

Re: Shape and Time Distortion Loss for Training Deep Time Series Forecasting Models

Manjot Singh Yiyu Wang
{manjot.singh, yiyu.wang}@epfl.ch

Abstract

In this paper, some results from DILATE(Distortion Loss including shape and time) (3) are reproduced and additional studies on hyperparameter explorations, loss function and different neural network models are done as a part of NeurIPS Reproducibility Challenge. The comparison between DILATE, Euclidean Loss (MSE) and Huber Loss has been shown when training Deep Neural Networks on non-stationary time series forecasting problems.

1 Introduction

Time series forecasting for non-stationary signals and multi-future-step predictions has been a challenging task. Deep Neural Networks such as RNNs, LSTMs and GRUs have been intensively studied in related domains because of their abilities to model complex non-linear time dependencies. The authors of (3) introduced DILATE(Distortion Loss including shape and time) for training Deep Neural Networks in the context of multi-step and non-stationary time series forecasting. DILATE is a differentiable loss function which penalizes two errors, shape and temporal localisation errors of change detection (2). Mostly, mean squared error (MSE) or its variants are used as a loss function for training Deep Neural Nets. In (3), the authors have shown that in some cases, MSE or its variants doesn't capture a sharp drop or is not a better fit for regulation purposes whereas DILATE reflects on the sharp changes of regime although with a slight delay or with a slight inaccurate amplitude(refer to Figure1:b,c in (3)).

The source code provided by the authors was used and it includes the implementation of DILATE loss function and seq2seq model. We tried to make contact with the authors for supplementary material and data preprocessing guidelines but didn't get a response back. We followed the directions as mentioned in (3) for data preprocessing and also evaluated the performance of DILATE on additional dataset section 3.1. In this paper, the main results from (3) were reproduced to evaluate the performance of DILATE. We propose Convolutional-LSTM model to make certain that DILATE can be used with other Neural Network architectures specifically designed for multi-step and non-stationary forecasting. Inspired by the robust regression, we also used Huber Loss to train deep neural nets in addition to DIALTE and MSE. In addition to the above, several tests were done to study the impact of α and γ on the overall performance of the model and results are shown in 3.4. We were given some compute resources from the CodeOcean(1) to run our experiments.

2 DILATE

A very brief introduction to the working of DILATE is given in this section. The framework proposed for multi-step time series forecasting is depicted in Figure2 in (3). For each input example of length n , the forecasting model predicts the output \hat{y} . The DILATE objective function (3) is composed of two terms balanced by the hyperparameter $\alpha \in [0, 1]$ which compares the prediction \hat{y}_i with the ground truth value y_i^* :

$$\mathcal{L}_{DILATE}(\hat{y}_i, y_i^*) = \alpha \mathcal{L}_{shape}(\hat{y}_i, y_i^*) + (1 - \alpha) \mathcal{L}_{temporal}(\hat{y}_i, y_i^*) \quad (1)$$

Shape term The shape loss function is based on Dynamic Time Warping (7) . DTW is non-differentiable but the smooth minimum operator γ proposed in (2) defines the shape term \mathcal{L}_{shape} as:

$$\mathcal{L}_{shape}(\hat{y}_i, y_i^*) = DTW_r(\hat{y}_i, y_i^*) := -\gamma \log\left(\sum_{A \in \mathcal{A}_{k,k}} \exp\left(-\frac{\langle \mathbf{A}, \Delta(\hat{y}_i, y_i^*) \rangle}{\gamma}\right)\right) \quad (2)$$

Temporal Term The temporal term penalizes the temporal irregularities between \hat{y}_i and y_i^* (3). The loss function is inspired from computing the Time Distortion Index (TDI) for temporal misalignment estimation (5). The loss function is given as:

$$\mathcal{L}_{temporal}(\hat{y}_i, y_i^*) := \langle \mathbf{A}_\gamma, \Omega \rangle \frac{1}{Z} \sum_{A \in \mathcal{A}_{k,k}} \langle \mathbf{A}, \Omega \rangle \exp\left(-\frac{\langle \mathbf{A}, \Delta(\hat{y}_i, y_i^*) \rangle}{\gamma}\right) \quad (3)$$

For more details on the above algorithm, please refer to (3), (2), (5).

3 Experiments

The results were reproduced from section 4.2 in (3) on 3 non-stationary time series datasets from different domains as implemented in (3). We did an evaluation on one more additional dataset to analyse the relevance of DILATE against Huber and Euclidean loss (MSE). The multi-step evaluation consists in forecasting the future trajectory on k future time steps (3).

3.1 Experimental Setup

Synthetic and ECG5000 Dataset description of these 2 datasets is same as in (3). For ECG5000, we first did the MinMax normalization for the data to narrow down the values from 0 to 1.

Traffic(k=24) dataset corresponds to the road occupancy rates (between 0 and 1) from the California Department of Transportation (48 months from 2015-2016) measured every 1h (3). For this, there were no data pre-processing instructions, this allowed us to work on univariate series of length 17544 with train/test on 75/25 weeks of data and predictions on 24 future points given the past 168 points (past week).

Wafer (k=62) comes from the UCR Time Series classification Archive (8) with sequence length 152. It consists of the collection of inline process control measurements (1000/6000: train/test) recorded from various sensors during the processing of silicon wafers for semiconductor fabrication.

Network Architectures and Training For multi-step forecasting, we used the Seq2Seq model with Gated Recurrent Units (GRU) as provided by the authors of (3). In addition to Seq2Seq, an implementation of CNN-LSTM model (section 3.3) is carried out so as to compare the performance of DILATE, Huber and MSE on different Neural Network architectures. The models are trained in Pytorch(6) for a maximum of 1000 epochs with the ADAM optimizer.

3.2 Evaluating DILATE Forecasting

The similar procedure is used as mentioned in Section4.2 in (3) to evaluate the forecasting performance of DILATE against Huber loss, Euclidean Loss (MSE) and DTW_γ . The results were averaged over 10 runs, same as in (3). The same experiment was repeated with CNN-LSTM Neural Network as well. The results are evaluated using three metrics: MSE, DTW(shape) and TDI(temporal) (3). Overall results are presented in Table1. For ECG dataset when evaluated on MSE and DTW, our results don't quite match with that of the authors in (3) and are highlighted in bold in Table1. Some of the results also differ for traffic dataset(seq2seq) and this could be due to the difference in ways of processing the data.

We display a few qualitative examples for Synthetic and Wafer dataset in Figure 1 and 2 when applied Seq2Seq and CNN-LSTM model. As in paper (3), our results for Synthetic dataset matches with that of the authors where DILATE is better than MSE in predicting sharp changes in regime. DTW_γ leads to very sharp predictions in shape, but with a large temporal misalignment. In the case of Wafer dataset, the performance of MSE and Huber loss is better than DILATE when evaluated on MSE.

Table 1: Forecasting results evaluated with MSE (x100), DTW (x100) and TDI (x10) metrics, averaged over 10 runs.

Dataset	Eval	Convolutional-LSTM Network				Recurrent Neural Network			
		MSE	DTW $_{\gamma}$	DILATE	Huber	MSE	DTW $_{\gamma}$	DILATE	Huber
Synth	MSE	1.64 ± 0.05	5.12 ± 0.33	1.80 ± 0.12	1.6 ± 0.04	1.06 ± 0.08	1.74 ± 0.26	1.32 ± 0.07	1.11 ± 0.05
	DTW	32.72±0.22	28.15±0.25	28.72±0.65	32.19±0.62	22.71±1.42	17.07±5.29	19.17±1.15	22.80±1.16
	TDI	12.46±0.53	30.16±1.96	14.17±0.56	12.79±0.36	18.37±1.77	20.64±1.30	14.78±1.57	17.67±0.94
ECG	MSE	0.57 ± 0.02	1.3 ± 0.37	0.56 ± 0.01	0.53 ± 0.02	0.60 ± 0.04	1.0 ± 0.09	1.9 ± 1.1	0.6 ± 0.1
	DTW	29.3 ± 0.5	0.6 ± 0.36	29.4 ± 0.35	28.7 ± 0.34	30.9 ± 0.95	31.3 ± 0.69	55 ± 22.5	31.6 ± 2.4
	TDI	1.7 ± 0.05	3.1 ± 0.66	1.68 ± 0.08	1.66 ± 0.09	2.1 ± 0.14	2.9 ± 0.41	5.8 ± 4.1	2.4 ± 0.6
Traffic	MSE	0.12±0.01	0.12±0.01	0.11±0.01	0.11 ± 0.01	1.64 ± 0.04	5.12 ± 0.33	1.80 ± 0.12	1.60 ± 0.04
	DTW	9.66±4.26	11.87±0.01	11.49±0.54	12.03±0.35	32.72±0.22	28.15±1.25	28.72±0.65	32.19±0.62
	TDI	0.62±0.05	0.67±0.03	1.46±0.33	0.71 ± 0.08	12.46±0.53	30.1 ± 1.9	14.17±0.56	12.79±0.36
Wafer	MSE	0.64±0.03	23.74±1.19	15.8 ± 3.8	0.65 ± 0.03	0.54 ± 0.10	23.9 ± 2.01	2.8 ± 0.99	0.58 ± 0.09
	DTW	45.0 ± 1.35	17.87 ± 2.0	36.6 ± 5.96	45.6 ± 1.4	33.37 ± 2.4	20.67 ± 1.5	22.6 ± 3.1	33.5 ± 2.4
	TDI	3.9 ± 0.27	82.94±8.06	23.75±9.45	7.79 ± 10.5	6.5 ± 1.34	88.8 ± 10.2	6.8 ± 2.8	6.5 ± 1.4

Huber Loss Comparison: From Table 1, it can be seen that MSE and Huber loss perform almost the same when evaluated on MSE, shape(DTW) and time(TDI). So here, we only compare DILATE against Huber loss. DILATE is better than Huber loss when evaluated on shape(DTW) in 5/8 experiments and 2 equivalent performances. On evaluation on time(TDI), Huber performs much better than DILATE in 3/8 experiments and 4 similar outcomes. When evaluated on MSE, DILATE is equivalent to Huber loss except for Wafer dataset.

DTW $_{\gamma}$ Comparison: On evaluation with shape(DTW), DILATE performs equivalently to DTW_{γ} (3 reductions and 5 similar performances). When evaluated on time, DILATE outperforms DTW_{γ} in 6/8 experiments except for Traffic dataset. For MSE evaluation, DILATE is significantly better than DTW_{γ} in all experiments.

In conclusion, DILATE performs better than DTW_{γ} . The performance of DILATE is better than MSE/Huber loss mostly when predicting sharp changes of regime and when evaluated on MSE, MSE/Huber outperforms DILATE.

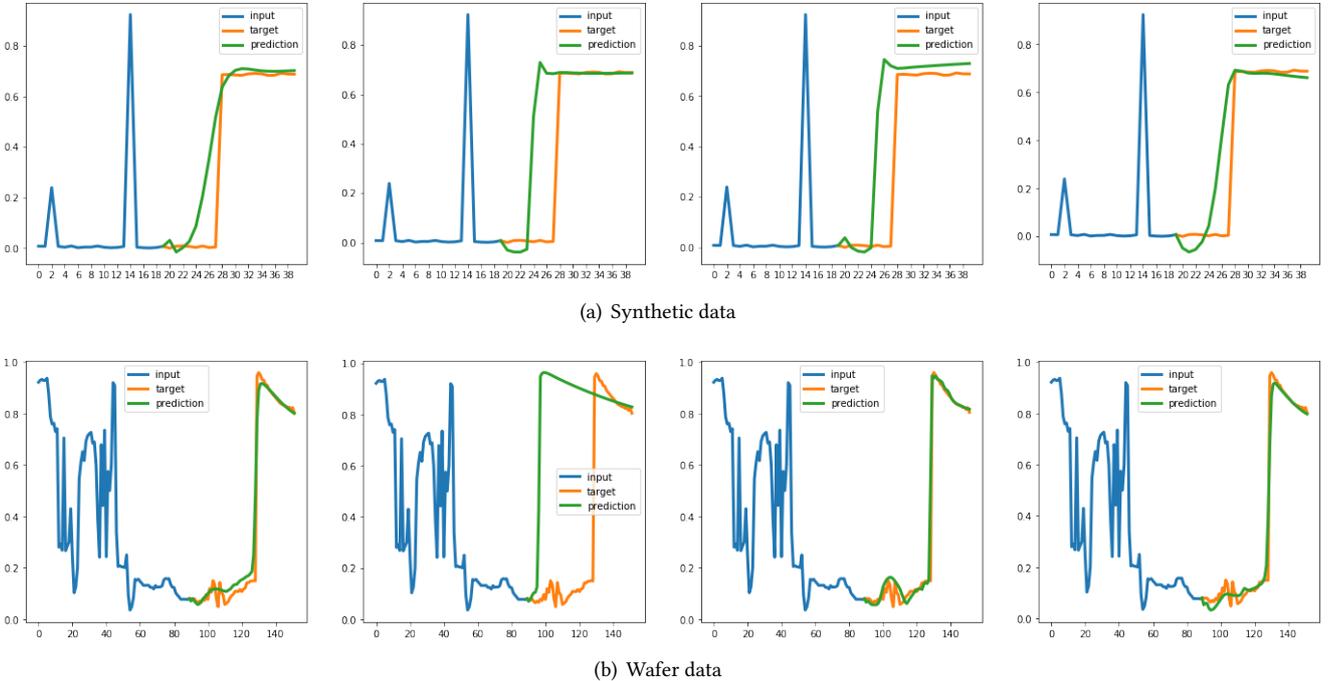


Figure 1: Forecasting results for Synthetic and Wafer data based on Seq2Seq model. **From left to right:** MSE loss, DTW $_{\gamma}$ loss, DILATE loss and Huber loss.

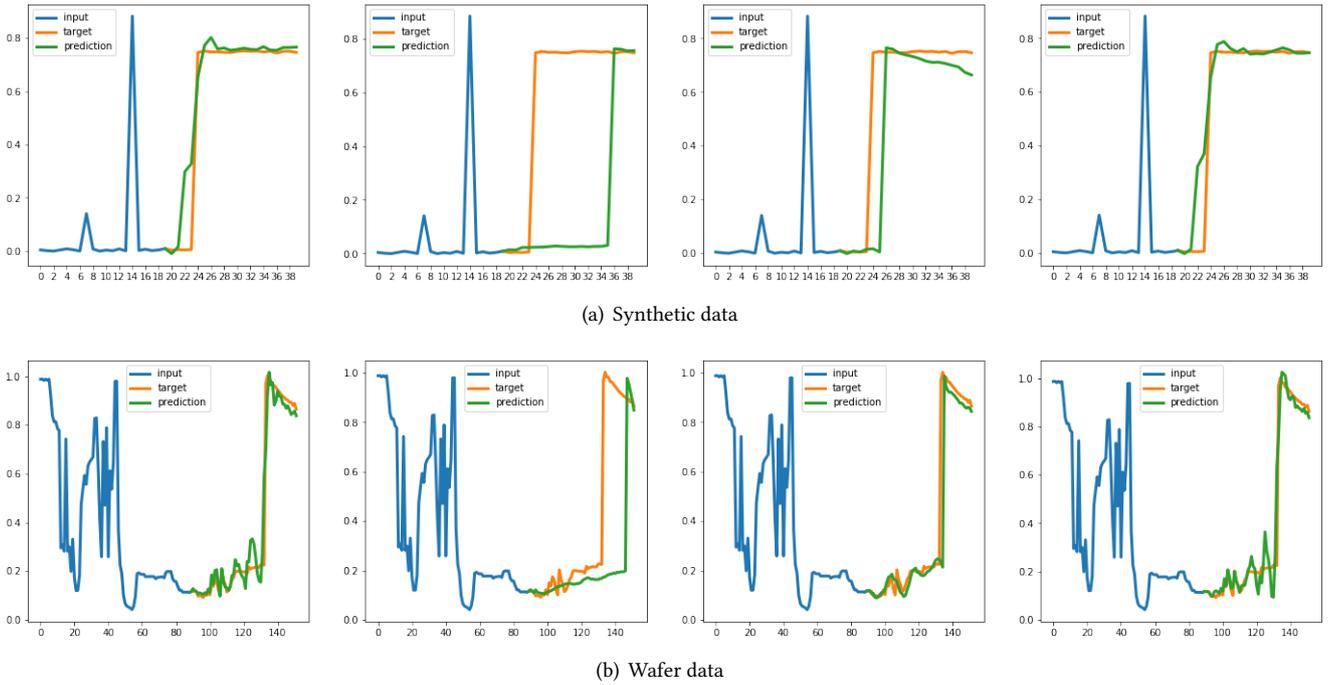


Figure 2: Forecasting results for Synthetic and Wafer data based on CNN_LSTM model. **From left to right:** MSE loss, DTW_γ loss, DILATE loss and Huber loss.

3.3 Comparison with CNN-LSTM

We compared DILATE performance on Seq2Seq model and a basic CNN-LSTM. CNN-LSTM model is developed with the goal to perform 1-D convolution to capture local patterns in the multi-step time series and then LSTM for capturing the complex nonlinear time dependencies. The Seq2Seq DILATE performs similarly to CNN-LSTM when evaluated on all the three metrics and four datasets used in this study. This highlights the relevance of DILATE loss function, which reaches better performances with simpler architectures. We also reproduced and evaluated results with fully connected neural network and the plots displaying forecasting results can be found in the supplementary material. If DILATE and MSE/Huber perform similar when used in complex Neural Network architectures (4), then MSE or Huber Loss seems to be a better choice because of computational reasons.

3.4 DILATE Analysis

Impact of γ : We carried out a thorough analysis of γ on the overall performance of the model using Synthetic dataset. We chose from 6 log-spaced values between 10^{-3} and 10. We observed that for low values of γ , DILATE training loss converges to an acceptable minimum value but as seen from the Figure 3, it has a possibility to get stuck in a very bad local minima. On the other hand, for high values of γ , loss converges smoothly to a reasonable solution. As γ decreases, DTW_γ achieves low loss values which verifies that DILATE is more compliant to optimisation by gradient-descent methods. The convergence of training loss for 3 different values of gamma is shown in the Figure 3. The same analysis can be carried out on different datasets for more validity.

Impact of α : We ran experiments on 10 different values of α ranging from 0 to 1 and observe its impact on the performance of DILATE. We carried out this analysis on Wafer and ECG5000 dataset. When $\alpha = 0$, temporal loss is minimized without any shape constraint. Both MSE and shape errors rise up in this case illustrating the fact that $\mathcal{L}_{temporal}$ is only meaningful in conjunction with \mathcal{L}_{shape} (3). For ECG5000 dataset, we observed the similar performance to that of Synthetic dataset as in (3) between $\mathcal{L}_{temporal}$ and \mathcal{L}_{shape} but MSE in this case seems to be independent of alpha as seen from Figure 4. We observed similar results in case of the Wafer dataset as well. This suggests that MSE loss does not depend on the value of α . For more insights on this, similar analysis can be carried out on different datasets.

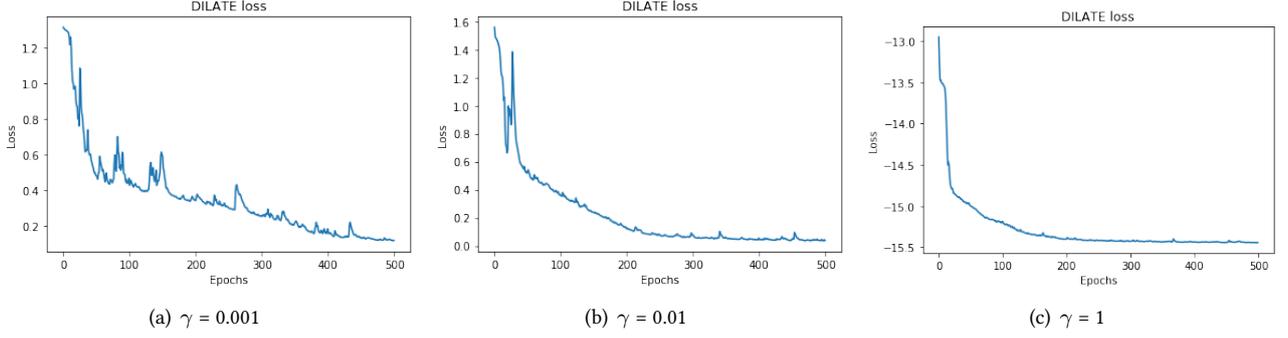


Figure 3: DILATE training loss on different γ values.

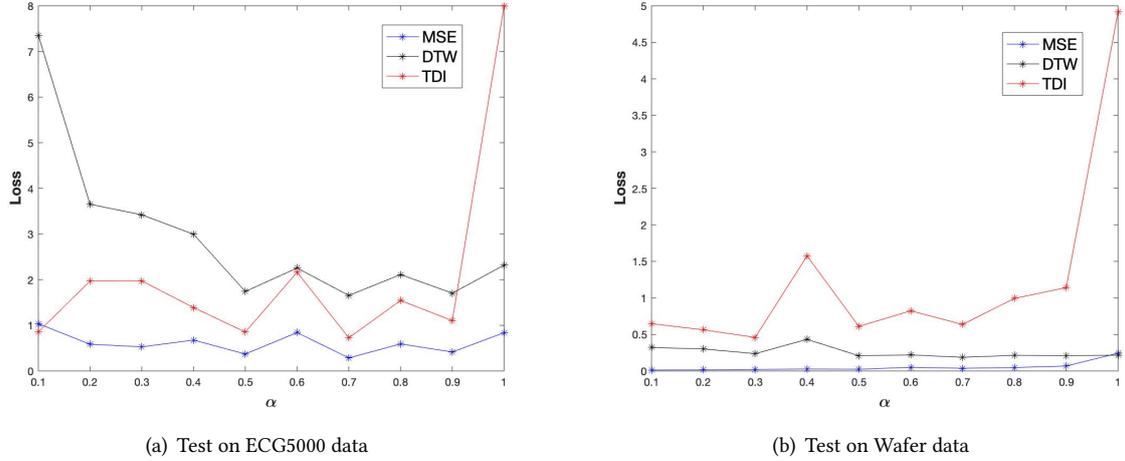


Figure 4: α on ECG and Wafer data

Custom backward implementation speedup: In order to verify the speedup, we apply the Pytorch build in function `TORCH.AUTOGRAD.PROFILER` to check the runtime of the back propagation of DILATE and MSE loss. According to section 3.2 in paper(3), the time complexity of DILATE is $\mathcal{O}(k^2)$ and the time complexity of MSE algorithm is known i.e. $\mathcal{O}(k)$. We tested different input length and different models, then estimated the speedup by computing $\frac{runtime_{mse}^2}{runtime_{dilate}}$. We verified that the custom backward implementation in (3) is very effective, but DILATE is still computationally expensive than MSE.

4 Conclusion

We reproduced some results which are similar to that of the paper (3). We did additional experiments using different datasets, loss function and models, and also studied the effect of two hyperparameters α and γ on the model performance. The results were easy to reproduce given one has to run the tests many times for quantitative reasons. We could not reproduce the Hausdorff distance and Ramp score as there was no access to the supplementary material of the paper. From the above experiments, DILATE is comparable to MSE/Huber. For more understanding, one might compare the performance of DILATE and MSE/Huber loss on transformer like models(4) with respect to better predictions and computational complexity.

References

- [1] April Clyburne-Sherin, Xu Fei, and Seth Ariel Green. Computational reproducibility via containers in social psychology. *Meta-Psychology*, 3, 2019.
- [2] Marco Cuturi and Mathieu Blondel. Soft-dtw: a differentiable loss function for time-series. In *International Conference on Machine Learning*.
- [3] Vincent Le Guen and Nicolas Thome. Shape and time distortion loss for training deep time series forecasting models. In *NIPS Neural Information Processing Systems*.
- [4] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyu Zhou, Wenhui Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. In *NIPS Neural Information Processing Systems*.
- [5] Laura Frías Paredes, Fermín Mallor, Teresa Leon, and Martín Gaston Romeo. Introducing the temporal distortion index to perform a bidimensional analysis of renewable energy forecast. *Elsevier*.
- [6] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W, 2017*.
- [7] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. In *IEEE Transactions on Acoustics, Speech, and Signal Processing*, pages 43 – 49. IEEE, 1990.
- [8] Chen Yanping, Keogh Eamonn, Hu Bing, Begum Nurjahan, Bagnall Anthony, Mueen Abdullah, and Gustavo Batista. Ucr time series classification archive.

5 Supplementary Material

The following are some more figures from the reproduction work.

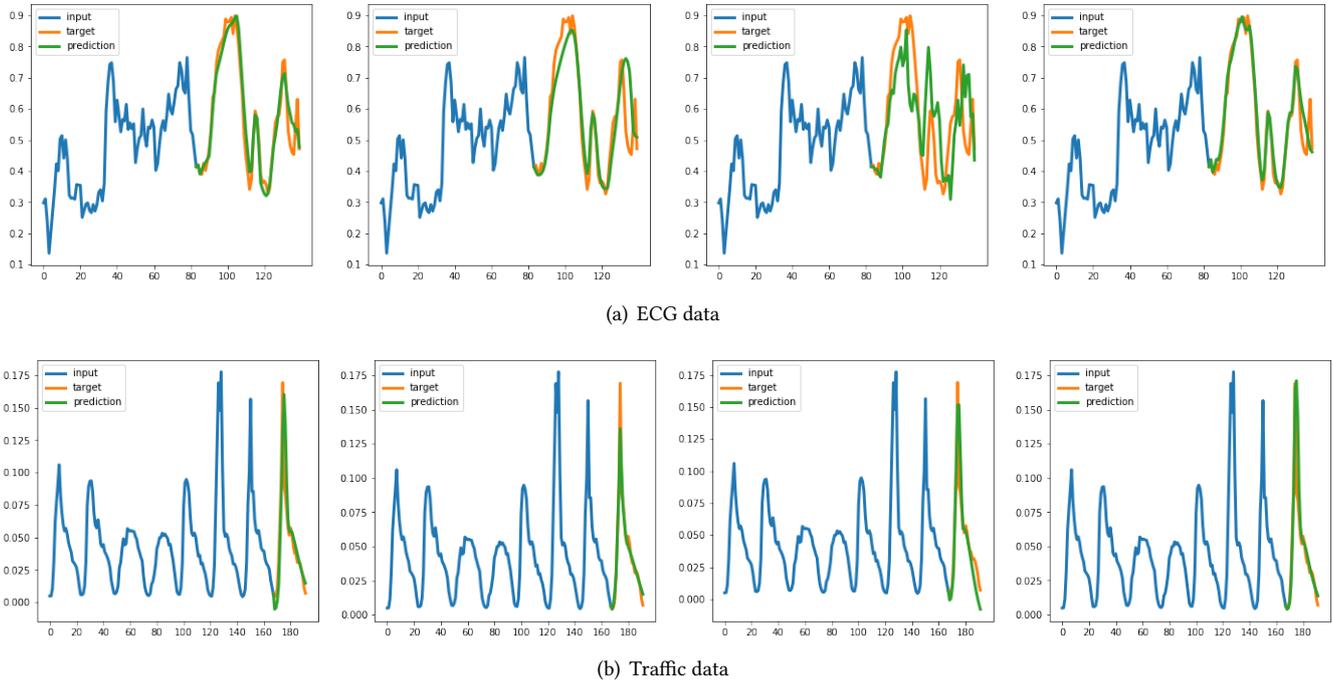


Figure 5: Forecasting results for ECG and Traffic data based on Seq2Seq model. **From left to right:** MSE loss, DTW γ loss, DILATE loss and Huber loss.

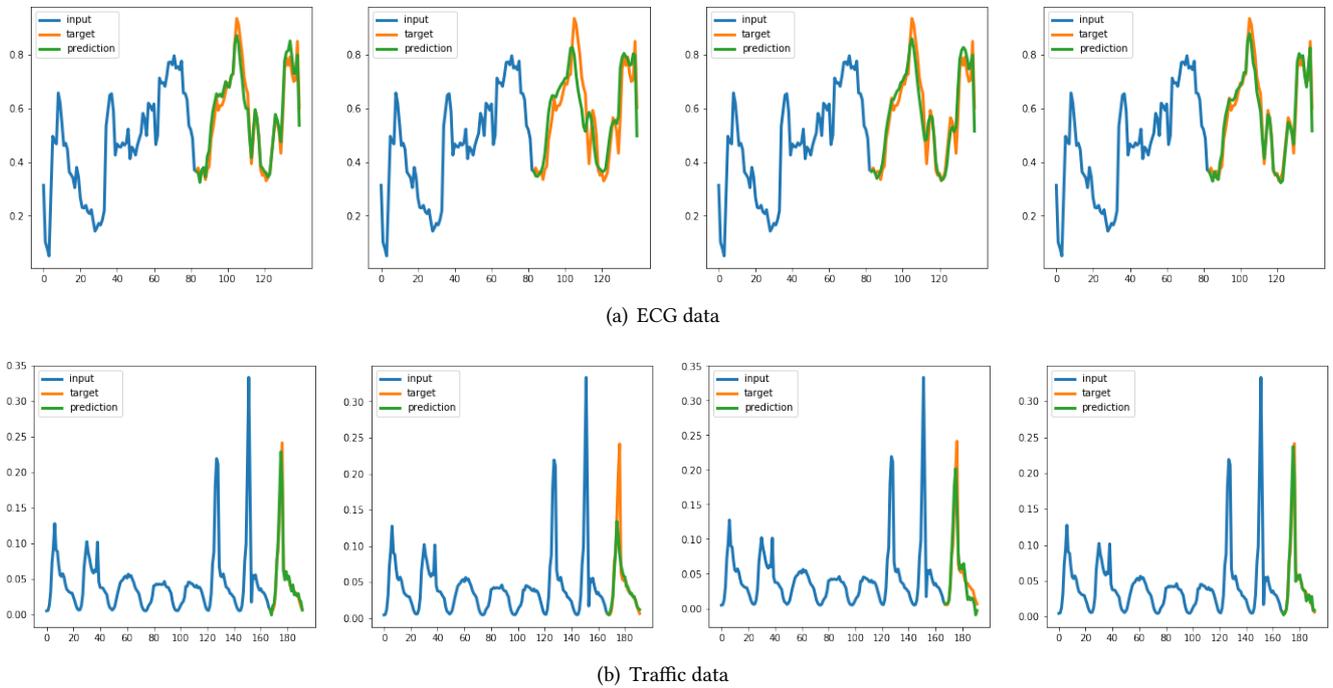
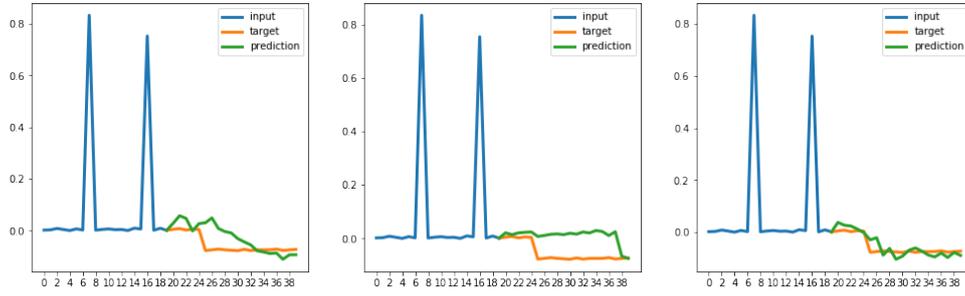
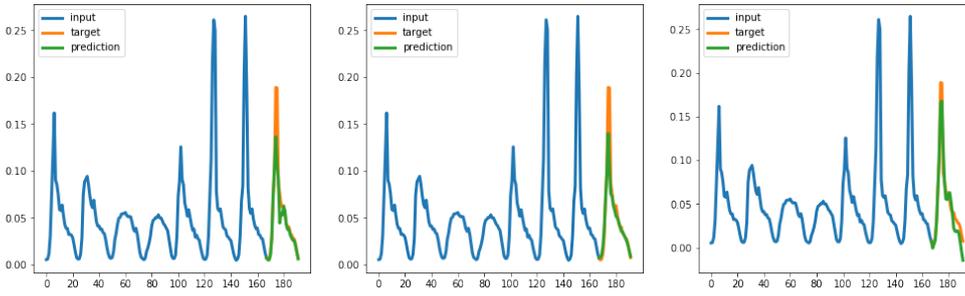


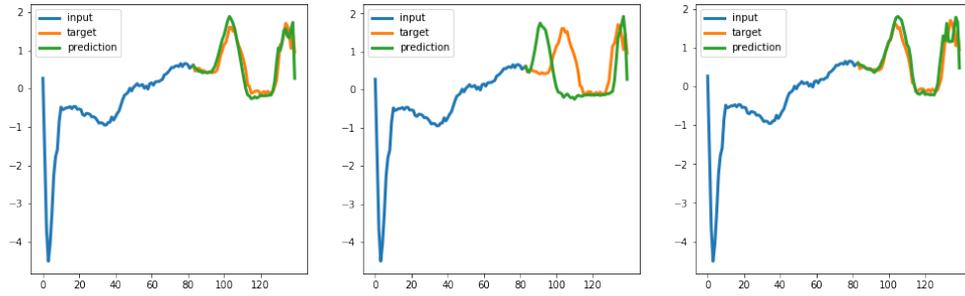
Figure 6: Forecasting results for ECG and Traffic data based on CNN_LSTM model. **From left to right:** MSE loss, DTW γ loss, DILATE loss and Huber loss.



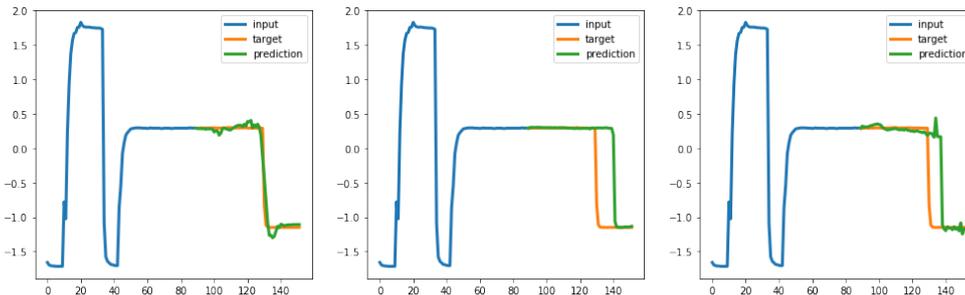
(a) Synthetic data



(b) Traffic data



(c) ECG data



(d) Wafer data

Figure 7: Forecasting results for 4 datasets based on Fully connected network(MLP) model. **Leftmost:** MLP with MSE loss, **Middle:** MLP with DTW γ loss, **Rightmost:** MLP with DILATE loss.