

TabMixer: Noninvasive Estimation of the Mean Pulmonary Artery Pressure via Imaging and Tabular Data Mixing

Michał K. Grzeszczyk¹, Przemysław Korzeniowski¹, Samer Alabed², Andrew J. Swift², Tomasz Trzcíński^{3,4,5}, and Arkadiusz Sitek⁶

¹ Sano Centre for Computational Medicine, Cracow, Poland

m.grzeszczyk@sano-science.org

² University of Sheffield, Sheffield, England

³ Warsaw University of Technology, Warsaw, Poland

⁴ IDEAS NCBR, Warsaw, Poland

⁵ Tooploox, Wrocław, Poland

⁶ Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA

Abstract. Right Heart Catheterization is a gold standard procedure for diagnosing Pulmonary Hypertension by measuring mean Pulmonary Artery Pressure (mPAP). It is invasive, costly, time-consuming and carries risks. In this paper, for the first time, we explore the estimation of mPAP from videos of noninvasive Cardiac Magnetic Resonance Imaging. To enhance the predictive capabilities of Deep Learning models used for this task, we introduce an additional modality in the form of demographic features and clinical measurements. Inspired by all-Multilayer Perceptron architectures, we present TabMixer, a novel module enabling the integration of imaging and tabular data through spatial, temporal and channel mixing. Specifically, we present the first approach that utilizes Multilayer Perceptrons to interchange tabular information with imaging features in vision models. We test TabMixer for mPAP estimation and show that it enhances the performance of Convolutional Neural Networks, 3D-MLP and Vision Transformers while being competitive with previous modules for imaging and tabular data. Our approach has the potential to improve clinical processes involving both modalities, particularly in noninvasive mPAP estimation, thus, significantly enhancing the quality of life for individuals affected by Pulmonary Hypertension. We provide a source code for using TabMixer at <https://github.com/SanoScience/TabMixer>.

Keywords: MLP · Pulmonary Hypertension · Tabular Data.

1 Introduction

Pulmonary Hypertension (PH) is a severe disease characterised by elevated pressure in the Main Pulmonary Artery. The gold standard for detecting PH is Right Heart Catheterization (RHC), an invasive procedure enabling the assessment of pulmonary circulation characteristics including mean Pulmonary Artery Pressure (mPAP), cardiac output, or stroke volume. PH is diagnosed when mPAP

exceeds 20 mmHg [12]. Regrettably, as an invasive procedure, RHC is prone to complications as well as it requires specialized equipment, trained staff and patient preparation. As a consequence, researchers have adopted Machine Learning (ML) models to diagnose PH from demographics data [16], measurements from cardiac computed tomography [14], physics-based models parameters [19] and Magnetic Resonance Imaging-derived (MRI) measurements [1,9,24]. In this study, we take a new route and we explore the potential of estimating mPAP directly from Cardiac MRI (CMR) videos of one cardiac cycle collected at two planes: 4 Chamber (4CH) and Short-Axis (SA) using Deep Learning.

In recent years, 3D Convolutional Neural Networks (CNNs) [3,28] and Vision Transformers [2,17] were the go-to standard for video processing. More recently the emergence of Multi-layer Perceptrons (MLP)-based models [5,26,27] has led to the development of their 3D counterparts [23]. In the medical domain, the common approach of collecting imaging and tabular data resulted in the introduction of fusion methods. Traditional methods involve concatenating tabular features in the network’s final layers [13], but this restricts the knowledge transfer with imaging data. To enhance the interaction between imaging and tabular data Duanmu *et al.* [7] proposed a method in which imaging features are channel-wise multiplied with tabular data at different stages of CNN. Pölsterl *et al.* [22] introduced Dynamic Affine Feature Map Transform (DAFT) module, an extension of the Feature-wise Linear Modulation (FiLM) [21] layer for scaling and shifting feature maps based on tabular data. Grzeszczyk *et al.* [10] proposed TabAttention, a module for learning channel, spatial and temporal attention conditioned on tabular data. In this work, we enhance the predictive capabilities of imaging models by introducing tabular data in the form of demographic features and measurements from MRI (e.g. right ventricular ejection fraction).

Traditionally, MLPs were used for tabular data processing and more recently for image analysis, however, no approach combines both modalities. Inspired by all-MLP architectures, we present TabMixer, a novel module facilitating the mixing of tabular data with imaging features. TabMixer enriches the video processing backbones by enabling interaction between tabular embeddings and imaging features across spatial, temporal and channel dimensions. Unlike convolutional or self-attention-based methods, TabMixer relies solely on MLPs, allowing the fusion of information stored in tabular data with imaging features during spatial, temporal and channel mixing. We evaluate TabMixer integrated with different backbones (CNNs, Transformer, all-MLP) on the mPAP prediction task and show its competitiveness with existing methods. Our contributions are three-fold: (1) we introduce TabMixer, a first module akin to MLP designed for the integration of imaging feature maps with tabular data, (2) we provide a comprehensive evaluation and comparison of methods that utilize imaging and/or tabular data for predicting mPAP, and (3) to our knowledge, this is the first demonstration of noninvasive mPAP estimation directly from CMR videos. In addition, we present a comparative analysis of the 4 Chamber (4CH) and Short-Axis (SA) planes for this estimation.




Fig. 1. The overview of the proposed TabMixer. TabMixer applies mixing operations over spatial, temporal and channel dimensions. Leveraging MLPs, TabMixer can interchange information between imaging and tabular features.

2 Method

In this section, we introduce the building blocks of the TabMixer module. We present the rationale behind employing MLP architecture for imaging and tabular data. Finally, we describe the integration of TabMixer with vision backbones.

The overview of TabMixer is presented in Fig. 1. Let $X_0 \in \mathbb{R}^{1 \times T_0 \times H_0 \times W_0}$ denote the grayscale CMR video comprising T_0 frames with H_0 height and W_0 width. The vision backbone generates intermediate temporal feature maps $X \in \mathbb{R}^{C \times T \times H \times W}$ where C is the number of channels. TabMixer refines these temporal feature maps through spatial, temporal and channel mixing with the aid of tabular data $Tab \in \mathbb{R}^D$. We follow the architectural principles of all-MLP and Transformer networks in the design of this module [29].

Initially, to reduce dimensionality and the number of parameters, the intermediate temporal feature maps X are embedded via average pooling and reshaping operations into a features cube $X' \in \mathbb{R}^{C \times T \times S}$ where S is $\frac{(H \times W)}{4}$. This structure of X' enables the interchange of spatial, temporal and channel information with tabular data via three sub-layers L_s , L_t and L_c , sequentially operating across every dimension. Each L_i component comprises an Affine transformation, MLP_i layer and a skip-connection followed by permutation operation to facilitate processing over the next dimension. Finally, the processed features are reshaped and linearly upsampled back to the input shape, thereby returning refined temporal feature maps integrated with tabular information.

Mixing with Multi-Layer Perceptron. In many networks, a normalization layer precedes model-specific layers (e.g. attention) to stabilize training. Inspired by [27], instead of Layer Normalization, we perform Affine transformation in L_i , to rescale and shift inputs element-wise without depending on batch statistics:

$$Aff_{\alpha, \beta}(X') = Diag(\alpha)X' + \beta \quad (1)$$

After the affine transformation, the three-dimensional cube is processed with MLP_i block consisting of two MLP layers and GELU [11] activation:

$$Y = MLP_{i,2}(GELU(MLP_{i,1}(Z))) \quad (2)$$

Here Z is the input to the MLP_i and Y is the output of the same size. The MLP_i block enables mixing information across the input’s last dimension and is shared across all vectors in this dimension as in [26]. The first MLP compresses the last dimension by a factor of two to create a bottleneck, while the second MLP expands it back to the initial size (as shown in Fig. 1). Subsequently, we employ the skip connection with the L_i input. As MLP_{i+1} operates on different dimension, the final operation within L_i entails axis permutation. The overall dimensionality evolves in TabMixer as follows:

$$\mathbb{R}^{C \times T \times H \times W}_{\text{input}} \xrightarrow{L_s} \mathbb{R}^{C \times T \times S} \xrightarrow{L_t} \mathbb{R}^{C \times S \times T} \xrightarrow{L_c} \mathbb{R}^{S \times T \times C} \xrightarrow{\text{output}} \mathbb{R}^{C \times T \times H \times W} \quad (3)$$

Tabular data within TabMixer. To incorporate tabular information within spatial, temporal and channel processing we embed tabular data with MLP_{tab} ($Tab \in \mathbb{R}^D \rightarrow Tab' \in \mathbb{R}^D$) duplicating the structure of MLP_i blocks. Inspired by the use of MLPs in tabular Deep Learning models and all-MLP networks for vision, we concatenate tabular embedding with the X' features cube before processing it via MLP_i . This concatenation involves repeating tabular embedding over the first and second axes of X' , augmenting the input to MLP_i by D values from tabular embedding (e.g. $\mathbb{R}^{C \times T \times S+D}$ for MLP_s). This approach enables the fusion of tabular features with imaging ones across all dimensions, as shown in Fig. 1. The processing of feature maps with tabular data can be expressed as:

$$\begin{aligned} TabMixer(X, Tab) &= U(L_c(L_t(L_s(R(X), Tab'), Tab'), Tab')) \\ Tab' &= MLP_{tab}(Tab) \end{aligned} \quad (4)$$

Here, R and U denote the reshaping with pooling and upsampling for simplicity. The design of TabMixer enables its integration with networks based on a hierarchical structure (e.g. ResNet-18, 3D-MLP) yielding intermediate feature maps.

3 Experiments and results

This section describes the PH dataset used for mPAP estimation and outlines the implementation details of TabMixer. We benchmark our method against various state-of-the-art modules and compare its performance with imaging-only and tabular-only ML and Deep Learning models. Additionally, we conduct an ablation study and demonstrate the importance of TabMixer’s main components.

Dataset. This study was approved by the Ethics Committee. The dataset employed comprises 1821 CMR studies (3642 videos of SA and 4CH planes) matched with invasively measured mPAP and is part of the ASPIRE Registry (Assessing the Severity of Pulmonary Hypertension In a Pulmonary Hypertension REferral




Fig. 2. Short-Axis (top) and corresponding 4 Chamber (bottom) frames from CMR depicting a cardiac cycle from left to right.

Centre) [15]. It contains data from 1708 patients (1045 females and 663 males, 64 ± 14 years old) suspected of PH who underwent both MRI and RHC within 48 hours (some patients had repeated the procedures over the years). To ensure a fair comparison across planes and methods based on tabular data we extract data samples from the ASPIRE Registry containing SA and 4CH CMR videos, and select tabular features (demographics and measurements from MRI) with no more than 200 empty values. We remove data samples without tabular features. The MRI videos, capturing one cardiac cycle (up to 50 frames) were acquired during a breath hold in the supine position, utilizing devices from various vendors, including GE, Philips and Siemens. Exemplary frames from corresponding SA and 4CH planes are shown in Fig. 2.

Implementation details. We implement all Deep Learning models using PyTorch and train them on NVIDIA A100 80 GPU for 100 epochs. To minimize the Mean Squared Error loss of the mPAP prediction we use AdamW optimizer [18] with a cosine annealing rate scheduler. An initial learning rate is set to 1×10^{-4} , L2 regularization to 1×10^{-5} and batch size to 8. We split the dataset into training, validation and testing sets with 1299, 217 and 305 samples respectively, ensuring that each patient’s data is present in only one set. We stratify the split based on the four mPAP bins (≤ 20 mmHg, ≤ 25 mmHg, ≤ 30 mmHg, > 30 mmHg). We resample the pixel spacing of all frames to $0.9375 \text{ mm} \times 0.9375 \text{ mm}$, pad them to 512×512 pixels (the maximum frame size in the dataset) and resize entire videos to 192×192 pixels with 16 frames. During training, we apply data augmentation methods, including translation, rotation, contrast adjustment, gaussian noise, intensity scaling and shifting. We one-hot categorical tabular features and standardize the numerical ones to a mean value of 0 and a standard deviation of 1. We retain tabular features exhibiting statistical significance measured with f-regression [20] and evaluate the performance in estimating mPAP using Mean Absolute Error [mmHg] (MAE), Root Mean Squared Error [mmHg] (RMSE) and Mean Absolute Percentage Error [%] (MAPE). Due to space limitations, we present mean results here, while standard deviations and the description of 29 tabular features are in the supplementary material.

Table 1. The performance of imaging (I.) and/or tabular (T.) methods. We **highlight** best and underline second MAE [mmHg], RMSE [mmHg], MAPE [%]. We indicate the statistically significant difference (two-tailed paired t-test p-value < 0.01) with TabMixer over the baseline on both 4CH and SA planes with *.

Method	I.	T.	MAE	RMSE	MAPE	MAE	RMSE	MAPE
LR [20]	✗	✓	8.50	11.01	23.35	—	—	—
MLP[20]	✗	✓	8.25	10.87	22.51	—	—	—
Trompt [4]	✗	✓	8.12	10.61	22.38	—	—	—
XGBoost [6]	✗	✓	7.82	10.11	21.84	—	—	—
ResNet _{tab} [8]	✗	✓	7.65	9.96	21.10	—	—	—
FT-Transformer [8]	✗	✓	7.62	10.03	<u>20.89</u>	—	—	—
GBDT [20]	✗	✓	<u>7.58</u>	<u>9.78</u>	20.85	—	—	—
RF [20]	✗	✓	7.47	9.61	20.97	—	—	—
			Short-Axis			4 Chamber		
Video Swin [17]	✓	✗	11.65	14.10	34.82	12.01	14.47	36.21
+ Concat	✓	✓	10.09	12.34	29.99	10.17	12.43	30.53
+ FiLM [21]	✓	✓	<u>7.98</u>	10.42	<u>22.62</u>	8.10	<u>10.43</u>	<u>22.80</u>
+ DAFT [22]	✓	✓	11.88	14.41	36.64	<u>8.09</u>	10.51	22.96
+ TabAttention [10]	✓	✓	8.02	<u>10.38</u>	23.07	8.60	11.11	25.12
+ TabMixer*	✓	✓	7.86	10.22	22.25	7.92	10.34	22.30
MLP-3D [23]	✓	✗	11.84	14.24	35.72	12.04	14.51	36.09
+ Concat	✓	✓	9.59	11.97	27.87	10.23	12.45	30.13
+ FiLM [21]	✓	✓	<u>8.13</u>	10.67	<u>22.62</u>	<u>8.01</u>	<u>10.54</u>	22.28
+ DAFT [22]	✓	✓	11.41	13.83	33.48	8.95	11.54	24.51
+ TabAttention [10]	✓	✓	8.29	<u>10.62</u>	24.55	8.48	10.72	24.48
+ TabMixer*	✓	✓	7.80	10.20	22.02	7.97	10.32	<u>22.76</u>
ResNet _{DAFT} [22]	✓	✗	12.09	14.53	35.48	12.24	14.67	34.91
+ Concat	✓	✓	10.50	12.82	31.97	9.79	11.81	28.93
+ FiLM [21]	✓	✓	<u>8.24</u>	<u>10.69</u>	23.38	8.07	10.67	22.41
+ DAFT [22]	✓	✓	8.25	10.98	<u>22.06</u>	<u>7.87</u>	<u>10.51</u>	<u>22.10</u>
+ TabAttention [10]	✓	✓	8.54	11.19	23.65	8.26	10.83	23.29
+ TabMixer*	✓	✓	7.82	10.22	22.04	7.68	10.24	21.50
ResNet-18 [28]	✓	✗	7.85	10.21	21.71	8.05	10.47	23.34
+ Concat	✓	✓	7.22	9.36	19.92	8.05	10.43	22.74
+ FiLM [21]	✓	✓	9.05	11.63	24.58	8.91	11.44	24.40
+ DAFT [22]	✓	✓	8.93	11.53	25.43	8.14	10.39	23.00
+ TabAttention [10]	✓	✓	7.92	10.03	22.77	<u>7.72</u>	<u>9.85</u>	<u>21.76</u>
+ TabMixer	✓	✓	<u>7.56</u>	<u>9.94</u>	<u>21.57</u>	7.53	9.91	21.75
I3D [3]	✓	✗	7.21	9.39	<u>19.11</u>	8.07	10.29	23.05
+ Concat	✓	✓	7.56	9.72	20.76	7.86	9.83	20.86
+ FiLM [21]	✓	✓	8.53	11.11	22.48	9.10	11.66	24.34
+ DAFT [22]	✓	✓	<u>7.18</u>	<u>9.25</u>	19.21	<u>7.74</u>	9.83	<u>20.56</u>
+ TabAttention [10]	✓	✓	7.34	9.49	21.34	7.77	<u>9.76</u>	<u>21.98</u>
+ TabMixer*	✓	✓	6.66	8.64	18.90	7.19	9.31	20.01

Table 2. The ablation study of the key components of the proposed TabMixer. The first row is the baseline I3D model. The next rows refer to I3D combined with TabMixer without tabular data or one of the mixing components and the last row is full TabMixer.

Method	I.	T.	MAE	RMSE	MAPE
I3D [3]	✓	✗	7.21 ± 6.01	9.39	19.11 ± 19.62
+ TM w/o tabular data	✓	✗	7.17 ± 5.71	9.17	19.97 ± 22.84
+ TM w/o channel mixing	✓	✓	7.08 ± 5.73	9.11	19.80 ± 22.72
+ TM w/o spatial & temporal mixing	✓	✓	7.01 ± 6.01	9.24	19.07 ± 20.00
+ TM w/o spatial mixing	✓	✓	7.00 ± 5.78	9.08	19.11 ± 20.12
+ TM w/o temporal mixing	✓	✓	6.89 ± 5.83	9.02	19.60 ± 21.72
+ TabMixer (TM)	✓	✓	6.66 ± 5.51	8.64	18.90 ± 20.83

Comparison with state-of-the-art methods. To compare TabMixer with several methods for imaging and tabular data (Concatenation, FiLM [21], DAFT [22], TabAttention [10]) we use five vision backbones: 3D CNNs (ResNet-18 [28], I3D [3], ResNet_{DAFT} [22]), Transformer (Video Swin Transformer [17]) and all-MLP network (MLP-3D [23]). In all backbones, we insert tabular modules before the global average pooling layer. We conduct the mPAP prediction on SA and 4CH planes to find the better one for this task. Additionally, we test the performance of mPAP estimation by methods based solely on tabular data. We experiment with Deep Learning models (Trompt [4], ResNet_{tab} [8], FT-Transformer [8], MLP) and ML - Linear Regression (LR), XGBoost [6], Gradient Boosting Decision Trees (GBDT), Random Forest (RF). Results are presented in Table 1. RF is the best-performing tabular-only method with MAE of 7.47 ± 6.05 , while I3D trained on the SA plane achieves the lowest error out of imaging-only methods (7.21 ± 6.01). TabMixer outperforms all modules when tested on the 4CH plane and has the best results for four out of five vision backbones on the SA plane (the improvement over the four baselines on both planes is statistically significant with a p-value below 0.002). Notably, the combination of I3D with TabMixer and SA plane reaches the lowest error with MAE of 6.66 ± 5.51 .

Ablation study. We conduct ablation experiments to validate the key components of TabMixer (Table 2). We perform validation of our proposed method when integrated with the I3D model. The incorporation of spatial, temporal and channel mixing components provides the lowest prediction error. Conversely, removing any of those blocks or tabular data mixing increases the error while still retaining performance gain over the baseline I3D.

4 Discussion and Conclusions

To the best of our knowledge, we are the first to directly estimate mPAP from CMR videos. Notably, our analysis indicates the superiority of the SA plane over 4CH in predicting mPAP which is the same CMR plane used for estimation of ejection fraction or stroke volume. Inspired by prior works combining imaging

Table 3. The parameters count and inference time of modules assuming feature maps shape from the I3D (1024, 4, 6, 6 - C, T, H, W) and the number of tabular features of 29. We also provide measures for TabMixer without channel mixing (TM w/o CM).

Method	FiLM	DAFT	TabAttention	TabMixer	TM w/o CM
Params [M]	0.015	0.022	0.202	1.070	0.001
Inference time [ms]	0.04	0.35	7.57	1.41	0.97




Fig. 3. The comparison of resistance to noisy imaging (left), tabular (middle) data and both (right). We compare the performance of I3D+TabMixer and the second best-performing method I3D+DAFT.

and tabular data, as well as all-MLP networks, we introduced TabMixer - a module for combining CMR and tabular data in 3D networks. TabMixer enhances the predictive capabilities of all investigated networks. It worked best with the I3D model (the improvement over the baseline is statistically significant with $p\text{-value} \approx 0.002$). However, while other methods exhibited improvement, they did not surpass the RF method. RF outperformed Deep Learning methods for tabular data like Prompt or FT-Transformer which is consistent with the challenges Deep Learning models face in effectively leveraging tabular data [25]. This highlights the importance of the right combination of the vision backbone and tabular module to the task.

To better understand the impact of imaging and tabular data, we experiment with noisy data applied to two best-performing methods I3D+TabMixer and I3D+DAFT. As shown in Fig. 3, both methods exhibit resilience to noise in tabular data with I3D+TabMixer showcasing greater robustness to noisy videos owing to its superior representation learning facilitated by tabular and imaging data mixing.

Our method exhibits certain limitations. Firstly, it was not tested on the external dataset and to prove its generalizability, further study is needed. Secondly, we experimented with estimations based on single CMR planes. The combination of both 4CH and SA videos could yield an additional improvement in the quality of predictions. Finally, TabMixer is not a lightweight module, due to the extensive use of MLP layers. As the spatial, temporal, and channel dimensions increase, the number of trainable parameters within the MLP layers escalates quadratically. For example, for I3D and 1024 channels in the final temporal feature maps the number of trainable parameters of TabMixer is higher than those

of FiLM or DAFT as shown in Table 3. However, TabMixer without channel mixing has much less trainable parameters but also outperforms all other modules in terms of MAE and RMSE as presented in the ablation study (Table 2).

In summary, this paper addresses the challenge of noninvasive mPAP estimation. We have introduced the TabMixer module, which facilitates the integration of tabular and imaging data. To our knowledge, this represents the first application of MLP models to merge these two data modalities. Our approach holds considerable promise for clinical workflows, particularly in mPAP estimation. By offering a noninvasive alternative, it has the potential to transform the field, conceivably obviating the need for high-risk invasive RHC procedures.

Acknowledgments. This research was funded in whole or in part by National Science Centre, Poland 2023/49/N/ST6/01841. For the purpose of Open Access, the author has applied a CC-BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission. This work is supported by the EU’s Horizon 2020 programme (grant no. 857533, Sano) and the Foundation for Polish Science’s International Research Agendas programme, co-financed by the EU under the European Regional Development Fund and the Polish Ministry of Science and Higher Education (contract no. MEiN/2023/DIR/3796).

Disclosure of Interests. The authors have no conflicts of interest to declare.

References

1. Alabed, S., Alandejani, F., Dwivedi, K., Karunasaagarar, K., Sharkey, M., Garg, P., de Koning, P.J., Tóth, A., Shahin, Y., Johns, C., et al.: Validation of artificial intelligence cardiac mri measurements: relationship to heart catheterization and mortality prediction. *Radiology* **305**(1), 68–79 (2022)
2. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 6836–6846 (October 2021)
3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6299–6308 (2017)
4. Chen, K.Y., Chiang, P.H., Chou, H.R., Chen, T.W., Chang, T.H.: Trompt: Towards a better deep neural network for tabular data. *arXiv preprint arXiv:2305.18446* (2023)
5. Chen, S., Xie, E., GE, C., Chen, R., Liang, D., Luo, P.: CycleMLP: A MLP-like architecture for dense prediction. In: *International Conference on Learning Representations* (2022)
6. Chen, T., Guestrin, C.: XGBoost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 785–794. *KDD ’16*, ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2939672.2939785>
7. 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: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23. pp. 242–252. Springer (2020)
8. Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems* **34**, 18932–18943 (2021)
 9. Grzeszczyk, M.K., Siatkawa, T., Lungu, A., Swift, A., Narracott, A., Hose, R., Trzcinski, T., Sitek, A.: Noninvasive estimation of mean pulmonary artery pressure using mri, computer models, and machine learning. In: International Conference on Computational Science. pp. 14–27. Springer (2022)
 10. Grzeszczyk, M.K., et al.: Tabattention: Learning attention conditionally on tabular data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 347–357. Springer (2023)
 11. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
 12. Hoepfer, M.M., et al.: Pulmonary hypertension. *Dtsch Arztebl Int* **114**, 73–84 (2017). <https://doi.org/10.3238/arztebl.2017.0073>
 13. Holste, G., Partridge, S.C., Rahbar, H., Biswas, D., Lee, C.I., Alessio, A.M.: End-to-end learning of fused image and non-image features for improved breast cancer classification from mri. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3294–3303 (2021)
 14. Huang, L., Li, J., Huang, M., Zhuang, J., Yuan, H., Jia, Q., Zeng, D., Que, L., Xi, Y., Lin, J., Dong, Y.: Prediction of pulmonary pressure after glenn shunts by computed tomography-based machine learning models. *European Radiology* **30**, 1369–1377 (2020). <https://doi.org/10.1007/s00330-019-06502-3>
 15. Hurdman, J., Condliffe, R., Elliot, C., Davies, C., Hill, C., et al.: Aspire registry: Assessing the spectrum of pulmonary hypertension identified at a referral centre. *European Respiratory Journal* **39**, 945–955 (4 2012). <https://doi.org/10.1183/09031936.00078411>
 16. Leha, A., Hellenkamp, K., Unsöld, B., Mushemi-Blake, S., Shah, A.M., Hasenfuß, G., Seidler, T.: A machine learning approach for the prediction of pulmonary hypertension. *PLoS ONE* **14** (10 2019). <https://doi.org/10.1371/journal.pone.0224453>
 17. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3202–3211 (2022)
 18. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
 19. Lungu, A., Swift, A.J., Capener, D., Kiely, D., Hose, R., Wild, J.M.: Diagnosis of pulmonary hypertension from magnetic resonance imaging-based computational models and decision tree analysis. *Pulmonary Circulation* **6**, 181–190 (6 2016). <https://doi.org/10.1086/686020>
 20. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
 21. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)

22. Pölsterl, S., Wolf, T.N., Wachinger, C.: Combining 3d image and tabular data via the dynamic affine feature map transform. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24. pp. 688–698. Springer (2021)
23. Qiu, Z., Yao, T., Ngo, C.W., Mei, T.: Mlp-3d: A mlp-like 3d architecture with grouped time mixing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3062–3072 (2022)
24. Reiter, U., Kovacs, G., Reiter, C., Kräuter, C., Nizhnikava, V., Fuchsjäger, M., Olschewski, H., Reiter, G.: Mr 4d flow-based mean pulmonary arterial pressure tracking in pulmonary hypertension. *European Radiology* **31**, 1883–1893 (2021)
25. Schwartz-Ziv, R., Armon, A.: Tabular data: Deep learning is not all you need. *Information Fusion* **81**, 84–90 (2022)
26. Tolstikhin, I.O., et al.: Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems* **34**, 24261–24272 (2021)
27. Touvron, H., et al.: Resmlp: Feedforward networks for image classification with data-efficient training. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 5314–5321 (2022)
28. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6450–6459 (2018)
29. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10819–10829 (2022)