and Vinyals 2018)). In the absence of label noise, it surpassed

Panel A. Datasets with more categorical variables								
model	Diabetes	Insurance	adult	bank	cmc	credit-approval	credit-g	
GFTab	0.3879 ±0.0013	0.4765 ±0.0022	0.7680 ±0.0015	0.6972 ±0.0040	0.4598 ±0.0284	0.7853 ±0.0054	0.5376 ±0.0144	
GRANDE	0.3730 ±0.0046	0.4594 ±0.0034	0.7329 ±0.0098	0.6772 ±0.0073	0.4316 ±0.0337	0.7330 ±0.0228	0.4958 ±0.0359	
TabPFN	0.2335 ±0.0120	0.4312 ±0.0000	0.7034 ±0.0029	0.4689 ±0.0012	0.4868 ±0.0070	0.7603 ±0.0010	0.4118 ±0.0100	
SCARF	0.2335 ±0.0010	0.4312 ±0.0000	0.4222 ±0.0116	0.4710 ±0.0035	0.3010 ±0.0310	0.6312 ±0.0854	0.4372 ±0.0374	
SubTab	0.2450 ±0.0021	0.4328 ±0.0011	0.7325 ±0.0039	0.6136 ±0.0059	0.4551 ±0.0143	0.3796 ±0.0237	0.6007 ±0.0144	
VIME	0.2219 ±0.0097	0.4101 ±0.0001	0.7142 ±0.0082	0.6755 ±0.0106	0.4307 ±0.0691	0.6794 ±0.1110	0.5629 ±0.0671	
XGBoost	0.3503 ±0.0020	0.4325 ±0.0008	0.7277 ±0.0004	0.5405 ±0.0049	0.4156 ±0.0087	0.7665 ±0.0121	0.5133 ±0.0100	
CatBoost	0.3388 ±0.0006	0.4312 ±0.0000	0.7328 ±0.0012	0.4806 ±0.0102	0.2565 ±0.0494	0.7738 ±0.0175	0.5106 ±0.0377	
model	dresses-sales	fars	jasmine	kick	okcupid-stem	online-shoppers	shruntime	
GFTab	0.4252 ±0.1100	0.6039 ±0.0021	0.7133 ±0.0194	0.5180 ±0.0066	0.4122 ±0.0210	0.7366 ±0.0042	0.6641 ±0.0016	
GRANDE	0.4824 ±0.0656	0.5575 ±0.0296	0.7361 ±0.0212	0.5175 ±0.0006	0.4783 ±0.0050	0.6635 ±0.0112	0.6866 ±0.0263	
TabPFN	0.3671 ±0.0002	0.3828 ±0.0137	0.7183 ±0.0031	0.4749 ±0.0020	0.3336 ±0.0052	0.7570 ±0.0033	0.5744 ±0.0101	
SCARF	0.4245 ±0.1140	0.1366 ±0.0211	0.5925 ±0.0394	0.4751 ±0.0006	0.2922 ±0.0117	0.4631 ±0.0044	0.4439 ±0.0014	
SubTab	0.6184 ±0.0105	0.4203 ±0.0009	0.5973 ±0.0069	0.4901 ±0.0013	0.3728 ±0.0316	0.6404 ±0.0068	0.6419 ±0.0037	
VIME	0.4388 ±0.1270	0.5588 ±0.0148	0.7394 ±0.0108	0.4524 ±0.0015	0.3198 ±0.0292	0.4601 ±0.0033	0.5718 ±0.0397	
XGBoost	0.5257 ±0.0169	0.4195 ±0.0003	0.7357 ±0.0152	0.4748 ±0.0000	0.3063 ±0.0000	0.7791 ±0.0050	0.6830 ±0.0051	
CatBoost	0.3838 ±0.0290	0.3571 ±0.0001	0.7494 ±0.0016	0.4749 ±0.0000	0.3036 ±0.0002	0.7284 ±0.0020	0.6334 ±0.0045	

Panel B. Datasets with more continuous variables								
model	KDD	Shipping	churn	eye-movements	nomao	qsar	road-safety	
GFTab	0.7774 ±0.0164	0.6203 ±0.0023	0.6285 ±0.0041	0.5600 ±0.0129	0.8866 ±0.0006	0.7099 ±0.0064	0.7349 ±0.0019	
GRANDE	0.7578 ±0.0126	0.6036 ±0.0046	0.5806 ±0.0091	0.4737 ±0.0874	0.8602 ±0.0065	0.6714 ±0.0374	0.7330 ±0.0023	
TabPFN	0.7090 ±0.0034	0.6411 ±0.0028	0.4624 ±0.0000	0.5575 ±0.0058	0.7818 ±0.0131	0.7716 ±0.0122	0.6960 ±0.0018	
SCARF	0.4969 ±0.0429	0.5569 ±0.1136	0.4624 ±0.0000	0.4841 ±0.0166	0.5293 ±0.0149	0.6909 ±0.0366	0.4993 ±0.0115	
SubTab	0.6135 ±0.0057	0.5403 ±0.0011	0.5855 ±0.0156	0.5502 ±0.0029	0.8019 ±0.0082	0.6712 ±0.0218	0.6624 ±0.0007	
VIME	0.6760 ±0.0066	0.5941 ±0.0411	0.6422 ±0.0266	0.5260 ±0.0308	0.8929 ±0.0032	0.7852 ±0.0044	0.7168 ±0.0012	
XGBoost	0.7708 ±0.0099	0.5983 ±0.0028	0.5569 ±0.0252	0.5641 ±0.0112	0.8996 ±0.0012	0.7318 ±0.0259	0.7334 ±0.0128	
CatBoost	0.7982 ±0.0068	0.6317 ±0.0045	0.4720 ±0.0166	0.5921 ±0.0169	0.8803 ±0.0006	0.7830 ±0.0110	0.7233 ±0.0068	

Table 2: Comparison of F1 score between GFTab and baseline models on 21 tabular benchmark datasets in 20% labeled training with 20% label noise. The best performing method is highlighted in red and the second best in blue, while the third best is bold.

the other functions on at least 14 out of 21 datasets. Notably, GFTab outperformed these alternative loss functions regardless of label noise.

The superiority of GFTab loss function over traditional methods may be attributed to its ability to capture the geometric relationships between data representations in a more nuanced and comprehensive manner. Traditional methods often rely on Euclidean space metrics, which may not adequately capture the complex relationships present in high-dimensional tabular data. In contrast, the geodesic flow kernel provides a continuous and smooth transition between subspaces, allowing GFTab to capture subtle variations in the data that might be missed by traditional approaches. This nature may enable GFTab to more accurately model and learn from the intrinsic structure of tabular datasets.

How to choose the balance between \mathcal{L}_{ce} and \mathcal{L}_{sim} ?

The GFTab is based on semi-supervised learning so that it is trained to minimize $\mathcal{L}_{GFTab} = \mathcal{L}_{sim} + \beta \mathcal{L}_{ce}$. we compared the performance of the model across various ranges of β . As a result, $\beta = 1.0$ yielded the best balance. As shown in Figure E.3(Appendix), we conducted an analysis of the impact of different β values on GFTab’s performance across various labeling scenarios. The results show that $\beta = 1.0$ yields good performance across different settings.

In the 10% labeled setting, $\beta = 1.0$ outperformed other β values on a majority of datasets. This trend continued for the 20% labeling setting, although slightly weaker, with $\beta = 1.0$ maintaining competitive performance. Even in the presence of label noise, $\beta = 1.0$ showed robustness, particularly in the 20% labeled setting with noise. However, the 10% la-

beled setting did not show a significant difference between β , but performed a little better at $\beta = 0.4$. While the optimal β might vary slightly depending on specific dataset characteristics or labeling conditions, our experiments indicate that $\beta = 1.0$ is a reliable choice that performs well across a range of settings.

Conclusion

GFTab is a semi-supervised framework for mixed-variable tabular datasets. It incorporates three key components: variable-specific corruption methods, similarity loss based on geodesic flow kernel, and tree-based embedding. This approach enables effective learning from both labeled and unlabeled data, addressing challenges specific to tabular data structures. Experiments across 21 diverse benchmark datasets show GFTab consistently outperforming existing ML/DL models, particularly in scenarios with limited labeled data or label noise. GFTab’s ability to handle tabular data characteristics makes it a robust and versatile solution applicable to various domains.

Acknowledgments

This work was supported by National Research Foundation of Korea (NRF) grant (No. NRF-2022R1I1A4069163 and No. RS-2024-00340339) and Institute of Information & Communications Technology Planning & Evaluation(IITP) grant (No. RS-2020-II201336, Artificial Intelligence Graduate School Program(UNIST)) funded by the Korea government(MSIT)

References

- Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; and Koyama, M. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2623–2631.
- Arik, S. Ö.; and Pfister, T. 2021. Tabnet: Attentive interpretable tabular learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 6679–6687.
- Bahri, D.; Jiang, H.; Tay, Y.; and Metzler, D. 2022. Scarf: Self-supervised contrastive learning using random feature corruption. *Proceedings of the 38th International Conference on Machine Learning*.
- Borisov, V.; Leemann, T.; Sessler, K.; Haug, J.; Pawelczyk, M.; and Kasneci, G. 2022. Deep Neural Networks and Tabular Data: A Survey. *IEEE Transactions on Neural Networks and Learning Systems*, 1–21.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Gallivan, K. A.; Srivastava, A.; Liu, X.; and Van Dooren, P. 2003. Efficient algorithms for inferences on Grassmann manifolds. In *IEEE Workshop on Statistical Signal Processing, 2003*, 315–318. IEEE.
- Gijsbers, P.; Bueno, M. L. P.; Coors, S.; LeDell, E.; Poirier, S.; Thomas, J.; Bischl, B.; and Vanschoren, J. 2022. AMLB: an AutoML Benchmark.
- Gong, B.; Shi, Y.; Sha, F.; and Grauman, K. 2012. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, 2066–2073. IEEE.
- Gorishniy, Y.; Rubachev, I.; Khrulkov, V.; and Babenko, A. 2021. Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 34: 18932–18943.
- Grinsztajn, L.; Oyallon, E.; and Varoquaux, G. 2022. Why do tree-based models still outperform deep learning on tabular data? *arXiv preprint arXiv:2207.08815*.
- Guo, X.; Quan, Y.; Zhao, H.; Yao, Q.; Li, Y.; and Tu, W. 2021. TabGNN: Multiplex graph neural network for tabular data prediction. *arXiv preprint arXiv:2108.09127*.
- Herzig, J.; Nowak, P. K.; Müller, T.; Piccinno, F.; and Eisen-schlos, J. M. 2020. TaPas: Weakly supervised table parsing via pre-training. *arXiv preprint arXiv:2004.02349*.
- Hollmann, N.; Müller, S.; Eggensperger, K.; and Hutter, F. 2023. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. In *The Eleventh International Conference on Learning Representations*.
- Huang, X.; Khetan, A.; Cvitkovic, M.; and Karnin, Z. 2020. Tabtransformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*.
- Iida, H.; Thai, D.; Manjunatha, V.; and Iyyer, M. 2021. Tabbie: Pretrained representations of tabular data. *arXiv preprint arXiv:2105.02584*.
- Jing, L.; and Tian, Y. 2020. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11): 4037–4058.
- Kim, S.; Tsai, Y.-C.; Singh, K.; Choi, Y.; Ibok, E.; Li, C.-T.; and Cha, M. 2020. DATE: Dual attentive tree-aware embedding for customs fraud detection. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2880–2890.
- Kossen, J.; Band, N.; Lyle, C.; Gomez, A. N.; Rainforth, T.; and Gal, Y. 2021. Self-attention between datapoints: Going beyond individual input-output pairs in deep learning. *Advances in Neural Information Processing Systems*, 34: 28742–28756.
- Lee, K.; Chang, M.-W.; and Toutanova, K. 2019. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*.
- Lim, B.; Arik, S. Ö.; Loeff, N.; and Pfister, T. 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4): 1748–1764.
- Mandolesi, A. L. 2019. Grassmann angles between real or complex subspaces. *arXiv preprint arXiv:1910.00147*.
- Marton, S.; Lüdtke, S.; Bartelt, C.; and Stuckenschmidt, H. 2024. GRANDE: Gradient-Based Decision Tree Ensembles. In *The Twelfth International Conference on Learning Representations*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Popov, S.; Morozov, S.; and Babenko, A. 2020. Neural oblivious decision ensembles for deep learning on tabular data. *ICLR*.
- Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A. V.; and Gulin, A. 2018. CatBoost: unbiased boosting with categorical features. *Advances in neural information processing systems*, 31.
- Qiu, X.; Sun, T.; Xu, Y.; Shao, Y.; Dai, N.; and Huang, X. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10): 1872–1897.
- Ruder, S.; and Plank, B. 2018. Strong baselines for neural semi-supervised learning under domain shift. *arXiv preprint arXiv:1804.09530*.
- Simon, C.; Koniusz, P.; and Harandi, M. 2021. On learning the geodesic path for incremental learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 1591–1600.
- Somepalli, G.; Goldblum, M.; Schwarzschild, A.; Bruss, C. B.; and Goldstein, T. 2021. Saint: Improved neural networks for tabular data via row attention and contrastive pre-training. *arXiv preprint arXiv:2106.01342*.

Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33: 16857–16867.

Song, W.; Shi, C.; Xiao, Z.; Duan, Z.; Xu, Y.; Zhang, M.; and Tang, J. 2019. AutoInt. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. ACM.

Ucar, T.; Hajiramezanali, E.; and Edwards, L. 2021. SubTab: Subsetting Features of Tabular Data for Self-Supervised Representation Learning. *Advances in Neural Information Processing Systems*, 34.

Van Loan, C. F. 1976. Generalizing the singular value decomposition. *SIAM Journal on numerical Analysis*, 13(1): 76–83.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 9929–9939. PMLR.

Yin, P.; Neubig, G.; Yih, W.-t.; and Riedel, S. 2020. TaBERT: Pretraining for joint understanding of textual and tabular data. *arXiv preprint arXiv:2005.08314*.

Yoon, J.; Zhang, Y.; Jordon, J.; and van der Schaar, M. 2020. Vime: Extending the success of self-and semi-supervised learning to tabular domain. *Advances in Neural Information Processing Systems*, 33: 11033–11043.

Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. *CoRR*, abs/2103.03230.

Zhai, X.; Oliver, A.; Kolesnikov, A.; and Beyer, L. 2019. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1476–1485.