



Temporal density extrapolation using a dynamic basis approach

G. Kremp1  · D. Lang² · V. Hofer³

Received: 16 September 2018 / Accepted: 20 May 2019 / Published online: 19 June 2019
© The Author(s) 2019

Abstract

Density estimation is a versatile technique underlying many data mining tasks and techniques, ranging from exploration and presentation of static data, to probabilistic classification, or identifying changes or irregularities in streaming data. With the pervasiveness of embedded systems and digitisation, this latter type of streaming and evolving data becomes more important. Nevertheless, research in density estimation has so far focused on stationary data, leaving the task of extrapolating and predicting density at time points outside a training window an open problem. For this task, temporal density extrapolation (TDX) is proposed. This novel method models and predicts gradual monotonous changes in a distribution. It is based on the expansion of basis functions, whose weights are modelled as functions of compositional data over time by using an isometric log-ratio transformation. Extrapolated density estimates are then obtained by extrapolating the weights to the requested time point, and querying the density from the basis functions with back-transformed weights. Our approach aims for broad applicability by neither being restricted to a specific parametric distribution, nor relying on cluster structure in the data. It requires only two additional extrapolation-specific parameters, for which reasonable defaults exist. Experimental evaluation on various data streams, synthetic as well as from the real-world domains of credit scoring and environmental health, shows that the model manages to capture monotonous drift patterns accurately and better than existing methods. Thereby, it requires not more than 1.5 times the run time of a corresponding static density estimation approach.

Keywords Density extrapolation · Density forecasting · Data streams · Concept drift · Non-stationary data · Compositional data

Responsible editor: Po-ling Loh, Evimaria Terzi, Antti Ukkonen, Karsten Borgwardt.

All authors contributed equally to this publication.

✉ G. Kremp1
g.m.kremp1@uu.nl

Extended author information available on the last page of the article

1 Introduction

The extent and scope of available data continues to grow, often comprising data that is continuously generated over longer time spans, in environments that are subject to changes over time (Reinsel et al. 2017). Volume and velocity of such data streams often require automated processing and analysis, while their dynamic nature and associated volatility needs to be considered as well (Fan and Bifet 2013). For example, the distribution of the data might change over time, a problem that is commonly denoted as concept drift (Widmer and Kubat 1996) or population drift (Kelly et al. 1999). In prediction tasks, such drift requires adaptation mechanisms, for example by forgetting outdated irrelevant data or models (Gama et al. 2014). However, merely attempting to avoid the negative influence that this dynamic nature can have on statistical models is leaving its potential unused. In contrast, aiming to understand and to incorporate the changes in data into the statistical model itself might provide additional value, by helping to perform more accurate predictions for the future and allowing a structured description of the occurring changes.

Several tasks and approaches have been proposed to identify change or irregularities in data, such as outlier detection [surveyed for example in Sadik and Gruenwald (2014)], anomaly detection [surveyed by Chandola et al. (2009)], or change detection [surveyed by Tran et al. (2014)]. Furthermore, some research has focused on the exploration, understanding, and exploitation of change: change diagnosis aims to estimate the spatio-temporal change rates in kernel density within a time window (Aggarwal 2005). This might be seen as an early proponent of the change mining paradigm, which was proposed afterwards by Böttcher et al. (2008) and calls for data mining approaches that provide understanding of change itself. Following this paradigm, drift mining techniques aim to provide explicit models of distributional changes over time, for example to transfer knowledge between different time points in data streams under verification latency (Hofer and Kreml 2013).

An essential technique in data science (Chacón and Duong 2018; Scott 2015) that is underneath many of the algorithms in the tasks above is density estimation (Rosenblatt 1956; Whittle 1958), which is also important for exploration and presentation of data in general (Silverman 1986), as it is for probabilistic classifiers such as Parzen Window Classifiers (Parzen 1962). Given a set of instances sampled within a training time window from a non-stationary (i.e., drifting) data stream, the provided density estimates typically correspond to the distribution over all instances within this window. However, these estimates might also be required for specific time points, rather than time windows. This time point might lie within the training time window, requiring an in-sample interpolation. Alternatively, it might lie outside, requiring an out-of-sample extrapolation to past or future time points.

Research in density estimation has so far focused on providing in-sample estimates, which are often adapted to the distribution of the most recently observed samples in the training time window. In contrast, density extrapolation beyond this window has only recently received attention (Kreml 2015), although it has some potential beyond being simply used in lieu of static density estimations in the applications above: Most importantly, extrapolating continuing trends allows to anticipate future distributional changes and to take preparatory measures. Examples are classifiers

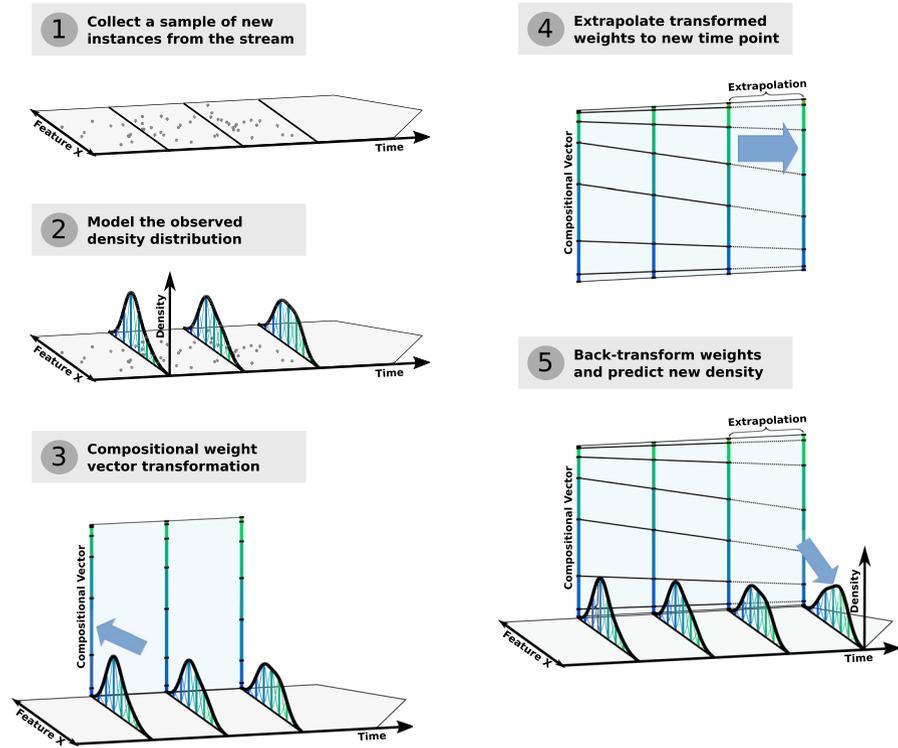


Fig. 1 Illustration of temporal density extrapolation, which aims to capture and predict the change in the distribution of instances over time. (1) A sample of instances X is observed at time points T in a data stream. (2) The observed density distribution is modelled as an expansion of density functions. (3) Their basis weights form a compositional vector. (4) The weight development over time is modelled, and extrapolated to new time points. (5) These extrapolated weights are back-transformed to predict the resulting new density distribution

or models that incorporate forthcoming distributional changes anticipatively; active learning and change detection approaches that proactively sample in regions where change is expected to happen; or in the identification of unexpected changes in drift.

We propose a novel approach to model and predict the density of a feature in a data stream. The underlying distribution is described as an expansion of Gaussian density functions, which are placed at fixed equidistant locations.¹ The basis weights of the Gaussian density functions are fitted functions of time. In this way the changes observed in the data over time are modelled as a change in the weighting of the basis expansion.

Figure 1 shows a feature X in a data stream, whose density at different points in time is described by an expansion of six Gaussian density functions. The weights associated to each of these basis functions form a composition (Aitchison 1982; Egozcue et al.

¹ This is related to Gaussian mixture models (GMM), with the subtle but important difference that in GMMs the location and variance parameters of the components are in general estimated from data and not fixed a priori. Note that the model is not restricted to Gaussian basis functions.

2003) that sums to one, their proportions are illustrated on the back end of the figure. The change in the density at the second and the third time point is modelled as a change in the distribution of the weights, as visible from the shift in the composition proportions. It is the core of the approach's fitting process to model these weights as functions of time, which will be elaborated on in Sect. 3.

Modelling the data stream in this way entails less computation compared to kernel density estimators, since we do not need to complete a full kernel matrix but only evaluate the small number of basis functions at the sample positions. This approach also allows a straight-forward way of forecasting the density at future time points, since the basis weight functions only need to be extrapolated to the desired time. Furthermore, the model delivers an easily interpretable description of drift by means of the basis weight functions.

The remainder of this article is organised as follows: in the Sect. 2, we review the related work, before presenting and discussing our temporal density extrapolation approach in detail in Sect. 3. An experimental evaluation of this approach is given in Sect. 4, concluding remarks in Sect. 5.

2 Related work

There is a vast literature on estimating the density within a training sample (Chacón and Duong 2018; Scott 2015), This includes approaches that provide spatio-temporal density estimates for time points within the training time window, e.g., by using time-varying mixture models (Lawlor and Rabbat 2016) or spatio-temporal kernels (Aggarwal 2005). However, our focus is on the task of density extrapolation in time-evolving data streams, as recently formulated in Kreml (2015). Thus, we will restrict the following review to this density extrapolation, then review the related problem of out-of-sample density forecasting, and finally discuss the broader context within the literature of concept drift in data streams.

Density extrapolation is described in Kreml (2015), together with a sketch of the general idea to extend kernel density estimation techniques for this task. In Lampert (2015), “extrapolating distribution dynamics” (EDD) is proposed, although this approach aims for predicting the distribution in the *immediate* future one-step ahead. EDD models the transformation between previous time points and applies this transformation to the most recent sample, to obtain a one-step ahead extrapolation of distribution dynamics.

A related problem is density forecasting [see, e.g., the survey in Tay and Wallis (2000)], where the realisations of a random variable are predicted. Applications exist in macroeconomics and finance (Tay and Wallis 2000; Tay 2015), as well as in specialised fields like energy demand forecasting (He and Li 2018; Mokilane et al. 2018). Forecast are either within (in-sample) or outside (out-of-sample) an observed sample and training time window. In our context, only out-of-sample forecasting is relevant. In Gu and He (2015), the ‘dynamic kernel density estimation’ in-sample forecasting approach by Harvey and Oryshchenko (2012) is extended to out-of-sample forecasting. Their method models a time-varying probability density function nonparametrically using kernel density estimation and schemes for observation weighting that are derived

from time series modelling. The resulting approach provides directional forecasting signals, specifically for the application of predicting the direction of stock returns. Another direction in out-of-sample forecasting is the use of histograms as approximations of the underlying distribution. Arroyo and Maté (2009) use time series of histograms, with a histogram being available for each observed time point. They propose a kNN-based method to forecast the distribution at a future time point based on the previously observed histograms. However, the method is limited to only being able to forecast an already previously observed histogram, making it more suited for context with recurring patterns. Furthermore, motivated by symbolic data analysis (Noirhomme-Fraiture and Brito 2011), Dias and Brito (2015) proposed an approach that uses linear regression to forecast the density of one variable based on observed histogram data from another variable. However, our objective is an extrapolation of the same variable, but to future time points. Another direction is the direct modelling of the probability density. Motivated by applications in energy markets, several approaches have been proposed for forecasting energy supply (e.g., wind power) and demand (power consumption). Bessa et al. (2012) developed a kernel density estimation model based on the Nadaraya–Watson estimator in the context of wind power forecasting. The employed kernels include the beta and gamma kernels as well as the von Mises distribution, underlining the very specialised nature of the approach for use in wind power forecasting. The work of He and Li (2018) is also targeted towards use in the context of wind power forecasting, for which they propose a hybrid model consisting of a quantile regression neural network and a Epanechnikov kernel density estimator. Quantile regression is also used in the work Mokilane et al. (2018) to predict the electricity demand in south Africa for the purpose of long-term planning, while the work of Bikcora et al. (2015) combines an ARMAX and a GARCH model to forecast the density of electricity load in the context of smart loading electric vehicles. However, the approaches to modelling the probability density presented above all share a strong specificity to their application.

Density extrapolation and forecasting are part of the more general topic of handling non-stationarity in streaming, time-evolving data. This problem has gained particular attention in the data stream mining community, where changes in the distribution are commonly denoted as *concept drift* by Widmer and Kubat (1996) or *population drift* by Kelly et al. (1999). As discussed in the taxonomy of drift given in Webb et al. (2016), the drifting subject might be the distribution of features X conditioned on the class label Y . Such drift in the class-conditional feature distribution $P(X|Y)$ might result in drift of the posterior distribution $P(Y|X)$. Of particular relevance for our work is *gradual* drift of $P(X|Y)$ and $P(Y|X)$, where the distribution slowly and continuously changes over time, as opposed sudden shift, where it changes abruptly. Many drift-adaptation techniques have been proposed that aim to incorporate the distribution of the most recently observed labelled data (X, Y) into a machine learning model, as surveyed in Gama et al. (2014). However, a challenge in some applications is that no such recent data is available for model update (Krempel et al. 2014). For example, in stream classification true labels Y might arrive only with a considerable delay [so-called *verification latency* (Marrs et al. 2010) or label delay (Plasse and Adams 2016)], or might only be available during the initial training (Dyer et al. 2014). As discussed in Hofer and Krempel (2013), this requires adaptation mechanisms that

Table 1 Overview of the notation used

X	Observed sample	R	Order of polynomial
x_i	i -th observation in sample	D	Dimensionality of feature space
T	Time attribute of observed sample	γ_j	Weight of j -th basis function
τ_i	Time value of i -th observation	$\boldsymbol{\gamma}$	Vector of all M basis weights
N	Observed instance sample size	B	Matrix of regression coefficients
M	Number of basis functions	ϕ_j	j -th basis function
h	Bandwidth	$\boldsymbol{\phi}$	Vector of all M basis functions

use the limited available data, which is either recent but unlabelled, or labelled but old.

These adaptation mechanisms build on *drift models* (Kreml and Hofer 2011), which model the gradual drift over time in the posterior distribution $P(Y|X)$ or the class-conditional feature distribution $P(X|Y)$. Then, they use this for temporal transfer learning from previous time points (source domains) to current or future time points (target domains). This is part of the broader *change mining* paradigm, introduced by Böttcher et al. (2008). This paradigm aims to understand the changes in a time-evolving distribution, and to use this to describe and predict changes. Various drift models and mechanisms for adaptation under verification latency have been proposed. However, they all model the class-conditional distribution of instances, for example by employing clustering assumptions (Kreml and Hofer 2011; Dyer et al. 2014; Souza et al. 2015), or calculating changes of a prior (Hofer and Kreml 2013) or weights (Tasche 2014), or by matching labelled and unlabelled instances (Kreml 2011; Hofer 2015; Courty et al. 2014), or they directly adapt the classifier (Plasse and Adams 2016). Thus, they are not applicable to model the changes in a non-parametric, multimodal distribution of unlabelled data, as it is the objective of this work. Thus, the approach proposed in this work complements the existing change mining literature. By providing a method for extrapolating $P(X|Y)$, it complements the ex-post drift analysis in Webb et al. (2017) and addresses the calls for better understanding of drift. This might be useful to assess so-called predictability of drift (Webb et al. 2016), and to adapt a classification model in presence of concept drift and verification latency.

3 Method

In the following subsections, we first derive in Subsection 3.1 the basic model for temporal density extrapolation by using an expansion of density functions and the ILR-transformation. Then, in Subsection 3.2, we extend this model by mechanisms for instance weighting and regularisation. Finally, in Subsection 3.3, we discuss how to solve the resulting optimisation problem. For the reader's convenience, we have summarised our notation in Table 1.

3.1 Basic model

Let \mathcal{X} be a univariate feature for which a sample $X = \{x_1, \dots, x_N\}$ of size N with associated time values $T = \{\tau_1, \dots, \tau_N\}$ with $\tau_i \in [0, 1]$ for $i = 1, \dots, N$ is observed.

This forms a data stream segment. Since a data stream is a dynamic setting, it is possible that the distribution of X changes as time progresses. This results in the probability density of X at a given time point t to be unequal to that at a future time $t + 1$. This change, referred to as concept drift, is assumed to not being limited to the observed sample, potentially continuing with time.

To begin modelling such a possibly changing feature, a model of the density is required. In this work the density is assumed to have the form of an expansion of M normed basis functions $\phi_j(x)$ with $j \in \{1, \dots, M\}$ and $\int \phi_j(x)dx = 1$. The basis functions $\phi_j(x)$ have been chosen as Gaussian density functions located at evenly spaced positions μ_j throughout the range of X , with standard deviation $\sigma = h$ and h being a fixed bandwidth. Other choices of basis functions and placement strategies are also possible. The basis functions could for example also be placed at the location of instances in the observed sample X . Associated with these basis functions are the weights γ_j , with $\gamma_j \in [0, 1] \forall j$ and $\sum_{j=1}^M \gamma_j = 1$. The density then takes the form

$$f(x) = \sum_{j=1}^M \gamma_j \phi_j(x) = \boldsymbol{\phi}(x)^\top \boldsymbol{\gamma} \tag{1}$$

with $\boldsymbol{\phi}(x) = (\phi_1(x), \dots, \phi_M(x))^\top$ and $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_M)^\top$.

This formulation of $f(x)$ is insufficient for the setting of a data stream as it does not account for time and the changes that might occur. Based on the model of $f(x)$ as a basis expansion it is proposed to model drift as time-dependent changes in the basis weights, while the basis functions themselves remain unchanged in location and bandwidth. The static basis weight γ_j in Eq. 1 is substituted with $\gamma_j(t)$, a function of time that models the changing basis weight of the j -th basis

$$f(x|t) = \sum_{j=1}^M \gamma_j(t) \phi_j(x) = \boldsymbol{\phi}(x)^\top \boldsymbol{\gamma}(t). \tag{2}$$

To fit this model to observed data, which entails determining the basis weight functions $\gamma_j(t)$ for $j = 1, \dots, M$, the likelihood \mathcal{L} of the model

$$\mathcal{L} = \prod_{i=1}^N f(x_i | \tau_i) = \prod_{i=1}^N (\boldsymbol{\phi}(x_i)^\top \boldsymbol{\gamma}(\tau_i))$$

is maximised by maximising the log-likelihood

$$\log \mathcal{L} = \sum_{i=1}^N \log(f(x_i | \tau_i)) = \sum_{i=1}^N \log(\boldsymbol{\phi}(x_i)^\top \boldsymbol{\gamma}(\tau_i)). \tag{3}$$

For this it is necessary to consider the properties of the basis weights mentioned above, which translate into two constraints on the basis weight functions, a sum constraint $\sum_{j=1}^M \gamma_j(t) = 1 \forall t$ and an interval constraint $\gamma_j(t) \in [0, 1] \forall j \forall t$.

Such constraints are a defining characteristic of what is referred to as compositional data (Aitchison 1982), i.e. data that describes the composition of a whole of M components. It is useful to approach the modelling of the basis weight functions as a compositional data problem, because among the methods developed for the analysis of such data there is one that enables the elegant incorporation of these constraints into the model—the isometric log-ratio (ilr) transformation (Egozcue et al. 2003). The ilr-transformation was proposed by Egozcue et al. to transform compositional data from its original feature space, which for a composition of M components is a M -dimensional simplex (Aitchison 1982), to \mathbb{R}^{M-1} .

By applying the ilr-transformation to the basis weight functions we acquire $\mathbf{v}(t)$, their representation in \mathbb{R}^{M-1}

$$ilr(\boldsymbol{\gamma}(t)) = \mathbf{U}^\top \log(\boldsymbol{\gamma}(t)) = \mathbf{v}(t)$$

where \mathbf{U} is a $M \times M - 1$ matrix that is defined as $\mathbf{U} = \tilde{\mathbf{U}} \cdot \mathbf{D}_2$ with

$$\tilde{\mathbf{U}} = \begin{pmatrix} -1 & -1 & -1 & \cdots & -1 \\ 1 & -1 & -1 & \cdots & -1 \\ 0 & 2 & -1 & \cdots & -1 \\ 0 & 0 & 3 & \cdots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & M-1 \end{pmatrix}, \quad \mathbf{D}_2 = \begin{pmatrix} \frac{1}{\|\tilde{\mathbf{u}}_1\|_2} & 0 & 0 & \cdots & 0 \\ 0 & \frac{1}{\|\tilde{\mathbf{u}}_2\|_2} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \cdots & \vdots \\ 0 & 0 & 0 & \cdots & \frac{1}{\|\tilde{\mathbf{u}}_{M-1}\|_2} \end{pmatrix}$$

$$\tilde{\mathbf{U}} = (\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{M-1}).$$

Due to this transformation we arrive at $\mathbf{v}(t) = (v_1(t), \dots, v_{M-1}(t))^\top$, so $M - 1$ functions that model the basis function weights in an unconstrained fashion. In this article it is assumed that these functions follow a polynomial regression model of order R

$$v_j(t) = \sum_{k=0}^R b_{j,k} t^k \quad j = 1, \dots, M - 1.$$

This allows the modelling of $\mathbf{v}(t)$ as a multivariate regression model

$$\mathbf{v}(t) = \mathbf{B}\mathbf{a}(t) \tag{4}$$

with

$$\mathbf{B} = \begin{pmatrix} b_{1,0} & b_{1,1} & \cdots & b_{1,R+1} \\ \vdots & \vdots & \vdots & \vdots \\ b_{M-1,0} & b_{M-1,1} & \cdots & b_{M-1,R+1} \end{pmatrix}, \quad \mathbf{a}(t) = (t^0, t^1, \dots, t^R)^\top.$$

Incorporating Eq. 4 into the log-likelihood function in Eq. 3 requires the inversion of the *ilr*-transformation

$$\boldsymbol{\gamma}(t) = \text{ilr}^{-1}(\mathbf{v}(t)) = \frac{\exp(\mathbf{U}\mathbf{B}\mathbf{a}(t))}{\mathbf{1}^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(t))}, \tag{5}$$

which is plugged into Eq. 3 to arrive at

$$\begin{aligned} \log \mathcal{L}(\mathbf{B}|X, T) &= \sum_{i=1}^N \log \left(\boldsymbol{\phi}(x_i)^\top \frac{\exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))}{\mathbf{1}^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))} \right) \tag{6} \\ &= \sum_{i=1}^N \log(\boldsymbol{\phi}(x_i)^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))) - \sum_{i=1}^N \log(\mathbf{1}^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))) \tag{7} \end{aligned}$$

with $\mathbf{1}$ being a vector of length M , each entry being 1.

3.2 Extension: instance weighting and regularisation

In this form a maximum likelihood estimation would result in fitting the centre of the training time window best. Since our objective is the extrapolation to future time points outside the training window, it is reasonable to promote a better fit to the data at the end of the training window, i.e. to the most recent data. To this end a time-based instance weighting is introduced to the above formulation in the form of multiplying with a weight vector $\mathbf{w} = (w_1, \dots, w_N)^\top$ that assigns a temporal weight to each observation based on its age. The age of an observation is considered the difference between the largest time value in the observed sample $\max(T)$, i.e. the most recent one, and the time value τ_i of the i -th observation. The weight assigned to the i -th observation is then defined as

$$w_i = \exp \left(\frac{\log(0.5)}{\kappa} \circ (\max(T) - \tau_i) \right) \tag{8}$$

where \circ denotes the Hadamard product, and κ is the age of an observation at which a weight of 0.5 is assigned. We recommend to set this value to half of the training window, i.e., $\kappa = 0.1$ in our experiments, but depending on the length of the time window, other values for κ are also possible. Including this temporal instance weighting into Eq. 7 we arrive at

$$\begin{aligned} \log \mathcal{L}(\mathbf{B}|X, T, \mathbf{w}) &= \sum_{i=1}^N w_i \log(\boldsymbol{\phi}(x_i)^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))) \\ &\quad - \sum_{i=1}^N w_i \log(\mathbf{1}^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(\tau_i))) \end{aligned} \tag{9}$$

which is the log-likelihood of the model on a biased sample.

Another issue that needs to be addressed is that the regression of the basis weights may suffer from overfitting the observed instances. This would compromise the generalisation to new, future data. To address this a regularisation term ζ is introduced to the above equation. This penalises the size of the coefficients in \mathbf{B} with exception of the offset. For this a regularisation strength parameter λ is included as well as a $R \times R - 1$ matrix \mathbf{C} so

$$\zeta = \lambda \operatorname{tr}(\mathbf{C}^\top \mathbf{B}^\top \mathbf{B} \mathbf{C}), \quad \mathbf{C} = \begin{pmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}, \quad (10)$$

which is equivalent to the sum of squares of the coefficient matrix and a common form of regularisation.

3.3 Optimisation problem and density model

Finally the objective function of the optimisation problem to be solved can be acquired by combining the log-likelihood function including instance weighting as seen in Eq. 9 with the regularisation term in Eq. 10. The maximum likelihood estimate of the coefficient matrix \mathbf{B} is then acquired by solving the unconstrained optimisation problem

$$\arg \max_{\mathbf{B}} \log \mathcal{L}(\mathbf{B}|X, T) - \zeta.$$

The objective gradient ∇l is supplied as

$$\nabla l = \left(\sum_{i=1}^N w_i \left(\frac{1}{\beta} \boldsymbol{\phi}(x_i) \mathbf{D} \mathbf{J} \right) - \sum_{i=1}^N w_i \left(\frac{1}{\beta} \mathbf{1} \mathbf{D} \mathbf{J} \right) \right) - \lambda \mathbf{B} \mathbf{C} \mathbf{C}^\top,$$

with $\mathbf{J} = \frac{\partial \mathbf{U} \mathbf{B} \mathbf{a}(\tau_i)}{\partial \mathbf{B}}$ and \mathbf{D} being a diagonal matrix with $\frac{\partial \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i))}{\partial \mathbf{B}}$ on the diagonal. The detailed derivation of the gradient can be found in the ‘‘Appendix’’.

To solve the optimisation problem we use MATLABs Global Optimization Toolbox and its implementation of the Quasi-Newton algorithm as provided by the function *fminunc*. Since only functionality for minimisation problems is provided, we redefine the optimisation problem stated above to

$$\arg \min_{\mathbf{B}} -(\log \mathcal{L}(\mathbf{B}|X, T) - \zeta)$$

and use $-\nabla l$ accordingly. Besides enabling the use of the supplied gradient and using a larger optimality tolerance of $1e^{-4}$ (instead of default $1e^{-6}$) we use the functions default parameters.

This solver configuration was used in a multiple starting point search executed via the *MultiStart* function provided by the previously mentioned toolbox. For this the ‘artificial bound’ parameter used for starting point generation is set to 2 and the number of start points was set to 4, the choice of the latter is a trade-off between higher optimality of the solution and shorter computational runtime.

After the optimisation problem is solved the coefficients in \mathbf{B} can be used to extrapolate the density of the model. For this the expression in Eq. 5 is plugged into Eq. 2, giving the models density estimate for a sample (x, t) as

$$f(x|t) = \phi(x)^\top \frac{\exp(\mathbf{U}\mathbf{B}\mathbf{a}(t))}{\mathbf{1}^\top \exp(\mathbf{U}\mathbf{B}\mathbf{a}(t))} \quad (11)$$

4 Experimental evaluation

The goals of the experimental evaluation are twofold. First, to assess the quality of the densities predicted by the models for future time points. This is done by comparing the proposed method to the other two methods that are available in the literature for this problem. Second, to investigate the behaviour of the proposed method in the form of an analysis of its sensitivity with respect to the model hyperparameters and its computational runtime. The experiments are conducted in the 2017b version of MATLAB with the exception of the EDD method, for which an implementation in Python 2.7 has graciously been provided by the author of the EDD method. All experiments are conducted on the ‘Gemini’ cluster of Utrecht University, using a PowerEdge R730 with 32 HT cores and 256 GB memory.

For the experiments a range of data sets have been selected, in part real and artificial, which show various different kinds of drift over time. These data sets will be discussed in detail in the following.

4.1 Data

Four different artificial data sets are generated to simulate different drift patterns. The use of artificial data in this context has the advantages that both the drift pattern is explicitly known as well as the data generating process in general, so the models’ estimates can be compared to the true density.

All four artificial data sets are generated by a mixture of three or four skew-normal distributions. To simulate different kinds of drift certain parameters of the mixture distribution are subjected to gradual changes over time, which is reflected in the names of the data sets:

- *meandrift*—four components with the location parameter of all but the first component changing over time.
- *weightdrift*—three components, mixture weight of the second decreases over time, weights of the other two increase.
- *sigmachange*—three components with scale parameter changing over time.
- *staticskewnormals*—four unchanging components.

All artificial data sets contain 25,000 instances in the range of $[0, 12]$ which are equally distributed over 120 unique time points in the interval $[0, 1]$.

The first real-world data set used in the experiments is taken from the website of the US peer-lending company ‘Lending Club’,² containing the accepted loan requests of the years 2007–2017. This data was then preprocessed by applying to each variable separately a Box–Cox transformation (Box and Cox 1964), as implemented in the R package MASS (Venables and Ripley 2002). Since the proposed method in this form can only handle univariate data the processed lending club data set is split. This results in each feature forming a separate data set, which is named after the feature. Each of these new data sets contains the feature values as well as the time stamps. In total the data contains 120 different time stamps, corresponding to the monthly data of 10 years. The features that have been selected from the original 75 are:

- *dti*—a ratio calculated using the borrower’s total monthly debt payments on the total debt obligations, excluding mortgage and the requested LC loan, divided by the borrower’s self-reported monthly income.
- *int_rate*—interest rate on the loan
- *loan_amnt*—the listed amount of the loan applied for by the borrower. If at some point in time the credit department reduces the loan amount, then it will be reflected in this value.
- *open_acc*—the number of open credit lines in the borrower’s credit file.
- *revol_util*—Revolving line utilisation rate, or the amount of credit the borrower is using relative to all available revolving credit.

One notable characteristic of the lending club data sets is the vastly bigger amount of data points at later time points in the stream. The 25th percentile lies at $t = 0.7$ and the 75th percentile at $t = 0.91$. This is due to the fact that the data set begins not long after the starting of the company and therefore also shows the increased amounts of credit applications.

Another data set that is taken from the real world is what will be referred to here as the ‘pollution’ data set. It is created from a Kaggle data set about air pollution in Skopje, Macedonia, between the years of 2008 to 2018.³ The original version of the data set contains measurements of carbon monoxide (CO), nitrogen dioxide (NO₂), ozone (O₃) as well as particulate matters smaller than 10 and 2.5 micrometers (PM10, PM25), measured at 7 different measurement stations in the city. While technically the measurements have been taken on an hourly basis within the time frame of 2008 to 2018, there are often missing values for certain stations or certain compounds. Since the original data represents time series of single measurements, it is not exactly fitting for our purposes, since we are interested in forecasting the distribution of the data.

To arrive at a more suitable data set the data was split into one subset per feature, resulting in 5 subsets with each a feature and the time variable. The separation by measuring station was removed, meaning that each hourly time stamp had at most 7 measurements associated to it. Each of these subsets underwent the same preprocessing. First, all entries whose feature value were missing have been discarded. Then

² <https://www.lendingclub.com/>.

³ See <https://www.kaggle.com/cokastefan/pm10-pollution-data-in-skopje-from-2008-to-2018>.

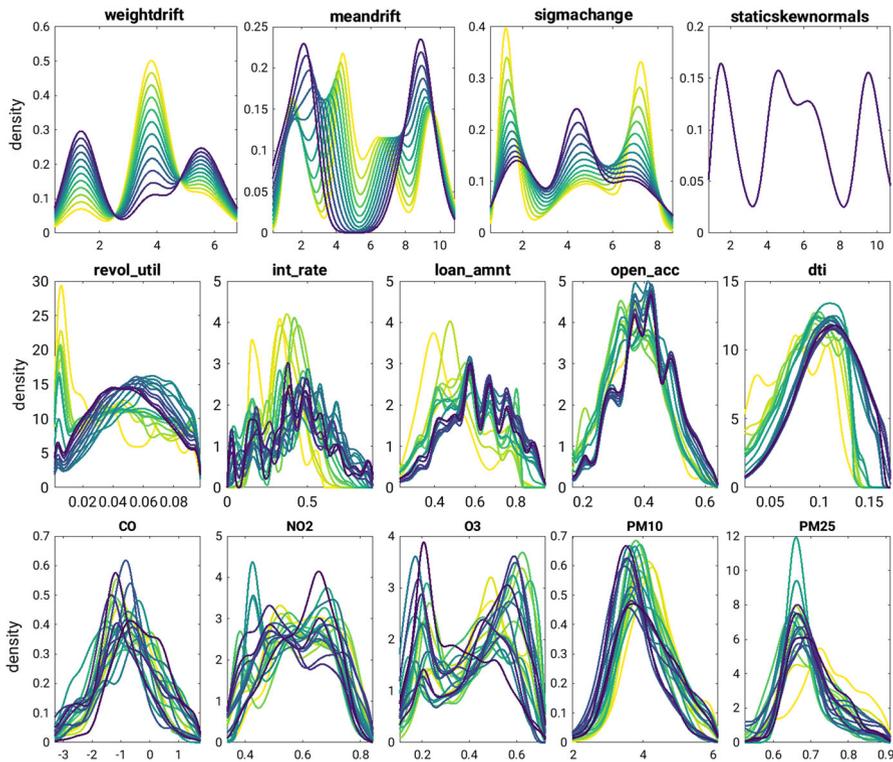


Fig. 2 True density of artificial data (top row), baseline density of lending club data (middle row) and baseline density of pollution data (bottom row) illustrated over time. Not all time points have been plotted for visibility's sake. Time progression colour-coded from light ($t = 0$) to dark ($t = 1$) (Color figure online)

all entries whose time stamp does not lie within the inclusive interval of 4 pm till 8 pm (considered the time frame of evening rush hour by common definition) were removed to eliminate the intra-day variation in the data and instead focusing on the trends on a monthly basis. To this end the time stamps were generalised to month-level by dropping the day component of the time stamp. Finally the time variable was normalised to a scale from 0 to 1 and the feature values were either box-transformed if the fitted λ value was not 0 and log-transformed if equal 0. In the end the processed data set represents the transformed compound measurements during evening rush hour in Skopje on a monthly basis, allowing for an analysis of the change in the distribution over time.

Figure 2 illustrates the changes in the density over time on both the real and artificial data sets. The X-axis shows the feature values (X), while the Y-axis shows the probability density of X . For the artificial data set this is the true density, for the real data it is an approximation as will be discussed later. The shade of the different lines indicates the time point associated to the density curve, with lighter shades representing earlier and darker shades representing later time points. The middle right plot in Fig. 2 shows the change in the approximated density on the lending club 'dti' data set, showing a

slow shift to the right that is fairly simple, while the plot of the ‘interest rate’ data set shows a more complex change. Many small changes at multiple locations make this pattern fairly complex and noisy. In contrast to this, the drift patterns of the artificial data sets are intentionally simple to clearly simulate certain causes for a changing distribution.

To give a comparison to the proposed method regarding the quality of the predicted density, two additional methods have been employed. The first of these is a static version of the proposed method, which fits the basis expansion model with static weights instead of modelling a time-dependent weight function. As such, this model does not account for the temporal dimension of the data. The second method used is the “Extrapolating the Distribution Dynamics” (EDD) approach proposed by Lampert (2015). Since EDD is based on kernel density estimates, the number of training instances needs to be kept in mind when tuning the bandwidth parameter of the model. If the number of training instances during model selection is very different from the number of instances in the experiments training window, the kernel bandwidth is likely to be unfit, resulting in poor performance. To account for this it was ensured in the experiments that the training set presented to EDD in the experiments comprises the same amount of instances as there are in the respective model selection training set. The procedure for this is as follows: if the training window of the model selection comprises less than 10% of the entire data set, then all instances in this window are used; if the model selection training window contains more than 10% of the overall instances, then a random subsample corresponding to 10% of the data is used. When training the model during the experiments it is ensured that the number of training instances does not exceed the number of instances used during model selection. If the experiment training window contains more instances than that, a random subsample is used—if it contains less, all samples are used.

The experiment setup comprises a model selection phase elaborated on in Sect. 4.2, a training phase using the selected hyperparameters and the application of the trained models to predict the density in a previously unseen segment of the data stream. Thus, on an infinite stream the model selection is done at the beginning, while training and prediction phases are repeated over time. The data sets have been partitioned into several time windows as illustrated in Fig. 3 for the experiments. Consider that $t = 0$ indicates the first time point in the data set and $t = 1$ the last time point. In this setting, we consider the data until $t = 0.8$ as available historic data and the data after that as entirely unseen. We chose the time frame $[0.0, 0.5[$ for model selection, because of the strongly skewed distribution of instances over time in the lending club data set. This way there is a sufficiently large initial sample for all algorithms. Therein, the interval $[0.0, 0.45[$ is used to train models with different hyperparameter combinations and the interval $[0.45, 0.5[$ is used to evaluate them. The segment of $[0.5, 0.8[$ was used to train the models with the previously selected hyperparameter combinations. In order to investigate the influence of the training window size on the model performance, the methods were trained once on each of the three sub-intervals $[0.5, 0.8[$, $[0.6, 0.8[$ and $[0.7, 0.8[$. These three training windows per method result in 9 final models per data set that are finally evaluated.

Each of the trained models is then applied to predict the out of time density, that is at the time points within the unseen testing window. This testing window spans the

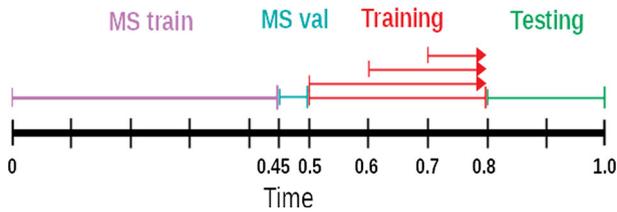


Fig. 3 Segmentation of the data for the experiments. Separate time intervals for training as part of the model selection ('MS train') and the validation of its models ('MS val'), as well as for the training of the final models with three different training window sizes ('Training') and their evaluation ('Testing')

remaining interval of $[0.8, 1.0]$ and is the same for all models. In the evaluation we will refer to the time difference between the end of the training window and the time point of the prediction as "latency".

The quality of a model's density estimate for a given time point is measured by the mean absolute error (MAE) of the prediction and the true density, evaluated at a previously defined set of points. This error measure was chosen because it is commonly used in forecasting and provides a straight-forward interpretation. Computing the error measure on the artificial data sets is simple, since the data generating process is entirely known and the true density can be computed. This however proves problematic with the real-world data, where the underlying process that generated the data is unknown and only random samples of varying size drawn from this process are available. To still evaluate the quality of the predicted density in this case, it is required to approximate the density based on the observed samples.

To approximate the true density at a time point t_i we consider the instances within a time window of size 4 around that time point to smooth potential noise, i.e. $\{t_{i-4}, \dots, t_i, \dots, t_{i+4}\}$. Based on this set of instances of size S an ensemble of 9 smoothed histograms is created, with each histogram using a different bin size b . The set of used bin sizes consists of 8 incremental integer steps, 4 in positive and 4 in negative direction, around the result of Sturges' formula $b_s = (\log_2 S) + 1$ (Sturges 1926). For each of these histograms we then compute the relative number of samples per bin with respect to the number of samples and divide it by the bin size. This density scaled, relative frequency is then associated to the bin centres and used to fit a cubic spline. Each smoothed histogram then approximates the density by the splines function value where it is greater than zero, while the density is regarded as equal to zero where the splines function value is smaller or equal zero. Each spline in the ensemble is then evaluated over the same set of points across the domain of X . The mean of the resulting 9 vectors then forms the approximation of the true density. This approximation is in the following only referred to as 'baseline density'.

4.2 Model selection

Both the method proposed in this article as well as EDD (Lampert 2015) require hyperparameters that need to be determined as part of the model selection step in the experiments. The proposed method requires four hyperparameters, namely the

number of basis functions M , their bandwidth h , the order of the polynomial R in the multivariate regression and the regularisation factor λ . EDD requires the Gaussian kernel bandwidth σ and also the regularisation factor λ .

In order to determine these hyperparameters a model selection step is incorporated into the experiment design, using the data designated for the model selection phase as discussed earlier. For the λ hyperparameter of EDD the author used $\frac{1}{N}$ as a default value for his experiments on the artificial data sets. We based the parameter search space for λ on this suggestion by multiplying it by a scaling factor $b \in \{0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5\}$. The parameter search space for σ consists of 20 evenly spaced numbers in the interval $[0.001, 0.3]$. For each possible combination of values in these two parameter search spaces an EDD model is trained and the parameter values of the model that scored the lowest MAE on the validation data is selected for use in the experiments for the given data set.

The model selection process for the proposed method is divided into two phases. First, M and h are optimised, since they determine the general fit of the basis expansion model. For M the values in $\{10, 12, 14\}$ are considered. Since M basis functions are spaced equidistantly and are fixed in their location, the coverage of the feature space by the basis functions is influenced by both M and h —e.g. a particular value for h might be a good choice for $M = 14$, but a poor choice for $M = 10$ for example by resulting in gaps between the bases. Because of this connection the limits of the searched interval $[h_{min}, h_{max}]$ are computed depending on M and the range of the available data as

$$h_{min} = \frac{P_{0.99}(X_{mst}) - P_{0.01}(X_{mst})}{M} \cdot 0.5$$

$$h_{max} = \frac{P_{0.99}(X_{mst}) - P_{0.01}(X_{mst})}{M} \cdot 1.2$$

with $P_z(\cdot)$ denoting the z -th percentile function and X_{mst} the data within the model selection training window.

Again, the hyperparameter combination that yields the smallest MAE is selected. This is then used in the second phase of the model selection in which R and λ are determined via a linear search of the parameter space for both parameters, with $R \in \{1, 2, 3\}$ and $\lambda \in \{1, 2, 3, 4, 5\}$. Note that the search spaces for these parameters are based on experience with experiments prior to those presented here. A detailed account of the models sensitivity to different hyperparameter values will be presented in Sect. 4.4.

4.3 Results

In this section we will first address the question of the models predictive quality, for which the three groups of data sets (artificial, lending club and pollution) will be regarded sequentially. Then the results of a series of experiments with the goal of investigating the sensitivity of the proposed method with respect to its hyperparameters will be presented. Finally the computational effort of the involved methods will be discussed.

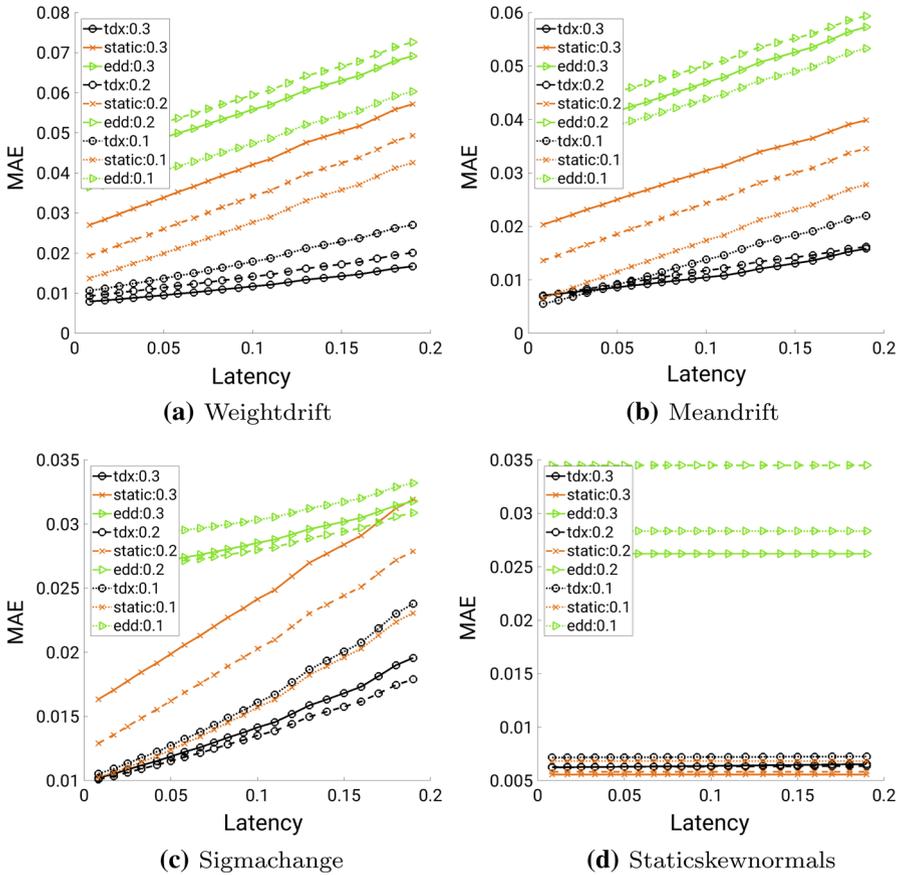


Fig. 4 MAE of each method (indicated by colour) for each training window length (indicated by linestyle) for different latency values on the artificial data sets (Color figure online)

4.3.1 Artificial data

The artificial data sets provide an important starting point for the evaluation of the proposed method since the drift patterns are known and clearly defined. This way one can evaluate whether a particular kind of pattern is recognised.

The weightdrift pattern matches the assumptions of the temporal density extrapolation (TDX) model the best, given that the weightdrift data is generated by a mixture distribution whose mixture weights are linearly changing over time. Figure 4a shows the MAE of the different models (tdx, static, edd) for the three different training window lengths (0.1, 0.2, 0.3) indicated by linestyle for a range of latency values (X-axis). As mentioned earlier, the latency is the time difference between the end of the training window and the time point of the forecast. TDX is performing best on this pattern for all three training window sizes and it can be seen that the length of the training window influences the error of the model and its change over time. The TDX model with the

longest training window shows both a lower error and a slower increase in error over time compared to the TDX models with window lengths of 0.2 and 0.1 respectively. The opposite effect can be observed with the static model, which performs better with a smaller training window.

On the meandrift pattern the location parameters of the sub-distributions in the mixture linearly change over time, which results in a less smooth drift compared to that of weightdrift as seen in Fig. 2. Nonetheless TDX also performs best on this pattern as can be seen in Fig. 4b. Here the difference between the different training window lengths for TDX is smaller. Although the model with the longest training window performs marginally worse for smaller latency values, it later consistently performs better than both models with shorter training windows.

The static models show a slightly higher error and increase in error, again with the smallest training window resulting in a lower error.

On the sigmachange data the drift is simulated by a linear change in the standard deviation of the components in the mixture distribution, resulting in the density change shown in Fig. 2. Figure 4c shows that all TDX models except the one with the shortest training window perform consistently better than the static model, showing lower error and a slower increase in error over time.

Finally the performance on the staticskewnormals data set is illustrated in Fig. 4d, which does not include any change in the distribution over time at all. Although all TDX and static models are fairly close in terms of error (< 0.0025) the best TDX model only scores the third lowest error. All TDX models show almost no increase in error over time, indicating that the model was able to recognise the absence of drift.

To statistically validate these results a series of two-sided Wilcoxon signed-rank tests was performed. For each data set, we selected the best performing model of each method with respect to the training window size based on the summed MAE over all forecasted timepoints. Based on this selection we considered two scenarios, "TDX vs EDD" and "TDX vs static". For each time point within the forecasting time window, the absolute deviations of the method's predicted density to the true/baseline density was computed at the same 200 equally spaced points within the domain of X . These error distributions were then tested with the two-sided Wilcoxon signed-rank test in order to determine whether the differences of these distributions are significant at a significance level of 0.01.

The results of these tests are illustrated in Fig. 5. The orientation of the triangular markers indicated whether the MAE of TDX is lower (downward) or larger (upward) than the other method in the test scenario, indicated in the labels of the Y-axis. If the difference in the error distributions at a given timepoint is significant according to the test mentioned above, the marker appears filled, otherwise empty.

These results show that the proposed method performs significantly better than EDD on all artificial data sets for all forecasted time points. Compared to the static model it also scores consistently lower MAE on weight- and meandrift with exception of the very first time point on meandrift. These results are significant on both data sets over all time points except for the first 3 on meandrift. On the sigmachange data TDX shows a consistently lower MAE than the static model, but the differences are not significant at any of these time points. Finally, although the MAE of TDX's predictions is only

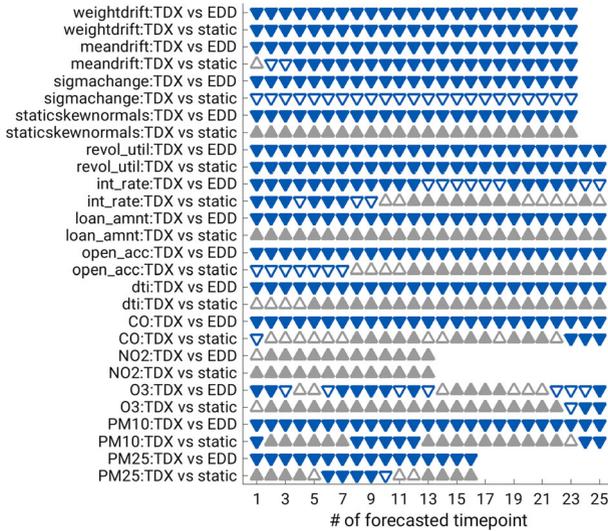


Fig. 5 Summary of TDXs estimation error in comparison to EDD and static. X-axis shows the forecasted timepoint, the markers on the Y-axis indicate TDXs relative error: triangle downward, if MAE of TDX is smaller than that of the other method, triangle upward if MAE of TDX is bigger. Filled out markers indicate that the difference is significant

slightly larger than that of the static model for all time points on staticskewnormals, the differences are significant for all forecasted timepoints.

In summary the proposed method manages to capture the drift patterns on weight-drift, meandrift and sigmachange well and performs better than the static model. On the staticskewnormals data set the TDX models consistently have higher error than the static model, showing that on an entirely static data set the static model is hard to surpass.

4.3.2 Lending club data

The lending club data sets are the first set of real-world data sets used in the experiments. Since the data contains 10 years worth of monthly data, the time variable is measured in months on these data sets as give a better understanding of the presented time frames.

As the first of these data sets the revolving line utilisation rate (short ‘revol_util’) is considered. A look at the evolution of the baseline density of this data set in Fig. 2 (second row, first from left) shows interesting, non-monotonous drift pattern. A peak on the left edge of the domain of X diminishes as the density shifts to form a new peak around $X \in [0.06, 0.08]$ that then shifts left towards $T \in [0.03, 0.06]$. The results in Fig. 6 show that on the revol_util data set TDX performs well with all three window sizes resulting in very similar error as well as almost identical, slight increase in error over time. Throughout the entire forecasting period all three TDX models achieve lower errors than the static models, while EDD scores an enormous error due to the bandwidth selection of the grid search being unfit for this later segment of the data set.

Fig. 6 MAE on *revol_util* for different training window sizes and latency values

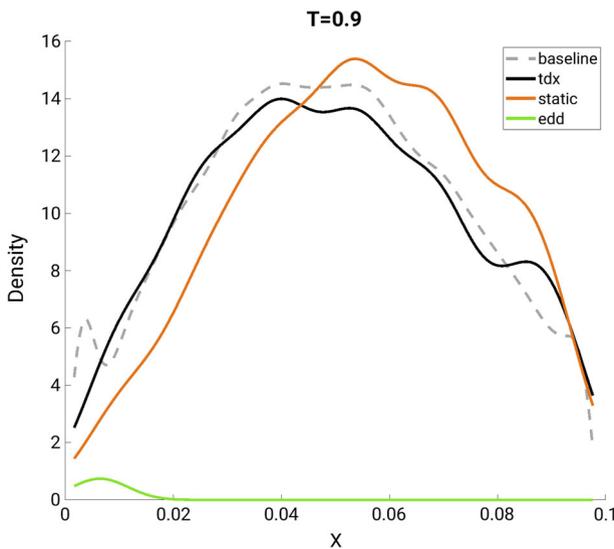
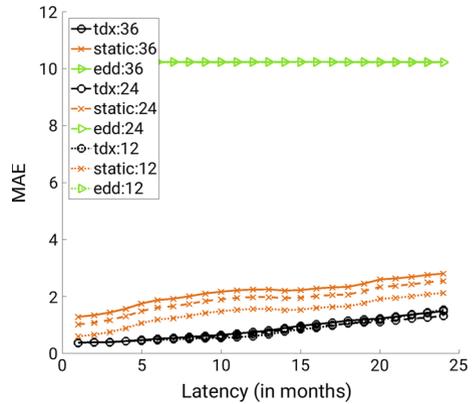
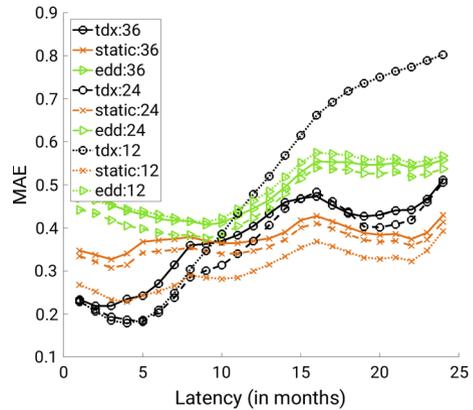


Fig. 7 Baseline and predicted density at $T = 0.9$ (12 months latency) of models with training window length of 12 months on the *revol_util* data

The evolution of the density on this data set as shown in Fig. 2 provides a hint as to why a poor choice of bandwidth resulted from the model selection, since the density distribution has a very different shape during the earlier time points. This indicates that EDD’s parameters might be very sensitive to distributional changes. Looking at the predicted density as well as the baseline density in Fig. 7 one notices that even with a latency of 12 months the TDX model manages to anticipate the change in the density, while the forecast of the static model still reflects the density at the end of the training window.

Figure 5 shows that TDX’s prediction not only result in a lower MAE than the other two methods, but that the distributions of the absolute error are also significantly different for all forecasted timepoints.

Fig. 8 MAE on `int_rate` for different training window sizes and latency values



A more volatile and mobile drift can be observed on the interest rate data set (short ‘`int_rate`’) in Fig. 2 (second row, second from left). Multiple smaller movements in the density make a difficult, non-monotonous pattern. Figure 8 shows that within the first 7 months of latency the TDX models with 12 and 24 months of training data reach the lowest and second-lowest error respectively. After this point the static model with a 12 month window surpasses them, remaining the lowest error model for the rest of the forecast period. For the TDX model with a 12 month training window a quick increase in error can be observed for latencies > 8 months, which is likely a result of the non-monotonous and often reversing trends within the data. Here TDX continues the drift pattern that has been learned within the training window, but the drift pattern of the data changed, and possibly even reversed.

As can be taken from Fig. 5 the MAE of the proposed method is only lower than that of the static model for the first nine forecasting timepoints, and only significant on 6 out of 9 compared to static. Furthermore TDX shows lower error than EDD for all forecasted time points with the differences being significant on 17 of the 25 time points.

Figure 9 shows the baseline and predicted density on the `int_rate` data with 5 months of latency. It is noticeable that the baseline density shows a lot of peaks and valleys compared to the `revol_util` data set, which none of the methods seem to capture entirely. Inspection of the corresponding empirical cdf in Fig. 10 shows that while there are deviations of the baseline from the ecdf, the baseline is reasonably close to the empirical distribution.

As shown in Fig. 5 the results for the other three data sets from the lending club series, namely `loan_amnt`, `open_acc` and `dti`, show worse performance than the two previously presented. While TDX shows lower and significantly different error compared to EDD on all three and all forecasted time points, the performance of TDX compared to static is worse. For `loan_amnt` and `dti` TDX shows consistently higher error than the static model. On `open_acc` TDXs error is lower for the first 7 time points but insignificantly so. The density plots of these data sets included in Fig. 2 show that the density at the later time points tends to only change slightly, as the darker lines are visibly close to each other. Similar to the results observed Fig. 4d, the static model

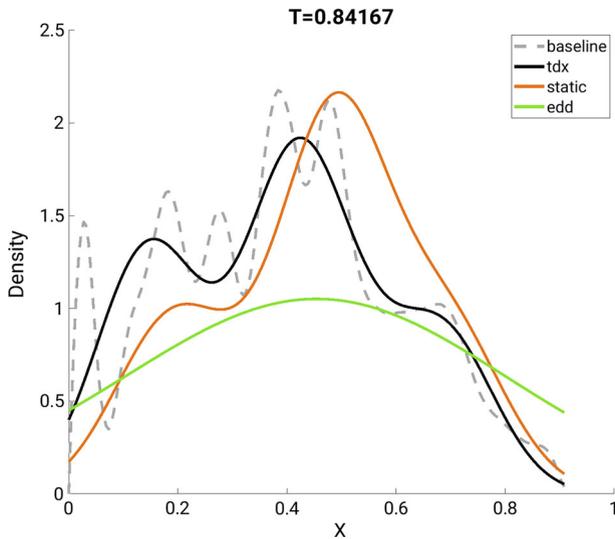
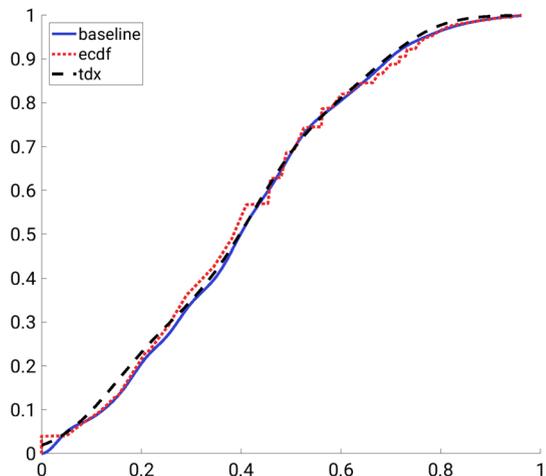


Fig. 9 Baseline and predicted density at $T = 0.841$ (5 months latency) of models with training window length of 24 months on *int_rate*

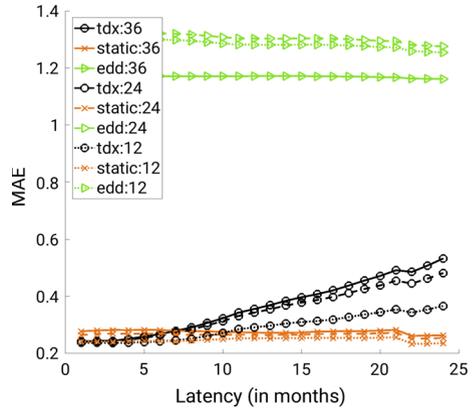
Fig. 10 CDF modelled by baseline (solid line), TDX (dashed line), as well as empirical CDF (dotted line)



proves as a tough competitor for TDX on these data sets. The MAE curves for these data sets are shown in Fig. 11 as well as Fig. 17a, b (see “Appendix”).

To summarise the results on the lending club data it can be said that TDX provides good forecasts on the two data sets that show a continuation of the drift pattern beyond the training window. If the drift does not continue or is not present at all, the static version of the model provides better estimates than the extrapolating one, which is to be expected.

Fig. 11 Error development on the open_acc data set



4.3.3 Pollution data

The second series of real-world data sets is the Skopje pollution data. These data sets are an interesting addition to the previously presented ones as they not only show non-monotonous drift patterns, but these are to some extent repeating. This proves particularly challenging for all models in the experiments, since none of them is equipped to adjust to such repeating patterns. An example of this can be seen in Fig. 12a that shows the error on the CO data, while Fig. 12b shows the development of the baseline density within the test timeframe. From the latter figure one can observe how the density around $x = -1.5$ increases, then decreases below the initial level, then increases again and again afterwards, before finally decreasing as the majority of the density shifts to the right. Matching this, an increase in error can be observed for all models around a latency of 0.03 and 0.13, right when the two major increases in density are observed. Due to the back-and-forth of this development, the TDX models struggle to anticipate this change, leading to the static model performing better overall by not trying to anticipate the movement of the drift. Nonetheless the TDX models manage to only score a slightly higher error for most of the forecasting time window and even managing to perform better at the end. The test results in Fig. 5 confirm this proximity, showing that the differences between TDX and static are only significant for only 14 of 25 time points.

While similar repeating patterns can be found on the other pollution data sets, these are not handled as well by TDX as on CO, resulting in the static method showing significantly lower error on large stretches of the forecasting time window on these data sets. Exceptions to this are observable on PM10 and PM25 between the time points 6 and 12, but it is nonetheless observable that the proposed method is not suited to handle this kind of drift patterns. The error curves for the O3, NO2, PM10 and PM25 data sets can be found in Fig. 17 in the “Appendix”.

4.4 Parameter sensitivity analysis

Since the proposed method has 4 hyperparameters we will discuss how sensitive the method is to changes in these parameters. To this end a series of experiments was

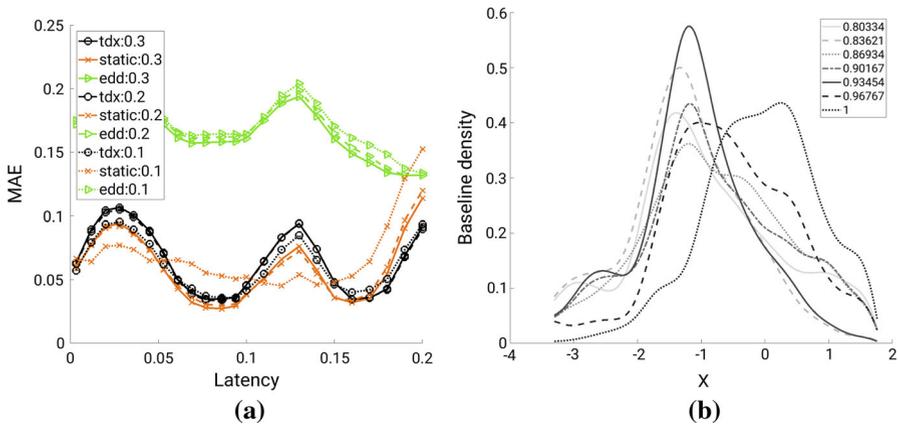


Fig. 12 **a** MAE of each method (indicated by colour) for each training window length (indicated by linestyle) for different latency values on the CO data set. **b** Baseline density of the CO data set within the test timeframe. Each line represents different time points as indicated in the legend (Color figure online)

conducted in the time interval of $t \in [0.3, 0.45]$ on the `int_rate` and meandrift data sets. In the first phase of these experiments the number of basis functions M and the bandwidth h was investigated by spanning a parameter grid of 30 bandwidth values and 9 different numbers of bases, while keeping the other two parameters fixed at $R = 2$ and $\lambda = 1$. To compare the resulting models, the MAE of the model's prediction and the true/baseline density for latency of 0.05 was used. In Fig. 13a the MAE for different h and M values on the `int_rate` data set is shown. While the extreme values for small M and h values make the surface difficult to read a region with lower error is visible for $h \leq 0.2$.

The surface in Fig. 13b shows a clearer image for the results on the meandrift data set. There is a region with lower error for M -values larger than 6 for h -values under 1, while very small h -values combined with small m -values result in higher error, reaching a maximum at the far corner of the surface at $m = 4$ and $h = 0.25$. In this case the combination of few bases and small bandwidth strongly impacts the model performance. For M -values greater than 6 the error increases quite similarly with increasing h , eventually reaching a plateau at $h = 3$ where the bandwidth is so large that the density estimate becomes so smooth that the influence of M is no longer visible. As it is also the case for the static density estimation, it can be seen that the choice of the M and h parameters is strongly dependent on the data so that it is best tuned on available data.

Proceeding with the best performing M and h values, the second phase of these experiments considers the influence of the remaining two hyperparameters, which are the order of the polynomial R and the regularisation strength λ . The heatmap in Fig. 14a shows the results for the `int_rate` data set where the minimum lies at $R = 6$ and $\lambda = 1$. The results for $R > 1$ and $\lambda = 1$ form the region with the lowest errors, ranging from 0.26 to 0.28, while $R > 1$ and $\lambda = 0$ shows the highest error. The exact opposite is the case on the meandrift data set, whose results are shown in Fig. 14b. Here, the region with lowest error is observed for $\lambda = 0$, with the error increasing with

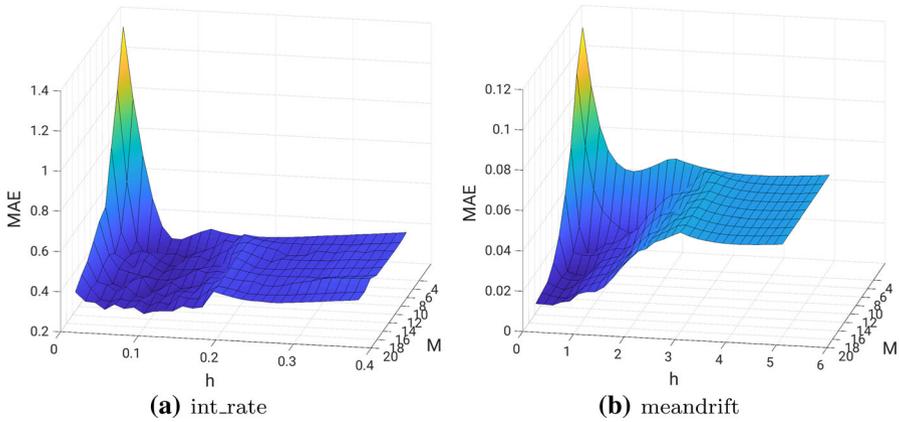


Fig. 13 Surface plots showing the MAE for different combinations of M and h values

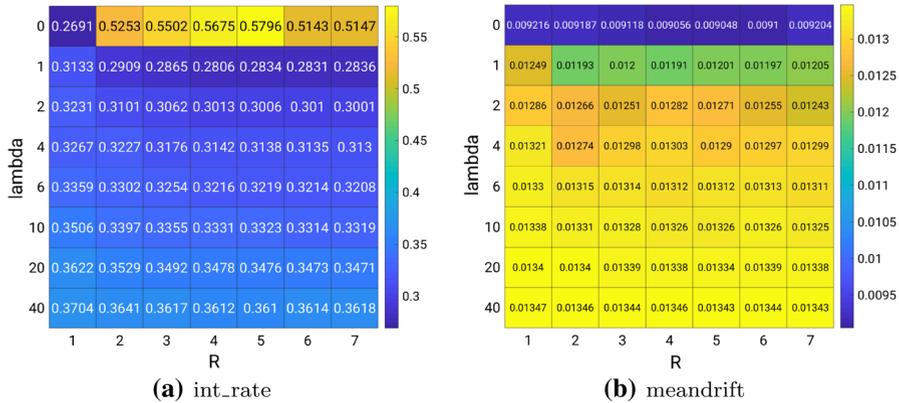


Fig. 14 Heatmaps showing the MAE for different combinations of R and λ values

increasing regularisation strength. In the case of meandrift, the drift pattern is very clear and shows little noise, leading to the unregularised models performing better. On real-world data like `int_rate` the models benefit from a certain amount of regularisation, so λ -values around 1 appear to be a good choice as a default.

During all experiments above the training window was sufficiently large, so the number of used training samples was no concern. Figure 15 shows how the error behaves on meandrift for different numbers of samples used within the training window, employing the best parameters from Figs. 13b and 14b. In these experiments only every n -th instance was used, with $n \in \{124, 62, 31, 16, 8, 4, 2, 1\}$. This results in sample sizes ranging from 31 to 3744, the latter being all instances with $t \in [0.3, 0.45]$. The results show that while using the entirety of the samples resulted in the smallest error, using only around 500 instances achieved a comparable result.

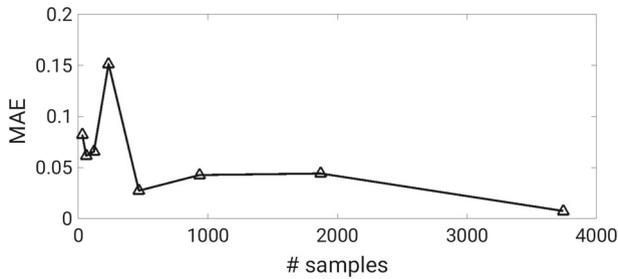
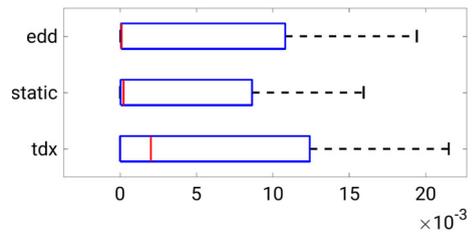


Fig. 15 MAE with varying number of samples taken from the training window

Fig. 16 Runtime in seconds required to fit the model divided by the number of training samples



4.5 Execution times

The execution time for fitting the model was measured programmatically during the experiments presented in Sects. 4.3.1–4.3.3. The boxplot in Fig. 16 shows the execution time in seconds the methods took to fit the model per observation in the training window. From the figure it is noticeable that TDX has the highest median fitting time, with the static model showing the lowest value for the 75th percentile (right end of box). Regarding the fitting time of TDX it has to be mentioned that during the experiments the optimisation problem is solved 4 times as part of a ‘multistart’ search as mentioned in Sect. 3.3. This configuration has been used as a reasonable default value and possibly has room for improvement.

5 Conclusion

In this article we presented a novel approach called temporal density extrapolation (TDX) with the goal of predicting the probability density of a univariate feature in a data stream. TDX models the density as an expansion of Gaussian density basis functions, whose weights are modelled as functions of time to account for drift in the data. Fitting these time-dependent weighting functions is approached by modelling the weights of the basis expansion as compositional data. For this purpose the isometric log-ratio transformation is used to ensure the compliance with the properties of the basis expansion weights. This approach allows to extrapolate the density model to time points outside the available training window while accounting for the changes over time that are often encountered in data streams.

The evaluation shows that the TDX manages to capture monotonous drift patterns, like changes in the component means or weights of a mixture distribution, better than a competing method (EDD) or a static version of the basis expansion model. Furthermore

the model also performs well on those two lending club data sets that exhibit very noticeable drift (`revol_util` and `int_rate`). All these data sets have in common that the drift in the data is both continuous and unidirectional. Data sets that show little to no drift however, like the static artificial data set or the remaining lending club data sets, prove challenging for TDX. In these cases the static version of the model that fits the weights of the basis expansion in a non-time-adaptive fashion performs better, as the data generating process aligns with its stationarity assumption.

Furthermore the results on the pollution data sets show that the method is currently not equipped to handle drift patterns that show seasonality. As the model is designed to handle monotonous drift, this is not surprising but shows another avenue for future work.

The analysis of the model's sensitivity with regard to its hyperparameters has shown that both the number of basis functions M and their bandwidth h has to be tuned on the data, as is necessary for the static density estimation model. For both extrapolation-specific parameters there exist reasonable defaults: for R , the order of the polynomial used to model the basis weights, and the regularisation strength λ . This reduces the effort of configuring TDX.

In summary, TDX handles its intended application, i.e., data with monotonous drift, better than the only other comparable density forecasting approach to date (EDD) and provides reliable density forecasts on these data sets. The next step in the models development will be the use of TDXs density forecasts for probabilistic classification in data streams. Furthermore we aim to extend the methods capabilities by improving the performance on data sets with little to no drift.

Acknowledgements We thank the Austrian National Bank for supporting our research Project No. 17028 as part of the OeNB Anniversary Fund. We thank Christoph Lampert for providing his implementation of the EDD approach and Utrecht University for providing the Gemini cluster for computations.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

This section contains additional material that was excluded from the main part of the work out of consideration of the available space as well as the information flow.

In the following the derivation of the objective gradient can be found. This is relevant because it is used to solve the optimization problem entailed by proposed method.

After that a table detailing the hyperparameters of both TDX and EDD that resulted from the model selection procedure can be found. Note that the static version of TDX is not listed separately, as it uses the same M and h parameters as the adaptive TDX version.

Model selection results

See Table 2.

Table 2 Hyperparameters of the models resulting from the model selection procedure

Data set	TDX				EDD	
	M	h	R	λ	σ	λ
meandrift	14	0.62	2	1	0.3	0.028
weightdrift	14	0.42	2	1	0.13	0.004
sigmachange	10	0.46	3	2	0.3	0.004
staticskewnormals	14	0.53	1	4	0.15	0.004
dti	10	0.007	3	1	0.095	0
int_rate	10	0.091	3	1	0.14	0
loan_amnt	12	0.056	3	3	0.048	0
open_acc	10	0.037	3	1	0.11	0
revol_util	14	0.007	3	1	0.23	0
CO	12	0.38	1	5	0.3	0.0014
NO2	12	0.043	3	1	0.0216	2.8e-5
O3	10	0.075	1	1	0.0629	1.4e-5
PM10	10	0.42	1	5	0.3	0.0012
PM25	10	0.044	1	5	0.15	4.6e-5

Additional figures

See Fig. 17.

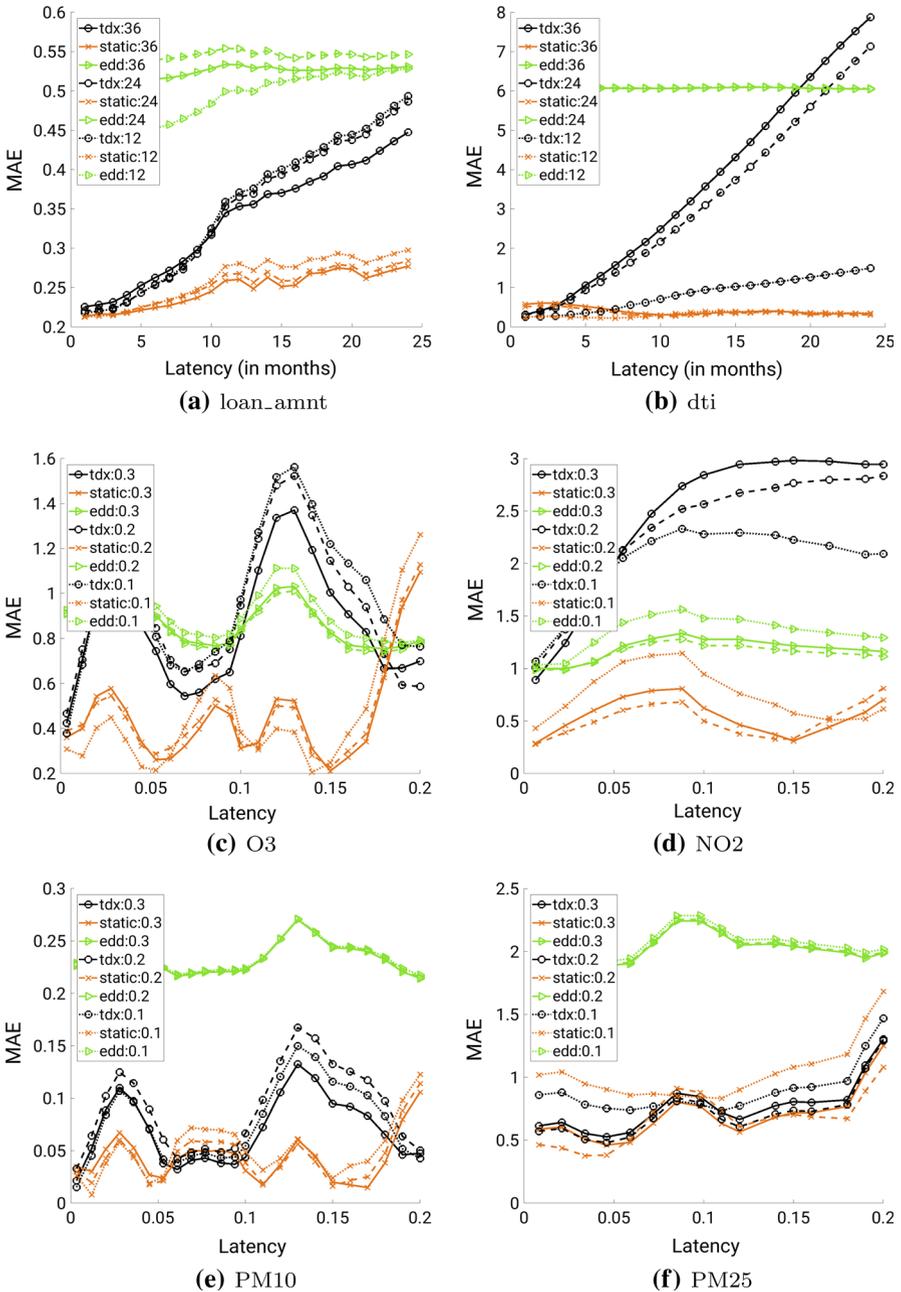


Fig. 17 MAE of each method (indicated by colour) for each training window length (indicated by linestyle) for different latency values on the remaining lending club and pollution data sets (Color figure online)

Derivation of the objective gradient

Consider the log-likelihood (l) equation of the temporal density extrapolation model, excluding instance weighting and regularization.

$$l = \sum_{i=1}^N \log(\boldsymbol{\phi}(x_i)^\top \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i))) - \sum_{i=1}^N \log(\mathbf{1}^\top \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i))) \tag{12}$$

Let us define a generalized version of the terms in Eq. 12 as $g(\mathbf{y})$

$$g(\mathbf{y}) = \sum_{i=1}^n \log(\mathbf{y}^\top \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i))). \tag{13}$$

Then the log-likelihood in Eq. 12 and its gradient ∇l can be expressed as

$$l = g(\boldsymbol{\phi}(x_i)) - g(\mathbf{1}), \tag{14}$$

$$\nabla l = \nabla g(\boldsymbol{\phi}(x_i)) - \nabla g(\mathbf{1}). \tag{15}$$

As such we will start by deriving the gradient of $g(\mathbf{y})$. Recall the matrix \mathbf{U} required for the ilr-transformation

$$\mathbf{U} = \tilde{\mathbf{U}} \cdot \mathbf{D}_2 \tag{16}$$

$$\tilde{\mathbf{U}} = \begin{pmatrix} -1 & -1 & -1 & \dots & -1 \\ 1 & -1 & -1 & \dots & -1 \\ 0 & 2 & -1 & \dots & -1 \\ 0 & 0 & 3 & \dots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & M-1 \end{pmatrix} = (\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{M-1}) \tag{17}$$

$$\mathbf{D}_2 = \begin{pmatrix} \frac{1}{\|\tilde{\mathbf{u}}_1\|_2} & 0 & 0 & \dots & 0 \\ 0 & \frac{1}{\|\tilde{\mathbf{u}}_2\|_2} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \dots & \vdots \\ 0 & 0 & 0 & \dots & \frac{1}{\|\tilde{\mathbf{u}}_{M-1}\|_2} \end{pmatrix} \tag{18}$$

which is part of the term $\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i)$ appearing in both parts of l . This term is then denoted as $\boldsymbol{\kappa}$ and its Jacobian matrix with respect to \mathbf{B} is denoted as \mathbf{J} .

$$\boldsymbol{\kappa} = \mathbf{U} \mathbf{B} \mathbf{a}(\tau_i) \tag{19}$$

$$\frac{\partial \boldsymbol{\kappa}}{\partial \mathbf{B}} = \mathbf{J} \tag{20}$$

As $U\mathbf{B}\mathbf{a}(\tau_i)$ appears in Eq. 13 in exponentiated form, we define

$$\mathbf{E} = \exp(\boldsymbol{\kappa}) = \begin{pmatrix} \exp(\kappa_1) \\ \vdots \\ \exp(\kappa_N) \end{pmatrix} = \begin{pmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_N \end{pmatrix} \tag{21}$$

and derive it in the following

$$\frac{\partial \varepsilon_i}{\partial \mathbf{B}_{ij}} = \exp(\kappa_i) \frac{\partial \kappa_i}{\partial \mathbf{B}_{ij}} \tag{22}$$

$$= \exp(\kappa_i) \mathbf{U}_{ij} \mathbf{a}_k \tag{23}$$

$$\mathbf{D} = \text{diag}(\mathbf{E}) = \begin{pmatrix} \varepsilon_1 & 0 & \cdots & 0 \\ 0 & \varepsilon_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & \cdots & \varepsilon_N \end{pmatrix} \tag{24}$$

$$\frac{\partial \mathbf{E}}{\partial \mathbf{B}} = \mathbf{D} \cdot \mathbf{J} \tag{25}$$

Then we reintroduce the placeholder variable y and define β as the expression inside the logarithm in Eq. 13

$$\mathbf{y}^\top \mathbf{E} = \mathbf{y}^\top \cdot \exp(\boldsymbol{\kappa}) \tag{26}$$

$$= \mathbf{y}^\top \cdot \exp(\mathbf{U}\mathbf{B}\mathbf{a}) = \beta \tag{27}$$

and take the derivative of this expression.

$$\frac{\partial \beta}{\partial \mathbf{B}_{ij}} = \mathbf{y}^\top \cdot \frac{\partial \mathbf{E}}{\partial \mathbf{B}_{ij}} \tag{28}$$

$$\frac{\partial \beta}{\partial \mathbf{B}} = \mathbf{y}^\top \mathbf{D}\mathbf{J} \tag{29}$$

Returning to the formulation in Eq. 13, we arrive at its gradient ∇g by substituting with the derivatives of its parts shown above.

$$g = \log(\beta) = \log(\mathbf{y}^\top \mathbf{E}) \tag{30}$$

$$\frac{\partial g}{\partial \mathbf{B}} = \frac{1}{\beta} \mathbf{y} \mathbf{D}\mathbf{J} \tag{31}$$

$$\nabla g = \frac{1}{\beta} \mathbf{y} \mathbf{D}\mathbf{J} \tag{32}$$

Applying this to the original formulation in Eq. 15 we get

$$\nabla l = \sum_{i=1}^N (\nabla g(\boldsymbol{\phi}(x_i)) - \nabla g(\mathbf{1})) \tag{33}$$

$$\nabla l = \sum_{i=1}^N \left(\left(\frac{1}{\beta} \boldsymbol{\phi}(x_i) \mathbf{D} \mathbf{J} \right) - \left(\frac{1}{\beta} \mathbf{1} \mathbf{D} \mathbf{J} \right) \right). \quad (34)$$

The inclusion of the temporal instance weighting via a weight vector \mathbf{w} is then straight-forward.

$$l = \sum_{i=1}^N w_i \log \left(\boldsymbol{\phi}(x_i)^\top \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i)) \right) - \sum_{i=1}^N w_i \log \left(\mathbf{1}^\top \exp(\mathbf{U} \mathbf{B} \mathbf{a}(\tau_i)) \right) \quad (35)$$

$$\nabla l = \sum_{i=1}^N w_i \left(\frac{1}{\beta} \boldsymbol{\phi}(x_i) \mathbf{D} \mathbf{J} \right) - \sum_{i=1}^N w_i \left(\frac{1}{\beta} \mathbf{1} \mathbf{D} \mathbf{J} \right) \quad (36)$$

This leaves only the regularization term ζ out, which was defined as

$$\zeta = \lambda \operatorname{tr}(\mathbf{C}^\top \mathbf{B}^\top \mathbf{B} \mathbf{C}) \quad (37)$$

$$\mathbf{C} = \begin{pmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix} \quad (38)$$

whose derivative then is $\lambda \mathbf{B} \mathbf{C} \mathbf{C}^\top$, arriving at the final form of the gradient

$$\nabla l = \left(\sum_{i=1}^N w_i \left(\frac{1}{\beta} \boldsymbol{\phi}(x_i) \mathbf{D} \mathbf{J} \right) - \sum_{i=1}^N w_i \left(\frac{1}{\beta} \mathbf{1} \mathbf{D} \mathbf{J} \right) \right) - \lambda \mathbf{B} \mathbf{C} \mathbf{C}^\top \quad (39)$$

References

- Aggarwal CC (2005) On change diagnosis in evolving data streams. *IEEE Trans Knowl Data Eng* 17(5):587–600
- Aitchison J (1982) The statistical analysis of compositional data. *J R Stat Soc Ser B Methodol* 44:139–177
- Arroyo J, Maté C (2009) Forecasting histogram time series with k-nearest neighbours methods. *Int J Forecast* 25(1):192–207
- Bessa RJ, Miranda V, Botterud A, Wang J, Constantinescu EM (2012) Time adaptive conditional kernel density estimation for wind power forecasting. *IEEE Trans Sustain Energ* 3(4):660–669
- Bikcora C, Verheijen L, Weiland S (2015) Semiparametric density forecasting of electricity load for smart charging of electric vehicles. In: 2015 IEEE conference on control applications (CCA), IEEE, pp 1564–1570
- Böttcher M, Höppner F, Spiliopoulou M (2008) On exploiting the power of time in data mining. *ACM SIGKDD Explor Newsl* 10(2):3–11
- Box GEP, Cox DR (1964) An analysis of transformations. *J R Stat Soc Ser B* 26:211–243
- Chacón JE, Duong T (2018) Multivariate kernel smoothing and its applications. CRC, Boca Raton
- Chandola V, Banerjee A, Kumar V (2009) Anomaly detection: a survey. *ACM Comput Surv* 41(3):15
- Courty N, Flamary R, Tuia D (2014) Domain adaptation with regularized optimal transport. In: Calders T, Esposito F, Hüllermeier E, Meo R (eds) Proceedings of the European conference on machine learning

- and knowledge discovery in databases (ECMLPKDD 2014), Springer, Lecture Notes in Artificial Intelligence, vol 8724, pp 370–385
- Dias S, Brito P (2015) Linear regression model with histogram-valued variables. *Stat Anal Data Min* 8(2):75–113. <https://doi.org/10.1002/sam.11260>
- Dyer KB, Capo R, Polikar R (2014) Compose: A semisupervised learning framework for initially labeled nonstationary streaming data (special issue on learning in nonstationary and dynamic environments). *IEEE Trans Neural Netw Learn Syst* 25(1):12–26
- Egozcue JJ, Pawlowsky-Glahn V, Mateu-Figueras G, Barcelo-Vidal C (2003) Isometric logratio transformations for compositional data analysis. *Math Geol* 35(3):279–300
- Fan W, Bifet A (2013) Mining big data: current status, and forecast to the future. *SIGKDD Explor Newsl* 14(2):1–5. <https://doi.org/10.1145/2481244.2481246>
- Gama J, Zliobaite I, Bifet A, Pechenizkiy M, Bouchachia A (2014) A survey on concept drift adaptation. *ACM Comput Surv* 46(4):1–44
- Gu W, He J (2015) A forecasting model based on time-varying probability density. In: Li M, Zhang Q, Zhang R, Shi X (eds) *Proceedings of 2014 1st International conference on industrial economics and industrial security*. Springer, Berlin, pp 519–525. https://doi.org/10.1007/978-3-662-44085-8_75
- Harvey A, Oryshchenko V (2012) Kernel density estimation for time series data. *Int J Forecast* 28(1):3–14
- He Y, Li H (2018) Probability density forecasting of wind power using quantile regression neural network and kernel density estimation. *Energy Convers Manag* 164:374–384
- Hofer V (2015) Adapting a classification rule to local and global shift when only unlabelled data are available. *Eur J Oper Res* 243(1):177–189
- Hofer V, Kreml G (2013) Drift mining in data: a framework for addressing drift in classification. *Comput Stat Data Anal* 57(1):377–391
- Kelly MG, Hand DJ, Adams NM (1999) The impact of changing populations on classifier performance. In: *Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining*, pp 367–371. <https://doi.org/10.1145/312129.312285>
- Kreml G (2011) The algorithm APT to classify in concurrence of latency and drift. In: Gama J, Bradley E, Hollmén J (eds) *Advances in intelligent data analysis X*. Lecture notes in computer science, vol 7014. Springer, Berlin, pp 222–233
- Kreml G (2015) Temporal density extrapolation. In: Douzal-Chouakria A, Vilar JA, Marteau PF, Maharaj A, Alonso AM, Otranto E, Nicolae MI (eds) *Proceedings of the 1st international workshop on advanced analytics and learning on temporal data (AALTD) co-located with ECML PKDD 2015, CEUR workshop proceedings*, vol 1425. <http://ceur-ws.org/Vol-1425/paper12.pdf>. Accessed 6 June 2049
- Kreml G, Hofer V (2011) Classification in presence of drift and latency. In: Spiliopoulou M, Wang H, Cook D, Pei J, Wang W, Zaiane O, Wu X (eds) *Proceedings of the 11th IEEE international conference on data mining workshops (ICDMW 2011)*, IEEE. <https://doi.org/10.1109/ICDMW.2011.47>
- Kreml G, Zliobaite I, Brzeziński D, Hüllermeier E, Last M, Lemaire V, Noack T, Shaker A, Sievi S, Spiliopoulou M, Stefanowski J (2014) Open challenges for data stream mining research (special issue on big data). *SIGKDD Explor* 16(1):1–10. <https://doi.org/10.1145/2674026.2674028>
- Lampert CH (2015) Predicting the future behavior of a time-varying probability distribution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 942–950. <http://pub.ist.ac.at/~chl/erc/papers/lampert-cvpr2015.pdf>. Accessed 6 June 2049
- Lawlor SF, Rabbat MG (2016) Estimation of time-varying mixture models: an application to traffic estimation. In: *Proceedings of the IEEE statistical signal processing workshop*, pp 1–5
- Marrs G, Hickey R, Black M (2010) The impact of latency on online classification learning with concept drift. In: Bi Y, Williams MA (eds) *Knowledge science, engineering and management. Lecture notes in computer science*, vol 6291. Springer, Berlin, pp 459–469
- Mokilane P, Galpin J, Sarma Yadavalli V, Debba P, Koen R, Sibiyi S (2018) Density forecasting for long-term electricity demand in South Africa using quantile regression. *S Afr J Econ Manag Sci* 21(1):1–14
- Noirhomme-Fraiture M, Brito P (2011) Far beyond the classical data models: symbolic data analysis. *Stat Anal Data Min ASA Data Sci J* 4(2):157–170
- Parzen E (1962) On estimation of a probability density function and mode. *Ann Math Stat* 33:1065–1076
- Plasse J, Adams N (2016) Handling delayed labels in temporally evolving data streams. In: *Big Data, IEEE*, pp 2416–2424
- Reinsel D, Gantz J, Rydning J (2017) Data age 2025: the evolution of data to life-critical. Technical report, IDC. <https://www.seagate.com/files/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf>. Accessed 6 June 2049

- Rosenblatt M (1956) Remarks on some non-parametric estimates of a density function. *Ann Math Stat* 27(3):832–837
- Sadik S, Gruenwald L (2014) Research issues in outlier detection for data streams. *ACM SIGKDD Explor Newsl* 15(1):33–40
- Scott DW (2015) *Multivariate density estimation: theory, practice, and visualization*. Wiley Online Library, 2nd edn. Wiley, Hoboken. <https://doi.org/10.1002/9781118575574.fmatter>
- Silverman BW (1986) *Density estimation for statistics and data analysis*. Monographs on statistics and applied probability. Chapman and Hall. <http://nedwww.ipac.caltech.edu/level5/March02/Silverman/paper.pdf>. Accessed 6 June 2049
- Souza VM, Silva DF, Gama J, Batista GE (2015) Data stream classification guided by clustering on nonstationary environments and extreme verification latency. In: *Proceedings of the 2015 SIAM international conference on data mining*. SIAM, pp 873–881
- Sturges HA (1926) The choice of a class interval. *J Am Stat Assoc* 21:65–66. <https://doi.org/10.1080/01621459.1926.10502161>
- Tasche D (2014) Exact fit of simple finite mixture models. *J Risk Financ Manag* 7:150–164
- Tay, AS (2015) A brief survey of density forecasting in macroeconomics. *Macroeconomic Review*. pp 92–97. Research Collection School Of Economics. https://ink.library.smu.edu.sg/soe_research/1901
- Tay AS, Wallis KF (2000) Density forecasting: a survey. *Companion Econ Forecast* 19:45–68
- Tran DH, Gaber MM, Sattler KU (2014) Change detection in streaming data in the era of big data: models and issues. *ACM SIGKDD Explor Newsl* 16(1):30–38
- Venables WN, Ripley BD (2002) *Modern applied statistics with S-PLUS*. Springer, Berlin (pub-SV:adr)
- Webb G, Lee LK, Goethals B, Petitjean F (2017) Understanding concept drift. *ArXiv preprint arXiv:1704.00362v1*
- Webb GI, Hyde R, Cao H, Nguyen HL, Petitjean F (2016) Characterizing concept drift. *Data Min Knowl Discov* 30(4):964–994
- Whittle P (1958) On the smoothing of probability density functions. *J R Stat Soc Ser B Methodol* 20:334–343
- Widmer G, Kubat M (1996) Learning in the presence of concept drift and hidden context. *Mach Learn* 23:2369–101

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

G. Krempl¹  · D. Lang² · V. Hofer³

D. Lang
dominik.lang@ovgu.de

V. Hofer
vera.hofer@uni-graz.at

¹ Utrecht University, Utrecht, The Netherlands

² Otto-von-Guericke University Magdeburg, Magdeburg, Germany

³ Karl-Franzens-University Graz, Graz, Austria