



**HAL**  
open science

## **LMT: Longitudinal mixing training, a framework for the prediction of disease progression using a single image**

Rachid Zeghlache, Pierre-Henri Conze, Mostafa El Habib Daho, Yihao Li, Hugo Le Boité, Pascale Massin, Ramin Tadayoni, Béatrice Cochener, Ikram Brahim, Gwenolé Quéllec, et al.

### ► **To cite this version:**

Rachid Zeghlache, Pierre-Henri Conze, Mostafa El Habib Daho, Yihao Li, Hugo Le Boité, et al.. LMT: Longitudinal mixing training, a framework for the prediction of disease progression using a single image. International Workshop on Machine Learning in Medical Imaging, Oct 2023, Vancouver, Canada. hal-04183655

**HAL Id: hal-04183655**

**<https://imt-atlantique.hal.science/hal-04183655>**

Submitted on 26 Jan 2024

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

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

# LMT: Longitudinal Mixing Training, a framework to predict disease progression from a single image

Rachid Zeghlache<sup>1,2</sup>, Pierre-Henri Conze<sup>1,3</sup>, Mostafa El Habib Daho<sup>1,2</sup>, Yihao Li<sup>1,2</sup>, Hugo Le boîte<sup>5</sup>, Ramin Tadayoni<sup>5</sup>, Pascal Massin<sup>5</sup>, Béatrice Cochener<sup>1,2,4</sup>, Ikram Brahim<sup>1,6</sup>, Gwenolé Quéllec<sup>1</sup>, and Mathieu Lamard<sup>1,2</sup>

<sup>1</sup> LaTIM UMR 1101, Inserm, Brest, France

<sup>2</sup> University of Western Brittany, Brest, France

<sup>3</sup> IMT Atlantique, Brest, France

<sup>4</sup> Ophthalmology Department, CHRU Brest, Brest, France

<sup>5</sup> Lariboisière Hospital, AP-HP, Paris, France

<sup>6</sup> LBAI UMR 1227, Inserm, Brest, France

**Abstract.** Longitudinal imaging is able to capture both static anatomical structures and dynamic changes in disease progression toward earlier and better patient-specific pathology management. However, conventional approaches rarely take advantage of longitudinal information for detection and prediction purposes, especially for Diabetic Retinopathy (DR). In the past years, Mix-up training and pretext tasks with longitudinal context have effectively enhanced DR classification results and captured disease progression. In the meantime, a novel type of neural network named Neural Ordinary Differential Equation (NODE) has been proposed for solving ordinary differential equations, with a neural network treated as a black box. By definition, NODE is well suited for solving time-related problems. In this paper, we propose to combine these three aspects to detect and predict DR progression. Our framework, Longitudinal Mixing Training (LMT), can be considered both as a regularizer and as a pretext task that encodes the disease progression in the latent space. Additionally, we evaluate the trained model weights on a downstream task with a longitudinal context using standard and longitudinal pretext tasks. We introduce a new way to train time-aware models using  $t_{mix}$ , a weighted average time between two consecutive examinations. We compare our approach to standard mixing training on DR classification using OPHDIAT a longitudinal retinal Color Fundus Photographs (CFP) dataset. We were able to predict whether an eye would develop a severe DR in the following visit using a single image, with an AUC of 0.798 compared to baseline results of 0.641. Our results indicate that our longitudinal pretext task can learn the progression of DR disease and that introducing  $t_{mix}$  augmentation is beneficial for time-aware models.

**Keywords:** Disease progression · mix-up training · diabetic retinopathy · time-aware model · predictive medicine

## 1 Introduction

According to the International Diabetes Federation, by 2045, diabetes will impact 700 million individuals globally, with over one-third suffering from Diabetic Retinopathy (DR) [20]. DR is the leading cause of vision loss worldwide and is caused by high blood sugar damaging retinal vessels, leading to swelling and leakage [15]. Color fundus photographs (CFP) are used clinically to detect DR. Severity is graded in five classes (0, 1, 2, 3 and 4) using the International Clinical DR (ICDR) scale: 0 is no apparent DR, 1 is mild non-proliferative DR (NPDR), 2 is moderate NPDR, 3 is severe NPDR, and 4 is proliferative DR (PDR). Early detection and treatment, particularly in mild to moderate NPDR, may slow DR progression and reduce blindness incidence. Very few papers try to predict the progression of DR using a single CFP [18,2]. Despite its difficulty, this task is crucial for better patient follow-up management.

Recently, longitudinal pretext tasks (LPT) have emerged to encode disease progression, such as Longitudinal Self-Supervised Learning (LSSL) introduced by Rivail et al. [17] using a Siamese network to predict time lapses between consecutive retinal Optical Coherence Tomography (OCT) scans. Zhao et al. [27] proposed a theoretical framework for LSSL using an auto-encoder and an alignment term that forces the topology of the latent space to change in the direction of longitudinal changes. An extension was proposed in [16] to create a smooth trajectory field and a dynamic graph was computed to connect nearby subjects and enforce maximally aligned progression directions. Authors in [26] successfully used LSSL in the context of DR to predict the change from grades  $\{0,1\}$  to  $\{2,3,4\}$ , referred to as Moderate+, between two consecutive CFPs. However, since these approaches are self-supervised, it is unclear whether they learn the disease progression or the longitudinal changes.

Neural Ordinary Differential Equation (NODE) is a new type of neural networks that parameterizes the continuous dynamics of ordinary differential equations. It is ideally suited for solving time-related problems. Time-aware models, including NODEs, have achieved state-of-the-art performance in various tasks related to irregularly-sampled time series data or disease progression [1,19,9,25]. However, NODEs remain challenging to train. Authors in [8] proposed a simple yet efficient technique to regularize NODEs by randomly solving ODE for longer time points.

In recent years, mixing augmentation training has been successful in computer vision [22,24,12]. Mix-up [22] performs on the training set linearly, mixing a random pair of examples and their corresponding labels. Manifold Mix-up [24] extends such principle to linear interpolation to the hidden representation. Based on two one-hot labels, the mix-up training generates soft labels that model the relationship between two classes. One-hot labels describe the intra-class relationship, while Mix-up labels describe the inter-class relationship.

These soft labels modulate the learned decision boundaries, providing the model with more information. In this sense, Mix-up training can be seen as a pretext task. Going further, longitudinal-based Mix-up can be considered as a pretext task that captures the disease’s (presumed) linear progression. Motivated

by that, we propose a LPT using Manifold Mix-up (MM) [24], which we consider to be more suitable form of Mix-up training for our objective because manifold Mix-up has proven to provide more informative hidden representations. In order to learn feature representations embedded with disease progression that can be reapplied in tackling longitudinal-based problems, especially for the challenging task of prediction the disease progression based on a single image. We introduce  $t_{mix}$  an intermediate time in the course of the disease progression,  $t_{mix}$  is the weighted average time between two consecutive examinations.  $t_{mix}$  is used to obtain soft labels based on a severity profile. Additionally, it acts as a data augmentation for training time-aware models. To the best of our knowledge, this work is the first to automatically assess DR and forecast its progression using Mix-up training and time-aware models, using only one CFP.

## 2 Methods

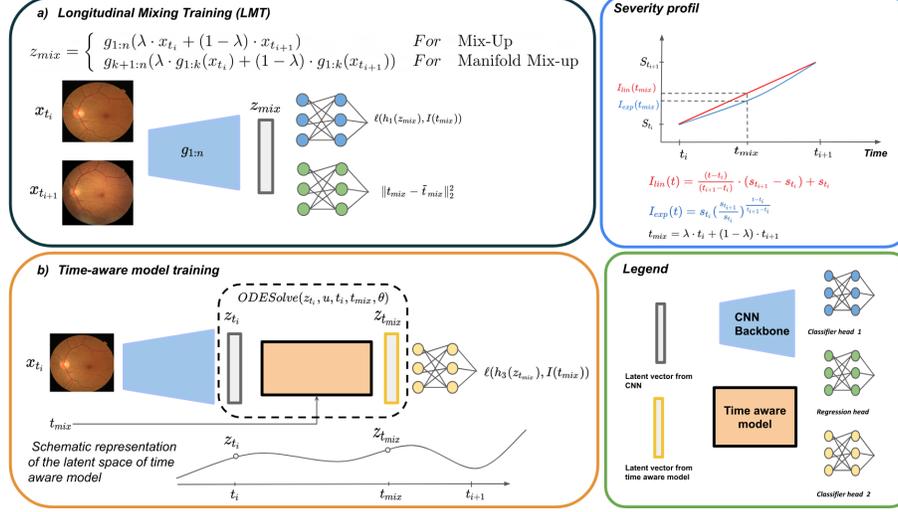
### 2.1 Preliminary

Let  $\mathcal{V}$  be the set of consecutive patient-specific image pairs for the collection of all CFP images.  $\mathcal{V}$  contains all  $(x_{t_i}, x_{t_{i+1}})$  that are from the same patient where  $x_{t_i}$  is scanned before  $x_{t_{i+1}}$  with  $i \in [0, m - 2]$ ,  $m$  being the number of visits for a given eye. Figure 1 displays a simplified architecture for clarity. We define backbone  $g_{1:n}$  with  $n$  layers, where  $g_{1:k}$  denotes the part of the neural network mapping the input data to the hidden representation at layer  $k$ .  $h_l$  represents a classification or regression head  $l$ ,  $(y, y')$  one-hot labels,  $\text{Beta}(\alpha, \alpha)$  the Beta distribution and  $\ell(\cdot)$  the Binary Cross-Entropy (BCE) loss. We define the mixing operator by  $\text{Mix}_\lambda(a, b) = \lambda \cdot a + (1 - \lambda) \cdot b$  with  $\lambda \sim \text{Beta}(\alpha, \alpha)$  where  $\lambda \in [0, 1]$ . **Mix-up** was introduced in [22] as a simple regularization method to minimize overfitting in deep neural networks. It linearly interpolates a mini-batch of random examples and their labels to transform the training set.

**Manifold Mix-up** is an extension of Mix-up to hidden representations [24]. During training, a random layer  $k$  from a set of eligible layers  $S$  in a neural network is selected. It processes two random data mini-batches  $(x, y)$  and  $(x', y')$ , until reaching layer  $k$ . The Mix-up is then performed on these two intermediate mini-batches  $(g_k(x), y)$  and  $(g_k(x'), y')$ , continuing the forward pass with the mixed representation until the ending layer  $n$ . These mixed representations are then fed to the classification head  $h_l$  and projected to the number of classes.

**Neural Ordinary Differential Equations (NODEs)** approximate unknown ordinary differential equations by a neural network [5] that parameterizes the continuous dynamics of hidden units  $\mathbf{z} \in \mathbb{R}^n$  over time with  $\mathbf{t} \in \mathbb{R}$ . NODEs are able to model the instantaneous rate of change of  $\mathbf{z}$  with respect to  $\mathbf{t}$  using a neural network  $u$  with parameters  $\theta$ .

$$\lim_{h \rightarrow 0} \frac{\mathbf{z}_{t+h} - \mathbf{z}_t}{h} = \frac{d\mathbf{z}}{dt} = u(t, \mathbf{z}, \theta) \quad (1)$$



**Fig. 1:** Illustration of Longitudinal Mixing Training in a) and time-aware model training using  $t_{mix}$  in b). Fig.1.a) and Fig.1.b) can be trained simultaneously or independently.

The analytical solution of Eq.1 is given by:

$$\mathbf{z}_{t_1} = \mathbf{z}_{t_0} + \int_{t_0}^{t_1} u(t, \mathbf{z}, \theta) dt = \text{ODESolve}(\mathbf{z}(t_0), u, t_0, t_1, \theta) \quad (2)$$

where  $[t_0, t_1]$  represents the time horizon for solving the ODE,  $u$  being a neural network, and  $\theta$  is the trainable parameters of  $u$ . By using a black-box ODE solver introduced in [5], we can solve the Initial Value Problem (IVP) and calculate the hidden state at any desired time using Eq.2. We can differentiate the solutions of the ODE solver with respect to the parameters  $\theta$ , the initial state  $\mathbf{z}_{t_0}$  at initial time  $t_0$ , and the solution at time  $t$ . This can be achieved by using the adjoint sensitivity method [5]. Through the latent representation of a given image, we define an IVP that aims to solve the ODE from  $t_i$  to a terminal time  $t_{i+1}$ :

$$\dot{z}(t) = u(z(t), t, \theta), \text{ with the initial value } z(t_i) = z_{t_i} \quad (3)$$

## 2.2 Longitudinal Mixing Training (LMT)

We denote  $s_{t_i}$  the severity grade of image  $x_{t_i}$ .  $I(t)$  is the severity interpolation function between two consecutive longitudinal pairs. In the linear case, between  $x_{t_i}$  and  $x_{t_{i+1}}$  we have  $I_{lin}(t) = \frac{(t-t_i)}{(t_{i+1}-t_i)} \cdot (s_{t_{i+1}} - s_{t_i}) + s_{t_i}$ . In conventional Mix-up training, the labels are mixed. Instead, we propose to mix the time between consecutive pairs  $t_{mix} = \text{Mix}_\lambda(t_i, t_{i+1})$ , then used this  $t_{mix}$  to evaluate  $I(t)$  and uses this signal as supervision. Motivated by the assumption that the progression of DR is a slow process, we tested another monotonic disease progression profile

expressed as  $I_{exp}(t) = s_{t_i} \left(\frac{s_{t_{i+1}}}{s_{t_i}}\right)^{\frac{t-t_i}{t_{i+1}-t_i}}$ . During training, we one-hot encode the value of interpolation at  $t_{mix}$  to get our soft label. Depending on the mix-up method, we have as latent representation of the mixed pair:

$$z_{mix} = \begin{cases} g_{1:n}(\text{Mix}_\lambda(x_{t_i}, x_{t_{i+1}})) & \text{for Mix-Up} \\ g_{k+1:n}(\text{Mix}_\lambda(g_{1:k}(x_{t_i}), g_{1:k}(x_{t_{i+1}}))) & \text{for Manifold Mix-up} \end{cases} \quad (4)$$

During the training of our LMM, we add a *time consistency* loss, as follows:

$$L_{t_{mix}} = \|t_{mix} - \tilde{t}_{mix}\|_2^2 \quad (5)$$

with  $\tilde{t}_{mix} = h_2(z_{mix})$ , where  $h_2$  is regression head that predicts the value of the current  $t_{mix}$  for a given pair. This term is inspired by [17] and motivated by authors in [13], who used Manifold Mix-up coupled with SSL loss to enhance the quality of the feature extraction. For the training of the LMM, the total loss is:

$$L = \mathbb{E}_{(x_{t_i}, x_{t_{i+1}}) \sim \mathcal{V}} \mathbb{E}_{\lambda \sim \text{Beta}(\alpha, \alpha)} \mathbb{E}_{k \sim \mathcal{S}} \ell(h_1(z_{mix}), I(t_{mix})) + L_{t_{mix}} \quad (6)$$

Concerning the use of  $t_{mix}$  for the NODE; instead of solving the ODE from  $t_i$  to a terminal time  $t_{i+1}$  using Eq.3, we solve to the intermediate time  $t_{mix}$  (see Fig.1.b) then use this  $t_{mix}$  to evaluate  $I(t)$  and take this signal as supervision for training. Note that this approach could be applied to any time-aware model.

### 3 Experiments and results

**Dataset.** The proposed models were trained and evaluated on OPHDIAT [14], a large CFP database collected from the Ophthalmology Diabetes Telemedicine network consisting of examinations acquired from 101,383 patients between 2004 and 2017. Out of 763,848 interpreted CFP images, 673,017 were assigned a DR severity grade, and the others were non-gradable. Patients range in age from 9 to 91, and image sizes vary from  $1440 \times 960$  to  $3504 \times 2336$  pixels. Each examination includes at least two eye images. To limit consecutive pairs without progression, 10412 patients were selected with at least one severity change. Each patient had 2–5 scans, averaging 3.43, spanning an average interval of 4.86 years. This dataset was further divided into training (60%), validation (20%), and test (20%) based on patients. We randomly selected one image per eye for each examination, resulting in 49578 pairs. Our longitudinal downstream task is to predict whether an eye without DR at the initial visit was later graded as having Mild+ DR within two years. 8,111 patients and 13,936 eyes fit this criterion for the training. For the DR assignment, a specific test set consisted of patients assigned the same grade from two ophthalmologists, resulting in 9,734 eyes of 4,996 patients. Except for the registration, we followed the same image processing performed in [26]. All the timestamp were normalized by  $2 \times 365$ .

**Implementation details.** In our basic architecture, we employed a stack of 2 pre-activated residual blocks (ReLU+BN). In each residual block, the residual

feature map was calculated using a series of three  $3 \times 3$  convolutions, the first of which always halves the number of the feature maps employed at the present scale. Our encoder comprised seven levels; the first six levels are composed of two residual blocks and the latter deals with only one residual block. This provides a final latent representation of size  $64 \times 4 \times 4$ . The last three layers of our backbone were used for the eligible layers S. The different networks were trained for 200 epochs by the AdamW optimizer, OneCycleLR as a scheduler, weight decay of  $10^{-4}$ , and a batch size of 128, using an NVIDIA A6000 GPU with the PyTorch framework. A grid search was performed for several key hyper-parameters of all Mix-up algorithms and SSL, including  $\alpha \in \{0.2, 0.5, 1.0, 2.0, 3.0, 5.0, 10.0\}$ , learning rate  $\in \{10^{-2}, 10^{-3}, 10^{-4}\}$  and three different seed values. Concerning the NODE, we used the Pytorch package TorchdiffEq [4]. This library provides ODE solvers, and backpropagation through ODE solutions and a support of the adjoint method for constant memory cost. Our NODE is a combination of dense layers followed by the tanh activation function and the adjoint method with "dopri5" as a solver. The loss in each task was used to monitor the model’s performance on the validation set, the best one was kept.

**Experiments for DR severity assessment using mixing training.** Usually, a new permutation is applied in mixing training at each batch. Since we only use fixed consecutive pairs, the network looks at fewer examples. To perform fair comparisons, we tried multiple permutations that matched the distribution of longitudinal pairs. We report the Quadratic-weighted Kappa for different scenarios of mixing training in Tab.1.

**Experiments to evaluate the quality of feature extraction.** For the downstream task, we provided the results from both linear evaluation and fine-tuning. The performance was evaluated with the Area Under the receiver operating characteristic Curve (AUC) and reported Tab.2. The linear evaluation was conducted by training a linear layer on top of the pre-trained and frozen encoder and trained using the same set-up that was tried with LMM. For the classical feature extractor, we used AE, VAE [11], and SimCLR [6]. For longitudinal SSL, we used longitudinal Siamese [17], LSSL [27], and LNE [16].

	Kappa	$\alpha$
Mix-up *	0.7646	0.2
Manifold Mix-up *	0.7747	2.0
Mix-up	0.7314	0.2
Manifold Mix-up	0.7342	2.0
Longitudinal Mix-up (LM) + $I_{lin}(t)$	0.7339	0.2
Longitudinal Manifold Mix-up (LMM) + $I_{lin}(t)$	<b>0.7511</b>	2.0
Longitudinal Mix-up (LM) + $I_{exp}(t)$	0.5595	0.5
Longitudinal Manifold Mix-up (LMM) + $I_{exp}(t)$	<b>0.7350</b>	2.0

**Table 1:** Comparison of the best Kappa for

Mix-up training for DR severity assessment. **Table 2:** Results on linear evaluation and fine-tuning of pre-trained model using AUC

Weights	AUC ( Mild+ DR within 2 years)	
	Frozen	Fine-tuned
Random	-	0.584
MM [24] ( $\alpha = 2.0$ )	0.564	0.595
AE	0.531	0.569
VAE [11]	0.510	0.575
SimCLR [6]	0.544	0.558
L-Siamese [17]	0.562	0.593
LSSL [27]	0.579	0.602
LNE [16]	0.570	0.595
Ours (LMM $\alpha = 2.0$ )	<b>0.613</b>	<b>0.627</b>

**Experiments using  $t_{mix}$  for time-aware models.** Another time-aware model T-LSTM, introduced in [1], was used for this experiment in order to demonstrate the effectiveness of  $t_{mix}$ . Both time-aware model take as input the latent representation  $z_{t_i}$  and the time difference between  $x_{t_i}$  and  $x_{t_{i+1}}$  ( $\Delta_t$ ) in order to

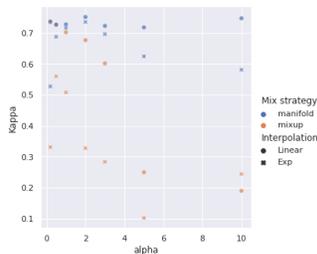
predict the latent representation of  $z_{t_{i+1}}$ . For the NODE, it was performed by the mean of IVP defined in Eq.3 while for T-LSTM [1], the LSTM gates were modulated by  $\Delta_t$  to produce the future  $z_{t_{i+1}}$ . We tested three training set-ups:

1. Using one image and the time of the next examination, we trained the model presented in Fig.1.b) with the loss  $\ell(h_3(z_{t_{i+1}}, s_{t+1}))$ .
2. Similarly to (1), we used image  $x_t$  but instead of giving the time of the next examination,  $t_{mix}$  was used and trained with the loss  $\ell(h_3(z_{t_{mix}}, I(t_{mix})))$ .
3. We used (2) with our LMM, i.e., we use (a) and (b) simultaneously (Fig.1).

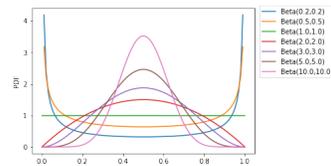
	AUC Mild+DR	AUC moderate+DR	AUC severe+DR	Best $\alpha$
(1) NODE [5]	0.584	0.617	0.641	-
(1) T-LSTM [1]	0.608	0.646	0.677	-
(2) NODE + $t_{mix}$ (ours)	0.632	0.695	0.725	2.0
(2) T-LSTM + $t_{mix}$ (ours)	0.610	0.661	0.725	0.5
(3) NODE+LMM (ours)	<b>0.657</b>	<b>0.721</b>	<b>0.798</b>	2.0

**Table 3:** Comparison of AUCs for the next visit for time-aware model training with and without  $t_{mix}$  with linear progression assumption.

The best results for the longitudinal task are obtained with LMM (Tab.2), indicating that it effectively captures disease progression. Moreover, the longitudinal task performed better than classical feature extraction methods, which is aligned with [7,16,17,27,26,25]. Concerning the use of  $t_{mix}$ , results in Tab.3 indicate that it is beneficial for both time-aware models. We believe  $t_{mix}$  plays the role of data augmentation in the context of disease progression for Time-Aware models and is regarded as a method to regularize the training of NODE like [8]. The fact that the set-up (3) performs better than other configurations supports the idea that LMM and  $t_{mix}$  are beneficial to solve time-related problems.



**Fig. 2:** Best value of Kappa when alpha varies for LM and LMM for the two profiles.



**Fig. 3:** Beta distribution for alpha values used.

Only LMM adapts  $I_{exp}(t)$  (Fig.2), indicating that with LMM, other types of severity profiles could be used. We suspect that LM suffers from manifold intrusion [10]. Mixing longitudinal pairs create an existing severity grade in the dataset, making training more challenging when using Mix-up. LMM performs better

than regular MM, according to Tab.1, when the profile of severity progression was supposed to be linear between two consecutive exams, suggesting that longitudinal pairs are more informative than random pairs with the same label distribution. The more  $\alpha$  increases, the more the severity interpolation is taken into consideration because when  $\lambda \leq 0.1 \Rightarrow t_{mix} \simeq t_1$  and  $\lambda = 0.5 \Rightarrow t_{mix} = \frac{t_0+t_1}{2}$  and finally  $\lambda \geq 0.9 \Rightarrow t_{mix} \simeq t_0$ , as illustrated by the Beta distribution in Fig.3. In Fig.2 for MML, we found that as alpha increases, Kappa differences between severity profiles increase. This could indicate that  $I_{exp}(t)$  is sub-optimal for DR progression. In addition, for  $\alpha = 10$ , the LMM is almost trained with the center of the severity interpolation  $I(t)$  as a label and yet, according to Fig.2, is able to assign DR with a Kappa of 0.75 (second best value of Kappa in all experiments). This could suggest, like in the original MM [24], that the LMM is able to disentangle factors of variations, such as the one responsible for encoding the disease progression. When alpha is low, there is a higher chance of sampling values of  $\lambda$  closer to either 0 or 1 from the Beta distribution. Time-aware models are then practically trained for two tasks: 1- predict the severity grade of the current image ( $t_{mix} \simeq t_i \Rightarrow I_{lin}(t_{mix}) \simeq s_{t_i}$ ) and 2- predict the next visit grade based on the last exam ( $t_{mix} \simeq t_{i+1} \Rightarrow I_{lin}(t_{mix}) \simeq s_{t_{i+1}}$ ). As a result, the model receives more information, which could explain the increase in performance. However, according to our experiments, T-LSTM does not perform well when  $\alpha$  is high. Only the NODE gains from having diverse time point  $t_{mix} \in [t_i, t_{i+1}]$  during training, showing that it can successfully change its hidden dynamic when time varies, in line with the conclusions of [5,19].

## 4 Discussion and conclusion

In this paper, we proposed straightforward modifications to Manifold Mix-up training. This adaptation aims to enhance training of time-aware models for disease progression by introducing  $t_{mix}$ . The results are encouraging and may help clinicians to choose the best DR screening intervals. Our framework is general and could be easily extended to other Mix-up training [12,23] and time-aware models [19,3] or to other diseases. However our work has some limitations. We did not register images between consecutive examinations. Image registration is a critical step, as mentioned in [21], and could enhance our results. To overcome the lack of grade diversity for a given pair during training, we could train with all potential pairs of a patient in a follow-up, as done in [7]. We made a strong assumption on the disease progression by supposing one common severity profile, yet observing good results. Since we can access the full examination, we could use a more accurate interpolation function to better fit the DR progression. We hope this work will benefit the fields of longitudinal analysis and disease progression.

**Acknowledgments** The work takes place in the framework of the ANR RHU project Evired. This work benefits from State aid managed by the French National Research Agency under the ‘‘Investissement d’Avenir’’ program bearing the reference ANR-18-RHUS-0008.

## References

1. Baytas, I., Xiao, C., Zhang, X., Wang, F., Jain, A., Zhou, J.: Patient subtyping via time-aware lstm networks. pp. 65–74 (08 2017). <https://doi.org/10.1145/3097983.3097997>
2. Bora, A., Balasubramanian, S., Babenko, B., Virmani, S., Venugopalan, S., Mitani, A., de Oliveira Marinho, G., Cuadros, J., Ruamviboonsuk, P., Corrado, G.S., Peng, L., Webster, D.R., Varadarajan, A.V., Hammel, N., Liu, Y., Bavishi, P.: Predicting the risk of developing diabetic retinopathy using deep learning. *The Lancet Digital Health* **3**, e10–e19 (1 2021). [https://doi.org/10.1016/S2589-7500\(20\)30250-8](https://doi.org/10.1016/S2589-7500(20)30250-8)
3. Brouwer, E.D., Simm, J., Arany, A., Moreau, Y.: Gru-ode-bayes: Continuous modeling of sporadically-observed time series (2019)
4. Chen, R.T.Q.: torchdiffeq (2018), <https://github.com/rtqichen/torchdiffeq>
5. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations (2018). <https://doi.org/10.48550/ARXIV.1806.07366>, <https://arxiv.org/abs/1806.07366>
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations (2020). <https://doi.org/10.48550/ARXIV.2002.05709>, <https://arxiv.org/abs/2002.05709>
7. Emre, T., Chakravarty, A., Rivail, A., Riedl, S., Schmidt-Erfurth, U., Bogunović, H.: Tinc: Temporally informed non-contrastive learning for disease progression modeling in retinal oct volumes. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022* p. 625–634 (2022). [https://doi.org/10.1007/978-3-031-16434-7\\_60](https://doi.org/10.1007/978-3-031-16434-7_60), [http://dx.doi.org/10.1007/978-3-031-16434-7\\_60](http://dx.doi.org/10.1007/978-3-031-16434-7_60)
8. Ghosh, A., Behl, H.S., Dupont, E., Torr, P.H.S., Namboodiri, V.: Steer: Simple temporal regularization for neural odes (2020). <https://doi.org/10.48550/ARXIV.2006.10711>, <https://arxiv.org/abs/2006.10711>
9. Gruffaz, S., Poulet, P.E., Maheux, E., Jedynek, B., DURRLEMAN, S.: Learning riemannian metric for disease progression modeling. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems*. vol. 34, pp. 23780–23792. Curran Associates, Inc. (2021), <https://proceedings.neurips.cc/paper/2021/file/c7b90b0fc23725f299b47c5224e6ec0d-Paper.pdf>
10. Guo, H., Mao, Y., Zhang, R.: Mixup as locally linear out-of-manifold regularization (2018). <https://doi.org/10.48550/ARXIV.1809.02499>, <https://arxiv.org/abs/1809.02499>
11. Kingma, D.P., Welling, M.: Auto-encoding variational bayes (2013). <https://doi.org/10.48550/ARXIV.1312.6114>, <https://arxiv.org/abs/1312.6114>
12. Liu, Z., Li, S., Wu, D., Liu, Z., Chen, Z., Wu, L., Li, S.Z.: Automix: Unveiling the power of mixup for stronger classifiers (2021). <https://doi.org/10.48550/ARXIV.2103.13027>, <https://arxiv.org/abs/2103.13027>
13. Mangla, P., Singh, M., Sinha, A., Kumari, N., Balasubramanian, V.N., Krishnamurthy, B.: Charting the right manifold: Manifold mixup for few-shot learning (2020)
14. Massin, P., Chabouis, A., Erginay, A., Viens-Bitker, C., Lecleire-Collet, A., Meas, T., Guillausseau, P., Choupot, G., André, B., Denormandie, P.: Ophdiat©: A telemedical network screening system for diabetic retinopathy in the ile-de-france. *Diabetes & metabolism* **34**, 227–34 (07 2008). <https://doi.org/10.1016/j.diabet.2007.12.006>

15. Ogurtsova, K., da Rocha Fernandes, J., Huang, Y., Linmenkamp, U., Guariguata, L., Cho, N.H., Cavan, D., Shaw, J., Makaroff, L.: IDF diabetes atlas: Global estimates for the prevalence of diabetes for 2015 and 2040. *Diabetes Research and Clinical Practice* **128**, 40–50 (2017)
16. Ouyang, J., Zhao, Q., Adeli, E., Sullivan, E.V., Pfefferbaum, A., Zaharchuk, G., Pohl, K.M.: Self-supervised longitudinal neighbourhood embedding
17. Rivail, A., Schmidt-Erfurth, U., Vogel, W.D., Waldstein, S.M., Riedl, S., Grechenig, C., Wu, Z., Bogunovic, H.: Modeling disease progression in retinal octs with longitudinal self-supervised learning. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **11843 LNCS**, 44–52 (2019). [https://doi.org/10.1007/978-3-030-32281-6\\_5](https://doi.org/10.1007/978-3-030-32281-6_5)
18. Rom, Y., Aviv, R., Ianchulev, S., Dvey-Aharon, Z.: Predicting the future development of diabetic retinopathy using a deep learning algorithm for the analysis of non-invasive retinal imaging. *medRxiv* (2022). <https://doi.org/10.1101/2022.03.31.22272079>, <https://www.medrxiv.org/content/early/2022/07/05/2022.03.31.22272079>
19. Rubanova, Y., Chen, R.T.Q., Duvenaud, D.: Latent odes for irregularly-sampled time series (2019). <https://doi.org/10.48550/ARXIV.1907.03907>, <https://arxiv.org/abs/1907.03907>
20. Saeedi, P., Petersohn, I., Salpea, P., Malanda, B., Karuranga, S., Unwin, N., Colagiuri, S., Guariguata, L., Motala, A.A., Ogurtsova, K., Shaw, J.E., Bright, D., Williams, R.: Global and regional diabetes prevalence estimates for 2019 and projections for 2030 and 2045: Results from the international diabetes federation diabetes atlas, 9th edition. *Diabetes Research and Clinical Practice* **157**, 107843 (Nov 2019). <https://doi.org/10.1016/j.diabres.2019.107843>, <https://doi.org/10.1016/j.diabres.2019.107843>
21. Saha, S.K., Xiao, D., Bhuiyan, A., Wong, T.Y., Kanagasingam, Y.: Color fundus image registration techniques and applications for automated analysis of diabetic retinopathy progression: A review. *Biomedical Signal Processing and Control* **47**, 288–302 (2019). <https://doi.org/10.1016/j.bspc.2018.08.034>, <https://doi.org/10.1016/j.bspc.2018.08.034>
22. Thulasidasan, S., Chennupati, G., Bilmes, J.A., Bhattacharya, T., Michalak, S.E.: On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In: *Neural Information Processing Systems* (2019)
23. Venkataramanan, S., Kijak, E., Amsaleg, L., Avrithis, Y.: Alignmixup: Improving representations by interpolating aligned features (2022)
24. Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: *International Conference on Machine Learning* (2018)
25. Zeghlache, R., Conze, P.H., Daho, M.E.H., Li, Y., Le Boité, H., Tadayoni, R., Massin, P., Cochener, B., Brahim, I., Quellec, G., Lamard, M.: Longitudinal self-supervised learning using neural ordinary differential equation. In: *International Workshop on Predictive Intelligence in Medicine*. Vancouver (Canada), Canada (Oct 2023), <https://imt-atlantique.hal.science/hal-04171357>
26. Zeghlache, R., Conze, P.H., Daho, M.E.H., Tadayoni, R., Massin, P., Cochener, B., Quellec, G., Lamard, M.: Detection of diabetic retinopathy using longitudinal self-supervised learning. In: Antony, B., Fu, H., Lee, C.S., MacGillivray, T., Xu, Y., Zheng, Y. (eds.) *Ophthalmic Medical Image Analysis*. pp. 43–52. Springer International Publishing, Cham (2022)
27. Zhao, Q., Liu, Z., Adeli, E., Pohl, K.M.: Longitudinal self-supervised learning. *Medical Image Analysis* **71** (2021). <https://doi.org/10.1016/j.media.2021.102051>