

TIP: Tabular-Image Pre-training for Multimodal Classification with Incomplete Data

Siyi Du^{*}, Shaoming Zheng, Yinsong Wang, Wenjia Bai, Declan P. O'Regan, and Chen Qin^{*}

Imperial College London, London, UK

{s.du23,s.zheng22,y.wang23,w.bai,declan.oregan,c.qin15}@imperial.ac.uk

Abstract. Images and structured tables are essential parts of real-world databases. Though tabular-image representation learning is promising for creating new insights, it remains a challenging task, as tabular data is typically heterogeneous and incomplete, presenting significant modality disparities with images. Earlier works have mainly focused on simple modality fusion strategies in complete data scenarios, without considering the missing data issue, and thus are limited in practice. In this paper, we propose TIP, a novel tabular-image pre-training framework for learning multimodal representations robust to incomplete tabular data. Specifically, TIP investigates a novel self-supervised learning (SSL) strategy, including a masked tabular reconstruction task to tackle data missingness, and image-tabular matching and contrastive learning objectives to capture multimodal information. Moreover, TIP proposes a versatile tabular encoder tailored for incomplete, heterogeneous tabular data and a multimodal interaction module for inter-modality representation learning. Experiments are performed on downstream multimodal classification tasks using both natural and medical image datasets. The results show that TIP outperforms state-of-the-art supervised/SSL image/multimodal methods in both complete and incomplete data scenarios. Our code is available at <https://github.com/siyi-wind/TIP>.

Keywords: Multimodal · Image-tabular Representation Learning · Missing Data · Self-supervised Learning

1 Introduction

While combining various modalities such as images and text to build a multimodal artificial intelligence (AI) system has achieved significant progress, integrating tabular data has been less explored [6, 9]. Tabular data, however, is increasingly accessible in multimodal datasets, and its integration is crucial in various applications [1, 15, 34, 36, 37]. For instance, in healthcare, rich tabular information, *e.g.*, demographics, lifestyle and laboratory tests (Fig. 1(a)), is commonly collected together with imaging data in hospitals, which are then used in

^{*}Corresponding authors.

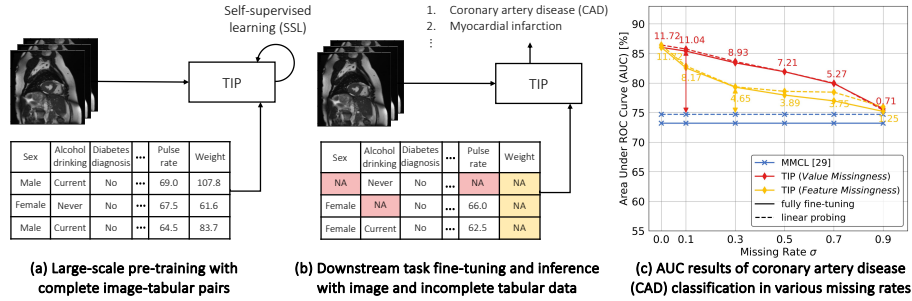


Fig. 1: The pipeline of TIP, which is pre-trained on large multimodal datasets (a) and can be deployed to downstream tasks with data missingness (b), *e.g.*, *Value Missingness* (red) and *Feature Missingness* (yellow). Results for coronary artery disease classification (c) show TIP’s superior performance over the SOTA multimodal pre-training method (numbers denote performance increase). Complete results in Fig. 4.

a joint way to inform clinical decision-making [5, 13, 16]. Large population studies [14, 42] such as the UK Biobank, have further enabled the wide availability of such multimodal resources to both machine learning and medical researchers. Despite these, current techniques for image-tabular data analysis are relatively limited. There is an increasingly strong interest in developing effective multimodal representation learning methods that can make the most of both image and tabular information to improve our understanding about human health.

Compared to vision-language modeling, incorporating image and tabular data in practical applications is a more difficult task with two main challenges. (1) Low-quality data: Despite the availability of large multimodal databases that allow pre-training, datasets for specific downstream tasks, *e.g.*, classification of rare diseases, can be limited and often suffer from data sparsity [10]. In Fig. 1(b), these datasets may inadvertently miss tabular values for some subjects, *i.e.*, *Value Missingness* (red table cells), or simply miss the entire features (columns), *i.e.*, *Feature Missingness* (yellow cells), due to diverse data collecting criteria across centers [7, 39, 49]. (2) Modality disparities: Unlike the homogeneous property of images and text, tabular data is heterogeneous with both dense numerical and sparse categorical features. The data columns exhibit varied value ranges, meanings and without clearly-defined inter-relationships [10]. Therefore, how to design a model that can effectively learn tabular and image representations to bridge the modality gap and address missing data is non-trivial.

Though there are some solutions for handling noisy and missing data [26, 38, 61], they mainly focus on unimodal tabular data analysis instead of multimodal tasks. Previous image-tabular models [37, 69, 77] are mostly trained and tested on relatively small labeled datasets (*e.g.*, 653 samples [69]) with a limited amount of tabular information (*e.g.*, 12 features [77]). These methods typically adopt shallow multi-layer perceptrons (MLPs) with simple modality fusion strategies and do not consider the challenges of incomplete data and modality disparities [38]. Recently, Hager *et al.* [29] proposed MMCL, the first SSL method that

jointly trains image and tabular encoders through multimodal contrastive learning. However, it only utilizes the image encoder for downstream tasks, neglecting the wealth of information in the tabular data for decision-making.

In this work, to address the above-identified two challenges, we propose TIP, a tabular-image pre-training framework based on a new multimodal representation learning network and a novel SSL pre-training strategy for managing small, incomplete downstream data (Fig. 1(a,b)). Specifically, we introduce a transformer-based tabular encoder with a versatile tabular embedding module, which serves two purposes: (1) it supports heterogeneous tabular inputs and diverse data missingness; (2) it captures inter-dependencies of tabular features and enhances representation learning. We further design a multimodal interaction module based on cross-modal attention to extract inter-modality information. To learn representations robust to missing data, we introduce three pre-training tasks: (1) masked tabular reconstruction to extract intra- and inter-modality relations from randomly masked data; (2) image-tabular contrastive learning to improve unimodal and multimodal representation learning; (3) image-tabular matching to obtain joint image-tabular representations for downstream tasks. Experiments on two representative datasets, *i.e.*, cardiac data from the UK Biobank [14] and natural image data from the DVM car advertisement dataset [36], demonstrate TIP’s superior performance. In particular, as illustrated in Fig. 1(c), even with 50% tabular values missing, TIP has 7.21% higher AUC than MMCL, the SOTA multimodal pre-training method.

Our contributions can be summarized as follows. (1) To the best of our knowledge, we are the first to propose SSL image-tabular pre-training to tackle the challenges of low-quality data and modality disparities and investigate various data missingness in multimodal tasks. (2) We introduce TIP, featuring a transformer-based multimodal architecture for enhanced representation learning and a novel SSL pre-training strategy for tackling tabular missingness. (3) Experiments on both natural and medical datasets demonstrate that TIP substantially surpasses SOTA supervised/SSL image/multimodal methods in both complete and incomplete data scenarios with various missing rates.

2 Related Work

Self-supervised Learning (SSL) approaches aim to acquire useful intermediate representations by pre-training models on unlabeled datasets with various intra- or inter-modal pretext tasks [41, 53]. Two groups of intra-modal tasks are currently popular: (1) contrastive learning that models similarity (and dissimilarities) between multiple input views [4, 18, 19, 61, 68]; and (2) generation-based learning that predicts the values of missing/corrupted input [3, 12, 20, 32, 54, 65, 71, 73]. With the increasing availability of multimodal datasets, inter-modal pretext tasks are gaining more attention and have demonstrated remarkable performance [17, 30]. Radford *et al.* [57] introduced CLIP, which performs image-text contrastive pre-training on massive web data and exhibits notable zero-shot performance. Follow-up research [45, 46, 74] added tasks such as masked language

modeling for more intricate cross-modal interactions. Nevertheless, few works have explored image-tabular pre-training [2, 29, 44]. Some generative image-tabular models [2, 44] were proposed, but were limited to using two or four tabular features. Though MMCL [29] used 117 tabular features, it only supported unimodal downstream tasks, ignoring the usefulness of multimodal information in fine-tuning and inference time. We are the first to handle multimodal downstream tasks with incomplete data using tabular-image SSL pre-training.

Deep Learning (DL) with Tabular Data has gained much interest [10], due to their ability to achieve an end-to-end multimodal data learning [10, 29]. Most existing works rely on MLPs and perform SSL tasks to learn representations [4, 65, 73]. Recently, transformers have been introduced to handle more challenging cases [50], *e.g.*, missing and noisy data [38], column permutation bias [70], and cross-table learning [68, 71]. These recent developments have inspired us to adapt the powerful transformer architecture for image-tabular learning.

Multimodal Image-Tabular Learning exploits tabular data as a complement to facilitate visual task learning, which is especially popular in the medical field [8, 35, 37] and has achieved improved results compared to pure image models [56, 66]. Previous works typically extracted image and tabular features through two separate encoders that are fused through various methods [11, 24, 24, 62, 66, 69, 77]. However, these methods mostly tested on small datasets with limited tabular features and did not consider missing data. Some works [29, 40] transferred tabular knowledge into image models, but used image features alone for downstream tasks, ignoring potentially helpful information in tabular data.

Missing Tabular Data is common problem in scientific data analysis, for which many solutions have been proposed [52, 60]. Some statistical approaches fill in missing values using column-wise mean or median [52]. An alternative popular method is iterative imputation, where each column with missing values is modeled as a function of other columns, employing a round-robin imputation process until convergence [39, 58, 59, 63]. With the emergence of DL, imputation algorithms based on deep generative models were introduced [51, 72], although they are limited to pre-processing steps and only support *Value Missingness* [26]. Some algorithms utilized SSL pre-training to make the model more robust to noisy or incomplete tabular data through reconstruction [38, 73], contrastive learning [4], or denoising [61]. In contrast, our study is the first to investigate various tabular missingness in a multimodal setting.

3 Methodology

In this section, we introduce our TIP, a tabular-image pre-training framework that is pre-trained on large multimodal datasets and then fine-tuned on downstream tasks, *e.g.*, classification with complete/incomplete data. To encode incomplete, heterogeneous tabular data and enhance representation learning, we propose a tailored tabular encoder with versatile tabular embedding and transformer layers and a multimodal interaction module based on cross-modal attention. Moreover, we design a novel SSL pre-training strategy for multimodal

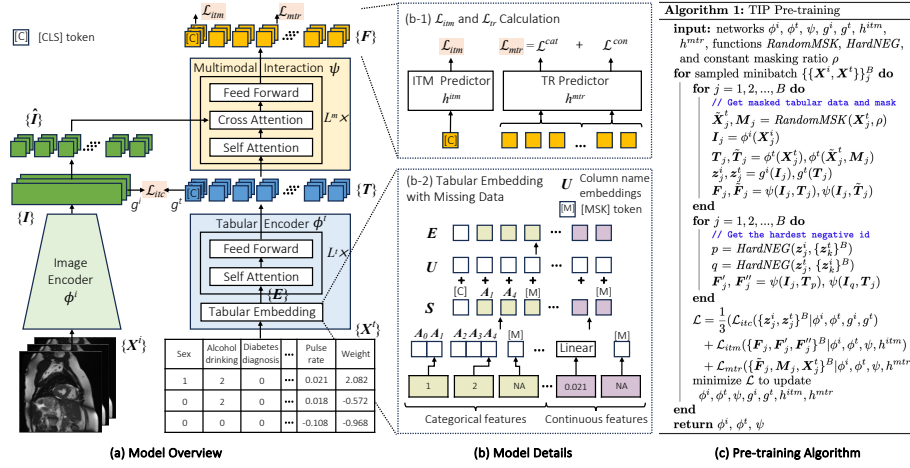


Fig. 2: Model architecture and algorithm of TIP: (a) Model overview with its image encoder, tabular encoder, and multimodal interaction module, which are pre-trained using 3 SSL losses: \mathcal{L}_{itm} , \mathcal{L}_{itm} , and \mathcal{L}_{mtr} . (b) Model details for (b-1) \mathcal{L}_{itm} and \mathcal{L}_{mtr} calculation and (b-2) tabular embedding with missing data. (c) Pre-training algorithm.

information extraction and for addressing potential data missingness. The overall framework of TIP is shown in Fig. 2. We describe TIP’s model architecture in Sec. 3.1 and then discuss its SSL pre-training strategy in Sec. 3.2.

3.1 TIP Model Architecture

Let $(\mathbf{X}^i \in \mathbb{R}^{H \times W \times 3}, \mathbf{X}^t = [x_1^t, \dots, x_N^t] \in \mathbb{R}^N)$ be an image-tabular pair, where N is the number of tabular features. Assume that each tabular input contains N_a categorical features, $x_1^t, \dots, x_{N_a}^t$, and $N - N_a$ continuous features, $x_{N_a+1}^t, \dots, x_N^t$. We convert categorical data into ordinal numbers and standardize continuous data as [29]. As shown in Fig. 2(a), TIP involves a convolutional neural network (CNN) based image encoder ϕ^i , a tabular encoder ϕ^t , and a multimodal interaction module ψ . The image representation $\mathbf{I} \in \mathbb{R}^{H' \times W' \times C}$ and the tabular representation $\mathbf{T} \in \mathbb{R}^{N \times D}$ are extracted by the image and tabular encoders, respectively, where C and D are their corresponding channel dimensions. We transform and project \mathbf{I} into a sequence of embeddings $\hat{\mathbf{I}} \in \mathbb{R}^{(H'W') \times D}$. Based on that, the multimodal interaction module then receives input of $\hat{\mathbf{I}}$ and \mathbf{T} to perform inter-modality learning, yielding a multimodal representation $\mathbf{F} \in \mathbb{R}^{N \times D}$. **Tabular Encoder:** To tackle heterogeneous data with potential missingness and to extract rich contextual information from tabular features, we propose to treat each tabular feature as a basic element and convert it to a token embedding through a versatile tabular embedding module (Fig. 2(b-2)). We then enable the embedded tokens to attend to related tokens through transformer layers. Our tabular embedding module contains 3 parts: *heterogeneous data processing*, *missing data processing*, and *column diversity integration*.

(1) *Heterogeneous data processing*: Tabular data usually contains categorical and continuous variables, which are very different attribute types and cannot be embedded using a single function [10, 27]. Therefore, our *tabular embedding* module processes these two types of features independently. In particular, we convert each categorical feature into a token embedding using a learnable embedding matrix $\mathbf{A} \in \mathbb{R}^{\bar{N}_a \times D}$, where \bar{N}_a is a summation of the number of unique values in N_a categorical feature. Meanwhile, We employ a shared linear layer to project each continuous feature into the D -dimensional space.

(2) *Missing data processing*: To handle incomplete data and inform our model which data is missing during training, we propose to embed each missing value using a special trainable D -dimensional [MSK] token embedding. This method does not require data imputation before inputting and can handle various types of missing data scenarios, which is more flexible than previous techniques [4, 39, 73] that only support random *Value Missingness* [26]. The embedded features are concatenated with a special tunable [CLS] token to generate a sequence of embeddings \mathbf{S} . The [CLS] token’s state at the last transformer layer serves as a learnt representation for downstream classification tasks as in [45, 46].

(3) *Column diversity integration*: Each column in tabular data generally has a different meaning, thus it is not suitable to treat them uniformly as pixels in images [44, 70]. To distinguish different columns and capture inter-column relationship, we propose to integrate column diversity through a sequence of learnable column name embeddings $\mathbf{U} \in \mathbb{R}^{(N+1) \times D}$. Instead of using pre-trained language models to tokenize column names into fixed textual embeddings [68, 70, 71], our data-driven strategy can dynamically capture hidden column dependencies existing in the training data. For example, the ‘weight’ and ‘alcohol drinking’ columns may have little semantic similarity but can present strong clinical association. The final embeddings are formulated as: $\mathbf{E} = \mathbf{S} + \mathbf{U}$.

Driven by transformers’ ability to capture long-range dependencies through self-attention [67] and to embed knowledge from large databases [30, 43], we utilize L_t transformer layers [67] to encode tabular information and extract a high-level tabular representation \mathbf{T} . To avoid the potential negative effect of missing data on model learning [32], we employ a self-attention mask, which restricts each token to only attend to itself and non-missing tokens, thus ensuring the tabular representation learning to be more robust and stable.

Multimodal Interaction Module: To enhance multimodal representation learning and capture cross-modal relationships, we propose to leverage the cross-attention mechanism in a transformer decoding module [67] to enable the [CLS] token and each tabular token to cross-attend to relevant image information. The interaction module consists of L^m layers, each including self-attention, cross-modal attention, an MLP feed-forward module, and layer normalization. The cross-modal attention in the l th layer can be written as:

$$\text{CrossAttention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}(\mathbf{Q}\mathbf{K}^T / \sqrt{d_k})\mathbf{V}, \quad (1)$$

where $\mathbf{Q} = \mathbf{F}_{l-1}\mathbf{W}_l^Q$, $\mathbf{K} = \hat{\mathbf{I}}\mathbf{W}_l^K$, $\mathbf{V} = \hat{\mathbf{I}}\mathbf{W}_l^V$, and $\mathbf{F}_0 = \mathbf{T}$.

3.2 SSL Pre-training Strategy

To enable the model to be robust to incomplete downstream data while improving representation learning, we propose to pre-train our model with three objectives, including masked tabular reconstruction, image-tabular contrastive learning, and image-tabular matching, as shown in Fig. 2(a,c).

Image-Tabular Contrastive Learning (ITC): We design the ITC task to capture better unimodal representations and align their feature spaces before modality fusion. This shares a similar motivation with image-text contrastive learning, which has demonstrated the ability to extract transferable representations for downstream tasks and to facilitate cross-modal learning [45,46,57]. ITC encourages image and tabular representations from a matched sample to be close compared to those from unmatched ones. We utilize two projection heads g^i and g^t to bring \mathbf{I} and \mathbf{T} to a shared low-dimensional hidden space and calculate the image-to-tabular and tabular-to-image similarities, *i.e.*, s_j^{i2t} and s_j^{t2i} , as [46]. The image-tabular contrastive loss can then be computed as $\mathcal{L}_{itc} = -(s_j^{i2t} + s_j^{t2i})/2$.

Image-Tabular Matching (ITM): We propose the ITM task to enable TIP’s multimodal interaction module to capture inter-modality relations and generate a joint multimodal representation, inspired by the success of image-text matching [45,46]. Our ITM aims to predict whether a pair of imaging and tabular input is positive (matched) or negative (unmatched). As displayed in Fig. 2(b-1), we feed the [CLS] embedding of \mathbf{F} , which captures a joint representation of an image-tabular pair, into the ITM predictor h^{itm} (a linear layer) for matching prediction based on a binary cross-entropy loss \mathcal{L}_{itm} . To enhance representation learning and capture discriminative features, we expose the model to more informative negative pairs using the hard negative mining strategy (*HardNEG*) proposed in [46]. Specifically, for each image/tabular representation, we select one unmatched tabular/image representation from the mini-batch using the similarity calculated in ITC as the sampling weight.

Masked Tabular Reconstruction (MTR): This task aims to learn multimodal representations that can be robust to missing tabular data in downstream tasks. Previous studies have found that reconstructing masked data can help mitigate noisy or missing data issues and produce promising performance in SSL representation learning [21,38,50,55,64,76]. We therefore propose MTR for multimodal representation learning and require the model to predict the masked parts in tabular data based on both image and unmasked tabular information. Specifically, we apply random masking (*RandomMSK*) on a tabular input based on a masking ratio ρ , to generate a masked version $\tilde{\mathbf{X}}^t$, *i.e.*, treating masked values as missing data, and a mask matrix \mathbf{M} used to record masking positions. The resulting $\tilde{\mathbf{X}}^t$ is then fed into the model to produce the masked multimodal representation $\tilde{\mathbf{F}} = [\tilde{\mathbf{f}}_1, \dots, \tilde{\mathbf{f}}_N]$. The generated $\tilde{\mathbf{F}}$ serves as the input to a MTR predictor h^{mtr} for reconstructing the missing values. Since reconstructing categorical data is a classification task, whereas reconstructing continuous data is a regression task, h^{mtr} has two distinct linear layers, processing N_a categorical

and $N - N_a$ continuous features correspondingly:

$$h^{mtr}(\tilde{\mathbf{f}}_n) = \begin{cases} \mathbf{W}_1 \tilde{\mathbf{f}}_n + \mathbf{b}_1, \mathbf{W}_1 \in \mathbb{R}^{\tilde{N}_a \times D} & 0 < n \leq N_a, \\ \mathbf{W}_2 \tilde{\mathbf{f}}_n + \mathbf{b}_2, \mathbf{W}_2 \in \mathbb{R}^{1 \times D} & N_a < n \leq N. \end{cases} \quad (2)$$

We formulate a reconstruction loss based on masked features only, $\mathcal{L}_{mtr} = \mathcal{L}^{cat} + \mathcal{L}^{con}$, where \mathcal{L}^{cat} represents the cross-entropy loss for categorical features and \mathcal{L}^{con} is a mean squared error loss for continuous features. Compared to previous tabular techniques that fill the masked cell with a randomly selected value from the same column [38, 73], our MTR task with the random masking strategy enables the model to learn a mask token for missing data in downstream tasks, so that the model can fully understand which values are missing and handle diverse data missingness, even if a whole column is missing. Ultimately, the overall pre-training loss function is formulated as:

$$\mathcal{L} = (\mathcal{L}_{itc} + \mathcal{L}_{itm} + \mathcal{L}_{mtr})/3. \quad (3)$$

Ensemble Learning during Fine-tuning: After pre-training, we can add linear classifiers after the feature extractor for downstream classification tasks. Given that our pre-training strategy enables the image encoder, tabular encoder, as well as multimodal interaction module to learn rich representations beneficial to downstream tasks, as well as motivated by ensemble learning’s ability to boost models’ generalizability and results [22, 25], we propose to build an ensemble model to further enhance the model’s performance. Specifically, we incorporate a linear classifier after each of the three modules and average the predictions from all three classifiers to generate the final output.

4 Experiment

Datasets and Metrics: Similar to [29], we experiment on two large datasets: a medical dataset – UK Biobank (UKBB) [14] and a natural image dataset – Data Visual Marketing (DVM) [36]. UKBB contains rich cardiac imaging and clinical tabular data gathered from individuals in the United Kingdom [48]. We carry out two cardiac disease classification tasks: coronary artery disease (CAD) and myocardial infarction (Infarction), using 2D short-axis cardiac magnetic resonance (MR) images and 75 disease-related tabular features (details in Sec. A of the supplementary material (supp.)). The dataset contains 36,167 image-tabular pairs, split into training (26,040), validation (6,510), and test (3,617) sets. Due to the low disease prevalence (3% for Infarction and 6% for CAD), we use balanced training datasets for fine-tuning, comprising 3,482 for CAD and 1,552 for Infarction, and evaluate all models with area under the curve (AUC). DVM [36] is a publicly available dataset for automotive applications, including 2D car images and car-related tabular data. We obtain 176,414 image-tabular pairs (17 tabular features, details in Sec. A of supp.) and implement a car model classification task with 283 classes. We split this dataset into training (70,565), validation (17,642), and test (88,207) sets and use accuracy for evaluation.

Table 1: Results of DVM, CAD, and Infarction classification tasks comparing TIP with supervised/SSL image/multimodal techniques on complete data. \ast denotes linear probing, *i.e.*, the feature extractors are frozen, and only the linear classifiers of the pre-trained models are tuned. \odot means fully fine-tuning, *i.e.*, all parameters are trainable. For supervised methods, all parameters are trainable in both \ast and \odot columns.

Model	DVM Accuracy (%) \uparrow		CAD AUC (%) \uparrow		Infarction AUC (%) \uparrow	
	\ast	\odot	\ast	\odot	\ast	\odot
(a) Supervised Image and Multimodal Methods						
ResNet-50 [33]	87.68	87.68	63.11	63.11	59.48	59.48
Concat Fuse (CF) [62]	94.60	94.60	85.76	85.76	85.05	85.04
Max Fuse (MF) [66]	94.39	94.39	85.31	85.31	84.75	84.75
Interact Fuse (IF) [24]	96.24	96.24	84.89	84.89	81.91	81.91
DAFT [69]	96.60	96.60	86.21	86.21	56.27	56.27
(b) SSL Image Pre-training Methods						
SimCLR [18]	61.06	87.65	68.42	72.58	68.86	75.07
BYOL [28]	56.26	88.64	65.67	69.18	66.63	70.12
SimSiam [19]	23.14	78.62	57.77	67.71	53.83	64.79
BarlowTwins [75]	53.60	88.36	55.64	61.68	50.01	60.14
(c) SSL Multimodal Pre-training Methods						
MMCL [29]	91.66	93.27	74.71	73.21	76.79	76.46
TIP	99.72	99.56	86.43	86.03	84.46	85.58

Implementation Details: We utilized a ResNet-50 [33] as the image encoder. Our tabular encoder and multimodal interaction module both have 4 transformer layers, with 8 attention heads and a hidden dimension of 512. We used an MLP with a hidden size of 2048 for g^i , and a hidden size of 512 for g^t in ITC. Both MLPs have an output size of 128. The temperature parameter τ for ITC is 0.1, and the masking ratio ρ for MTR is 0.5. The images are resized to 128×128 . During pre-training, we conducted tabular augmentation for ITC and ITM, and image augmentation for 3 pre-training tasks. Note that the CAD and Infarction tasks utilize the same pre-trained model during fine-tuning. Additional implementation details for TIP and other comparing models are in Sec. B of supp..

4.1 Comparison with SOTAs on Complete Downstream Data

We first investigate the performance of TIP in a complete downstream data regime by comparing it with other supervised and SSL pre-training algorithms. For supervised learning, we trained a fully supervised image model, ResNet-50, and reproduced 4 image-tabular learning strategies: concatenation fusion (CF) [62], maximum fusion (MF) [66], interactive fusion through channel-wise multiplication (IF) [24], and dynamic affine feature map transform (DAFT) [69]. For fair comparison, the image encoder used in all these approaches is ResNet-50. For SSL image pre-training, we tested 4 popular contrastive learning solutions: SimCLR [18], BYOL [28], SimSiam [19], and BarlowTwins [75]. We also compared our TIP with MMCL [29], a recent multimodal image-tabular pre-training method. We evaluated all pre-trained models using linear probing, which only tunes linear classifiers, and fully fine-tuning, which trains all parameters.

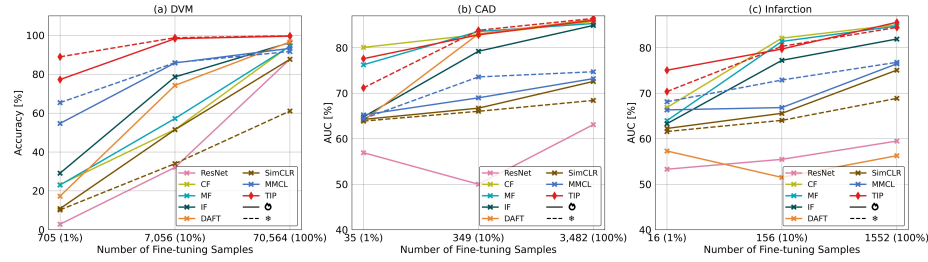


Fig. 3: Result comparison with supervised/SSL image/multimodal methods on various number of fine-tuning samples. \bullet denotes fully fine-tuning, and $*$ means linear probing.

As shown in Tab. 1(a,b), TIP outperforms supervised/SSL image-only models in linear probing and fully fine-tuning by a large margin, which indicates that integrating multiple modalities in pre-training improves the representation learning and that tabular information facilitates our classification tasks. Moreover, TIP significantly surpasses MMCL, *e.g.*, in linear probing, boosting the accuracy by 8.06% on DVM and AUC by 11.72% on CAD. While MMCL transfers tabular information related to visual features into the image branch and discards the tabular branch during fine-tuning, tabular data often contains task-related complementary information that is not visible in images [1,5]. Our results showcase TIP can exploit tabular information that is visible or non-visible in images to improve downstream tasks. Finally, compared to supervised multimodal methods (Tab. 1(a)), TIP achieves higher performance on DVM, *e.g.*, raising accuracy by 3.12% in linear probing. On CAD and Infarction tasks, TIP performs competitively against supervised multimodal methods, indicating the usefulness of features learnt via self-supervised pre-training and a requirement for a larger pre-training dataset (70,565 in DVM *vs.* 26,040 in UKBB).

Robustness to Low-data Regimes: As data annotation for downstream tasks is often costly, we propose to assess the performance of TIP and other SOTA methods on low-data regimes (10% and 1% of the original training data size). We used 7,056 (10%) and 705 (1%) training samples for DVM, 349 and 35 for CAD, and 156 and 16 for Infarction. For SSL image approaches, only SimCLR’s results are presented since it showed the best performance among them (complete results in Sec. C of supp.). Fig. 3 shows that TIP is more robust at low-data regimes and outperforms other SOTAs on DVM. For TIP, 10% of the training data can already achieve a performance close to that of 100%, indicating the potential to use fewer data for fine-tuning. Only for two cases, 1% CAD and 10% Infarction, TIP slightly underperforms CF, a supervised multimodal method, possibly due to the relatively small pre-training datasets of CAD and Infarction used by TIP.

4.2 Comparison with SOTAs on Incomplete Downstream Data

We conduct a study to assess the model performance on tackling tabular missingness. To achieve that, we introduce 4 types of missing scenarios: (a) random

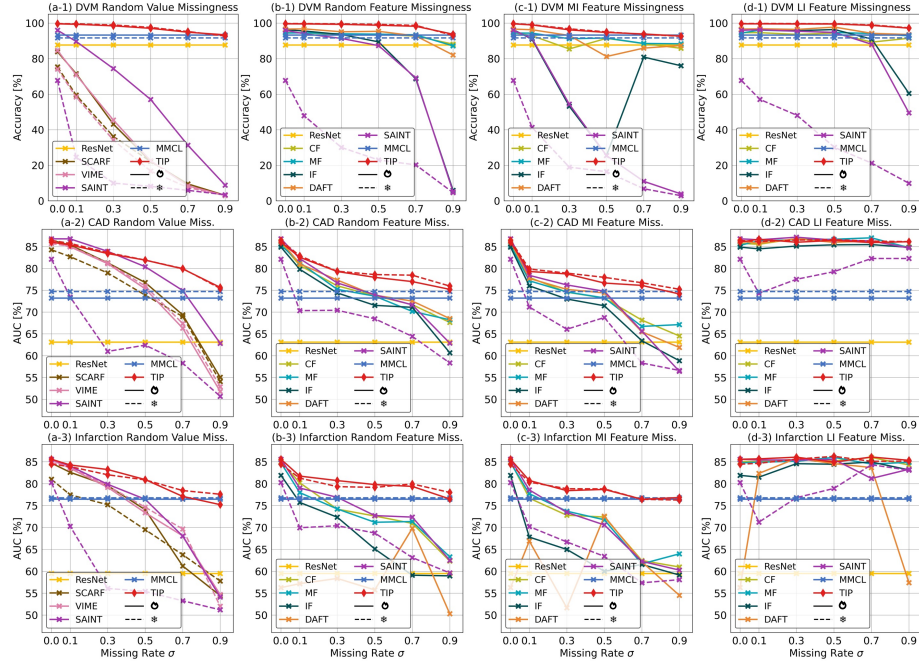


Fig. 4: Results of 4 missing scenarios: (a) RVM, (b) RFM, (c) MIFM, and (d) LIFM, on DVM (1st row), CAD (2nd row), and Infarction (3rd row) tasks with different missing rates σ . \bullet denotes fully fine-tuning, and \ast means linear probing.

value missingness (RVM), where tabular values (cells) are randomly missing; (b) random feature missingness (RFM), where a random set of features (columns) is missing; (c) most important feature missingness (MIFM), where the most important features for the prediction task are eliminated in descending order; (d) least important feature missingness (LIFM), where the least important features are removed first. The importance of features is determined using a random forest algorithm [47] trained on downstream training datasets. To showcase TIP’s capability in handling missing data using multimodal information, we compared it with 3 SSL tabular pre-training methods: VIME [73], SCARF [4], and SAINT [61]. MLP-based VIME and SCARF do not support feature missingness, while transformer-based SAINT can handle all missing scenarios. Hence, we only evaluated VIME and SCARF in RVM by filling missing positions with randomly chosen values from the same column, as in [4, 73]. Supervised multimodal approaches cannot address random value missingness. Thus, they are compared in RFM, MIFM, and LIFM. SSL image techniques and MMCL are not affected by missing tabular data. For complete comparison, we include the highest MMCL’s result among them. We executed each scenario with 6 missing rates.

As depicted in Fig. 4, the most challenging scenario is MIFM, where most models’ performance drops significantly. Besides, supervised multimodal meth-

Table 2: Result comparison of TIP and data imputation methods for reconstructing missing continuous features across various missing rates on DVM and UKBB test sets.


Model	DVM RMSE ↓			UKBB RMSE ↓		
	0.3	0.5	0.7	0.3	0.5	0.7
Mean [31]	0.9621	0.9783	0.9733	1.0162	1.0191	1.0070
MissForest [63]	0.6700	0.7653	0.8833	0.7516	0.7754	0.8177
GAIN [72]	1.0447	0.9428	2.9705	0.7920	2.0039	2.8130
MIWAE [51]	1.0105	1.0265	1.0218	1.0644	1.0680	1.0557
Hyperimpute [39]	0.6329	0.9428	0.9793	0.6803	0.7242	0.8060
TIP	0.3899	0.4651	0.5055	0.6039	0.6460	0.7106




ods outperform ResNet and MMCL in MIFM and RFM when $\sigma = 0.1$ and achieve improved results across various missing rates in LIFM. This indicates even if downstream data is incomplete, they can still provide useful information, especially when MI features are not missing. We notice that some models show an increase with higher missing rates in MIFM and attribute this to their feature importance not being exactly the same as identified by the random forest model.

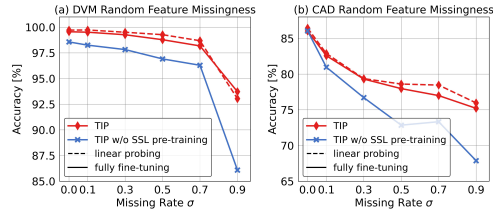
In comparison to other approaches, TIP can cope with all types of data missingness scenarios and significantly surpasses other methods. For supervised multimodal models (Fig. 4(b,c,d)), missing data can significantly reduce their performance, especially when the missing rate is high. However, TIP remains robust across different missing rates and exhibits improved performance. For instance, in RFM ($\sigma = 0.5$), TIP increases accuracy by 3.9% on DVM and AUC by 4.7% on CAD compared to DAFT. This implies that our pre-training strategy allows the model to capture valuable multimodal embeddings and their relationships from unlabeled image-tabular pairs.

Moreover, TIP performs better than tabular pre-training techniques (VIME and SCARF in Fig. 4(a)), especially with a high missing rate, *e.g.*, it increases accuracy by 75.02% on DVM when $\sigma = 0.5$. When compared to SAINT, which also uses learnable mask tokens during fine-tuning, TIP’s superior performance indicates that our SSL strategy is more effective for incomplete downstream data and allows the model to be able to integrate visual and tabular information for predicting missing components. Finally, compared to the multimodal pre-training model MMCL, TIP outperforms it by a large margin, even with a high missing rate of 0.7, in all scenarios. This shows that TIP can fully leverage tabular information in incomplete data. Even when certain tabular features in downstream tasks are inaccessible, the intra- and inter-modality relations learnt during pre-training enable TIP to generate promising outcomes.

Missing Value Reconstruction: We further assess TIP’s missing data reconstruction performance by comparing with 5 data imputation methods: column-wise mean substitution (Mean) [31], MissForest [63], GAIN [72], MIWAE [51], and HyperImpute [39]. Since GAIN and MIWAE are hard to apply to mixed-type data containing both continuous and categorical features, our experiments focus on the continuous data within DVM and UKBB test sets, using root mean squared error (RMSE) for evaluation. In this case, we masked the categorical

Table 3: Experiments using different image encoder backbones and ablation study of TIP. * means linear probing, and  represents fully fine-tuning.

Model	DVM Accuracy (%) \uparrow		CAD AUC (%) \uparrow		Infarction AUC (%) \uparrow	
	*		*		*	
(a) Applicability to Various Image Encoder Backbones						
TIP (ViT-S [23])	99.67	99.40	85.85	86.94	83.83	86.16
TIP (ViT-B [23])	99.40	99.28	84.90	86.93	83.15	85.76
(b) Ablation Study						
TIP w/o SSL pre-training	98.57	98.57	86.04	86.04	84.19	84.19
TIP w/o column name emb.	97.38	97.21	79.40	81.12	82.00	75.15
TIP w/o ensemble	99.63	99.35	86.00	86.97	84.43	84.00
TIP	99.72	99.56	86.43	86.03	84.46	85.58

**Fig. 5:** Results comparing TIP with or without the proposed SSL pre-training strategy on the DVM and the CAD RFM scenario.**Table 4:** TIP’s RMSE results on the DVM test set for missing continuous feature reconstruction. σ denotes missing rate, and ρ means masking ratio.

σ	0.3	0.5	0.7
$\rho = 0.1$.5349	.6752	.7871
$\rho = 0.3$.4110	.5128	.5924
$\rho = 0.5$.3899	.4651	.5055
$\rho = 0.7$.3986	.4612	.4733
$\rho = 0.9$.4279	.4800	.4816

data input for TIP. As shown in Tab. 2, TIP exceeds those imputation algorithms across varying missing rates, which showcases that our pre-training task and multimodal architecture have enabled the model to capture the relations with multimodal data and thus predict missing information more accurately.

4.3 Ablation Study and Visualization

Applicability to Different Image Encoder Backbones: We propose to vary the image encoder backbone to showcase the general applicability of the proposed method. Specifically, we utilized two vision transformers (ViTs): ViT-S/16 and ViT-B/16 [23], as the variations for the image encoder. Since ViTs output a sequence representation, we directly project it into the same hidden dimension as the tabular representation. Tab. 3(a) exemplifies that using ViTs results in similar outcomes to using ResNet-50, and large ViT-B does not perform better than small ViT-S. We suspect this is due to ViTs performing better than CNNs when pre-trained on much larger datasets, as found in [23].

Ablation Study on Key Model Components: We directly trained TIP’s model architecture on downstream tasks in a supervised manner to evaluate the efficacy of our SSL pre-training strategy. Tab. 3(b) and Fig. 5 demonstrate that our pre-training strategy improves the performance of downstream tasks, especially on incomplete data with a high missing rate. In addition, we conducted experiments that removed the column name embeddings in our tabular encoder

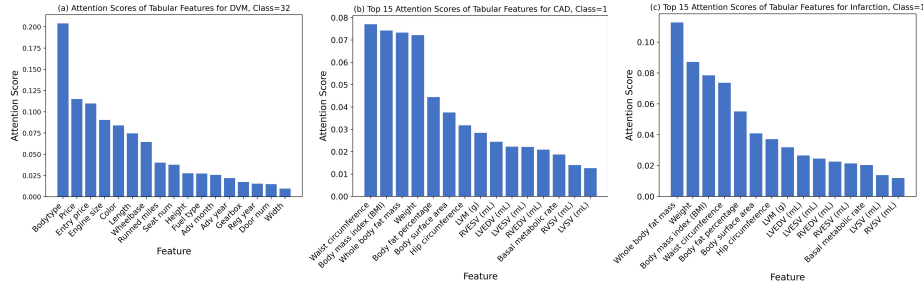


Fig. 6: The [CLS] token’s attention scores to tabular features for a particular class from the last layer of TIP’s tabular encoder.

or the ensemble learning during fine-tuning. Tab. 3(b) shows that subtracting any of those techniques leads to inferior performance. Additional ablation studies on each pre-training task and TIP’s tabular encoder in Sec. C.2-3 of supp..

Sensitivity Analysis of the Masking Ratio: We study the impact of different masking ratios ρ in the MTR task. Tab. 4 shows that moderate masking ratios $\rho \in (0.5, 0.7)$ achieve the best performance, whereas too high (0.9) or too low (0.1) ratios adversely affect model learning. More analysis in Sec C.4 of supp..

Visualization of TIP’s Tabular Feature Attention: In Fig. 6, we visualize TIP’s attention to different tabular features when predicting a specific class in downstream tasks. To achieve this, we average the self-attention map of samples belonging to the same class in test sets and present the attention scores of the [CLS] token. We observe that TIP attends to not only image-related features, *e.g.*, color in DVM, but also to features that are not directly visible in the image, *e.g.*, price in DVM. This showcases the importance of integrating multimodal data, as it can provide additional complementary information during downstream tasks, which can be difficult to obtain with image data alone. More visualization on cross-attention and case studies in Sec. C.5 of supp..

5 Conclusion

We have proposed TIP, a novel tabular-image pre-training framework for multimodal representation learning. TIP is a transformer-based multimodal network with a versatile tabular encoder and a multimodal interaction module, which are trained by a novel self-supervised pre-training strategy. In particular, TIP accounts for tabular data missingness, which makes it applicable to real-world datasets. Experiments on natural and medical image datasets have showcased TIP’s SOTA performance in various missing data scenarios and the efficacy of the proposed model components. The current work utilizes simulated missing data and 2D images. Future works will incorporate real-world incomplete data and extend to higher-dimensional images, *e.g.*, 3D and temporal imaging data. Potential societal impact is discussed in the supplementary material.

Acknowledgements

This research has been conducted using the UK Biobank Resource under Application Number 40616. The MR images presented in the figures are reproduced with the kind permission of UK Biobank ©. We also thank Paul Hager from the Lab for AI in Medicine at the Technical University of Munich for providing the pre-processing code for the UKBB dataset. DO'R is supported by the Medical Research Council (MC_UP_1605/13); National Institute for Health Research (NIHR) Imperial College Biomedical Research Centre; and the British Heart Foundation (RG/19/6/34387, RE/24/130023, CH/P/23/80008).

References

1. Acosta, J.N., Falcone, G.J., Rajpurkar, P., Topol, E.J.: Multimodal biomedical AI. *Nature Medicine* **28**(9), 1773–1784 (2022)
2. Antelmi, L., Ayache, N., Robert, P., Ribaldi, F., Garibotto, V., Frisoni, G.B., Lorenzi, M.: Combining multi-task learning and multi-channel variational auto-encoders to exploit datasets with missing observations-application to multi-modal neuroimaging studies in dementia. hal preprint hal-03114888v2 (2021)
3. Assran, M., Duval, Q., Misra, I., et al.: Self-supervised learning from images with a joint-embedding predictive architecture. In: CVPR. pp. 15619–15629 (2023)
4. Bahri, D., Jiang, H., Tay, Y., Metzler, D.: SCARF: Self-supervised contrastive learning using random feature corruption. In: ICLR (2022)
5. Bai, W., Suzuki, H., Huang, J., Francis, C., Wang, S., Tarroni, G., et al.: A population-based phenome-wide association study of cardiac and aortic structure and function. *Nature Medicine* **26**(10), 1654–1662 (2020)
6. Baltrušaitis, T., Ahuja, C., Morency, L.P.: Multimodal machine learning: A survey and taxonomy. *IEEE TPAMI* **41**(2), 423–443 (2018)
7. Barnard, J., Meng, X.L.: Applications of multiple imputation in medical studies: from AIDS to NHANES. *Statistical Methods in Medical Research* **8**(1), 17–36 (1999)
8. Bayasi, N., Hamarneh, G., Garbi, R.: Continual-Zoo: Leveraging zoo models for continual classification of medical images. In: CVPRW. pp. 4128–4138 (2024)
9. Bayoudh, K., Knani, R., Hamdaoui, F., Mtibaa, A.: A survey on deep multimodal learning for computer vision: advances, trends, applications, and datasets. *The Visual Computer* **38**(8), 2939–2970 (2022)
10. Borisov, V., Leemann, T., Sekler, K., Haug, J., Pawelczyk, M., Kasneci, G.: Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
11. Borsos, B., Allaart, C.G., van Halteren, A.: Predicting stroke outcome: A case for multimodal deep learning methods with tabular and CT perfusion data. *Artificial Intelligence in Medicine* **147**, 102719 (2024)
12. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., et al.: Language models are few-shot learners. *NIPS* **33**, 1877–1901 (2020)
13. Buntin, M.B., Burke, M.F., Hoaglin, M.C., Blumenthal, D.: The benefits of health information technology: a review of the recent literature shows predominantly positive results. *Health Affairs* **30**(3), 464–471 (2011)
14. Bycroft, C., Freeman, C., Petkova, D., et al.: The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**(7726), 203–209 (2018)

15. Cai, Q., Wang, H., et al.: A survey on multimodal data-driven smart healthcare systems: approaches and applications. *IEEE Access* **7**, 133583–133599 (2019)
16. Chaudhry, B., Wang, J., Wu, S., Maglione, M., Mojica, W., Roth, E., et al.: Systematic review: impact of health information technology on quality, efficiency, and costs of medical care. *Annals of Internal Medicine* **144**(10), 742–752 (2006)
17. Chen, F.L., Zhang, D.Z., Han, M.L., Chen, X.Y., et al.: VLP: A survey on vision-language pre-training. *Machine Intelligence Research* **20**(1), 38–56 (2023)
18. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *ICML*. pp. 1597–1607. PMLR (2020)
19. Chen, X., He, K.: Exploring simple siamese representation learning. In: *CVPR*. pp. 15750–15758 (2021)
20. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: *ACL*. pp. 4171–4186 (2019)
21. Dong, H., Cheng, Z., He, X., Zhou, M., Zhou, A., Zhou, F., Liu, A., Han, S., Zhang, D.: Table pre-training: A survey on model architectures, pre-training objectives, and downstream tasks. *arXiv preprint arXiv:2201.09745* (2022)
22. Dong, X., Yu, Z., Cao, W., Shi, Y., Ma, Q.: A survey on ensemble learning. *Frontiers of Computer Science* **14**, 241–258 (2020)
23. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2020)
24. Duanmu, H., Huang, P.B., Brahmavar, S., Lin, S., Ren, T., Kong, J., Wang, F., Duong, T.Q.: Prediction of pathological complete response to neoadjuvant chemotherapy in breast cancer using deep learning with integrative imaging, molecular and demographic data. In: *MICCAI*. pp. 242–252. Springer (2020)
25. Ganaie, M.A., Hu, M., Malik, A., et al.: Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence* **115**, 105151 (2022)
26. Ghorbani, A., Zou, J.Y.: Embedding for informative missingness: Deep learning with incomplete data. In: *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. pp. 437–445. IEEE (2018)
27. Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting deep learning models for tabular data. *NIPS* **34**, 18932–18943 (2021)
28. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Dohersch, C., Avila Pires, B., et al.: Bootstrap your own latent-a new approach to self-supervised learning. *NIPS* **33**, 21271–21284 (2020)
29. Hager, P., Menten, M.J., Rueckert, D.: Best of both worlds: Multimodal contrastive learning with tabular and imaging data. In: *CVPR*. pp. 23924–23935 (2023)
30. Han, X., Wang, Y.T., Feng, J.L., Deng, C., et al.: A survey of transformer-based multimodal pre-trained modals. *Neurocomputing* **515**, 89–106 (2023)
31. Hawthorne, G., Hawthorne, G., Elliott, P.: Imputing cross-sectional missing data: Comparison of common techniques. *Australian & New Zealand Journal of Psychiatry* **39**(7), 583–590 (2005)
32. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *CVPR*. pp. 16000–16009 (2022)
33. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. pp. 770–778 (2016)
34. He, R., McAuley, J.: Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In: *WWW*. pp. 507–517 (2016)
35. Heiliger, L., Sekuboyina, A., Menze, B., et al.: Beyond medical imaging-a review of multimodal deep learning in radiology. *Authorea Preprints* (2023)

36. Huang, J., Chen, B., Luo, L., et al.: DVM-CAR: A large-scale automotive dataset for visual marketing research and applications. In: 2022 IEEE International Conference on Big Data (Big Data). pp. 4140–4147. IEEE (2022)
37. Huang, S.C., Pareek, A., Seyyedi, S., Banerjee, I., Lungren, M.P.: Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ Digital Medicine* **3**(1), 136 (2020)
38. Huang, X., Khetan, A., Cvitkovic, M., Karnin, Z.: TabTransformer: Tabular data modeling using contextual embeddings. arXiv preprint arXiv:2012.06678 (2020)
39. Jarrett, D., Cebere, B.C., Liu, T., Curth, A., van der Schaar, M.: HyperImpute: Generalized iterative imputation with automatic model selection. In: ICML. pp. 9916–9937. PMLR (2022)
40. Jiang, J.P., Ye, H.J., Wang, L., Yang, Y., Jiang, Y., Zhan, D.C.: On transferring expert knowledge from tabular data to images. In: NIPSW (2023)
41. Jing, L., Tian, Y.: Self-supervised visual feature learning with deep neural networks: A survey. *IEEE TPAMI* **43**(11), 4037–4058 (2020)
42. Johnson, A.E., Pollard, T.J., Shen, L., Lehman, L.w.H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Anthony Celi, L., Mark, R.G.: MIMIC-III, a freely accessible critical care database. *Scientific Data* **3**(1), 1–9 (2016)
43. Kalyan, K.S., Rajasekharan, A., Sangeetha, S.: AMMU: a survey of transformer-based biomedical pretrained language models. *Journal of Biomedical Informatics* **126**, 103982 (2022)
44. Ko, W., Jung, W., Jeon, E., Suk, H.I.: A deep generative–discriminative learning for multimodal representation in imaging genetics. *IEEE Transactions on Medical Imaging* **41**(9), 2348–2359 (2022)
45. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML. pp. 12888–12900. PMLR (2022)
46. Li, J., Selvaraju, R., Gotmare, A., et al.: Align before fuse: Vision and language representation learning with momentum distillation. *NIPS* **34**, 9694–9705 (2021)
47. Liaw, A., Wiener, M., et al.: Classification and regression by randomforest. *R News* **2**(3), 18–22 (2002)
48. Littlejohns, T.J., Holliday, J., Gibson, L.M., Garratt, S., et al.: The UK Biobank imaging enhancement of 100,000 participants: rationale, data collection, management and future directions. *Nature Communications* **11**(1) (2020)
49. Mackinnon, A.: The use and reporting of multiple imputation in medical research—a review. *Journal of Internal Medicine* **268**(6), 586–593 (2010)
50. Majmundar, K.A., Goyal, S., Netrapalli, P., Jain, P.: MET: Masked encoding for tabular data. In: NIPSW (2022)
51. Mattei, P.A., Frellsen, J.: MIWAE: Deep generative modelling and imputation of incomplete data sets. In: ICML. pp. 4413–4423. PMLR (2019)
52. Miao, X., Wu, Y., et al.: An experimental survey of missing data imputation algorithms. *IEEE Transactions on Knowledge and Data Engineering* (2022)
53. Min, B., Ross, H., Sulem, E., Veyseh, A.P.B., Nguyen, T.H., et al.: Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys* **56**(2), 1–40 (2023)
54. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *NIPS* **35**, 27730–27744 (2022)
55. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. pp. 2536–2544 (2016)

56. Pölsterl, S., Wolf, T.N., Wachinger, C.: Combining 3D image and tabular data via the dynamic affine feature map transform. In: MICCAI. pp. 688–698. Springer (2021)
57. Radford, A., Kim, J.W., Hallacy, C., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763. PMLR (2021)
58. Raghunathan, T.E., Lepkowski, J.M., Van Hoewyk, J., Solenberger, P., et al.: A multivariate technique for multiply imputing missing values using a sequence of regression models. *Survey Methodology* **27**(1), 85–96 (2001)
59. Royston, P., White, I.R.: Multiple imputation by chained equations (MICE): implementation in stata. *Journal of Statistical Software* **45**, 1–20 (2011)
60. Schafer, J.L., Graham, J.W.: Missing data: our view of the state of the art. *Psychological Methods* **7**(2), 147 (2002)
61. Somepalli, G., Goldblum, M., Schwarzschild, A., Bruss, C.B., Goldstein, T.: SAINT: Improved neural networks for tabular data via row attention and contrastive pre-training. arXiv preprint arXiv:2106.01342 (2021)
62. Spasov, S., Passamonti, L., Duggento, A., Lio, P., Toschi, N., et al.: A parameter-efficient deep learning approach to predict conversion from mild cognitive impairment to alzheimer’s disease. *NeuroImage* **189**, 276–287 (2019)
63. Stekhoven, D.J., Bühlmann, P.: MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics* **28**(1), 112–118 (2012)
64. Sun, K., Luo, X., Luo, M.Y.: A survey of pretrained language models. In: International Conference on Knowledge Science, Engineering and Management. pp. 442–456. Springer (2022)
65. Ucar, T., Hajiramezanali, E., Edwards, L.: SubTab: Subsetting features of tabular data for self-supervised representation learning. *NIPS* **34**, 18853–18865 (2021)
66. Vale-Silva, L.A., Rohr, K.: Long-term cancer survival prediction using multimodal deep learning. *Scientific Reports* **11**(1), 13505 (2021)
67. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *NIPS* **30** (2017)
68. Wang, Z., Sun, J.: TransTab: Learning transferable tabular transformers across tables. *NIPS* **35**, 2902–2915 (2022)
69. Wolf, T.N., Pölsterl, S., et al.: DAFT: a universal module to interweave tabular data and 3D images in CNNs. *NeuroImage* **260**, 119505 (2022)
70. Yang, J., Gupta, A., Upadhyay, S., He, L., Goel, R., Paul, S.: TableFormer: Robust transformer modeling for table-text encoding. In: ACL. pp. 528–537 (2022)
71. Ye, C., Lu, G., Wang, H., et al.: CT-BERT: learning better tabular representations through cross-table pre-training. arXiv preprint arXiv:2307.04308 (2023)
72. Yoon, J., Jordon, J., Schaar, M.: Gain: Missing data imputation using generative adversarial nets. In: ICML. pp. 5689–5698. PMLR (2018)
73. Yoon, J., Zhang, Y., et al.: VIME: Extending the success of self-and semi-supervised learning to tabular domain. *NIPS* **33**, 11033–11043 (2020)
74. Yu, J., Wang, Z., Vasudevan, V., et al.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
75. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow Twins: Self-supervised learning via redundancy reduction. In: ICML. pp. 12310–12320. PMLR (2021)
76. Zhang, C., Zhang, C., Song, J., Yi, J.S.K., Kweon, I.S.: A survey on masked auto-encoder for visual self-supervised learning. In: IJCAI. pp. 6805–6813 (2023)
77. Zheng, H., Lin, Z., Zhou, Q., Peng, X., Xiao, J., Zu, C., Jiao, Z., Wang, Y.: Multi-transSP: Multimodal transformer for survival prediction of nasopharyngeal carcinoma patients. In: MICCAI. pp. 234–243. Springer (2022)