
Extended Deep Adaptive Input Normalization for Preprocessing Time Series Data for Neural Networks

Marcus A. K. September
Department of Mathematics
Imperial College London

Francesco Sanna Passino
Department of Mathematics
Imperial College London

Leonie Goldmann
Decision Science, Credit & Fraud Risk
American Express

Anton Hinel
Decision Science, Credit & Fraud Risk
American Express

Abstract

Data preprocessing is a crucial part of any machine learning pipeline, and it can have a significant impact on both performance and training efficiency. This is especially evident when using deep neural networks for time series prediction and classification: real-world time series data often exhibit irregularities such as multi-modality, skewness and outliers, and the model performance can degrade rapidly if these characteristics are not adequately addressed. In this work, we propose the EDAIN (Extended Deep Adaptive Input Normalization) layer, a novel adaptive neural layer that learns how to appropriately normalize irregular time series data for a given task in an end-to-end fashion, instead of using a fixed normalization scheme. This is achieved by optimizing its unknown parameters simultaneously with the deep neural network using back-propagation. Our experiments, conducted using synthetic data, a credit default prediction dataset, and a large-scale limit order book benchmark dataset, demonstrate the superior performance of the EDAIN layer when compared to conventional normalization methods and existing adaptive time series preprocessing layers.

1 INTRODUCTION

There are many steps required when applying deep neural networks or, more generally, any machine learning model, to a problem. First, data should be gathered, cleaned, and formatted into machine-readable values. Then, these values need to be preprocessed to facilitate learning. Next, features are designed from the processed data, and the model architecture and its hyperparameters are chosen. This is followed by parameter optimisation and evaluation using suitable metrics. These steps may be iterated several times.

A step that is often overlooked in the literature is *pre-processing* (see, for example, Koval, 2018), which consists in operations used to transform raw data in a format that is suitable for further modeling, such as detecting outliers, handling missing data and normalizing features. Applying appropriate preprocessing to the data can have significant impact on both performance and training efficiency (Cao et al., 2018; Nawi et al., 2013; Passalis et al., 2020; Sola and Sevilla, 1997; Tran et al., 2021). However, determining the most suitable preprocessing method usually requires a substantial amount of time and relies on iterative training and performance testing. Therefore, the main objective of this work is to propose a *novel efficient automated data preprocessing method for optimising the predictive performance of neural networks*, with a focus on normalization of multivariate time series data.

1.1 Preprocessing multivariate time series

Let $\mathcal{D} = \{\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}, i = 1, \dots, N\}$ denote a dataset containing N time series, where each time series $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$ is composed of T d -dimensional feature vectors. The integers d and T refer to the feature and temporal dimensions of the data, respectively.

Also, we use $\mathbf{x}_t^{(i)} \in \mathbb{R}^d$, $t = 1, \dots, T$ to refer to the d features observed at timestep t in the i -th time series. Before feeding the data into a model such as a deep neural network, it is common for practitioners to perform z -score normalization (see, for example, Koval, 2018) on $\mathbf{x}_t^{(i)} = (x_{t,1}^{(i)}, \dots, x_{t,d}^{(i)}) \in \mathbb{R}^d$, obtaining

$$\tilde{x}_{t,k}^{(i)} = \frac{x_{t,k}^{(i)} - \mu_k}{\sigma_k}, \quad k = 1, \dots, d,$$

where μ_k and σ_k denote the mean and standard deviation of the measurements from the k -th predictor variable. Another commonly used method is min-max scaling, where the observations for each predictor variable are transformed to fall in the value range $[0, 1]$ (see, for example, Koval, 2018). Other common preprocessing methods are winsorization (see, for example, Nyitrai and Virág, 2019) and the Yeo-Johnson power transform (see, for example, Yeo and Johnson, 2000). In this work, we refer to these conventional methods as *static preprocessing* methods, as the transformation parameters are fixed statistics that are computed through a single sweep of the training data. Most of these transformations only change the location and scale of the observations, but real-world data often contains additional irregularities such as skewed distributions, outliers, extreme values, heavy tails and multiple modes (Cao et al., 2023; Nawi et al., 2013), which are not mitigated by transformations such as z -score and min-max normalization. Employing static normalization in such cases may lead to sub-optimal results, as demonstrated in Passalis et al. (2020, 2021); Tran et al. (2021) and in the experiments on real and synthetic data in Section 4 of this work.

In contrast, better results are usually obtained by employing *adaptive* preprocessing methods (see, for example Lubana et al., 2021; Passalis et al., 2020, 2021; Tran et al., 2021), where the preprocessing is integrated into the deep neural network by augmenting its architecture with additional layers. Both the transformation parameters and the neural network model parameters are then jointly optimised in an end-to-end fashion as part of the objective function of interest. The main contribution of this this paper belongs to this class of methods: we propose a novel adaptive normalization approach, called EDAIN (Extended Deep Adaptive Input Normalization), which can appropriately handle irregularities in the input data, without making any assumption on their distribution.

1.2 Contributions

The main contribution of our work is EDAIN, displayed in Figure 1, a neural layer that can be added to any neural network architecture for preprocessing

multivariate time series. This method complements the shift and scale layers proposed in DAIN (Passalis et al., 2020, described in details in Section 2) with an adaptive outlier mitigation sublayer and a power transform sublayer, used to handle common irregularities observed in real-world data, such as outliers, heavy tails, extreme values and skewness.

Additionally, our EDAIN method can be implemented in two versions, named *global-aware* and *local-aware*, suited to unimodal and multi-modal data respectively. Furthermore, we propose a computationally efficient variation of EDAIN, trained via the Kullback-Leibler divergence, named EDAIN-KL. Like EDAIN, this method can normalize skewed data with outliers, but in an unsupervised fashion, and can be used in conjunction with non-neural-network models.

EDAIN is described in details in Section 3, after a discussion on related methods in Section 2. The proposed methodology is extensively evaluated on synthetic and real-world data in Section 4, followed by a discussion on its performance and a conclusion. Also, an open-source implementation of the EDAIN layer, along with code for reproducing the experiments, is available in the GitHub repository [marcusGH/edain_paper](https://github.com/marcusGH/edain_paper).

2 RELATED WORK

Several works consider adaptive normalization methods, but they all apply the normalization transformation to the outputs of the inner layers within the neural network, known as activations in the literature. A well-known example of these transformations is batch normalization (Ioffe and Szegedy, 2015), which applies z -score normalization to the output of each inner layer, but several alternatives and extensions exist (see, for example, Huang et al., 2023; Lubana et al., 2021; Yu and Spiliopoulos, 2023). To the best of our knowledge, there are only three other methods where the deep neural network is augmented by inserting the adaptive preprocessing layer as the first step, transforming the data before it enters the network: the Deep Adaptive Input Normalization (DAIN) layer (Passalis et al., 2020), the Robust Deep Adaptive Input Normalization (RDAIN) layer (Passalis et al., 2021), and the Bilinear Input Normalization (BIN) layer (Tran et al., 2021).

The DAIN layer normalizes each time series $\mathbf{X}^{(i)}$ using three sublayers: each time series is first shifted, then scaled, and finally passed through a gating layer that can suppress irrelevant features. The unknown parameters are the weight matrices $\mathbf{W}_a, \mathbf{W}_b, \mathbf{W}_c \in \mathbb{R}^{d \times d}$, and the bias term $\mathbf{d} \in \mathbb{R}^d$, and are used for the shift, scale, and gating sublayer, respectively. The first two

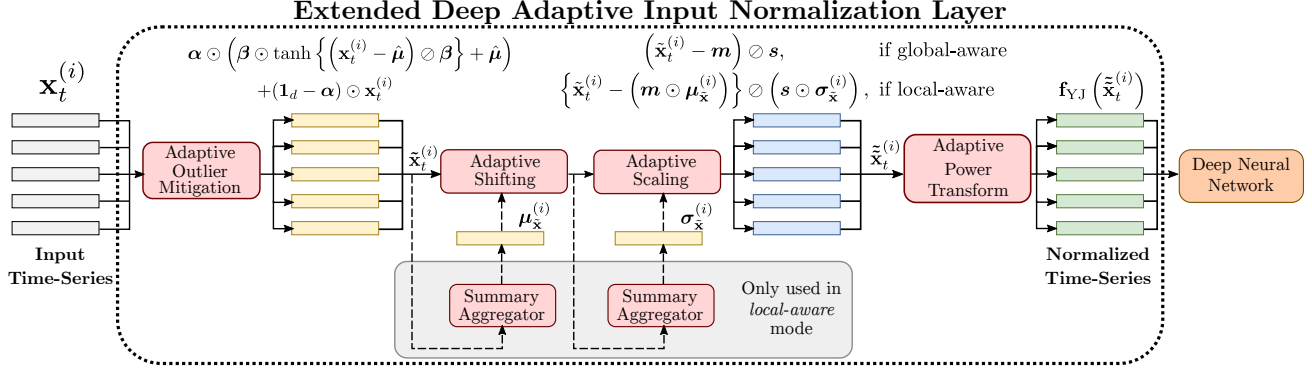


Figure 1: Architecture of the proposed EDAIN (Extended Deep Adaptive Input Normalization) layer. The layout and color choices of the diagram are based on Figure 1 from Passalis et al. (2020).

layers together, perform the operation

$$\tilde{\mathbf{x}}_t^{(i)} = \left(\mathbf{x}_t^{(i)} - \mathbf{W}_a \mathbf{a}^{(i)} \right) \oslash \mathbf{W}_b \mathbf{b}^{(i)},$$

where $\mathbf{x}_t^{(i)} \in \mathbb{R}^d$ is the input feature vector at timestep t of time series i , \oslash denotes element-wise division, and $\mathbf{a}^{(i)} \in \mathbb{R}^d$ and $\mathbf{b}^{(i)} \in \mathbb{R}^d$ are summary statistics that are computed for the i -th time series as follows:

$$\mathbf{a}^{(i)} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t^{(i)}, \quad \mathbf{b}^{(i)} = \sqrt{\frac{1}{T} \sum_{t=1}^T \left(\mathbf{x}_t^{(i)} - \mathbf{W}_a \mathbf{a}^{(i)} \right)^2}. \quad (1)$$

In Equation (1), the power operations are applied element-wise. The third sublayer, the gating layer, performs the operation

$$\tilde{\tilde{\mathbf{x}}}_t^{(i)} = \tilde{\mathbf{x}}_t^{(i)} \odot S \left(\mathbf{W}_c \mathbf{c}^{(i)} + \mathbf{d} \right).$$

Here, \odot is the element-wise multiplication operator, $S : \mathbb{R}^d \rightarrow \mathbb{R}^d$ denotes the logistic function applied element-wise, and $\mathbf{c}^{(i)}$ is the summary statistic

$$\mathbf{c}^{(i)} = \frac{1}{T} \sum_{t=1}^T \tilde{\tilde{\mathbf{x}}}_t^{(i)}.$$

The final output of the DAIN layer is thus

$$\tilde{\tilde{\mathbf{X}}}^{(i)} = \left[\tilde{\tilde{\mathbf{x}}}_1^{(i)}, \dots, \tilde{\tilde{\mathbf{x}}}_T^{(i)} \right] \in \mathbb{R}^{d \times T}.$$

In the RDAIN layer proposed by Passalis et al. (2021), a similar 3-stage normalization pipeline as that of the DAIN layer is used, but a residual connection across the shift and scale sublayers is also introduced. The BIN layer (Tran et al., 2021) has two sets of linear shift and scale sublayers that work similarly to the DAIN layer, which are applied across columns and rows of each time series $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$. The output of the BIN layer is a trainable linear combination of the two.

In addition to the described adaptive preprocessing methods, using a combination of static preprocessing methods has been proposed. In particular, McCarter proposes the Kernel Density Integral Transformation (KDJIT). In KDJIT, the data distribution is estimated via a Gaussian kernel density estimate with bandwidth depending on a parameter $\alpha \in \mathbb{R}_+$. The estimated density is then used to construct an estimate of the cumulative distribution function, which is used in turn to standardise the data to the range $[0, 1]$ (McCarter, 2023). For $\alpha \rightarrow \infty$, KDJIT converges to min-max scaling, whereas $\alpha \rightarrow 0$ corresponds to a quantile transformation (McCarter, 2023).

In real world applications, data often present additional irregularities, such as outliers, extreme values, heavy tails and skewness, which the aforementioned adaptive preprocessing methods are not designed to handle. Therefore, this work proposes EDAIN, a layer which comprises two novel sublayers that can appropriately treat skewed and heavy-tailed data with outliers and extreme values, resulting in significant improvements in performance metrics on real and simulated data. Also, the DAIN, RDAIN and BIN adaptive preprocessing methods are only primarily designed to handle multi-modal and non-stationary time series, which are common in financial forecasting tasks (Passalis et al., 2020). They do this by making the shift and scale parameters a parameterised function of each $\mathbf{X}^{(i)}$, allowing a transformation specific to each time series data point, henceforth referred to as *local-aware* preprocessing. However, these normalization schemes do not necessarily preserve the relative ordering between time series data points, which can degrade performance on unimodal datasets. As discussed in the next section, we address this drawback by proposing a novel *global-aware* version for our proposed EDAIN layer, which preserves ordering by learning a monotonic transformation. It must be remarked that the EDAIN layer can also be fitted in local-aware fashion,

to address multi-modality when present, providing additional flexibility when modelling real-world data.

3 EXTENDED DEEP ADAPTIVE INPUT NORMALIZATION

In this section, we describe in detail the novel EDAIN preprocessing layer, which can be added to any deep learning architecture for time series data. EDAIN adaptively applies local transformations specific to each time series, or global transformations across all the observed time series $\mathbf{X}^{(i)}$, $i = 1, \dots, N$ in the dataset \mathcal{D} . These transformations are aimed at appropriately preprocessing the data, mitigating the effect of skewness, outliers, extreme values and heavy-tailed distributions. This section discusses in details the different sublayers of EDAIN, its global-aware and local-aware versions, and training strategies via stochastic gradient descent or Kullback-Leibler divergence minimization.

An overview of the EDAIN layer’s architecture is shown in Figure 1. Given some input time series $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$, each feature vector $\mathbf{x}_t^{(i)} \in \mathbb{R}^d$ is independently transformed sequentially in four stages: an outlier mitigation operation $\mathbf{h}_1 : \mathbb{R}^d \rightarrow \mathbb{R}^d$, a shift operation: $\mathbf{h}_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$, a scale operation: $\mathbf{h}_3 : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and a power transformation operation: $\mathbf{h}_4 : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

Outlier mitigation sublayer. In the literature, it has been shown that an appropriate treatment of outliers and extreme values can increase predictive performance (Yin and Liu, 2022). The two most common ways of doing this are omission and winsorization (Nyitrai and Virág, 2019). The former corresponds to removing the outliers from further analysis, whereas the latter seeks to replace outliers with a censored value corresponding to a given percentile of the observations. In this work we propose the following smoothed winsorization operation obtained via the $\tanh(\cdot)$ function:

$$\check{\mathbf{x}}_t^{(i)} = \boldsymbol{\beta} \odot \tanh \left\{ \left(\mathbf{x}_t^{(i)} - \hat{\boldsymbol{\mu}} \right) \oslash \boldsymbol{\beta} \right\} + \hat{\boldsymbol{\mu}}, \quad (2)$$

where the parameter $\boldsymbol{\beta} \in [\beta_{\min}, \infty)^d$ controls the range to which the measurements are restricted to, and $\hat{\boldsymbol{\mu}} \in \mathbb{R}^d$ is the global mean of the data, considered as a fixed constant. In this work, we let $\beta_{\min} = 1$. Additionally, we consider a ratio of winsorization to apply to each predictor variable, controlled by an unknown parameter vector $\boldsymbol{\alpha} \in [0, 1]^d$, combined with the smoothed winsorization operator (2) via a residual connection. This gives the following adaptive outlier mitigation operation for an input time series $\mathbf{x}_t^{(i)} \in \mathbb{R}^d$:

$$\mathbf{h}_1 \left(\mathbf{x}_t^{(i)} \right) = \boldsymbol{\alpha} \odot \check{\mathbf{x}}_t^{(i)} + (\mathbf{1}_d - \boldsymbol{\alpha}) \odot \mathbf{x}_t^{(i)},$$

where $\mathbf{1}_d$ is a d -dimensional vector of ones. Both $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are considered as unknown parameters as part of the full objective function optimised during training.

Shift and scale sublayers. The adaptive shift and scale layer, combined, perform the operation

$$\mathbf{h}_3 \left\{ \mathbf{h}_2 \left(\mathbf{x}_t^{(i)} \right) \right\} = \left(\mathbf{x}_t^{(i)} - \mathbf{m} \right) \oslash \mathbf{s},$$

where the unknown parameters are $\mathbf{m} \in \mathbb{R}^d$ and $\mathbf{s} \in (0, \infty)^d$. Note that the EDAIN scale and shift sublayers generalise z -score scaling, which does not treat \mathbf{m} and \mathbf{s} as unknown parameters, but it sets them to the mean and standard deviation instead:

$$\mathbf{m} = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{x}_t^{(i)},$$

and

$$\mathbf{s} = \sqrt{\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left(\mathbf{x}_t^{(i)} - \mathbf{m} \right)^2},$$

where the power operations are applied element-wise.

Power transform sublayer. Many real-world datasets exhibit significant skewness, which is often corrected using power transformations (Schroth and Muma, 2021), such as the commonly used Box-Cox transformation (Box and Cox, 1964). One of the main limitations of the Box-Cox transformation is that it is only valid for positive values. A more general alternative which is more widely applicable is the Yeo-Johnson (YJ) transform (Yeo and Johnson, 2000):

$$f_{\text{YJ}}^\lambda(x) = \begin{cases} \frac{(x+1)^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0, x \geq 0, \\ \log(x+1) & \text{if } \lambda = 0, x \geq 0, \\ \frac{(1-x)^{2-\lambda} - 1}{\lambda - 2} & \text{if } \lambda \neq 2, x < 0, \\ -\log(1-x) & \text{if } \lambda = 2, x < 0. \end{cases} \quad (3)$$

The transformation f_{YJ}^λ only has one unknown parameter, $\lambda \in \mathbb{R}$, and it can be applied to any $x \in \mathbb{R}$, not just positive values (Yeo and Johnson, 2000). The power transform sublayer of EDAIN simply applies the transformation in Equation (3) along each dimension of the input time series $\mathbf{X}^{(i)}$. That is, for each $i = 1, \dots, N$ and $t = 1, \dots, T$, the sublayer outputs

$$\mathbf{h}_4 \left(\mathbf{x}_t^{(i)} \right) = \left[f_{\text{YJ}}^{\lambda_1} \left(x_{t,1}^{(i)} \right), \dots, f_{\text{YJ}}^{\lambda_d} \left(x_{t,d}^{(i)} \right) \right],$$

where the unknown quantities to be optimised are the power parameters $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d) \in \mathbb{R}^d$.

3.1 Global- and local-aware normalization

For highly multi-modal and non-stationary time series data, Passalis et al. (2020, 2021) and Tran et al.

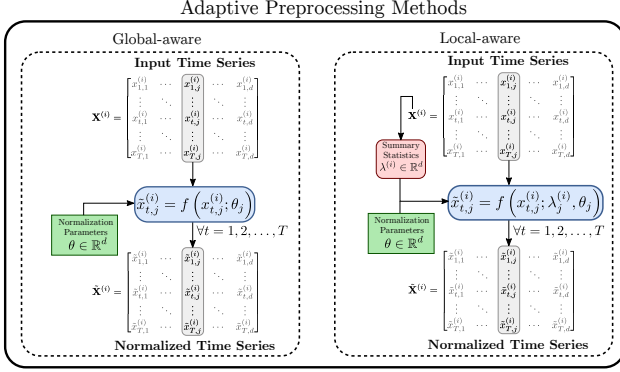


Figure 2: Visual comparison of the local- and global-aware versions of adaptive preprocessing schemes.

(2021) observed significant performance improvements when using local-aware preprocessing methods, as these allow forming a unimodal representation space from predictor variables with multi-modal distributions. Therefore, we also propose a *local-aware* version of the EDAIN layer in addition to the *global-aware* version we presented earlier. In the local-aware version of EDAIN, the shift and scale operations also depend on a summary representation of the current time series $\mathbf{X}^{(i)}$ to be preprocessed:

$$\mathbf{h}_3 \left\{ \mathbf{h}_2 \left(\mathbf{x}_t^{(i)} \right) \right\} = \left\{ \mathbf{x}_t^{(i)} - \left(\mathbf{m} \odot \boldsymbol{\mu}_{\mathbf{x}}^{(i)} \right) \right\} \odot \left(\mathbf{s} \odot \boldsymbol{\sigma}_{\mathbf{x}}^{(i)} \right).$$

The summary representations $\boldsymbol{\mu}_{\mathbf{x}}^{(i)}, \boldsymbol{\sigma}_{\mathbf{x}}^{(i)} \in \mathbb{R}^d$, are computed through a reduction along the temporal dimension of each time series $\mathbf{X}^{(i)}$ (cf. Figure 2):

$$\boldsymbol{\mu}_{\mathbf{x}}^{(i)} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t^{(i)}, \quad \boldsymbol{\sigma}_{\mathbf{x}}^{(i)} = \sqrt{\frac{1}{T} \sum_{t=1}^T \left(\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_{\mathbf{x}}^{(i)} \right)^2}, \quad (4)$$

where the power operations are applied element-wise in (4). The outlier mitigation and power transform sublayers remain the same for the local-aware version, except the $\hat{\boldsymbol{\mu}}$ statistic in Equation (2) is no longer fixed, but rather the mean of the input time series:

$$\hat{\boldsymbol{\mu}}^{(i)} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t^{(i)}.$$

Order preservation. Adaptive local-aware preprocessing methods work well on multi-modal data such as financial forecasting datasets (Passalis et al., 2020, 2021; Tran et al., 2021). However, as we will experimentally demonstrate in the next section, local-aware preprocessing methods might perform worse than conventional methods such as z -score normalization when applied to unimodal data. This is because local-aware methods are not *order preserving*, since the shift and

scale amount is different for each time series. In many applications in which the features have a natural and meaningful ordering, it would be desirable to have

$$\tilde{x}_{t,k}^{(i)} < \tilde{x}_{t,k}^{(j)} \text{ if } x_{t,k}^{(i)} < x_{t,k}^{(j)}, \quad (5)$$

for all $i, j \in \{1, \dots, N\}$, $t = 1, \dots, T$ and $k = 1, \dots, d$, where $\tilde{\mathbf{X}}^{(i)}$ denotes the output of a transformation with $\mathbf{X}^{(i)}$ as input. This property does not necessarily hold for the local-aware methods (DAIN, RDAIN, BIN, and local-aware EDAIN). For unimodal features, the qualitative interpretation of the predictor variables might dictate that property (5) should be maintained (for example, consider the case of credit scores for a default prediction application, cf. Section 4.2).

As a solution, the proposed *global-aware* version of the proposed EDAIN layer does not use any time series specific summary statistics, which makes each of the four sublayers monotonically non-decreasing. This ensures that property (5) is maintained by the global-aware EDAIN transformation, providing additional flexibility for applications on real world data where ordering within features should be preserved.

3.2 Optimising the EDAIN layer

The output of the proposed EDAIN layer is obtained by feeding the input time series $\mathbf{X}^{(i)}$ through the four sublayers in a feed-forward fashion, as shown in Figure 1. The output is then fed to the deep neural network used for the task at hand. Letting \mathbf{W} denote the weights of the deep neural network model, the weights of both the deep model and the EDAIN layer are simultaneously optimised in an end-to-end manner using stochastic gradient descent, with the update equation:

$$\begin{aligned} \Delta(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{m}, \mathbf{s}, \boldsymbol{\lambda}, \mathbf{W}) = \\ - \eta \left(\eta_1 \frac{\partial \mathcal{L}}{\partial \boldsymbol{\alpha}}, \eta_1 \frac{\partial \mathcal{L}}{\partial \boldsymbol{\beta}}, \eta_2 \frac{\partial \mathcal{L}}{\partial \mathbf{m}}, \eta_3 \frac{\partial \mathcal{L}}{\partial \mathbf{s}}, \eta_4 \frac{\partial \mathcal{L}}{\partial \boldsymbol{\lambda}}, \frac{\partial \mathcal{L}}{\partial \mathbf{W}} \right), \end{aligned}$$

where $\eta \in \mathbb{R}_+$ is the base learning rate, whereas $\eta_1, \dots, \eta_4 \in \mathbb{R}_+$ correspond to sublayer-specific corrections to the global learning rate η . As Passalis et al. (2020) observed when training their DAIN layer, the gradients of the unknown parameters for the different sublayers might have vastly different magnitudes, which prevents a smooth convergence of the preprocessing layer. Therefore, they proposed using separate learning rates for the different sublayers. We therefore introduce corrections η_ℓ , $\ell = \{1, 2, 3, 4\}$ as additional hyperparameters that modify the learning rates for each of the four different EDAIN sublayers.

Furthermore, note that computing the fixed constant $\hat{\boldsymbol{\mu}}$ in the outlier mitigation sublayer (2) would require a sweep on the entire dataset before training the

EDAIN-augmented neural network architecture, which could be computationally extremely expensive. As a solution to circumvent this issue, we propose to calculate $\hat{\boldsymbol{\mu}}$ iteratively during training, updating it using a cumulative moving average estimate at each forward pass of the sublayer. We provide more details on this in Appendix A.

3.3 EDAIN-KL

In addition to the EDAIN layer, we also propose another novel preprocessing method, named EDAIN-KL (Extended Deep Adaptive Input Normalization, optimised with Kullback-Leibler divergence). This approach uses a similar neural layer architecture as the EDAIN method, but modifies it to ensure the transformation is invertible. Its unknown parameters are then optimised with an approach inspired by normalizing flows (see, for example, Kobyzev et al., 2021).

The EDAIN-KL layer is used to transform a Gaussian base distribution $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, I_{dT})$ via a composite function $\mathbf{g}_{\boldsymbol{\theta}} = \mathbf{h}_1^{-1} \circ \mathbf{h}_2^{-1} \circ \mathbf{h}_3^{-1} \circ \mathbf{h}_4^{-1}$ comprised of the inverses of the operations in the EDAIN sublayers, applied sequentially with parameter $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{m}, \mathbf{s}, \boldsymbol{\lambda})$. The parameter $\boldsymbol{\theta}$ is chosen to minimize the KL-divergence between the resulting distribution $\mathbf{g}_{\boldsymbol{\theta}}(\mathbf{Z})$ and the empirical distribution of the dataset \mathcal{D} :

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \operatorname{KL} \{ \mathcal{D} \parallel \mathbf{g}_{\boldsymbol{\theta}}(\mathbf{Z}) \}.$$

Note that we apply all the operations in reverse order, compared to the EDAIN layer, because we use $\mathbf{g}_{\boldsymbol{\theta}}$ to transform a base distribution \mathbf{Z} into a distribution that resembles the training dataset \mathcal{D} . To normalize the dataset after fitting the EDAIN-KL layer, we apply

$$\mathbf{g}_{\hat{\boldsymbol{\theta}}}^{-1} = \mathbf{h}_4 \circ \mathbf{h}_3 \circ \mathbf{h}_2 \circ \mathbf{h}_1$$

to each $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$, similarly to the EDAIN layer. The main advantage of the EDAIN-KL approach over standard EDAIN is that it allows training in an unsupervised fashion, separate from the deep model. This enables its usage for preprocessing data in a wider set of tasks, including non-deep-neural-network models. An exhaustive description of the EDAIN-KL method is provided in Appendix B.

4 EXPERIMENTAL EVALUATION

For evaluating the proposed EDAIN layer we consider a synthetic dataset, a large-scale default prediction dataset, and a large-scale financial forecasting dataset. We compare the two versions of the EDAIN layer (global-aware and local-aware) and the EDAIN-KL layer to the DAIN (Passalis et al., 2020) layer,




Figure 3: Histogram across timesteps $t = 1, \dots, T$ of the $d = 3$ predictor variables from the synthetic data.

to the BIN (Tran et al., 2021) layer and to the KDIT method. We also consider a statistical baseline consisting of different combinations of z -score scaling, win-sorization and a Yeo-Johnson power transformation. For all experiments, we use a recurrent neural network (RNN) model composed of gated recurrent unit (GRU) layers, followed by a classifier head with fully connected layers. Categorical features, when present, are passed through an embedding layer, whose output is combined with the output of the GRU layers and then fed to the classifier head. Full details on the model architectures, optimization procedures, including learning rates and number of epochs, can be found in Appendix C and in the code repository associated with this work.

4.1 Synthetic Datasets

Before considering real-world data, we evaluate our method on synthetic data, where we have full control over the data generating process. To do this, we develop a synthetic time series data generation algorithm which allows specifying arbitrary unnormalized probability density functions (PDFs) for each of the d predictor variables. It then generates N time series of the form $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$, along with N binary response variables $y^{(i)} \in \{0, 1\}$. We present a detailed description of the algorithm in Appendix D.

For our experiments, we generated $N_{\mathcal{D}} = 100$ datasets, each with $N = 50\,000$ time series of length $T = 10$ and dimensionality $d = 3$. The three predictor variables were configured to be distributed as follows:

$$f_1(x) = 10 \cdot \Phi_{\mathcal{N}} \{ 10(x + 4) \} \cdot p_{\mathcal{N}}(x + 4) + \mathbb{I}_{(8,9.5)}(x) \cdot e^{x-8}/10, \quad (6)$$

$$f_2(x) = \begin{cases} 20 \cdot p_{\mathcal{N}}(x - 20), & \text{if } x > \pi, \\ e^{x/6} \cdot \{ 10 \sin(x) + 10 \}, & \text{if } x \leq \pi, \end{cases} \quad (7)$$

$$f_3(x) = 2 \cdot \Phi_{\mathcal{N}} \{ -4(x - 4) \} \cdot p_{\mathcal{N}}(x - 4), \quad (8)$$

where $p_{\mathcal{N}}(\cdot)$ and $\Phi_{\mathcal{N}}(\cdot)$ denote the PDF and cumulative distribution function (CDF) of the standard normal distribution, and $\mathbb{I}_{\mathcal{A}}(\cdot)$ is the indicator function on the set \mathcal{A} . Samples from the dataset are visualised in Figure 3. We train and evaluate a RNN model

with the architecture described earlier on each of the $N_{\mathcal{D}}$ datasets using a 80%-20% train-validation split. Our results are presented in Table 1, where the binary cross-entropy (BCE) loss and the accuracy on the validation set are used as evaluation metrics.

From our experiments on the synthetic datasets, we observe that the model performance is more unstable when no preprocessing is applied, as seen from the increased variance in Table 1. We also observe that z -score normalization only gives minor performance improvements when compared to no preprocessing, aside from reducing the variance. As we have perfect information about the underlying data generation mechanism from Equations (6), (7) and (8), we also compared our methods to what we refer to as *CDF inversion*, where each observation is transformed by the CDF of its corresponding distribution, and then transformed via $\Phi_{\mathcal{N}}^{-1}(\cdot)$, giving predictor variables with standard normal distributions. We also apply this method to the real-world datasets in Section 4.2 and Section 4.3, but since the true PDFs are unknown in those settings, we estimate the CDFs using quantiles from the distribution of the training data. The CDF inversion method resembles KDIT for low α , since KDIT becomes a quantile transformation in the limiting case $\alpha \rightarrow 0$ (McCarter, 2023). The main difference is that the KDIT method does not involve the Gaussianization step, implying that the preprocessing matches the data generation mechanism, representing a gold-standard for this example. As discussed in more details in Appendix D, the synthetic responses are generated from a linear combination of uniform random variables, so the KDIT gives the model a stronger signal than that of the CDF inversion method. Out of all the methods not exploiting the mechanics of the underlying data generation mechanism, the global-aware version of EDAIN demonstrates superior performance. It also almost performs as well as CDF inversion, which is able to perfectly normalize each predictor variable via its data generation mechanism. Finally, we observe that the local-aware methods perform even worse than no preprocessing: this might be due to the ordering not being preserved, as discussed in Section 3.1.

4.2 Default Prediction Dataset

The first real-world dataset we consider is the publicly available default prediction dataset published by American Express (Howard et al., 2022), which contains data from $N = 458\,913$ credit card customers. For each customer, a vector of $d = 188$ aggregated profile features has been recorded at $T = 13$ different credit card statement dates, producing a multivariate time series of the form $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$, $i = 1, 2, \dots, N$. Given $\mathbf{X}^{(i)}$, the task is to predict a binary label

Preprocessing method	BCE loss	Binary accuracy (%)
No preprocessing	0.1900 ± 0.0036	0.9168 ± 0.0017
z -score	0.1871 ± 0.0011	0.9176 ± 0.0007
z -score + YJ	0.1789 ± 0.0009	0.9211 ± 0.0006
Winsorize + z -score	0.1876 ± 0.0010	0.9172 ± 0.0006
Winsorize + z -score + YJ	0.1788 ± 0.0009	0.9211 ± 0.0006
<u>CDF inversion</u>	<u>0.1627 ± 0.0009</u>	<u>0.9289 ± 0.0006</u>
BIN	0.2191 ± 0.0010	0.9036 ± 0.0006
DAIN	0.2153 ± 0.0015	0.9048 ± 0.0008
EDAIN (local-aware)	0.2099 ± 0.0010	0.9071 ± 0.0006
EDAIN (global-aware)	0.1636 ± 0.0009	0.9283 ± 0.0005
EDAIN-KL	0.1760 ± 0.0009	0.9224 ± 0.0006
KDIT ($\alpha = 0.1$)	<u>0.1532 ± 0.0011</u>	<u>0.9329 ± 0.0006</u>

Table 1: Experimental results on synthetic data, with 95% normal confidence intervals $\mu \pm 1.96\sigma/\sqrt{K}$ calculated across $K = 100$ datasets. The gold-standard CDF-based transformations are underlined.

Preprocessing method	BCE loss	AMEX metric
No preprocessing	0.3242 ± 0.0066	0.6430 ± 0.0087
z -score	0.2213 ± 0.0017	0.7872 ± 0.0030
Winsorize + z -score	0.2217 ± 0.0018	0.7867 ± 0.0025
z -score + YJ	0.2224 ± 0.0014	0.7846 ± 0.0022
Winsorize + z -score + YJ	0.2926 ± 0.1229	0.6321 ± 0.2679
CDF inversion	0.2215 ± 0.0018	0.7861 ± 0.0032
EDAIN-KL	0.2218 ± 0.0018	0.7858 ± 0.0027
EDAIN (local-aware)	0.2245 ± 0.0015	0.7813 ± 0.0025
EDAIN (global-aware)	0.2199 ± 0.0015	0.7890 ± 0.0035
DAIN	0.2224 ± 0.0016	0.7847 ± 0.0024
BIN	0.2237 ± 0.0017	0.7829 ± 0.0029
KDIT ($\alpha = 1$)	0.2258 ± 0.0014	0.7791 ± 0.0024

Table 2: Experimental results on the default prediction dataset, with 95% normal asymptotic confidence intervals $\mu \pm 1.96\sigma/\sqrt{K}$ calculated across $K = 5$ folds.

$y^{(i)} \in \{0, 1\}$ indicating whether the i -th customer defaulted at any point within 18 months after the last observed data point in the time series. A default event is defined as not paying back the credit card balance amount within 120 days after the latest statement date (Howard et al., 2022). Out of the $d = 188$ features, only the 177 numerical variables are preprocessed.

To evaluate the different preprocessing methods, we perform 5-fold cross validation, which produces evaluation metrics for five different 20% validation splits. The evaluation metrics we consider are the validation BCE loss and a metric that was proposed by Howard et al. (2022) for use with this dataset, and we refer to it as the *Amex metric*. This metric is calculated as $M = 0.5 \cdot (G + D)$, where D is the default rate captured at 4% (corresponding to the proportion of positive labels captured within the highest-ranked 4% of the model predictions) and G is the normalized Gini coefficient (see Appendix E and, for example, Lerman and Yitzhaki, 1989).

Our results are reported in Table 2. From the table, it can be inferred that neglecting the preprocessing step

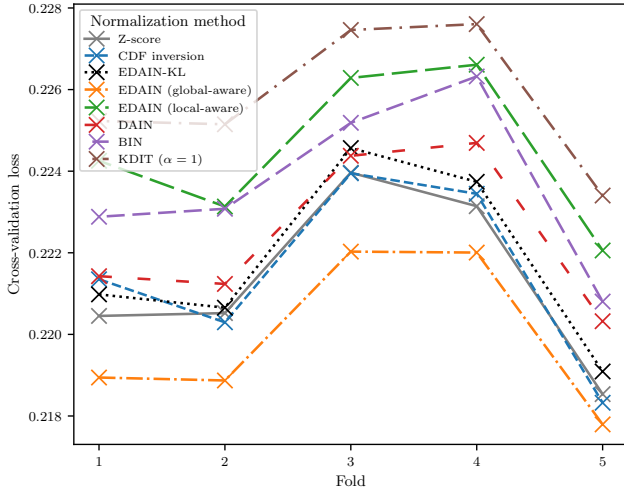


Figure 4: BCE cross-validation loss across different folds in the Amex default prediction dataset.

deteriorates the performance significantly. Moreover, we observe that the local-aware methods (local-aware EDAIN, BIN, and DAIN) all perform worse than z -score normalization. This might be because the data in the default prediction dataset is mostly distributed around one central mode, and the local-aware methods discard this information in favour of forming a common representation space. Another possible reason is that the local-aware methods do not preserve the relative ordering between data points $\mathbf{X}^{(i)}$ as per Equation (5), which might be detrimental for these types of datasets. As discussed in the previous section, predictor variables such as credit scores may be present in the dataset, which should only be preprocessed via monotonic transformations. From the results presented in Table 2, we can also conclude that the proposed global-aware version of the EDAIN layer shows superior average performance when compared to all alternative preprocessing methods. Additionally, it must be remarked that most of the variance in the methods’ performance arise from the folds themselves, as seen in Figure 4, and the global-aware EDAIN method consistently shows superior performance across all cross-validation folds when compared to the other methods. This was further confirmed via a paired sign test between EDAIN and BIN (the second best performing method), which returned p -value 0.015, indicating a significant difference in the performance of the two preprocessing methods.

4.3 Financial Forecasting Dataset

For evaluating the proposed methods, we also considered a limit order book (LOB) time series dataset (FI-2010 LOB Ntakaris et al., 2018), used as a benchmark dataset in Passalis et al. (2020) and Tran et al.

Preprocessing method	Cohen’s κ	Macro- F_1 -score
No preprocessing	0.0035 ± 0.0016	0.2859 ± 0.0076
z -score	0.2777 ± 0.0185	0.5052 ± 0.0137
Winsorize + z -score	0.2928 ± 0.0205	0.5166 ± 0.0168
CDF inversion	0.3618 ± 0.0199	0.5798 ± 0.0124
BIN	0.3670 ± 0.0213	0.5889 ± 0.0160
DAIN	0.3588 ± 0.0169	0.5776 ± 0.0114
EDAIN (local-aware)	0.3836 ± 0.0185	0.5946 ± 0.0144
EDAIN (global-aware)	0.2820 ± 0.0235	0.5111 ± 0.0216
EDAIN-KL	0.2870 ± 0.0214	0.5104 ± 0.0173
KDIT ($\alpha = 0.1$)	0.2974 ± 0.0225	0.5260 ± 0.0170

Table 3: Experimental results on the FI-2010 LOB dataset, with 95% normal asymptotic confidence intervals $\mu \pm 1.96\sigma/\sqrt{K}$ calculated across $K = 9$ anchored folds.

(2021). The data was collected across 10 business days in June 2010 from five Finnish companies (Ntakaris et al., 2018), and it was cleaned and features were extracted based on the pipeline of Kercheval and Zhang (2015). This resulted in $N = 453\,975$ vectors of dimensionality $d = 144$. The task is predicting whether the mid price will increase, decrease or remain stationary with a prediction horizon of $H = 10$ timesteps, where a stock is labelled as stationary if the mid price changes by less than 0.01%. More details on the FI-2010 benchmark dataset can be found in Ntakaris et al. (2018).

For training, we use the anchored cross-validation scheme of Ntakaris et al. (2018). We sequentially increase the size of the training set, starting with a single day and extending it to nine days, while reserving the subsequent day for evaluation in each iteration, obtaining nine different evaluation folds. For evaluating the model performance, we look at the Cohen’s κ (Artstein and Poesio, 2008) and macro- F_1 score, obtained by averaging class-specific F_1 scores across the three possible outcomes (increase, decrease or stationary).

We can draw several conclusions from the results reported in Table 3. Firstly, employing some form of preprocessing is essential, as not applying any preprocessing gives κ values close to 0, which is what is expected to occur by random guessing. Secondly, applying a local-aware normalization scheme (local-aware EDAIN, BIN and DAIN) greatly improves performance when compared to conventional preprocessing methods such as z -score scaling. Our third observation is that the proposed local-aware EDAIN method outperforms both BIN and DAIN on average. Finally, we note that most of the variability in the evaluation metrics arises from the folds themselves, similarly to the default prediction example (*cf.* Figure 4). Again, this was further assessed via a paired sign test between EDAIN and the second best performing method, resulting in a p -value 0.00097, which confirms the significance of our results.

Preprocessing method	BCE loss	AMEX metric
<i>z</i> -score	0.2214 ± 0.0018	0.7873 ± 0.0029
scale	0.2214 ± 0.0016	0.7863 ± 0.0029
shift	0.2213 ± 0.0016	0.7871 ± 0.0035
shift+scale	0.2212 ± 0.0018	0.7872 ± 0.0035
shift+scale+PT	0.2207 ± 0.0016	0.7885 ± 0.0026
OM+shift+scale	0.2203 ± 0.0015	0.7884 ± 0.0031
OM+shift+scale+PT	0.2199 ± 0.0015	0.7889 ± 0.0031

Table 4: Ablation study on the default prediction, with 95% normal asymptotic confidence intervals $\mu \pm 1.96\sigma/\sqrt{K}$ calculated across $K = 5$ folds. In all but the first row, the global-aware EDAIN layer is used with the corresponding subset of its sublayers enabled.

4.4 Ablation Study

To identify the effect of each of the four sublayers on the final predictive performance, we conducted an ablation study. The results are reported in Table 4, where OM and PT refer to the outlier mitigation and power transform sublayers, respectively. These experiments were conducted on the default prediction dataset using the same model architecture and evaluation procedure as described in Section 4.2. We observe that just using adaptive shift and scale sublayers provides marginal gains over the *z*-score scaling baseline. Meanwhile, introducing either the adaptive power transform or the outlier mitigation sublayers reduces the loss to a greater extent, with the outlier mitigation sublayer proving most effective.

5 CONCLUSION

In this work, we proposed EDAIN, an adaptive data preprocessing layer which can be used to augment any deep neural network architecture that takes multivariate time series as input. EDAIN has four adaptive sublayers (outlier mitigation, shift, scale, and power transform). It also has two versions (local-aware and global-aware), which apply local or global transformations to each time series. Also, we proposed a computationally efficient variant of EDAIN, optimised via the Kullback-Leibler divergence, named EDAIN-KL. The EDAIN layer’s ability to increase the predictive performance of the deep neural network was evaluated on a synthetic dataset, a default prediction dataset, and a financial forecasting dataset. On all datasets considered, either the local-aware or global-aware version of the proposed EDAIN layer consistently demonstrated superior performance.

In Section 4.3, we observed that the local-aware preprocessing methods gave significantly better performance than the global-aware version of EDAIN and *z*-score normalization. However, in Section 4.1 and Sec-

tion 4.2 the opposite is observed, with the global-aware version of EDAIN demonstrating superior performance and the local-aware methods being outperformed by *z*-score normalization. We hypothesize these differences occur because local-aware methods do not preserve the relative ordering between observations, while the global-aware EDAIN method does. In the financial forecasting dataset, which is highly multi-modal, it appears that the observations’ feature values relative to their mode are more important than their absolute ordering. Such considerations should be taken into account when deciding what adaptive preprocessing method is most suitable for application on new data.

There are several directions for future work. Passalis et al. (2021) observed that the performance improvements from DAIN differed greatly between different deep neural network architectures. Therefore, the effectiveness of EDAIN with other architectures could be further explored as only GRU-based RNNs were considered in this work. Additionally, with the proposed EDAIN method, one has to manually decide whether to apply local-aware or global-aware preprocessing. This drawback could be eliminated by extending the proposed neural layer to apply both schemes to each feature and adaptively learning which version is most suitable. Another common irregularity in real-world data is missing values (Nawi et al., 2013; Cao et al., 2018). A possible direction would be to extend EDAIN with an adaptive method for treating missing values that makes minimal assumptions on the data generation mechanism (for example, Cao et al., 2018).

References

- Artstein, R. and Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Atkinson, K. E. (1989). *An Introduction to Numerical Analysis*. John Wiley & Sons, New York, 2nd edition.
- Box, G. E. P. and Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2):211–252.
- Cao, W., Wang, D., Li, J., Zhou, H., Li, L., and Li, Y. (2018). Brits: Bidirectional recurrent imputation for time series. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Cao, Z., Gao, X., Chang, Y., Liu, G., and Pei, Y. (2023). Improving synthetic CT accuracy by combining the benefits of multiple normalized preprocesses. *Journal of Applied Clinical Medical Physics*, page e14004.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and

- Stein, C. (2001). *Introduction to Algorithms*. The MIT Press, 2nd edition.
- Higham, N. J. (1988). Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra and its Applications*, 103:103–118.
- Hinton, G. and Tieleman, T. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, (4):26–31.
- Howard, A., Bera, A., Xu, D. W., Vashani, H., Hayatbini, N., and Dane, S. (2022). American Express – Default Prediction. <https://kaggle.com/competitions/amex-default-prediction>. Accessed: 2023-06-09.
- Huang, L., Qin, J., Zhou, Y., Zhu, F., Liu, L., and Shao, L. (2023). Normalization Techniques in Training DNNs: Methodology, Analysis and Application. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 45(08):10173–10196.
- Hyndman, R. J. and Athanasopoulos, G. (2018). *Forecasting: principles and practice, 2nd edition*. OTexts, Melbourne, Australia. Accessed: 29-08-2023.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Bach, F. and Blei, D., editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 448–456, Lille, France. PMLR.
- Kercheval, A. N. and Zhang, Y. (2015). Modelling high-frequency limit order book dynamics with support vector machines. *Quantitative Finance*, 15(8):1315–1329.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA*.
- Kobyzev, I., Prince, S. J., and Brubaker, M. A. (2021). Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979.
- Koval, S. I. (2018). Data preparation for neural network data analysis. In *2018 IEEE Conference of Russian Young Researchers in Electrical and Electronic Engineering (EIConRus)*, pages 898–901.
- Lerman, R. I. and Yitzhaki, S. (1989). Improving the accuracy of estimates of Gini coefficients. *Journal of Econometrics*, 42(1):43–47.
- Lubana, E. S., Dick, R., and Tanaka, H. (2021). Beyond batchnorm: Towards a unified understanding of normalization in deep learning. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc.
- McCarter, C. (2023). The kernel density integral transformation. *Transactions on Machine Learning Research*.
- Nawi, N. M., Atomi, W. H., and Rehman, M. (2013). The effect of data pre-processing on optimized training of artificial neural networks. *Procedia Technology*, 11:32–39. 4th International Conference on Electrical Engineering and Informatics, ICEEI 2013.
- Ntakaris, A., Magris, M., Kannianen, J., Gabbouj, M., and Iosifidis, A. (2018). Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. *Journal of Forecasting*, 37(8):852–866.
- Nyitrai, T. and Virág, M. (2019). The effects of handling outliers on the performance of bankruptcy prediction models. *Socio-Economic Planning Sciences*, 67:34–42.
- Passalis, N., Kannianen, J., Gabbouj, M., Iosifidis, A., and Tefas, A. (2021). Forecasting financial time series using robust deep adaptive input normalization. *Journal of Signal Processing Systems*, 93(10):1235–1251.
- Passalis, N., Tefas, A., Kannianen, J., Gabbouj, M., and Iosifidis, A. (2020). Deep adaptive input normalization for time series forecasting. *IEEE Transactions on Neural Networks and Learning Systems*, 31(9):3760–3765.
- Paszke et al., A. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Schroth, C. and Muma, M. (2021). Real elliptically skewed distributions and their application to robust cluster analysis. *IEEE Transactions on Signal Processing*, 69:3947–3962.
- Sola, J. and Sevilla, J. (1997). Importance of input data normalization for the application of neural networks to complex industrial problems. *IEEE Transactions on Nuclear Science*, 44:1464 – 1468.
- Tran, D. T., Kannianen, J., Gabbouj, M., and Iosifidis, A. (2021). Data normalization for bilinear structures in high-frequency financial time-series. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7287–7292.
- Yeo, I.-K. and Johnson, R. A. (2000). A new family of power transformations to improve normality or symmetry. *Biometrika*, 87(4):954–959.
- Yin, S. and Liu, H. (2022). Wind power prediction based on outlier correction, ensemble reinforcement

ment learning, and residual correction. *Energy*, 250:123857.

Yu, J. and Spiliopoulos, K. (2023). Normalization effects on deep neural networks. *Foundations of Data Science*, 5(3):389–465.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model.
Yes: a clear description of the proposed neural layer is provided in Section 3, and the setting in which it is to be used is described in Section 2.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm.
Not applicable: the proposed EDAIN and EDAIN-KL methods are not algorithms, but rather neural layers.
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries.
Yes: this is linked at the end of Section 1.
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results.
Not applicable.
 - (b) Complete proofs of all theoretical results.
Not applicable.
 - (c) Clear explanations of any assumptions.
Not applicable.
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL).
Yes: the code, along with instructions on how to acquire the data and reproduce the results, is provided in the code repository hyperlinked at the end of Section 1.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen).
Yes: these details are provided in Appendix C.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times).
Yes: a clear description of the confidence intervals presented in Tables 1, 2 and 3 are given in the captions.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider).
Yes: this is described in Appendix C.
 - (a) Citations of the creator If your work uses existing assets.
Yes: the existing assets used (DAIN, BIN, default prediction dataset, and FI-2010 LOB dataset) are all cited when introduced.
 - (b) The license information of the assets, if applicable.
Yes: license information about the assets used (datasets for experimental evaluation) is provided in the code repository hyperlinked at the end of Section 1.
 - (c) New assets either in the supplemental material or as a URL, if applicable.
Yes: the new assets (implementations of EDAIN, EDAIN-KL, and the synthetic data generation algorithm) is contained in the code repository hyperlinked at the start of Section 1.
 - (d) Information about consent from data providers/curators.
Yes: additional consent from data providers is reported in the code repository hyperlinked at the end of Section 1.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content.
Not applicable.
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots.
Not applicable.
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable.
Not applicable.
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation.
Not applicable.

Supplementary Materials

Extended Deep Adaptive Input Normalization for Preprocessing Time Series Data for Neural Network

A MEAN ESTIMATION IN THE OUTLIER MITIGATION LAYER

The $\hat{\boldsymbol{\mu}} \in \mathbb{R}^d$ vector in Equation (2) is an estimate of the mean of the data, and it is used to ensure the smoothed winsorization operation is centered. When using the local-aware version of the EDAIN layer, it is the mean of the input time series data point:

$$\hat{\boldsymbol{\mu}}^{(i)} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t^{(i)}.$$

In the global-aware version of EDAIN, we consider $\hat{\boldsymbol{\mu}} \in \mathbb{R}^d$ a fixed quantity, estimating the global mean of the data. As mentioned in the main paper, the $\hat{\boldsymbol{\mu}}$ estimate is iteratively computed during training using a *cumulative moving average estimate* at each forward pass of the sublayer. To do this, we keep track of the current estimated average at forward pass $n = 1, 2, 3, \dots$, denoted $\hat{\boldsymbol{\mu}}^{(n)}$, and when we process an input time series $\mathbf{X}^{(i)}$, we apply the update:

$$\hat{\boldsymbol{\mu}}^{(n+1)} = \frac{nT \cdot \hat{\boldsymbol{\mu}}^{(n)} + \sum_{t=1}^T \mathbf{x}_t^{(i)}}{(n+1)T}.$$

We also initialise $\hat{\boldsymbol{\mu}}^{(0)} = \mathbf{0}$.

B EDAIN-KL

The EDAIN-KL layer has a very similar architecture to the EDAIN layer (*cf.* Section 3 in the main paper), but the unknown parameters are learned via a different optimization procedure. Unlike the EDAIN layer, the EDAIN-KL layer is not attached to the deep neural network during training, but rather training in isolation before training the neural network. This is done by establishing an invertible mapping to transform a standard normal distribution into a distribution that resembles that of our training dataset. Then, after the EDAIN-KL weights have been optimized, we use the layer in reverse to normalize the time series from the training dataset before passing them to the neural network.

B.1 Architecture

Aside from the outlier mitigation sublayer, the EDAIN-KL layer has an identical architecture to the global-aware EDAIN layer. The outlier mitigation transformation has been simplified to ensure its inverse is analytic. Additionally, the layer no longer supports local-aware mode, as this would make the inverse intractable. The EDAIN-KL transformations are:

$$\text{(outlier mitigation)} \quad \mathbf{h}_1 \left(\mathbf{x}_t^{(i)} \right) = \boldsymbol{\beta} \odot \tanh \left\{ \left(\mathbf{x}_t^{(i)} - \hat{\boldsymbol{\mu}} \right) \oslash \boldsymbol{\beta} \right\} + \hat{\boldsymbol{\mu}} \quad (9)$$

$$\text{(shift)} \quad \mathbf{h}_2 \left(\mathbf{x}_t^{(i)} \right) = \mathbf{x}_t^{(i)} - \mathbf{m} \quad (10)$$

$$\text{(scale)} \quad \mathbf{h}_3 \left(\mathbf{x}_t^{(i)} \right) = \mathbf{x}_t^{(i)} \oslash \mathbf{s} \quad (11)$$

$$\text{(power transform)} \quad \mathbf{h}_4 \left(\mathbf{x}_t^{(i)} \right) = \left[f_{\text{YJ}}^{\lambda_1} \left(x_{t,1}^{(i)} \right), \quad \dots, \quad f_{\text{YJ}}^{\lambda_d} \left(x_{t,d}^{(i)} \right) \right], \quad (12)$$

where $f_{\text{YJ}}^{\lambda_i}(\cdot)$ is defined in the main paper in Equation (3).

B.2 Optimisation through Kullback-Leibler divergence

The optimisation approach used to train the EDAIN-KL method is inspired by normalizing flows (see, for example, Kobyzev et al., 2021). Before describing the approach, we provide a brief overview of related notation and some background on the concept behind normalizing flows. After this, we describe how the EDAIN-KL layer itself can be treated as an invertible bijector to fit into the normalizing flow framework. In doing so, we derive analytic and differentiable expressions for certain terms related to the EDAIN-KL layer.

B.2.1 Brief background on normalizing flow

The idea behind normalizing flows is taking a simple random variable, such as a standard Gaussian, and transforming it into a more complicated distribution, for example, one that resembles the distribution of a given real-world dataset. Consider a random variable $\mathbf{Z} \in \mathbb{R}^d$ with a known and analytic expression for its PDF $p_{\mathbf{Z}} : \mathbb{R}^d \rightarrow \mathbb{R}$. We refer to \mathbf{Z} as the *base distribution*. We then define a parametrised invertible function $\mathbf{g}_{\theta} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, also known as a *bijector*, and use this to transform the base distribution into a new probability distribution: $\mathbf{Y} = \mathbf{g}_{\theta}(\mathbf{Z})$. By increasing the complexity of the bijector \mathbf{g}_{θ} (for example, by using a deep neural network), the transformed distribution \mathbf{Y} can grow arbitrarily complex as well. The PDF of the transformed distribution can then be computed using the change of variable formula (Kobyzev et al., 2021), where

$$p_{\mathbf{Y}}(\mathbf{y}) = p_{\mathbf{Z}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y})) \cdot |\det \mathbf{J}_{\mathbf{Y} \rightarrow \mathbf{Z}}(\mathbf{y})| = p_{\mathbf{Z}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y})) \cdot |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y}))|^{-1}, \quad (13)$$

and where $\mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}$ is the Jacobian matrix for the *forward mapping* $\mathbf{g}_{\theta} : \mathbf{z} \mapsto \mathbf{y}$. Taking logs on both sides of Equation (13), it follows that

$$\log p_{\mathbf{Y}}(\mathbf{y}) = \log p_{\mathbf{Z}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y})) - \log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y}))|. \quad (14)$$

One common application of normalizing flows is density estimation (Kobyzev et al., 2021). Given a dataset $\mathcal{D} = \{\mathbf{y}^{(i)}\}_{i=1}^N$ with samples from some unknown, complicated distribution, we want to estimate its PDF. This can be done with likelihood-based estimation, where we assume the data points $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(N)}$ come from, say, the parametrised distribution $\mathbf{Y} = \mathbf{g}_{\theta}(\mathbf{Z})$ and we optimise θ to maximise the data log-likelihood,

$$\log p(\mathcal{D}|\theta) = \sum_{i=1}^N \log p_{\mathbf{Y}}(\mathbf{y}^{(i)}|\theta) \stackrel{(14)}{=} \sum_{i=1}^N \log p_{\mathbf{Z}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y}^{(i)})) - \log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\mathbf{g}_{\theta}^{-1}(\mathbf{y}^{(i)}))|. \quad (15)$$

It can be shown that this is equivalent to minimising the KL-divergence between the empirical distribution \mathcal{D} and the transformed distribution $\mathbf{Y} = \mathbf{g}_{\theta}(\mathbf{Z})$ (see, for example, Kobyzev et al., 2021):

$$\operatorname{argmax}_{\theta} \log p(\mathcal{D}|\theta) = \operatorname{argmax}_{\theta} \sum_{i=1}^N \log p_{\mathbf{Y}}(\mathbf{y}^{(i)}|\theta) = \operatorname{argmin}_{\theta} \operatorname{KL} \{ \mathcal{D} \parallel \mathbf{g}_{\theta}(\mathbf{Z}) \}.$$

When training a normalizing flow model, we want to find the parameter values θ that minimize the above KL-divergence. This is done using back-propagation, where the criterion \mathcal{L} is set to be the negation of Equation (15). That is, the loss becomes the negative log likelihood of a batch of samples from the training dataset. To perform optimisation with this criterion, we need to compute all the terms in Equation (15), and this expression needs to be differentiable because the back-propagation algorithm uses the gradient of the loss with respect to the input data. We therefore need to find

- (i) an analytic and differentiable expression for the inverse transformation $\mathbf{g}_{\theta}^{-1}(\cdot)$,
- (ii) an analytic and differentiable expression for the PDF of the base distribution $p_{\mathbf{Z}}(\cdot)$, and
- (iii) an analytic and differentiable expression for the log determinant of the Jacobian matrix for \mathbf{g}_{θ} , that is, $\log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}|$.

We will derive these three components for our EDAIN-KL layer in the next section, where we extensively use the following proposition (Kobyzev et al., 2021, see, for example,):

Proposition B.1. Let $\mathbf{g}_1, \dots, \mathbf{g}_n : \mathbb{R}^d \rightarrow \mathbb{R}^d$ all be bijective functions, and consider the composition of these functions, $\mathbf{g} = \mathbf{g}_n \circ \mathbf{g}_{n-1} \cdots \circ \mathbf{g}_1$. Then, \mathbf{g} is a bijective function with inverse

$$\mathbf{g}^{-1} = \mathbf{g}_1^{-1} \circ \cdots \circ \mathbf{g}_{n-1}^{-1} \circ \mathbf{g}_n^{-1},$$

and the log of the absolute value of the determinant of the Jacobian is given by

$$\log |\det \mathbf{J}_{\mathbf{g}^{-1}}(\cdot)| = \sum_{i=1}^N \log |\det \mathbf{J}_{\mathbf{g}_i^{-1}}(\cdot)|.$$

B.2.2 Application to EDAIN-KL

Like with the EDAIN layer, we want to compose the outlier mitigation, shift, scale and power transform transformations into one operation, which we do by defining

$$\mathbf{g}_{\boldsymbol{\theta}} = \mathbf{h}_1^{-1} \circ \mathbf{h}_2^{-1} \circ \mathbf{h}_3^{-1} \circ \mathbf{h}_4^{-1}, \quad (16)$$

where $\boldsymbol{\theta} = (\boldsymbol{\beta}, \mathbf{m}, \mathbf{s}, \boldsymbol{\lambda})$ are the unknown parameters and $\mathbf{h}_1, \dots, \mathbf{h}_4$ are defined in Equations (9), (10), (11) and (12), respectively. Notice how we apply all the operations in reverse order, compared to the EDAIN layer. This is because we will use $\mathbf{g}_{\boldsymbol{\theta}}$ to transform our base distribution \mathbf{Z} into a distribution that resembles the training dataset, \mathcal{D} , not the other way around. Then, to normalize the dataset after fitting the EDAIN-KL layer, we apply

$$\mathbf{g}_{\boldsymbol{\theta}}^{-1} = \mathbf{h}_4 \circ \mathbf{h}_3 \circ \mathbf{h}_2 \circ \mathbf{h}_1 \quad (17)$$

to each time series data point, similar to the EDAIN layer. It can be shown that all the transformations defined in Equations (9), (10), (11) and (12) are invertible, of which a proof is given in the next subsection. Using Lemma B.1, it thus follows that $\mathbf{g}_{\boldsymbol{\theta}}$, as defined in Equation (16), is bijective and that its inverse is given by Equation (17). Since we already have analytic and differentiable expressions for \mathbf{h}_1 , \mathbf{h}_2 , \mathbf{h}_3 and \mathbf{h}_4 from Equations (9), (10), (11) and (12), it follows that $\mathbf{g}_{\boldsymbol{\theta}}^{-1}$, as defined in Equation (17), is an analytic and differentiable expression, so part (i) is satisfied.

We now move onto deciding what our base distribution should be. As we validated experimentally in Section 4.1 of the main paper, *Gaussianizing* the input data could increase the performance of deep neural networks (depending on the data generating process). Therefore, we let the base distribution be the standard multivariate Gaussian distribution

$$\mathbf{Z} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d),$$

whose PDF $p_{\mathbf{Z}}(\cdot)$ has an analytic and differentiable expression, so part (ii) is satisfied.

In the next subsection, we will derive part (iii): an analytic and differentiable expression for the log of the determinant of the Jacobian matrix of $\mathbf{g}_{\boldsymbol{\theta}}$, $\log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}|$. Once that is done, parts (i), (ii) and (iii) are satisfied, so $\boldsymbol{\theta}$ can be optimised using back-propagation using the negation of Equation (15) as the objective. In other words, we can optimise $\boldsymbol{\theta}$ to maximise the likelihood of the training data under the assumption that it comes from the distribution $\mathbf{Y} = \mathbf{g}_{\boldsymbol{\theta}}(\mathbf{Z})$. This is desirable, as if we can achieve a high data likelihood, the samples $\mathcal{D} = \{\mathbf{y}^{(i)}\}_{i=1,2,\dots,N}$ will more closely resemble a standard normal distribution after being transformed by $\mathbf{g}_{\boldsymbol{\theta}}^{-1}$. Also recall that multivariate time series data are considered in this work, so the “ \mathbf{y} ”-samples will be of the form $\mathbf{X}^{(i)} \in \mathbb{R}^{d \times T}$.

B.2.3 Derivation of the inverse log determinant of the Jacobian

Recall that the EDAIN-KL architecture is a bijector composed of four other bijective functions. Using the result in Equation (B.1), we get

$$\log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\cdot)| = \sum_{i=1}^4 \log |\det \mathbf{J}_{\mathbf{h}_i^{-1}}(\cdot)|.$$

Considering the transformations in Equations (9), (10), (11) and (12), we notice that all the transformations happen element-wise, so for $i \in \{1, 2, 3, 4\}$, we have $\left[\frac{\partial \mathbf{h}_i^{-1}(\mathbf{x})}{\partial x_k}\right]_j = 0$ for $k \neq j$. Therefore, the Jacobians are

diagonal matrices, implying that the determinant is the product of the diagonal entries, giving

$$\log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\mathbf{x})| = \sum_{i=1}^4 \log \left| \prod_{j=1}^d \left[\frac{\partial \mathbf{h}_i^{-1}(\mathbf{x})}{\partial x_j} \right]_j \right| = \sum_{i=1}^4 \sum_{j=1}^d \log \left| \frac{\partial h_i^{-1}(x_j; \theta_j^{(i)})}{\partial x_j} \right|, \quad (18)$$

where in the last step we used the fact that h_1, h_2, h_3 and h_4 are applied element-wise to introduce the notation $h_i(x_j; \theta_j^{(i)})$. This means, applying \mathbf{h}_i to some vector where the j th element is x_j , and the corresponding j th transformation parameter takes the value $\theta_j^{(i)}$. For example, for the scale function $\mathbf{h}_3(\mathbf{x}) = \mathbf{x} \oslash \mathbf{s}$, we have $h_3(x_j; \lambda_j) = \frac{x_j}{s_j}$. From Equation (18), we know that we only need to derive the derivatives for the element-wise inverses, which we will now do for each of the four transformations, also demonstrating that each transformation is bijective.

Shift. We first consider $h_2(x_j; m_j) = x_j - m_j$. Its inverse is $h_2^{-1}(z_j; m_j) = z_j + m_j$, and it follows that

$$\log \left| \frac{\partial h_2^{-1}(z_j; m_j)}{\partial z_j} \right| = \log 1 = 0.$$

Scale. We now consider $h_3(x_j; s_j) = \frac{x_j}{s_j}$, whose inverse is $h_3^{-1}(x_j; s_j) = z_j s_j$. It follows that

$$\log \left| \frac{\partial h_3^{-1}(z_j; s_j)}{\partial z_j} \right| = \log |s_j|.$$

Outlier mitigation. We now consider $h_1(x_j; \beta_j) = \beta_j \tanh \left\{ \frac{(x_j - \hat{\mu}_j)}{\beta_j} \right\} + \hat{\mu}_j$. Its inverse is

$$h_1^{-1}(z_j; \beta_j) = \beta_j \tanh^{-1} \left\{ \frac{z_j - \hat{\mu}_j}{\beta_j} \right\} + \hat{\mu}_j.$$

It follows that

$$\log \left| \frac{\partial h_1^{-1}(z_j; \beta_j)}{\partial z_j} \right| = \log \left| \frac{1}{1 - \left(\frac{z_j - \hat{\mu}_j}{\beta_j} \right)^2} \right| = -\log \left| 1 - \left(\frac{z_j - \hat{\mu}_j}{\beta_j} \right)^2 \right|.$$

Power transform. By considering Equation (12), it can be shown that for fixed λ , negative inputs are always mapped to negative values and vice versa, which makes the Yeo-Johnson transformation invertible. Additionally, in $\mathbf{h}_4(\cdot)$ the Yeo-Johnson transformation is applied element-wise, so we get

$$\mathbf{h}_4^{-1}(\mathbf{z}) = \left[\left[f_{\text{YJ}}^{\lambda_1} \right]^{-1}(z_1), \left[f_{\text{YJ}}^{\lambda_2} \right]^{-1}(z_2), \dots, \left[f_{\text{YJ}}^{\lambda_d} \right]^{-1}(z_d) \right],$$

where it can be shown that the inverse Yeo-Johnson transformation for a single element is given by

$$\left[f_{\text{YJ}}^{\lambda} \right]^{-1}(z) = \begin{cases} (z\lambda + 1)^{1/\lambda} - 1, & \text{if } \lambda \neq 0, z \geq 0; \\ e^z - 1, & \text{if } \lambda = 0, z \geq 0; \\ 1 - \{1 - z(2 - \lambda)\}^{1/(2-\lambda)}, & \text{if } \lambda \neq 2, z < 0; \\ 1 - e^{-z}, & \text{if } \lambda = 2, z < 0. \end{cases}$$

The derivative with respect to z then becomes

$$\frac{\partial \left[f_{\text{YJ}}^{\lambda} \right]^{-1}(z)}{\partial z} = \begin{cases} (z\lambda + 1)^{(1-\lambda)/\lambda}, & \text{if } \lambda \neq 0, z \geq 0; \\ e^z, & \text{if } \lambda = 0, z \geq 0; \\ \{1 - z(2 - \lambda)\}^{(\lambda-1)/(2-\lambda)}, & \text{if } \lambda \neq 2, z < 0; \\ e^{-z}, & \text{if } \lambda = 2, z < 0. \end{cases}$$

It follows that

$$\log \left| \frac{\partial [f_{YJ}^\lambda]^{-1}(z)}{\partial z} \right| = \begin{cases} \frac{1-\lambda}{\lambda} \log(z\lambda + 1), & \text{if } \lambda \neq 0, z \geq 0; \\ z, & \text{if } \lambda = 0, z \geq 0; \\ \frac{\lambda-1}{2-\lambda} \log \{1 - z(2-\lambda)\}, & \text{if } \lambda \neq 2, z < 0; \\ -z, & \text{if } \lambda = 2, z < 0, \end{cases}$$

which we use as the expression for $\log \left| \frac{\partial h_4^{-1}(z_j; \lambda_j)}{\partial z_j} \right|$ for $z = z_1, \dots, z_d$.

Combining these expressions, we get an analytical and differentiable form for $\log |\det \mathbf{J}_{\mathbf{Z} \rightarrow \mathbf{Y}}(\mathbf{x})|$, as required.

C EXPERIMENTAL EVALUATION DETAILS ON MODELS AND TRAINING

In this section, we provide details on the specific RNN model architectures used for evaluation in Section 4 of the main paper. We then cover the optimization procedures used for the three datasets, including details such as number of training epochs, learning rates and choice of optimizers. Then, we list the learning rate modifiers used for the different adaptive preprocessing layers, and explain how these were selected.

C.1 Deep Neural Network Model Architectures

Synthetic dataset. The GRU RNN architecture consisted of two GRU cells with a dimensionality of 32 and dropout layer with dropout probability $p = \frac{1}{5}$ between these cells. This was followed by a linear feed-forward neural network with 3 fully-connected layers, separated by ReLU activation functions, of 64, 32 and 1 units, respectively. The output was then passed through a sigmoid layer to produce a probability $p \in (0, 1)$.

Default prediction dataset. We use a RNN sequence model, with a classifier head, for all our experiments with the default prediction dataset. It consists of two stacked GRU RNN cells, both with a hidden dimensionality of 128. Between these cells, there is a dropout layer with the dropout probability of $p = 20\%$. For the 11 categorical features present in the dataset, we pass these through separate embedding layers, each with a dimensionality of 4. The outputs of the embedding layers are then combined with the output of the last GRU cell, after it has processed the numeric columns, and the result is passed to the linear classifier head. The classifier head is a conventional linear neural network consisting of 2 linear layers with 128 and 64 units each, respectively, and separated by ReLU activation functions. The output is then fed through a linear layer with a single output neuron, followed by a sigmoid activation function to constrain the output to be a probability in the range $(0, 1)$. The described architecture was chosen because it worked well in our initial experiments.

Financial forecasting dataset. For the FI-2010 LOB dataset, we use a similar GRU RNN model as with the Amex dataset, but change the architecture slightly to match the RNN model used by Passalis et al. (2020). This was done to make the comparison between the proposed EDAIN method and their DAIN method more fair, seeing as they also used the LOB dataset to evaluate DAIN. Instead of using two stacked GRU cells, we use one with 256 units. We also do not need any embedding layers because all the predictor variables are numeric. The classifier head that follows the GRU cells consists of one linear layer with 512 units, followed by a ReLU layer and a dropout layer with a dropout probability of $p = 0.5$. The output layer is a linear layer with 3 units, as we are classifying the multivariate time series into one of three classes, $\mathcal{C} = \{\text{decrease, stationary, increase}\}$. These outputs are then passed to a *softmax* activation function such that the output is a probability distribution over the three classes and sums to 1: $[\text{softmax}(\mathbf{y})]_j = e^{y_j} / \sum_{i=1}^k e^{y_i}$, $j = 1, 2, \dots, k$. In our case, we have $k = 3$.

C.2 Optimization Procedures

Synthetic dataset. For the synthetic dataset, all the $N_{\mathcal{D}}$ generated datasets were separated into 80%-20% train-validation splits, and the metrics reported in Table 1 of the main paper are based on metrics computed for model performance on the validation splits. The training was done using the Adam optimizer proposed by Kingma and Ba (2015) using a base learning rate of $\eta = 10^{-3}$ and the model was trained for 30 epochs. We also used a multi-step learning rate scheduler with decay $\gamma = \frac{1}{10}$ at the 4th and 7th epoch. Additionally, an early stopper was used on the validation loss with a patience of 5. The model was trained using binary cross-entropy loss, and the batch size was 128.

Default prediction dataset. To evaluate the different preprocessing methods on this dataset, we perform 5-fold cross validation, which gives evaluation metrics for five different 20% validation splits. For training, we used a batch size of 1024. This was chosen to give a good trade-off between the time required to train the model and predictive performance. This model was also optimised with the Adam optimizer proposed by Kingma and Ba (2015). The learning rate was also set to $\eta = 10^{-3}$ after testing what learning rate from the set $H = \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$ gave the most stable convergence. The momentum parameters β_1 and β_2 were set to their default values according to the PyTorch implementation of the optimizer (Paszke et al., 2019). We also used a multi-step learning rate scheduler, with the step size set to $\gamma = \frac{1}{10}$. For the milestones (corresponding to the epoch indices at which the learning rate changes), we first tuned the first step based on observing the learning rate curve during training; then, we set the step epoch to ensure the performance did not change too rapidly. This was done until we got milestones at 4 and 7. We also used an early stopper based on the validation loss, with the patience set to $p_{\text{patience}} = 5$ as this worked well in our initial experiments. The maximum number of training epochs was set to 40.

Financial forecasting dataset. As discussed in the main paper, we use an anchored cross-validation scheme for training, which gives 9 different train-validation splits based on the 10 days of training data available. The targets are ternary labels $y_1, y_2, \dots, y_N \in \{0, 1, 2\}$, denoting whether the mid-price decreased, remained stationary, or increased, and the output of the model is probability vectors $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N \in (0, 1)^3$. For optimising the model parameters, we use the cross-entropy loss function, defined as

$$\mathcal{L}(\mathbf{p}_i, y_i) = - \sum_{c=0}^2 \mathbb{I}_{\{y_i=c\}} \log(p_{i,c}),$$

where $p_{i,c}$ denotes the predicted probability of class c for the i th input sample. We used a batch size of 128, and used the RMSProp optimizer proposed by Hinton and Tieleman (2012) to be consistent with the experiments by Passalis et al. (2020). The base learning rate was set to $\eta = 10^{-4}$. No learning rate scheduler nor early stoppers were used: this was done to best reproduce the methodology used by Passalis et al. (2020). Despite not using any early stoppers, all the metrics were computed based on the model state at the epoch where the validation loss was lowest. This was because the generalization performance started to decline in the middle of training in most cases. At each training fold, the model was trained for 20 epochs.

C.3 Sublayer Learning Rate Modifiers

C.3.1 Synthetic dataset

For the synthetic data, we use the learning rate modifier $\eta' = 10^{-1}$ for all sublayers of all adaptive preprocessing layers (BIN, DAIN, and EDAIN), while as mentioned in Section C.2, the base learning rate is $\eta = 10^{-3}$.

C.3.2 Default prediction dataset

The BIN method proposed by Tran et al. (2021) and the DAIN method proposed by Passalis et al. (2020) have both never been applied to the American Express default prediction dataset before, so their hyperparameters are tuned for this dataset. We also tune the two versions of EDAIN. For all these experiments, we run a grid-search with different combinations of the hyperparameters, specifically the learning rate modifiers, using the validation loss from training a single model for 10 epochs on a 80%-20% split of the training dataset. We then pick the combination giving the lowest validation loss. We now provide details on the grids used for each of the different preprocessing layers, and what learning rate modifiers gave the best performance.

BIN. We used the grids:

$$H_\beta = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}\},$$

$$H_\gamma = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}\}, \text{ and}$$

$$H_\lambda = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}\}.$$

The combination giving the lowest average cross-validation loss was found to be $\eta_\beta = 10$, $\eta_\gamma = 1$, and $\eta_\lambda = 10^{-6}$, giving 0.2234.

DAIN. We used the grids:

$$H_{\text{shift}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\} \text{ and}$$

$$H_{\text{scale}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}.$$

The combination giving the lowest average cross-validation loss was found to be $\eta_{\text{shift}} = 1$ and $\eta_{\text{scale}} = 1$, giving 0.2216.

Global-aware EDAIN. We used the grids:

$$H_{\text{scale}} = H_{\text{shift}} = H_{\text{outlier}} = H_{\text{pow}} = \{100, 10, 1, 10^{-1}, 10^{-2}, 10^{-3}\}.$$

The combination giving the lowest average cross-validation loss was found to be $\eta_{\text{shift}} = 10^{-2}$, $\eta_{\text{scale}} = 10^{-2}$, $\eta_{\text{outlier}} = 10^2$, and $\eta_{\text{pow}} = 10$, giving 0.2190.

Local-aware EDAIN. We used the grids:

$$H_{\text{scale}} = H_{\text{shift}} = \{1, 10^{-1}, 10^{-2}\} \text{ and } H_{\text{outlier}} = H_{\text{pow}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}\}.$$

The combination giving the lowest average cross-validation loss was found to be $\eta_{\text{shift}} = 1$, $\eta_{\text{scale}} = 1$, $\eta_{\text{outlier}} = 10$, and $\eta_{\text{pow}} = 10$, giving 0.2243.

EDAIN-KL. Note that due to numerical gradient errors in the power transform layers occurring for some choices of power transform learning rates, the values considered are all low to avoid these errors. We used the grids

$$H_{\text{outlier}} = \{100, 10, 1, 10^{-1}, 10^{-2}, 10^{-3}\},$$

$$H_{\text{scale}} = \{100, 10, 1, 10^{-1}, 10^{-2}, 10^{-3}\},$$

$$H_{\text{shift}} = \{100, 10, 1, 10^{-1}, 10^{-2}, 10^{-3}\}, \text{ and}$$

$$H_{\text{pow}} = \{10^{-7}\}$$

The combination giving the lowest average cross-validation loss was found to be $\eta_{\text{shift}} = 10$, $\eta_{\text{scale}} = 10$, $\eta_{\text{outlier}} = 10^2$, and $\eta_{\text{pow}} = 10^{-7}$, giving 0.2208.

C.3.3 Financial forecasting dataset

The BIN method proposed by Tran et al. (2021) and the DAIN method proposed by Passalis et al. (2020) have already been applied to the LOB dataset before, but only DAIN has been applied with the specific GRU RNN architecture we are using. Therefore, the learning rate modifiers found by Passalis et al. (2020) will be used as-is ($\eta_{\text{shift}} = 10^{-2}$, $\eta_{\text{scale}} = 10^{-8}$), and the learning rate modifiers for the remaining methods will be tuned. The details of this are presented here. For all the learning rate tuning experiments, we used the first day of data from the FI-2010 LOB for training and the data from the second day for validation. We then pick the combination giving the highest validation Cohen’s κ -metric.

BIN. We used the grids:

$$H_{\beta} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}, 10^{-8}\},$$

$$H_{\gamma} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}, 10^{-8}\}, \text{ and}$$

$$H_{\lambda} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-6}, 10^{-8}\}.$$

The combination giving the highest κ was found to be $\eta_{\beta} = 1$, $\eta_{\gamma} = 10^{-8}$, and $\eta_{\lambda} = 10^{-1}$, giving $\kappa = 0.3287$.

Global-aware EDAIN. We used the grids:

$$H_{\text{scale}} = H_{\text{shift}} = H_{\text{outlier}} = H_{\text{pow}} = \{10, 10^{-1}, 10^{-3}, 10^{-6}\}.$$

The combination giving the highest κ was found to be $\eta_{\text{shift}} = 10$, $\eta_{\text{scale}} = 10$, $\eta_{\text{outlier}} = 10^{-6}$, and $\eta_{\text{pow}} = 10^{-3}$, giving $\kappa = 0.2788$.

Local-aware EDAIN. We used the grids:

$$H_{\text{scale}} = \{10^{-1}, 10^{-4}, 10^{-8}\},$$

$$H_{\text{shift}} = \{10^{-1}, 10^{-2}\},$$

$$H_{\text{outlier}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}\}, \text{ and}$$

$$H_{\text{pow}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}\}.$$

The combination giving the highest κ was found to be $\eta_{\text{shift}} = 10^{-2}$, $\eta_{\text{scale}} = 10^{-4}$, $\eta_{\text{outlier}} = 10$, and $\eta_{\text{pow}} = 10$, giving $\kappa = 0.3859$.

EDAIN-KL. $H_{\text{scale}} = \{10^{-1}, 10^{-4}, 10^{-7}\},$

$$H_{\text{shift}} = \{10^{-1}, 10^{-2}\},$$

$$H_{\text{outlier}} = \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}\}, \text{ and}$$

$$H_{\text{pow}} = \{10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}\}.$$

The combination giving the highest κ was found to be $\eta_{\text{shift}} = 10^{-2}$, $\eta_{\text{scale}} = 10^{-4}$, $\eta_{\text{outlier}} = 10$, and $\eta_{\text{pow}} = 10^{-3}$, giving $\kappa = 0.2757$.

C.4 Other Hyperparameters

For selecting the α parameter in the KDIT method, we tried all $\alpha \in \{0.1, 1, 10, 100\}$ and selected the value that gave the lowest average validation loss.

C.5 Description of Computing Infrastructure

All experiments were run on a *Asus ESC8000 G4* server with the following specifications:

- two 16-core CPUs of model *Intel(R) Xeon(R) Gold 6242 CPU @ 2.80GHz*,
- 896 GiB of system memory,
- eight GPUs of model *NVIDIA GeForce RTX 3090*, each with 24576 MiB of video memory (VRAM).

D SYNTHETIC DATA GENERATION ALGORITHM

To help with designing preprocessing methods that can increase the predictive performance as much as possible, we propose a synthetic data generation algorithm that gives full control over how the variables are distributed, through only needing to specify an unnormalized PDF function for each variable. This allows synthesizing data with distributions having the same irregularities as that seen in real-world data. Additionally, the covariance structure of the generated data can be configured to resemble that of real-world multivariate time series, which most often have significant correlation between the predictor variables. Along with each time series generated, we also create a response $y \in \{0, 1\}$ that is based on the covariates, allowing for supervised classification tasks.

The main input to the data generation procedure is the time series length, $T \in \mathbb{N}$, and the number of features, $d \in \mathbb{N}$. For each predictor variable $j = 1, 2, \dots, d$, we also specify an unnormalized PDF, $f_j : \mathbb{R} \rightarrow \mathbb{R}_+$. The data generation procedure then generates a multivariate time series covariate $\mathbf{X} \in \mathbb{R}^{d \times T}$ and a corresponding response $y \in \{0, 1\}$ in three steps. Note that this procedure is repeated N times to, say, generate a dataset of N time series. An overview of the the three steps of the data generation algorithm is shown in Figure 5. Each row in the three matrices corresponds to one predictor variable and the column specifies the timestep.

In the first step in Equation (19), we generate Gaussian random variables that have a similar covariance structure to a multivariate time series. This ensures the covariates' covariance more closely resemble that of real-world sequence data. In the second step, shown in Equation (20), we convert the Gaussian random variables into uniform random variables using the inverse normal CDF, after standardizing each variable. In this step, we also form the response through a linear combination of unknown parameters β and the uniform random variables.

$$\begin{array}{l} \text{Hidden correlated} \\ \text{Gaussian RVs} \end{array} \left\{ \begin{array}{l} N_{1,1} \quad N_{1,2} \quad \cdots \quad N_{1,T} \\ N_{2,1} \quad N_{2,2} \quad \cdots \quad N_{2,T} \\ \vdots \quad \vdots \quad \ddots \quad \vdots \\ N_{d,1} \quad N_{d,2} \quad \cdots \quad N_{d,T} \end{array} \right. \sim \mathcal{N}(\mathbf{0}, \Sigma'), \text{ where } \Sigma' \in \mathbb{R}^{(dT) \times (dT)} \quad (19)$$

$$\begin{array}{l} \text{Hidden correlated} \\ \text{uniform RVs} \end{array} \left\{ \begin{array}{l} U_{1,1} \quad U_{1,2} \quad \cdots \quad U_{1,T} \\ U_{2,1} \quad U_{2,2} \quad \cdots \quad U_{2,T} \\ \vdots \quad \vdots \quad \ddots \quad \vdots \\ U_{d,1} \quad U_{d,2} \quad \cdots \quad U_{d,T} \end{array} \right. \xrightarrow{\text{Form response}} Y = \mathbb{I} \left(\sum_{j=1}^d \sum_{t=1}^T \beta_{j,t} U_{j,t} + \zeta > \frac{1}{2} \right) \quad (20)$$

$$\begin{array}{l} \text{Output multivariate} \\ \text{time series} \end{array} \left\{ \begin{array}{l} X_{1,1} \quad X_{1,2} \quad \cdots \quad X_{1,T} \\ X_{2,1} \quad X_{2,2} \quad \cdots \quad X_{2,T} \\ \vdots \quad \vdots \quad \ddots \quad \vdots \\ X_{d,1} \quad X_{d,2} \quad \cdots \quad X_{d,T} \end{array} \right. \quad (21)$$

Figure 5: High-level overview of the proposed synthetic data generation algorithm.

This ensures there is some mutual information between the response and the covariates that are generated in the final step. In step 3, shown in Equation (21), we form the final covariates using the provided PDFs f_1, f_2, \dots, f_d . This is done by estimating each PDF's inverse CDF using numerical methods, and transforming the uniform random variables with these. This makes the samples come from a distribution matching that of the provided PDFs. Moreover, in Equations (19), (20), and (21) we use random variable notation for each of the steps, but in practice, all the transformations are applied on samples from these.

D.1 Step 1: Generating random variables with a time series covariance structure

One approach to reproducing the covariance structure of a multivariate time series is to assume that each of the d individual time series follow a *moving average* model, which is a common type of theoretical time series (Hyndman and Athanasopoulos, 2018). With this model, the covariate at timestep t takes the form

$$X_t = c - \sum_{j=0}^q \theta_j \epsilon_{t-j}, \quad (22)$$

where $c \in \mathbb{R}$ is a constant, $\epsilon_0, \epsilon_1, \dots$ are uncorrelated random variables with zero mean and finite variance $\sigma_\epsilon \in \mathbb{R}$. Also, $\theta_0 = -1$ and $\theta_1, \dots, \theta_q \in (-1, 1)$ are the unknown parameters. Under this model, Hyndman and Athanasopoulos (2018) state that the covariance between a sample from timestep t and a sample from $\tau \in \mathbb{Z}$

timesteps into the future is

$$s_\tau = \text{cov}\{X_t, X_{t+\tau}\} = \sigma_\epsilon^2 \sum_{j=0}^{q-\tau} \theta_j \theta_{j+\tau}. \quad (23)$$

We will not be generating our covariates using the model in Equation (22) as this would make it infeasible to get samples that are distributed according to arbitrary PDFs. However, we can use the covariance formula from Equation (23) to set the covariance between each pair of variables generated. To do this, we first specify the parameters q , σ_ϵ , and $\theta_0, \dots, \theta_q$ for each of the d predictor variables. Then, we stack the Gaussian random variables $N_{1,1}, N_{1,2}, \dots, N_{2,1}, N_{2,2}, \dots, N_{d,T}$ in Equation (19) row-wise so that they form a dT -long vector. Let $\Sigma \in \mathbb{R}^{dT \times dT}$ denote the covariance matrix of this dT -long Gaussian multivariate random variable. While still thinking of each T -length row as its own univariate time series, we fill out the entries in Σ based on Equation (23), using the parameters specified for each of the d time series. The remaining entries of Σ are randomly initialised with samples from $\mathcal{N}(\mu = 0, \sigma = \sigma_{\text{cor}})$, where σ_{cor} is a hyperparameter for the data synthesis, with the motivation being to create some cross-dependence between each time series. In order to use Σ as a valid covariance matrix for sampling from the dT -dimensional multivariate normal distribution, it needs to be *symmetric positive semi-definite*. The Σ matrix we have constructed so far has no guarantee of satisfying this. Therefore, we use the algorithm proposed by Higham (1988) to find the symmetric positive semi-definite matrix $\Sigma' \in \mathbb{R}^{dT \times dT}$ that is closest to Σ according to the Frobenius norm. More details on this procedure can be found in Higham (1988). After this, we generate a dT -dimensional sample $\mathbf{N} \sim \mathcal{N}(\mathbf{0}, \Sigma')$ and imagine “unrolling” this into a $d \times T$ matrix where we have a T -timestep-long time series in each row, just as in Equation (19).

D.2 Step 2: Forming the response

Before forming the response y , we need to convert the Gaussian random variables generated in step 1 into uniform random variables. By the probability integral transform, if a normal random variable is passed through its inverse CDF-function, a uniform random variable is obtained. Therefore, we do this for each of the normal random variables, as shown in the transition between Equation (19) and Equation (20), giving d time series of uniform random variables, each of length T .

To form the response, we randomly sample a noise term $\zeta \sim \mathcal{N}(0, \sigma_\zeta^2)$ and set

$$Y = \mathbb{I} \left(\sum_{j=1}^d \sum_{t=1}^T \beta_{j,t} U_{j,t} + \zeta > \frac{1}{2} \right).$$

The idea behind this is to make sure each variable contributes to the response, but the contribution of each variable might differ and some might be completely irrelevant, just like in real-world data. Note that the noise term ζ is regenerated for each multivariate time series $\mathbf{X} \in \mathbb{R}^{d \times T}$ we generate, while the parameters $\beta \in \mathbb{R}^{d \times T}$ are held fixed for each time series generated as part of a synthetic dataset. When synthesizing a new dataset of say N samples, we first generate *one set* of $\beta \in \mathbb{R}^{d \times T}$ unknown parameters with $\beta_{1,1}, \dots, \beta_{d,T} \stackrel{\text{iid}}{\sim} \mathcal{N}\left(\frac{1}{dT}, \sigma_\beta^2\right)$, where σ_β is another hyperparameter for the dataset synthesis. The mean $\frac{1}{dT}$ was set to ensure the dataset generated is balanced, that is, the ratio between true and false labels is equal.

D.3 Step 3: Transforming the variables based on provided PDFs

The third step of the data generation procedure involves taking samples from uniformly distributed random variables $U_{j,t} \sim \text{U}[0, 1]$ and transforming them using the inverse CDF of the corresponding specified PDF, f_j , which results in a sample from the distribution specified by f_j . To do this, we need to estimate the inverse CDF function $\widehat{F}_j^{-1}(\cdot)$ using only f_j . This is done by evaluating f_j on a fine-grid of values $\mathcal{X} = \{A_j, A_j + \delta, A_j + 2\delta, \dots, B_j\}$, where A_j , B_j , and δ are additional hyperparameters specified in conjunction with the PDF f_j . Then we use *trapezoidal numerical integration* (see, for example, Atkinson, 1989) to estimate $\widehat{F}_j(x) = \int_{-\infty}^x f_j(x') dx'$ for all $x \in \mathcal{X}$. Since the provided PDFs are unnormalized, we normalize the CDF estimates by dividing them by $\widehat{F}_j(B_j)$. To get $\widehat{F}_j^{-1}(\cdot)$, we create an inverse *look-up table* that maps the $\widehat{F}_j(x)$ -values to the corresponding x -values. When implementing this procedure, the look-up table is *cached* to ensure the integration only needs to be done once for each $x \in \mathcal{X}$. Then, to evaluate $\widehat{F}_j^{-1}(u)$ for some $u \in (0, 1)$, we perform a *binary search* (see,

for example, Cormen et al., 2001) on the look-up table to find the smallest $x \in \mathcal{X}$ such that $\widehat{F}_j(x) \geq u$, which can efficiently be done since the values are already in sorted order because CDFs are monotonically increasing.

D.4 Hyperparameters for Synthetic Dataset Generation

Here we present the hyperparameters used when synthesizing the irregular data we used for the experiments in Section 4.1 of the main paper. We used the following bounds for the three PDFs: $(A_1, B_1) = (-8, 10)$, $(A_2, B_2) = (-30, 30)$, and $(A_3, B_3) = (-1, 7)$. The θ s were all configured with $q = 3$, and we used

$$\Theta = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix} = \begin{bmatrix} -1 & \frac{1}{3} & -\frac{1}{5} & \frac{4}{5} \\ -1 & \frac{3}{10} & \frac{9}{10} & 0 \\ -1 & \frac{4}{5} & \frac{3}{10} & -\frac{9}{10} \end{bmatrix}.$$

The standard deviations were set to $\sigma_{\text{cov}} = 1.4$, $\sigma_\eta = \frac{1}{2}$, and $\sigma_\beta = 2$.

E THE AMEX METRIC

Here we explain how to compute the *Amex metric* introduced in Section 4.2 of the main paper, which is calculated as the mean of the default rate captured at 4%, denoted D , and the normalized Gini coefficient, referred to as G . Assume we have made predictions p_1, p_2, \dots, p_N for each of the N customers, and assume these predicted default probabilities have been sorted in non-increasing order. Also assume they have associated normalized weights w_1, w_2, \dots, w_N such that $\sum_{i=1}^N w_i = 1$. To compute D , we take the predictions $p_1, p_2, \dots, p_\omega$ captured within the highest-ranked 4% of our predictions considering the weights w_1, w_2, \dots, w_N . Then, we look at the default rate within these predictions, normalized by the overall default rate. In other words,

$$D = \frac{\sum_{i=1}^{\omega} y_i}{\sum_{i=1}^N y_i}, \text{ where } \omega \text{ is the highest integer such that } \sum_{i=1}^{\omega} w_i \leq 0.04.$$

We now describe how to compute G , which requires computing the Gini coefficient in two ways for $k \in \{0, 1\}$. From Lerman and Yitzhaki (1989), we know the Gini coefficient can be computed as

$$G_k = 2 \sum_{j=1}^N w_{i_j^{(k)}} \left(\frac{p_{i_j^{(k)}} - \bar{p}}{\bar{p}} \right) \left(\widehat{F}_{i_j^{(k)}} - \bar{F} \right), \quad (24)$$

where $\widehat{F}_{i_j^{(k)}} = w_{i_j^{(k)}}/2 + \sum_{\ell=1}^{j+1} w_{i_\ell^{(k)}}$ and $\bar{p} = \sum_{j=1}^N w_{i_j^{(0)}} p_{i_j^{(0)}} = \sum_{j=1}^N w_{i_j^{(1)}} p_{i_j^{(1)}}$ and $\bar{F} = \sum_{j=1}^N w_{i_j^{(0)}} \widehat{F}_{i_j^{(0)}} = \sum_{j=1}^N w_{i_j^{(1)}} \widehat{F}_{i_j^{(1)}}$. To compute the normalized Gini coefficient, we first sort the predictions in non-decreasing order by the *true labels* y_1, y_2, \dots, y_N , and denote this ordering $i_1^{(0)}, i_2^{(0)}, \dots, i_N^{(0)}$. Let G_0 denote the result of computing equation (24) with this sorting. Then we sort the values by the *predicted probabilities* p_1, p_2, \dots, p_N in non-decreasing order, denoting this ordering as $i_1^{(1)}, i_2^{(1)}, \dots, i_N^{(1)}$. Let the value of equation (24) computed with this ordering be denoted G_1 . The normalized Gini coefficient is then $G = G_1/G_0$, which is what we use in the final metric

$$\text{Amex metric} = \frac{1}{2} (G + D) = \frac{1}{2} \left(\frac{G_1}{G_0} + D \right).$$