



Concept Drift Monitoring and Diagnostics of Supervised Learning Models via Score Vectors

Kungang Zhang, Anh T. Bui & Daniel W. Apley

To cite this article: Kungang Zhang, Anh T. Bui & Daniel W. Apley (2023) Concept Drift Monitoring and Diagnostics of Supervised Learning Models via Score Vectors, Technometrics, 65:2, 137-149, DOI: [10.1080/00401706.2022.2124310](https://doi.org/10.1080/00401706.2022.2124310)

To link to this article: <https://doi.org/10.1080/00401706.2022.2124310>



View supplementary material [↗](#)



Published online: 31 Oct 2022.



Submit your article to this journal [↗](#)



Article views: 467



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 5 View citing articles [↗](#)



Concept Drift Monitoring and Diagnostics of Supervised Learning Models via Score Vectors

Kungang Zhang^a , Anh T. Bui^b , and Daniel W. Apley^a

^aDepartment of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL; ^bDepartment of Statistical Sciences and Operations Research, Virginia Commonwealth University, Richmond, VA

ABSTRACT

Supervised learning models are one of the most fundamental classes of models. Viewing supervised learning from a probabilistic perspective, the set of training data to which the model is fitted is usually assumed to follow a stationary distribution. However, this stationarity assumption is often violated in a phenomenon called concept drift, which refers to changes over time in the predictive relationship between covariates X and a response variable Y and can render trained models suboptimal or obsolete. We develop a comprehensive and computationally efficient framework for detecting, monitoring, and diagnosing concept drift. Specifically, we monitor the Fisher score vector, defined as the gradient of the log-likelihood for the fitted model, using a form of multivariate exponentially weighted moving average, which monitors for general changes in the mean of a random vector. In spite of the substantial performance advantages that we demonstrate over popular error-based methods, a score-based approach has not been previously considered for concept drift monitoring. Advantages of the proposed score-based framework include applicability to broad classes of parametric models, more powerful detection of changes as shown in theory and experiments, and inherent diagnostic capabilities for helping to identify the nature of the changes.

ARTICLE HISTORY

Received July 2021

Accepted July 2022

KEYWORDS

Control chart; Multivariate EWMA; Predictive model; Score function

1. Introduction

In supervised learning, models are trained to predict a response variable Y , given an observed set of covariates $X \in \mathbb{R}^p$. The desire is usually to make test prediction accuracy metrics (i.e., R^2 , classification accuracy, F1-score, etc.) as high as possible. From a probabilistic perspective, a training sample of observations, $\{x_i, y_i\}_{i=1}^n$, is usually assumed to be drawn from some joint distribution, $P(X, Y)$, such that the conditional distribution, $P(Y|X; \theta)$, is stationary across the n observations (Hulten, Spencer, and Domingos 2001). Here, $x_i \in \mathbb{R}^p$ and $y_i \in \mathbb{R}$ are the covariates and response variable in the i th observation, respectively (the case with scalar response can easily be extended to a vector response), and we have parameterized the conditional distribution via the vector θ . However, the stationarity assumption is often violated in real applications, a phenomenon known as concept drift (Moreno-Torres et al. 2012; Žliobaitė, Pechenizkiy, and Gama 2016). For example, models used to predict customers' probabilities of defaulting by credit-scoring agencies are usually fitted to training data collected over a three to five years period, so that changing financial environments may result in a partially outdated predictive relationship by the time the trained model is used to score new customers (Crook et al. 1992; Žliobaitė, Pechenizkiy, and Gama 2016). In this work, we follow what appears to be the most common terminology according to (Moreno-Torres

et al. 2012) and (Žliobaitė, Pechenizkiy, and Gama 2016). In general, any temporal drift in the joint distribution $P(X, Y)$ is called “dataset drift.” Decomposing the joint distribution as $P(X, Y) = P(Y|X)P(X)$, the term “concept drift” refers to temporal drift in $P(Y|X)$ (which is what the fitted supervised learning model approximates), while “covariate drift” refers to temporal drift in $P(X)$. Our approach is intended to detect drift in the predictive relationship $P(Y|X)$ and not in $P(X)$ (see Section 4 for elaboration). For detecting changes in $P(X)$ or in $P(X)$ and $P(Y|X)$ together, we refer readers to (Raza, Prasad, and Li 2015) and (Mejri, Limam, and Weihs 2018; Mejri et al. 2021).

Concept drift poses many challenges in constructing trustworthy supervised learning models. Concept drift during both training (for historical data) and model usage (for future data) can degrade performance of supervised learning models: during training, concept drift results in there being no single predictive model $P(Y|X; \theta)$, since the model is changing over time; while concept drift during usage degrades the accuracy of the predictive model, relative to what it could be with an updated model for $P(Y|X; \theta)$. In order to obtain and maintain the highest possible predictive performance of supervised learning models, *retrospective* concept drift analysis of the training data and *prospective* concept drift monitoring of the future data to which the model is applied are important and should be a standard component of the predictive modeling process.

Concept drift is common in practice and often originates from changes of some hidden effects in generating the response outcomes of interest (Widmer and Kubat 1996; Kukar 2003; Donoho 2004; Tsymbal et al. 2008; Carmona-Cejudo et al. 2010; Žliobaitė, Bakker, and Pechenizkiy 2012; Harel et al. 2014; Fong, Di, and Permar 2015; Žliobaitė, Pechenizkiy, and Gama 2016). For example, in the problem of antibiotic resistance in nosocomial infections, supervised learning models are trained and validated using a dataset, with response labels indicating whether the level of sensitivity of a pathogen to an antibiotic is sensitive, resistant, or intermediate, and the covariates being patients' demographical data and conditions of hospitalization (Pechenizkiy et al. 2005). After properly training and validating such a model, it can with high accuracy determine whether or not a pathogen is sensitive to an antibiotic or not. However, according to medical experts (Kukar 2003), hidden effects due to failures and/or replacements of some medical equipment, changes in personnel, and seasonal bacterial outbreaks could result in unexpected changes over time in resistance of new pathogen strains to antibiotics. In other words, there is likely a significant level of concept drift in nosocomial infections. Consequently, bacteria that are predicted to be sensitive to a particular antibiotic may now be resistant, and the errors of the predictive model become unacceptable.

The problem of concept drift is gaining increasing attention, because increasingly data are organized in the form of data streams, and it is often unrealistic to expect that data distributions remain stable for a long period of time. Most state-of-the-art concept drift detection and adaptation methods, like ensemble methods (Wang et al. 2003), neural networks (Calandra et al. 2012), and other nonparametric methods (Bifet and Gavalda 2007; Frías-Blanco et al. 2015), are based on monitoring the classification error or some metrics derived from it (Ross et al. 2012; Gonçalves Jr et al. 2014; Barros and Santos 2018).

The main purpose of this article is to introduce a new and comprehensive concept drift framework that is based on monitoring the Fisher scores of the individual observations over time, as opposed to metrics derived from their prediction errors. Existing concept drift methods that are based on monitoring the error rate may fail to detect concept drift, because a change in $P(Y|X; \theta)$ does not necessarily result in an increase (or a decrease) in the error rate (see Section 3.1). Note that a change in $P(Y|X; \theta)$ that does not change the error rate still means that, if the model is updated accordingly after concept drift detection, the predictive accuracy of the model can be improved (see Section B.1, supplementary materials). Thus, it is still desirable to detect such changes. The Fisher score, which we will refer to as simply the score, of an observation (X, Y) is defined as the gradient of the log-likelihood, $\log P(Y|X; \theta)$, with respect to the parameters, θ . As shown in the following sections, comparisons between the score-based and error-based methods from both theoretical and empirical perspectives provide strong justification for a score-based approach. The score-based framework that we develop and demonstrate is intended to detect, monitor, and diagnose concept drift with the following major advantages and attractive properties.

- *The approach is based on monitoring for changes in the mean of the score vector, which has strong theoretical justification*

since theory dictates that the mean of the score vector changes if and only if $P(Y|X; \theta)$ changes, under fairly general conditions (Section 3.1).

- *The score-based approach can detect concept drift even when the error rate doesn't change. Upon the detection of the concept drift, the predictive model can be updated giving a lower prediction error rate.*
- *The score-based approach is more sensitive to changes in $P(Y|X; \theta)$ and thus more quickly detects concept drift than error-based methods, as demonstrated empirically with a number of examples (Section 5 and Sections B.2 and B.4, supplementary materials).*
- *The score-based approach is applicable to broad classes of parametric classification or regression models that can be interpreted probabilistically in terms of $P(Y|X; \theta)$, parameterized by a vector of parameters θ .*
- *The score-based approach also naturally provides a convenient means of diagnosing the nature of the concept drift (e.g., which parameters have changed) as derived in Section 4 and demonstrated in Section 5 and Sections B.4 and C, supplementary materials.*
- *The same score-based concepts can be used in a consistent manner either retrospectively (e.g., after fitting a supervised learning model to a set of training data, to validate that the training data were stable and, if not, provide diagnostic information as to why) or prospectively (e.g., when using a fitted supervised learning model to predict new cases, to quickly signal when the predictive model has changed, indicating that it is time to update the model).*

We elaborate on these properties in Sections 3 and 4 and provide an illustrative example with real data (credit risk scoring during the subprime mortgage crisis) in Section 5. The supplementary materials sections provide additional real and simulation comparison examples and derivations of various results.

2. Relation to Prior Work

In this section, we briefly review relevant existing literature on monitoring different types of drift in data distributions, in terms of goals and methodologies.

In general, the methods in the concept drift literature can be categorized into two classes (Tsymbal 2004; Harel et al. 2014): (a) model adaptation (or online learning) methods and (b) concept drift detection methods. Model adaptation methods mainly focus on maintaining the performance of machine learning models in the presence of concept drift, without formally detecting or diagnosing the drift (Wang et al. 2003; Tsymbal et al. 2008; Gonçalves Jr et al. 2014; Barros and Santos 2018). To maintain a good prediction metric related to classification or regression error, models are automatically updated (i.e., adapted) online continuously as new observations are collected, which is sometimes called online or incremental learning. This class of methods is not particularly relevant to our work, since our goal is concept drift detection and diagnosis, and not model adaptation. In fact, we view our approach as something that can be used in conjunction with model adaptation methods to make them more efficient and interpretable. In particular, the

model adaptation could be turned on only when the concept drift detection component has indicated that $P(Y|X; \theta)$ has changed.

The concept drift detection methods are more relevant to our work. Although most of these methods were originally proposed in a continuous model adaptation setting, we focus on their use and performance in our setting, in which one has a single fixed model fitted to one set of training data and applies it to usage data without updating it, so that small and gradual concept drift can be more easily detected. The drift detection method (DDM) (Gama et al. 2004) monitors the accumulated classification error rate as it evolves over time. The early drift detection method (EDDM) (Baena-Garcia et al. 2006) instead monitors the intervals between consecutive errors. The exponentially weighted moving average (EWMA) for concept drift detection (ECDD) (Ross et al. 2012) method monitors the misclassification rate using a univariate EWMA control chart. In the Linear-4-rate method (Wang and Abraham 2015), four statistics are monitored simultaneously to detect concept drift in binary classification problems. Most of the methods are designed specifically for binary classification and based on simple error-based quantities, like classification error, precision, recall, and residuals of the classification model (Wang and Abraham 2015; Barros and Santos 2018). As we demonstrate throughout the article, our score-based approach is far more powerful than error-based approaches.

Although the concept drift community seems to be unaware of the potential of score-based approaches for concept drift (e.g., it is not mentioned in the recent surveys in Barros and Santos (2018) and Lu et al. (2018)), there have been some works in the econometrics literature that have used the score vector to test for changes in the parameters of regression models (Kuan and Hornik 1995; Zeileis 2005; Zeileis and Hornik 2007; Xia, Guo, and Zhao 2009). Our work differs from this prior work in that we develop a comprehensive framework for, and investigate issues more relevant in, the typical situations to which the so-called concept drift paradigm refers. The econometrics work is developed mainly around the change-point paradigm in which it is assumed there is a single point in time at which θ changes from some before-value to some after-value, and the goal is to identify the change point with formal hypotheses testing. In contrast, our approach is developed around the much more general and flexible situation in which θ can continuously drift and/or change abruptly at multiple change-points, which is far more common in typical concept drift applications. Ours is more of an exploratory approach to inform the predictive modeling process, as opposed to formal hypothesis testing of well-defined but restrictive hypotheses. Our approach is also more general in the sense that it applies to broad classes of parametric supervised learning models in which $\log P(Y|X; \theta)$ is differentiable in θ (e.g., generalized linear models, neural networks, Gaussian process models, etc.). Moreover, we develop and study a number of other aspects that are highly relevant to the concept drift setting, including a diagnostic framework for understanding the nature of the concept drift.

3. Monitoring the Score Vector for Concept Drift

To monitor supervised learning models for concept drifts, we propose a systematic framework based on the sample score vectors derived from parametric models. In Section 3.1, we present theoretical arguments for using the score vector as the basis for concept drift monitoring. In Section 3.2, we explain empirical counterparts to the theoretical arguments, including interpretations when the parametric model structure is only an approximation to the true $P(Y|X)$, and connections to stochastic gradient descent (SGD) algorithms. In Section 3.3, we develop the multivariate EWMA (MEWMA) procedure for monitoring the mean of the score vector, and in Section 3.4 we discuss various implementation issues and how to handle high-dimensional θ and regularized models. Algorithm 1 provides high-level pseudocode for the entire procedure, details of which are explained in the subsequent sections.

Algorithm 1 Score-Based Concept Drift Monitoring and Diagnostics

- 1: **Step 1: Retrospective Analysis**
 - 2: Set *Stable* = *False*
 - 3: **while** *Stable* = *False* **do**
 - 4: Fit a supervised learning model to training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$.
 - 5: Calculate score vectors (Equation (1)) for $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ and apply the MEWMA (Section 3.3).
 - 6: **if** the MEWMA indicates the training data are stable **then**
 - 7: Set *Stable* = *True*
 - 8: **else**
 - 9: Discard the portion of training data over which the drift occurred and denote the remaining subset of training data as $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ with n reduced.
 - 10: **end if**
 - 11: **end while**
 - 12: **Step 2: Concept Drift Monitoring**
 - 13: Partition the stable dataset from Step 1, $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, into $\mathcal{D}_1 := \{\mathbf{x}_i, y_i\}_{i=1}^{n_1}$ and $\mathcal{D}_2 := \{\mathbf{x}_i, y_i\}_{i=n_1+1}^n$.
 - 14: Fit a new parametric supervised learning model \mathcal{M} to \mathcal{D}_1 .
 - 15: Calculate the score vectors (Equation (1)) for \mathcal{D}_2 using \mathcal{M} .
 - 16: Compute MEWMA statistics $\{\mathbf{z}_t\}_{t=n_1+1}^n$ and $\{T_t^2\}_{t=n_1+1}^n$ (Equation (8)) for \mathcal{D}_2 using \mathcal{M} .
 - 17: For a desired false alarm rate α , set UCL as the $1 - \alpha$ sample quantile of $\{T_t^2\}_{t=n_1+1}^n$.
 - 18: For Phase-II “on-line” data $\tilde{\mathcal{D}} := \{\tilde{\mathbf{x}}_i, \tilde{y}_i\}_{i=1}^{\tilde{n}}$, calculate the score vectors and apply the MEWMA with the UCL from Line 17 to monitor and detect concept drift.
 - 19: **Step 3: Concept Drift Diagnosis After Concept Drift Detection**
 - 20: Plot the T_t^2 statistics over time for the data over which concept drift was detected.
 - 21: Decouple the concept drift in the score vectors by calculating the Fisher-decoupled score vectors $\tilde{\mathbf{s}}_t := \hat{\mathbb{I}}^{-1}(\theta^{(0)})\mathbf{s}_t$ (see line following Equation (12)).
 - 22: Plot the univariate EWMA charts for each component of Fisher-decoupled score vectors $\tilde{\mathbf{s}}_t$.
-

3.1. The Score Vector as a Basis for Concept Drift Monitoring

Supervised learning is used to approximate an underlying conditional response distribution, which, if parametric, can be denoted as $P(Y|X; \theta)$. For example, for a regression neural network, the conditional mean $E[Y|X; \theta]$ is modeled as a neural network, and Y is assumed to be its conditional mean plus (typically) a Gaussian error; or for a classification neural network, Y is multinomial with class probabilities that are modeled as a neural network. Fitting the model then entails estimating the parameters θ by maximizing the log-likelihood, which can be viewed as an empirical approximation to $E_{\theta^{(0)}}[\log P(Y|X; \theta)]$, where $\theta^{(0)}$ denotes the true parameters and the subscript $\theta^{(0)}$ on the expectation operator indicates that it is with respect to $Y|X$ following the distribution $P(Y|X; \theta^{(0)})$. This is generally valid because of Shannon's Lemma (Shannon 1948), which states that if the model is correct and identifiable, the true parameter vector $\theta^{(0)}$ uniquely maximizes the expected log-likelihood, $E_{\theta^{(0)}}[\log P(Y|X; \theta)]$. For ease of exposition, we develop the approach for the setting in which maximum likelihood estimation (MLE) is used to fit the model. However, for complex models some form of regularization is virtually always used as a modification of the log-likelihood function. With minor modifications discussed in Section 3.4, the approach still applies to regularized MLE.

Given a parametric distribution or marginal likelihood function $P(Y|X; \theta)$ for an individual observation (X, Y) , and assuming we have an iid training dataset $\{(x_i, y_i)\}_{i=1}^n$ drawn from this distribution, the score vector for (x_i, y_i) is defined as

$$s(\theta; (x_i, y_i)) = \nabla_{\theta} \log P(y_i|x_i; \theta) \quad (1)$$

where ∇_{θ} is the derivative operator with respect to the vector of parameters, θ . From fundamental theory (Proposition 3.4.4 from Bickel and Doksum 2015), under certain regularity conditions, if the model is correct and we have a true parameter vector $\theta^{(0)}$, the expected score vector evaluated at $\theta^{(0)}$ is zero:

$$\begin{aligned} E_{\theta^{(0)}}[s(\theta^{(0)}; (X, Y))|X] \\ = \int s(\theta^{(0)}; (X, Y = y))P(Y = y|X; \theta^{(0)})dy = \mathbf{0}. \end{aligned} \quad (2)$$

In other words, the conditional expectation of the score vector evaluated at $\theta^{(0)}$ is identically zero.

Concept drift is defined as a change in $P(Y|X)$, and in our parametric model setting it equates to a change in the parameters of the conditional distribution $P(Y|X; \theta)$. This and Equation (2) suggest that a general approach for monitoring for concept drift is to monitor for changes in the mean of the sample score vector (1). In particular, with no concept drift, we have $\theta = \theta^{(0)}$, so that the sample score vector $s(\theta^{(0)}; (X, Y))$ is zero-mean from the above discussion. In contrast, if there is concept drift, this means θ has changed from $\theta^{(0)}$ to some other value $\theta^{(1)} \neq \theta^{(0)}$, in which case the new mean of the score vector $E_{\theta^{(1)}}[s(\theta^{(0)}; (X, Y))|X] \neq \mathbf{0}$ is no longer zero (for small changes in θ and under certain identifiability assumptions). As a preview to how we will apply the above concepts in practice (see Section 3.2 for details), a set of estimated parameters from a training dataset will take the place of $\theta^{(0)}$, and to

whatever the current values of θ have drifted will take the place of $\theta^{(1)}$.

To illustrate the above arguments more concretely, consider the generalized linear model (GLM) (Nelder and Wedderburn 1972), which encompasses many models commonly used in applications. The canonical form of the GLM model and score vector are

$$\begin{aligned} P(Y|X; \theta, \phi) &= \exp \left\{ \frac{Y\psi(X; \theta) - b(\psi(X; \theta))}{a(\phi)} + c(Y; \phi) \right\}, \text{ and} \\ s(\theta; (X, Y)) &= \frac{1}{a(\phi)} (Y - b'(\psi(X; \theta))) \nabla_{\theta} \psi(X; \theta), \end{aligned} \quad (3)$$

for functions $b(\cdot)$, $\psi(\cdot)$, $a(\cdot)$, and $c(\cdot)$ and dispersion parameter ϕ . The conditional mean is $E[Y|X] = b'(\psi(X; \theta))$, where the derivative is with respect to ψ , and the canonical link function $g(\cdot)$ is defined such that $g(E[Y|X]) = \psi(X; \theta) (= X^T \theta)$ for the traditional GLM). Using a Taylor expansion, the expected score vector for $\psi(X; \theta) = X^T \theta$ with a small parameter change $\Delta \theta = \theta^{(1)} - \theta^{(0)}$ is

$$\begin{aligned} E_{\theta^{(1)}}[s(\theta^{(0)}; (X, Y))|X] &\approx \frac{1}{a(\phi)} b''(\psi(X; \theta^{(0)})) \nabla_{\theta} \psi(X; \theta^{(0)}) \\ &\quad \nabla_{\theta}^T \psi(X; \theta^{(0)}) \Delta \theta \\ &= \frac{1}{a(\phi)} b''(\psi(X; \theta^{(0)})) X X^T \Delta \theta. \end{aligned} \quad (4)$$

It is known that $\text{var}[Y|X] = a(\phi)b''(\psi(X; \theta))$, so that $b''(\psi(X; \theta)) > 0$ always holds if Y is not a deterministic function of X . If we further take the expectation of Equation (4) with respect to the distribution of X , then the resulting matrix that premultiplies $\Delta \theta$ in (4) is always positive definite if the distribution of X is not degenerate. In other words, whenever concept drift happens (i.e., $\Delta \theta \neq \mathbf{0}$), the mean of the score vector is nonzero.

Notice that this approximation (4) does not require $\psi(X; \theta)$ to be a linear function of θ , if certain identifiability conditions on θ are satisfied. For example, $\psi(X; \theta)$ can be a neural network, in which case θ is the vector of weights and bias parameters for the neural net. In later studies with real and simulated datasets in Section 5 and Section C, supplementary materials, we empirically show that our score-based approach is also effective for neural networks.

In summary, for GLM-type response models with linear or nonlinear $\psi(\cdot)$ that satisfy certain identifiability conditions on θ , the mean of the score vector changes if and only if $P(Y|X; \theta)$ changes, at least for sufficiently small changes in θ (for large changes in θ in most practical settings, the mean of the score vector is likely to change substantially, although it is not clear if any mathematical theory exists to guarantee this). This is an important property that provides the theoretical basis for our score-based concept drift monitoring and diagnosis. Note that error-based methods do not enjoy this property, that is, $P(Y|X; \theta)$ can change in ways that do not change the error rate (see Section B.1, supplementary materials for a detailed discussion and a concrete example).

3.2. Interpretations with Empirical Data and Incorrect Models

The zero-mean property $E_{\theta^{(0)}}[\log P(Y|X; \theta^{(0)})] = \mathbf{0}$ of the score vector and the uniqueness of the parameters θ that maximize the

expected log-likelihood $E_{\theta^{(0)}}[\log P(Y|X; \theta)]$ are guaranteed to hold only when the model is correct; that is, when the supervised learning model $P(Y|X; \theta)$ is of the same structure as the true predictive relationship $P(Y|X)$. Recalling the adage that “All models are wrong, but some are useful” (Box 1976), one might wonder to what extent the results in the previous section are applicable when the structure of the model $P(Y|X; \theta)$ differs from the true $P(Y|X)$. A related question is what should we take to be the empirical counterpart to $E_{\theta^{(0)}}[s(\theta^{(0)}; (X, Y))]$ when $\theta^{(0)}$ is replaced by its estimate from a sample of training data, and the expectation is replaced by a sample average over a set of new data or over the same training data. We address both of these issues in this section and also relate the empirical counterpart to SGD for computational reasons.

Regardless of whether the model structure is correct, in analogy with Equation (2) we always have

$$\begin{aligned} \hat{E}_{(0)}[s(\hat{\theta}^{(0)}; (X, Y))] &:= \frac{1}{n} \sum_{i=1}^n s(\hat{\theta}^{(0)}; (x_i, y_i)) = \mathbf{0}, \text{ where} \\ \hat{\theta}^{(0)} &:= \operatorname{argmax}_{\theta} \hat{E}_{(0)}[\log P(Y|X; \theta)] \\ &:= \operatorname{argmax}_{\theta} \frac{1}{n} \sum_{i=1}^n \log P(y_i|x_i; \theta), \end{aligned} \quad (5)$$

the operator $\hat{E}_{(0)}$ denotes a sample average over the training data $\{(x_i, y_i)\}_{i=1}^n$, and $\hat{\theta}^{(0)}$ is the MLE of $\theta^{(0)}$ for the training data. That is, when we fit a model using MLE, the gradient of the training log-likelihood is identically zero, that is, $\nabla_{\theta} \hat{E}_{(0)}[\log P(Y|X; \theta)]|_{\theta=\hat{\theta}^{(0)}} := \nabla_{\theta} \frac{1}{n} \sum_{i=1}^n \log P(y_i|x_i; \theta)|_{\theta=\hat{\theta}^{(0)}} = \frac{1}{n} \sum_{i=1}^n s(\hat{\theta}^{(0)}; (x_i, y_i)) = \hat{E}_{(0)}[s(\hat{\theta}^{(0)}; (X, Y))] = \mathbf{0}$, even if the model is not the correct structure. Thus, (5) is the empirical counterpart of (2) with the estimated $\hat{\theta}^{(0)}$ taking the place of the true $\theta^{(0)}$.

Now suppose the true predictive relationship $\tilde{P}(\tilde{Y}|\tilde{X})$ changes from $P(Y|X)$ over some new set of data $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^{\tilde{n}}$. In this case a different set of parameters $\hat{\theta}^{(1)} \neq \hat{\theta}^{(0)}$ for the same supervised learning model structure $P(Y|X; \theta)$ will generally provide a better fit to the new data than did $\hat{\theta}^{(0)}$, where

$$\begin{aligned} \hat{\theta}^{(1)} &:= \operatorname{argmax}_{\theta} \hat{E}_{(1)}[\log P(\tilde{Y}|\tilde{X}; \theta)] \\ &:= \operatorname{argmax}_{\theta} \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \log P(\tilde{y}_i|\tilde{x}_i; \theta), \end{aligned} \quad (6)$$

and the operator $\hat{E}_{(1)}$ denotes the sample average over the new data. Thus, the gradient $\nabla_{\theta} \hat{E}_{(1)}[\log P(\tilde{Y}|\tilde{X}; \theta)]|_{\theta=\hat{\theta}^{(0)}} := \nabla_{\theta} \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \log P(\tilde{y}_i|\tilde{x}_i; \theta)|_{\theta=\hat{\theta}^{(0)}} = \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} s(\hat{\theta}^{(0)}; (\tilde{x}_i, \tilde{y}_i)) = \hat{E}_{(1)}[s(\hat{\theta}^{(0)}; (\tilde{X}, \tilde{Y}))]$ of the log-likelihood for the new data (but with the gradient evaluated at the original estimate $\hat{\theta}^{(0)}$) will generally differ from zero. The more $\tilde{P}(\tilde{Y}|\tilde{X})$ changes from $P(Y|X)$, the more we expect $\hat{\theta}^{(1)}$ to differ from $\hat{\theta}^{(0)}$, and the more we expect the new average score vector $\hat{E}_{(1)}[s(\hat{\theta}^{(0)}; (\tilde{X}, \tilde{Y}))]$ to differ from $\mathbf{0}$. This provides the justification for our score-based

concept drift monitoring approach, which tracks the mean of the score vector $s(\hat{\theta}^{(0)}; (X_i, Y_i)) = \nabla_{\theta} \log P(Y_i|X_i; \theta)|_{\theta=\hat{\theta}^{(0)}}$ to detect and analyze changes in it. In Section B.3, supplementary materials, we provide an example investigating the effect of fitting an incorrect model structure on the performance of our score-based approach and demonstrate that it still performs substantially better than the residual-based EWMA in this case.

If the supervised learning model $P(Y|X; \theta)$ is of the same structure as the true predictive relationship $P(Y|X)$, both $P(Y|X)$ and $P(X)$ are constant across the training data $\{(x_i, y_i)\}_{i=1}^n$, and $n \rightarrow \infty$, then under some regularity conditions the MLE $\hat{\theta}^{(0)}$ is consistent and $\hat{E}_{(0)}[s(\hat{\theta}^{(0)}; (X, Y))] \rightarrow E_{\theta^{(0)}}[s(\hat{\theta}^{(0)}; (X, Y))] \rightarrow E_{\theta^{(0)}}[s(\theta^{(0)}; (X, Y))] = \mathbf{0}$. In this case, there is no distinction between the theoretical version of the score-based monitoring arguments and their empirical version discussed above. With large n , SGD is often used to fit models to the training data, which involves approximating the gradient of the log-likelihood function using individual training observations or mini-batches of training observations at each iteration of the optimization algorithm. The SGD estimator of $\theta^{(0)}$ converges to the batch MLE under certain conditions involving step size and other considerations (see, e.g., Theorem 4.7 of Bottou et al. 2018). In this case, under the same asymptotic conditions stated above, the SGD estimator $\hat{\theta}_{\text{SGD}}$ is also consistent and $\hat{E}_{(0)}[s(\hat{\theta}_{\text{SGD}}; (X, Y))] \rightarrow \mathbf{0}$ as $n \rightarrow \infty$.

The score-based approach does not add much extra cost to the current framework of training and using models. For retrospective analysis, the sample score vectors are a byproduct of the SGD (or related optimizers, e.g., ADAM (Kingma and Ba 2014)), since the mini-batch gradients are of the form $\nabla_{\theta} \sum_{i=1}^m \log P(y_i|x_i; \theta) = \sum_{i=1}^m s(\theta; (x_i, y_i))$, where m is the batch size. Even for prospective analysis, prediction of new data usually partially calculates the score vectors. For example, prediction for neural networks requires forward-propagation, and another backward-propagation in memory would generate the score vector. Thus, we do not need much extra computation to apply the score-based approach.

With finite training data size n and finite new sample sizes \tilde{n} for monitoring (including $\tilde{n} = 1$), noise in $s(\theta; (x_i, y_i))$ and its sample averages must be considered. In particular, we need to distinguish by how much $s(\hat{\theta}^{(0)}; (x_i, y_i))$ (or its sample average over some moving time window) should differ from $\mathbf{0}$ before we conclude that $P(Y|X)$ has changed. The MEWMA monitoring strategy in the next section is designed to distinguish this type of noise from legitimate changes in $P(Y|X)$. Moreover, for models fitted with regularization, the gradient of the log-likelihood itself is not zero over the training data, because the regularization penalty is included in the optimization objective function. Regardless, the score-based monitoring method can still be applied with the minor modification to the score vectors discussed in Section 3.4.

3.3. An MEWMA Approach for Monitoring the Score Vector

As discussed in the previous sections, monitoring for concept drift reduces to monitoring for changes over time in the mean

of the score vector. Among other challenges, this requires distinguishing between noise in the score vectors for individual observations versus an actual mean change. The MEWMA (Lowry et al. 1992) has emerged as one of the most effective techniques for detecting changes in the mean vectors in general, and in this section we develop it for monitoring the mean of the score vector.

The MEWMA at time t , denoted by \mathbf{z}_t , is defined recursively (for $t = 1, 2, \dots$) via:

$$\mathbf{z}_t = \lambda \mathbf{s}_t + (1 - \lambda) \mathbf{z}_{t-1}, \quad (7)$$

where \mathbf{s}_t is the to-be-monitored random vector at time t , which in our case is the score vector $\mathbf{s}_t := \mathbf{s}(\hat{\boldsymbol{\theta}}^{(0)}; (\mathbf{x}_t, y_t))$; λ is a weighting parameter; and \mathbf{z}_0 can be initialized as the sample mean of \mathbf{s}_t over some small initial batch of observations. An equivalent expression for the recursive relationship (7) is $\mathbf{z}_t = \lambda \sum_{\tau=1}^t (1 - \lambda)^{t-\tau} \mathbf{s}_\tau + (1 - \lambda)^t \mathbf{z}_0$, which gives exponentially decaying weights to the historic observations up to \mathbf{s}_t .

Since \mathbf{s}_t and \mathbf{z}_t are vectors, and we desire to detect changes in the mean of \mathbf{s}_t in any direction, our MEWMA approach monitors the Hotelling T^2 (Hotelling 1947) statistic

$$T_t^2 = (\mathbf{z}_t - \bar{\mathbf{s}})^T \widehat{\boldsymbol{\Sigma}}^{-1} (\mathbf{z}_t - \bar{\mathbf{s}}), \quad (8)$$

where $\widehat{\boldsymbol{\Sigma}}$ and $\bar{\mathbf{s}}$ are the covariance matrix and sample mean vector of \mathbf{s}_t , respectively, estimated from the training and Phase-I data, as described in later parts of this section. In the statistical process control (SPC) literature, Phase-I refers to a retrospective analysis of previously collected data, with the goals of statistically characterizing the process behavior and establishing control limits to be used in Phase-II. Phase-II refers to a prospective analysis in which one monitors (usually in real-time) new observations as they arrive with the goal of detecting any change in the process behavior as quickly as possible after it occurs. If T_t^2 at some t exceeds a specified upper control limit (UCL) determined as described later, this is taken to be an indication that $P(Y|X)$ in the time-vicinity of observation t has changed from what it was when $\hat{\boldsymbol{\theta}}^{(0)}$ was estimated. The convention that we recommend and that we have used in all of our later examples in the article is to estimate $\bar{\mathbf{s}}$ from only the Phase-I data and $\widehat{\boldsymbol{\Sigma}}$ from only the training data. The reason we do not use any Phase-I data to estimate $\widehat{\boldsymbol{\Sigma}}$ is that the control limits are determined from the Phase-I data, and estimating $\widehat{\boldsymbol{\Sigma}}$ from the same Phase-I data can result in a form of overfitting and a UCL that is too small, which results in too many false alarms. The reason we do not use any training data to estimate $\bar{\mathbf{s}}$ is that if there is concept drift between the training and Phase-I data and we use the training data to estimate $\bar{\mathbf{s}}$, then $\mathbf{z}_t - \bar{\mathbf{s}}$ will not be zero-mean over the Phase-I data, which can result in a UCL that is too large, which reduces the sensitivity to concept drift in Phase-II. Note that if there is concept drift between the training and Phase-I data, this can be detected via our retrospective analysis, as illustrated in the later examples.

The score-based concept drift monitoring framework is divided into three steps, which will be illustrated in more detail in the later examples. In Step 1, we collect a batch of data in time order, $\{\mathbf{x}_i, y_i\}_{i=1}^m$, where m denotes the sample size of this batch. Then, after fitting a preliminary supervised learning model to these data, a retrospective analysis is conducted by applying the

MEWMA to these same data. If the MEWMA indicates these data are stable and no significant concept drift is detected, we proceed to Step 2. On the other hand, if the MEWMA detects significant concept drift in these data then one can discard a portion of the data prior to the concept drift in an attempt to ensure that the remaining data are stable. For other kinds of concept drift, like gradual drifting or seasonality as in the bike-sharing example in Section B.4, supplementary materials, we can try to modify the model by incorporating other covariates to reduce the concept drift. This retrospective analysis would then be repeated (perhaps multiple times) on the remaining data to verify whether it was stable versus experienced concept drift. When the remaining data are deemed stable, then we proceed to Step 2.

Let $n \leq m$ denote the size of the stable data from Step 1 and denote these data by $\mathcal{D} := \{\mathbf{x}_i, y_i\}_{i=1}^n$. The purpose of Step 2 is to establish the UCL and then to prospectively monitor new data for concept drift, as the new data are collected. To establish the UCL, we divide \mathcal{D} into two parts: $\mathcal{D}_1 := \{\mathbf{x}_i, y_i\}_{i=1}^{n_1}$ and $\mathcal{D}_2 := \{\mathbf{x}_i, y_i\}_{i=n_1+1}^n$. The first part, \mathcal{D}_1 , is used to retrain the parametric supervised learning model of interest, and the second part, \mathcal{D}_2 , is used to establish the UCL in Phase-I. In the Phase-I analysis, the score vectors for the \mathcal{D}_2 data are computed based on the model fitted to the \mathcal{D}_1 data. The MEWMA statistics $\{\mathbf{z}_t\}_{t=n_1+1}^n$ and T^2 statistics $\{T_t^2\}_{t=n_1+1}^n$ are then computed for the \mathcal{D}_2 data. The user specifies a desired false alarm probability α (e.g., $\alpha = 0.0001$, $\alpha = 0.001$, etc.), and the UCL is set as the $1 - \alpha$ sample quantile of $\{T_t^2\}_{t=n_1+1}^n$. After computing the UCL in Phase-I, in Phase-II the MEWMA with this UCL is applied prospectively to the sample score vectors for a new set of “on-line” data $\tilde{\mathcal{D}} := \{\tilde{\mathbf{x}}_i, \tilde{y}_i\}_{i=1}^{\tilde{n}}$ (i.e., new data for which the fitted supervised learning model is to be used to predict the response, e.g., new credit card applicants who are being scored for credit risk by the fitted model). The purpose of the Phase-II analysis is to detect as quickly as possible if concept drift occurs in the new data, so that the supervised learning model can be updated accordingly.

The purpose of Step 3 is to conduct a diagnostic analysis to help determine the nature of the concept drift, if any drift is detected in Phase-II. This involves plotting the T_t^2 statistic versus t over the dataset of interest. Note that the T^2 statistic aggregates changes in any of the components of the score vector into a single scalar statistic. To provide richer diagnostic information and help understand which parameters have changed, analogous univariate EWMA control charts for (transformed) individual components of the score vector should also be constructed. We describe these univariate control charts in Section 4.

The Step 3 diagnostic procedure can also be used in a purely retrospective analysis following Step 1, if it is desired to understand the nature of the nonstationarity in $P(Y|X; \boldsymbol{\theta})$ over any set of training data to which a supervised learning model is fitted.

For the other methods based on scalar metrics (e.g., error rates) to which we compare our approach and for the transformed components of the score vector that we describe in Section 4, we use univariate EWMA (Roberts 1959) control charts because the EWMA in Equation (7) is a scalar in this case. One main difference between a univariate EWMA and an MEWMA

is that the MEWMA T^2 statistic aggregates the changes in the mean vector into a single scalar statistic, and only larger values of T^2 indicate a change in the mean from \bar{s} . Thus, only an UCL is needed in the MEWMA. In contrast, the univariate version of the EWMA in Equation (7) to detect changes in the mean of a scalar random variable are two-sided in nature and have a lower control limit (LCL) as well as a UCL. Changes in the mean are indicated by the univariate EWMA statistic falling either below the LCL or above the UCL. After concept drift is detected, users can also attempt to validate whether the alarm was truly due to concept drift or to some other phenomenon such as covariate drift using methods discussed in (Raza, Prasad, and Li 2015) and (Mejri, Limam, and Weihs 2018; Mejri et al. 2021), for example.

3.4. Handling High-Dimensional and Regularized Models

Many machine learning models are becoming increasingly complex, and some state-of-the-art models can have millions of parameters, for example, deep neural networks. With such high-dimensional parameters, the sample covariance matrix $\hat{\Sigma}$ in Equation (8), which must be inverted, is very likely to be singular or close to it. For example, when the sample size of our training data (denoted by n_1 as in Section 3.3) is smaller than the number of parameters, the sample covariance is always singular.

With very high-dimensional θ like in convolutional neural networks (CNNs), one simple solution is to monitor the elements of the score vector corresponding to only a specific subset of the parameters, which we further discuss in the next section. An alternative solution is to modify the covariance matrix to circumvent the problem of inverting a singular or poorly-conditioned matrix. One way to accomplish this is to add a nugget parameter (borrowing terminology from Gaussian process modeling) to all diagonal entries of $\hat{\Sigma}$. Specifically, for some $\delta > 0$, we substitute $\tilde{\Sigma} := \hat{\Sigma} + \delta \mathbf{I}$ for $\hat{\Sigma}$ in Equation (8). To understand the effect of this nugget parameter, denote the eigen-decomposition of the sample covariance matrix as $\hat{\Sigma} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$, where \mathbf{Q} is an orthogonal matrix of eigenvectors and $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_q\}$, (where $q = \dim(\theta)$), is a diagonal matrix of eigenvalues, and suppose the eigenvalues are arranged in nonincreasing order. Then, we can write the approximated sample covariance matrix as $\tilde{\Sigma} = \mathbf{Q}\tilde{\mathbf{\Lambda}}\mathbf{Q}^T$, where $\tilde{\mathbf{\Lambda}} := \mathbf{\Lambda} + \delta \mathbf{I}$. Using $\tilde{\Sigma}$ in place of $\hat{\Sigma}$ in Equation (8) suppresses unimportant directions of variation. The resulting MEWMA with this modified covariance matrix would not have the issue of ill-conditioning.

Another way to accomplish this is to use a pseudo-inverse of $\hat{\Sigma}$ in Equation (8) instead of its actual inverse, after dropping any eigenvalues λ_i with $\lambda_1/\lambda_i > \gamma$ for some specified maximum condition number γ . Define the diagonal matrix $\mathbf{\Lambda}^- = \text{diag}\{\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_k^{-1}, 0, \dots, 0\}$ where λ_k is the smallest eigenvalue with λ_1/λ_k no greater than γ . Then, the pseudo-inverse is defined as $\hat{\Sigma}^- = \mathbf{Q}\mathbf{\Lambda}^-\mathbf{Q}^T$. This is equivalent to applying principal component analysis to s_t and retraining only the principal component directions in which the variation in s_t is not negligible. In our approach we used the nugget parameter instead of the pseudo-inverse, so as to enable detection in changes in the mean of the score vector in all directions. For all

of our examples, we chose the nugget parameter as the smallest value of δ for which the condition number of $\hat{\Sigma} + \delta \mathbf{I}$ was no larger than 10^4 .

Another related issue is regularization of complex models, which is almost always required to combat overfitting. A regularization term $J(\theta)$ is used to penalize complex models and large parameters by fitting the model to minimize the regularized loss function $l(\theta) = -\sum_{i=1}^n \log P(y_i|x_i; \theta) + J(\theta) = -\sum_{i=1}^n \left\{ \log P(y_i|x_i; \theta) - \frac{J(\theta)}{n} \right\}$, instead of the MLE loss function $-\sum_{i=1}^n \log P(y_i|x_i; \theta)$. For example, for L_2 regularization, we use $J(\theta) = \frac{c}{2} \|\theta\|_2^2$, where $c > 0$ is the regularization parameter. The gradient of the regularized loss function becomes $\nabla_{\theta} l(\theta) = -\sum_{i=1}^n \left\{ s(\theta; (x_i, y_i)) - \frac{\nabla_{\theta} J(\theta)}{n} \right\}$. Suppose we redefine the score vector as $s(\theta; (x_i, y_i)) \leftarrow s(\theta; (x_i, y_i)) - \frac{\nabla_{\theta} J(\theta)}{n}$ and view this as the regularized score vector. Using the same arguments as in the unregularized situation, it follows that the regularized score vector is zero-mean when there is no concept drift, and so it still makes sense to monitor for changes in the mean of the regularized score vector. The MEWMA and Hotelling T^2 computations in the concept drift monitoring procedure are exactly the same but with the score vector replaced by the regularized score vector. For L_2 regularization, this amounts to replacing $s(\theta; (x_i, y_i)) \leftarrow s(\theta; (x_i, y_i)) - \frac{c}{n} \theta$. Since the score vector and the regularized score vector only differ by the constant vector $\frac{\nabla_{\theta} J(\theta)}{n}$, the monitoring statistic remains unchanged when regularization is used, because the centered MEWMA $z_t - \bar{s}$ and the covariance matrix $\hat{\Sigma}$ in Equation (8) are translation-invariant. Another way to view the above is that we are replacing the likelihood with its Bayesian counterpart $P(Y|X; \theta) \exp \left\{ -\frac{J(\theta)}{n} \right\}$, which incorporates a prior distribution (a Gaussian prior for L_2 regularization) on the parameters.

Note that for L_1 regularization, which is commonly used for linear and logistic regression modeling (i.e., LASSO regression), $J(\theta)$ is not differentiable with respect to the parameters that have been regularized to zero. Similar to what was described above for the case of high-dimensional θ , one could take the score vector for only the subset of parameters that are nonzero. In this case, it is not clear whether changes in the parameters that were regularized to zero would be reflected in a change in the mean of the partial score vector. This warrants future study. In the Section D, supplementary materials, we discuss additional extensions of the approach to supervised learning models with extremely high-dimensional parameters (e.g., deep learning networks) and to nondifferentiable models (e.g., tree-based models).

4. Diagnostics and Enhanced Monitoring of Individual Components

The Hotelling T^2 statistic aggregates mean shifts in the components of the score vector into a single scalar statistic. In order to provide diagnostic insight into the nature of the change in $P(Y|X; \theta)$ (e.g., which parameters have changed) and/or to enhance the ability of the procedure to detect some changes, it is helpful to monitor individual components of the score vector $s(\theta^{(0)}; (X, Y))$ or the transformed version described below. To construct a univariate EWMA chart for the j th component ($j = 1, 2, \dots, q$) of the score vector, denoted by

$s_{j,t} := [s(\hat{\theta}^{(0)}; \mathbf{x}_t, y_t)]_j$, the univariate counterpart of Equation (7) is

$$z_{j,t} = \lambda s_{j,t} + (1 - \lambda) z_{j,t-1}. \quad (9)$$

In this case, $z_{j,t}$ is plotted directly on the univariate EWMA chart with both a LCL and UCL. The chart signals a change in the mean of the j th component at observation t if $z_{j,t}$ falls either below the LCL or above the UCL. If λ is small, $z_{j,t}$ is approximately normal by the central limit theorem, and one can set $\{LCL_j, UCL_j\} = \bar{z}_j \pm z_{\alpha/2} SD[z_j]$, where $SD[z_j]$ and \bar{z}_j denote the standard deviation and sample average over the training data, $\{z_{j,t}\}_{t=1}^{n_1}$, and Phase-I data, of $\{z_{j,t}\}_{t=n_1+1}^n$, $z_{\alpha/2}$ is the upper $\alpha/2$ quantile of the standard normal distribution, and α is the desired false alarm rate. Alternatively, if the Phase-I sample size $n - n_1$ is sufficiently large, one can choose LCL_j and UCL_j directly as the lower and upper $\alpha/2$ quantiles of the empirical distribution of $\{z_{j,t}\}_{t=n_1+1}^n$.

If the goal is to diagnose and isolate which parameter(s) have changed, then it is preferable to replace the score components in the univariate EWMA by the components of a transformed version of the score vector discussed below, to effectively decouple the changes in the individual parameters. To illustrate this coupling, consider a linear Gaussian regression model $Y = \mathbf{X}^T \boldsymbol{\theta} + \epsilon$ with ϵ following a zero-mean Gaussian distribution with variance σ^2 , the score vector for which is $\mathbf{s}(\boldsymbol{\theta}; (\mathbf{X}, Y)) = \frac{(Y - \mathbf{X}^T \boldsymbol{\theta}) \mathbf{X}}{\sigma^2}$. With no concept drift (i.e., $\boldsymbol{\theta} = \boldsymbol{\theta}^{(0)}$), the mean of the score vector is $E_{\boldsymbol{\theta}^{(0)}}[\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))] = E[E_{\boldsymbol{\theta}^{(0)}}[\frac{(Y - \mathbf{X}^T \boldsymbol{\theta}^{(0)}) \mathbf{X}}{\sigma^2} | \mathbf{X}]] = E[\mathbf{0}] = \mathbf{0}$. After concept drift, suppose the parameters change to $\boldsymbol{\theta}^{(1)}$, and denote this change by $\Delta \boldsymbol{\theta} = \boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(0)}$. The mean of the score vector after the concept drift is $E_{\boldsymbol{\theta}^{(1)}}[\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))] = E[E_{\boldsymbol{\theta}^{(1)}}[\frac{(Y - \mathbf{X}^T \boldsymbol{\theta}^{(0)} + \mathbf{X}^T \Delta \boldsymbol{\theta}) \mathbf{X}}{\sigma^2} | \mathbf{X}]] = E[E_{\boldsymbol{\theta}^{(1)}}[\frac{\mathbf{X} \mathbf{X}^T \Delta \boldsymbol{\theta}}{\sigma^2} | \mathbf{X}]] = \frac{E[\mathbf{X} \mathbf{X}^T] \Delta \boldsymbol{\theta}}{\sigma^2}$. Thus, we can decouple the changes in the parameters by premultiplying the score vector mean by the inverse of the expected Fisher information matrix $\frac{E[\mathbf{X} \mathbf{X}^T]}{\sigma^2}$, that is, via the transformation $\Delta \boldsymbol{\theta} = \sigma^2 (E[\mathbf{X} \mathbf{X}^T])^{-1} E_{\boldsymbol{\theta}^{(1)}}[\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))]$.

For more general models and $P(Y|\mathbf{X}; \boldsymbol{\theta})$ we consider two different forms of decoupling, depending on whether one is interested in decoupling the components of $\boldsymbol{\theta}$ or (for regularized models) the components of the regularized version $\tilde{\boldsymbol{\theta}}$ of the true parameters, where $\tilde{\boldsymbol{\theta}}$ denotes the minimizer of the regularized version of the population loss function $E_{\boldsymbol{\theta}}[-\log P(Y|\mathbf{X}; \boldsymbol{\theta}) + J(\boldsymbol{\theta})/n]$. The appropriate decoupling equations are (see Section A, supplementary materials for derivations):

$$\Delta \boldsymbol{\theta} = \mathbb{I}^{-1}(\boldsymbol{\theta}^{(0)}) E_{\boldsymbol{\theta}^{(1)}}[\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))], \quad (10)$$

and

$$\Delta \tilde{\boldsymbol{\theta}} = \tilde{\mathbb{I}}^{-1}(\boldsymbol{\theta}^{(0)}) E_{\boldsymbol{\theta}^{(1)}}[\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))], \quad (11)$$

where $\mathbb{I}(\boldsymbol{\theta}^{(0)}) := E_{\boldsymbol{\theta}^{(0)}}[\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y)) \mathbf{s}^T(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))]$ is the expected Fisher information matrix (which can be estimated as the sample covariance matrix of $\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))$ over the training data prior to concept drift), $\tilde{\mathbb{I}}(\boldsymbol{\theta}^{(0)}) := -E_{\boldsymbol{\theta}^{(0)}}[\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T \{\log P(Y|\mathbf{X}; \boldsymbol{\theta}) - \frac{J(\boldsymbol{\theta})}{n}\}]_{\boldsymbol{\theta}=\tilde{\boldsymbol{\theta}}^{(0)}}$ is an alternative

expression for the expected Fisher information matrix with regularization which can be estimated as the sample average of the Hessian matrix $-\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T \{\log P(Y|\mathbf{X}; \boldsymbol{\theta}) - \frac{J(\boldsymbol{\theta})}{n}\}_{\boldsymbol{\theta}=\tilde{\boldsymbol{\theta}}^{(0)}}$ over the training data prior to concept drift. Here, $E_{\boldsymbol{\theta}^{(1)}}[\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))]$ can be estimated as the sample mean of $\mathbf{s}(\tilde{\boldsymbol{\theta}}^{(0)}; (\mathbf{X}, Y))$ over some relevant window of data following the concept drift. Alternatively, expressions for the population expectations of the Hessian matrices are available for some models like traditional GLMs (e.g., linear and logistic regression), in which case these expressions can be used instead of their sample versions. Notice that the two decoupling expressions (10) and (11) are different in general. However, if there is no regularization (i.e., if $J(\boldsymbol{\theta}) = 0$ as discussed in Section 3.4), then $\tilde{\boldsymbol{\theta}}^{(0)} = \boldsymbol{\theta}^{(0)}$, and it is well known that \mathbb{I} and $\tilde{\mathbb{I}}$ are equivalent expressions for the expected Fisher information matrix, in which case, (10) and (11) are the same. We will refer to either the transformation (10) or (11) as “Fisher decoupling.”

In light of the above Fisher decoupling results, as an alternative to the univariate EWMA (9) on the score components, we can construct univariate EWMA on the decoupled score components

$$z_{j,t} = \lambda \tilde{s}_{j,t} + (1 - \lambda) z_{j,t-1}, \quad (12)$$

where $\tilde{s}_{j,t}$ is the j th component of the decoupled score vector $\tilde{\mathbf{s}}_t := \tilde{\mathbb{I}}^{-1}(\boldsymbol{\theta}^{(0)}) \mathbf{s}_t$, where $\tilde{\mathbb{I}}(\boldsymbol{\theta}^{(0)})$ is an estimate of either $\mathbb{I}(\boldsymbol{\theta}^{(0)})$ or $\tilde{\mathbb{I}}(\boldsymbol{\theta}^{(0)})$. The LCL and UCL for (12) are determined analogously to those for (9), based on the empirical distribution of $\{\tilde{s}_{j,t}\}_{t=n_1+1}^n$.

In addition to providing diagnostic information on which parameter(s) have changed, and how they have changed, the decoupled univariate EWMA (12) have the following, additional benefit over just using the MEWMA. As discussed earlier, under fairly general conditions, the mean of $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))$ changes if and only if $P(Y|\mathbf{X}; \boldsymbol{\theta})$ changes. Thus, a change in $P(\mathbf{X})$ alone with no change in $P(Y|\mathbf{X}; \boldsymbol{\theta})$ will not cause the mean of $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))$, or the mean of \mathbf{z}_t in (7), to change. However, a change in $P(\mathbf{X})$ alone can cause the mean of the Hotelling T^2 statistic in (8) to increase (by changing the covariance of $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))$), thus, potentially leading to more frequent false alarms if the goal is to signal only when $P(Y|\mathbf{X}; \boldsymbol{\theta})$ changes. This is easy to see for the case of the linear Gaussian regression model $Y = \mathbf{X}^T \boldsymbol{\theta} + \epsilon$ with score vector $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y)) = \frac{(Y - \mathbf{X}^T \boldsymbol{\theta}^{(0)}) \mathbf{X}}{\sigma^2}$. If $P(Y|\mathbf{X}; \boldsymbol{\theta})$ does not change (i.e., $\boldsymbol{\theta}$ remains unchanged at $\boldsymbol{\theta}^{(0)}$), then we still have that $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y)) = \frac{\epsilon \mathbf{X}}{\sigma^2}$ is zero mean regardless of whether $P(\mathbf{X})$ changes, but the covariance of $\epsilon \mathbf{X}$ can obviously change.

The univariate EWMA are more robust to false alarms that are caused by a change in $P(\mathbf{X})$ and a resulting covariance change in $\mathbf{s}(\boldsymbol{\theta}^{(0)}; (\mathbf{X}, Y))$, because they chart the individual $z_{j,t}$ components directly and not some quadratic form like the T^2 statistic, and the mean of $z_{j,t}$ changes if and only if $P(Y|\mathbf{X}; \boldsymbol{\theta})$ changes. A variance change in $z_{j,t}$ typically would not increase the false alarm rate as much as if its mean changes. Moreover, by visual inspection of the univariate EWMA charts, it is easier to determine if an alarm was due to a mean change or to a variance change in $z_{j,t}$ than with the aggregated T^2 statistic in the MEWMA chart. Because a single aggregated monitoring

statistic has other advantages, in practice we recommend using both the MEWMA and the decoupled univariate EWMA, which we illustrate with the examples in the subsequent sections.

For classification problems (categorical Y), the fitted supervised learning model represents $P(Y|X; \theta)$ directly, and changes in θ equate to changes in $P(Y|X; \theta)$ if the parametric model structure is appropriate. However, for regression problems, the fitted supervised learning model typically (with squared error loss) represents only the conditional mean $E(Y|X; \theta)$, and so changes in θ equate to changes in $E(Y|X; \theta)$. If (for example) $\text{Var}(Y|X; \theta)$ changes but $E(Y|X; \theta)$ does not, then the mean of the score vector remains zero (since the same θ still provides optimal prediction of the response). However, similar to the situation described above in which $P(X)$ changes, changes in $\text{var}(Y|X; \theta)$ may result in changes in the covariance of the score vector and affect the false alarm rate of the multivariate EWMA. Developing diagnostic methods that can distinguish this is a topic that warrants future study.

5. An Illustrative Example: Credit Risk Modeling

In order to demonstrate the performance of our score-based concept drift approach and illustrate its usage, we now present a real data example, which involves credit risk modeling with data collected over the period 2003–2008 from a major financial company, during which time the subprime mortgage crisis happened. In Section B.1, supplementary materials we provide a simulation example where concept drift results in no change in expected error rate (and thus, error-based methods cannot detect it), but our score-based approach is able to effectively detect it. In Section B.2, supplementary materials, we provide more comprehensive simulation results comparing the drift detection performances of our score-based approach and alternative existing methods. The main conclusion is that our score-based approach achieves better performance across every example that we considered, often substantially so. In Section B.4, supplementary materials, we also provide another real data example, the Capital Bikeshare rental data from 2010 to 2020 during which time the “sharing economy” steadily expanded, to demonstrate retrospective and prospective analysis for a regression problem. In Section C, supplementary materials, we use simulation examples to investigate the performance of the Fisher decoupling approach in diagnosing which parameters have changed and how they have changed.

In the credit risk dataset, each row corresponds to a unique credit card customer of a major financial company, and we have 196,587 rows of data in total. The covariates for row/customer i include various customer information (x_i) available at the time the customer applies for the credit account, and the binary response (y_i) indicates whether the customer defaults within the first 9 months after opening the account. For the purposes of plotting various quantities over time, we associate each row i with a time stamp that is taken to be the day on which the response y_i first becomes available. Specifically, the set of all customers associated with a particular day, which we refer to as their “entry day” into the dataset, are those who opened their account within nine months of that day and defaulted on that day (in which case they are assigned $y_i = 1$), together with

those who opened their account exactly nine months prior to that day and did not default (in which case they are assigned $y_i = 0$). The data are imbalanced in the sense that there are 3536 observations with $y_i = 1$, which accounts for around 1.8% of the data. Consequently, when fitting the model to the training data, we assign a weight of 7.5 to the minority class (which corresponds to upsampling the minority class by a factor of 7.5) so that the weighted percentage of the minority class in training data is around 10%. For the training, validation, Phase-I, and Phase-II datasets, the classification errors and score vectors in Equation (1) and any quantities related to them are also using the same weighting scheme. That is, for each minority-class observation, its score vector in the MEWMA equation in Equation (7) is multiplied by 7.5, its Hessian and score vector are multiplied by 7.5 when computing the average Hessian (for Fisher decoupling) and average score vector (for use in the Hotelling T^2 statistic (8)), etc.

These data were originally considered in (Im et al. 2012), who focused on the same 10 covariates that we consider in this study: x_1 (credit risk score from an earlier model used by the company), x_2 (credit bureau risk score), x_3 (highest credit limit for open revolving credit accounts), x_4 (total balance on all open revolving credit accounts), x_5 (balance on the highest-utilization open revolving credit account), x_6 (credit limit on the highest-utilization open revolving credit account), x_7 (number of inquiries in the last 24 months), x_8 (balance on open mortgages), x_9 (categorical variable involving status of savings accounts), and x_{10} (number of inquiries in the last 24 months, excluding the last two weeks).

We applied concept drift detection to two supervised learning models: A logistic regression model and a neural network (with one hidden layer having 50 activation nodes). The main reason the company fits models of this nature was to score credit card applicants for risk (via their predicted probability of defaulting) at the time they apply. For the purpose of evaluating the concept drift detection, we trained and validated both models on the data from January 2003 to December 2005. The data from January 2006 to December 2006 are used as the Phase-I data to calculate control limits. The remainder of the data, from January 2007 to August 2008 are used as the Phase-II data to test the capability of detecting the concept drift due to the subprime mortgage crisis. Recall that the S&P 500 declined by more than 50% over a 15-month period between the end of December 2007 and the end of March 2009 (from 1478 to 683). The prevailing view is that the major root cause of the stock market decline was the subprime crisis, which had been gradually developing prior to that.

We applied our score-based MEWMA with the goal of detecting concept drift as far in advance of the beginning of the crash (December 2007) as possible. For comparison, we also applied an EWMA on the (weighted) binary classification error, which requires that the predicted probability $P[Y = 1|X]$ for each observation be converted to a classified label ($Y = 0$ or $Y = 1$). For this, we used the standard approach of classifying $Y = 1$ if the predicted probability exceeds some threshold, and $Y = 0$ otherwise. We chose this classification probability threshold to be 0.1057 to match the true positive rate and the true negative rate as closely as possible (alternatively, if costs of false positives and false negatives were known in advance, we

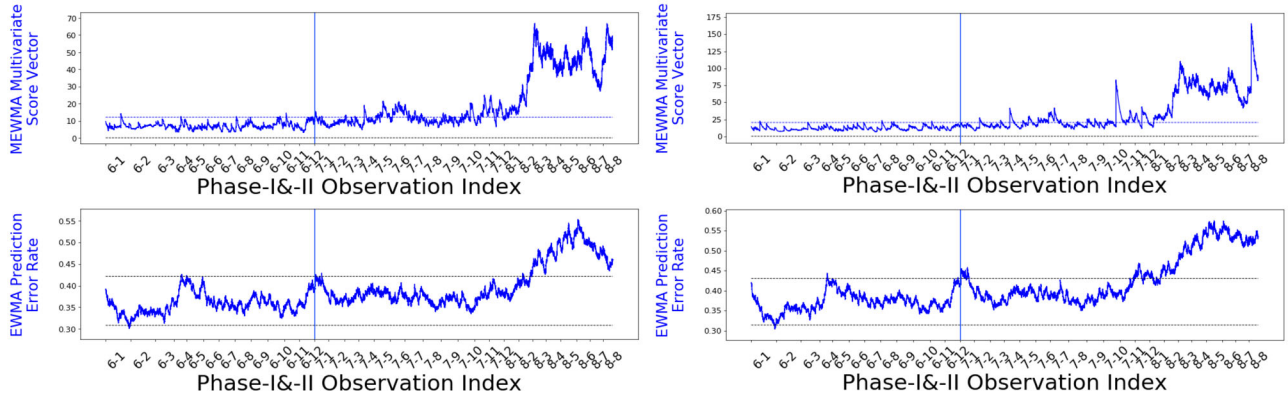


Figure 1. For the credit risk example, comparison of concept drift monitoring performance for our score-based MEWMA (the top plots) versus an EWMA on the classification error (the bottom plots). The left and right plots are for the logistic regression and neural network models, respectively. The horizontal dashed lines are the control limits. The blue vertical line is the boundary between the Phase-I and Phase-II data. The tilted numbers along the horizontal axes below each plot are the year-month indices (e.g., 6–1 stands for January 2006). The score-based MEWMA consistently signals beginning March 2007, which is nine months prior to the beginning of the stock market crash in December 2007. In contrast, the EWMA on the classification error does not consistently signal until around January 2008.

could have chosen the threshold to minimize the total misclassification cost). For both methods, we selected $\lambda = 0.001$, so that the effective window is $\frac{1}{\lambda} = 1000$ customers, which corresponds roughly to one week. Figure 1 shows the results of both concept drift monitoring methods for both models (logistic regression and neural network) over the Phase-I and Phase-II data.

As seen in Figure 1, for the logistic regression model, the score-based MEWMA consistently signals (i.e., the T^2 statistic consistently falls above the UCL) beginning around March 2007, which is about 9 months prior to the beginning of the economic crash in December 2007. In this sense, the score-based approach provides advanced warning that something has changed substantially in the predictive relationship between Y and X , which could have been indicative that serious economic changes were evolving. At the very least, it would have been an indication that the fitted credit risk scoring model was becoming obsolete and that more effective scoring could be achieved by updating the model. In contrast, the EWMA on the error rate does not consistently signal until much later, around January 2008, which is 10 months after the score-based MEWMA began to consistently signal and 1 month after the economic crash began. The results for the neural network model, which are shown in the right panels of Figure 1, are very similar.

Figure 2 shows several representative univariate component EWMA control charts (for $\theta_1, \theta_2, \theta_3, \theta_8$, and the intercept θ_0) for the logistic regression model over the Phase-I and Phase-II data. The MEWMA is also shown in the top row as a reference. The component EWMA control charts in the left column are for the original (coupled) score components (9), and those in the right column are for the Fisher-decoupled score components (12) using the Hessian matrix version of $\tilde{\mathbb{I}}(\theta^{(0)})$ for decoupling $\Delta\theta$ via (11). Rather than estimating the Fisher information matrix as the sample covariance matrix as in Equation (10), we used its theoretical Fisher information matrix expression for logistic regression.

Decoupling the score components reveals different patterns of concept drift than are seen in the original score components. In particular, the decoupled component EWMA for the intercept parameter (the bottom right plot) shows a clear upwards drift over the entire range of data, whereas this drift is not

apparent in the corresponding original component EWMA (the bottom left plot). The correlation between the intercept term and some of the other covariates has evidently conflated the drift in their corresponding parameters in the left column plots. The upwards drift in the decoupled intercept parameter is telling, as the intercept can be viewed as a regression-adjusted indicator of the overall default rate. Since the drift was upwards, this means that the intercept parameter increased over time, which means that the regression-adjusted likelihood of default (i.e., the default for applicants having the same covariate values) increased substantially over time. In fact, the concept drift may even be evident earlier in the decoupled component EWMA for θ_0 than in the MEWMA. For example, the decoupled component EWMA for θ_0 falls consistently above the center line (0) beginning back in the Phase-I data, around June 2006. Even though it is not above the UCL at this point, the fact that it is consistently above the center line is an indication that θ_0 has increased. In SPC control charting in general, users typically look for these types of patterns in the charts to signal changes in the mean of what is being charted (Montgomery 2007).

The decoupled component EWMA chart for θ_1 (the right column, the second from top) also shows a clear downwards drift in θ_1 , which means that x_1 (credit risk score from an earlier model) should be given less weight in predicting credit risk as time evolves. This makes sense, because the earlier model had been fitted to even earlier data. The decoupled component EWMA chart for θ_2 (the right column, the third from top) has a different trend that is also less extreme than for θ_1 . It is interesting that the component EWMA charts for θ_1 and θ_2 prior to decoupling (second and third plots in the left column) are far more similar than their decoupled counterparts. This is likely because x_1 and x_2 have a correlation coefficient of 0.45, and correlation between covariates conflates the concept drift in their coefficients, whereas the Fisher decoupling is intended to distinguish the concept drift in the coefficients.

6. Conclusions

Predictive models are trained on historical datasets, but due to potential changes in the conditional distribution $P(Y|X)$ (aka,

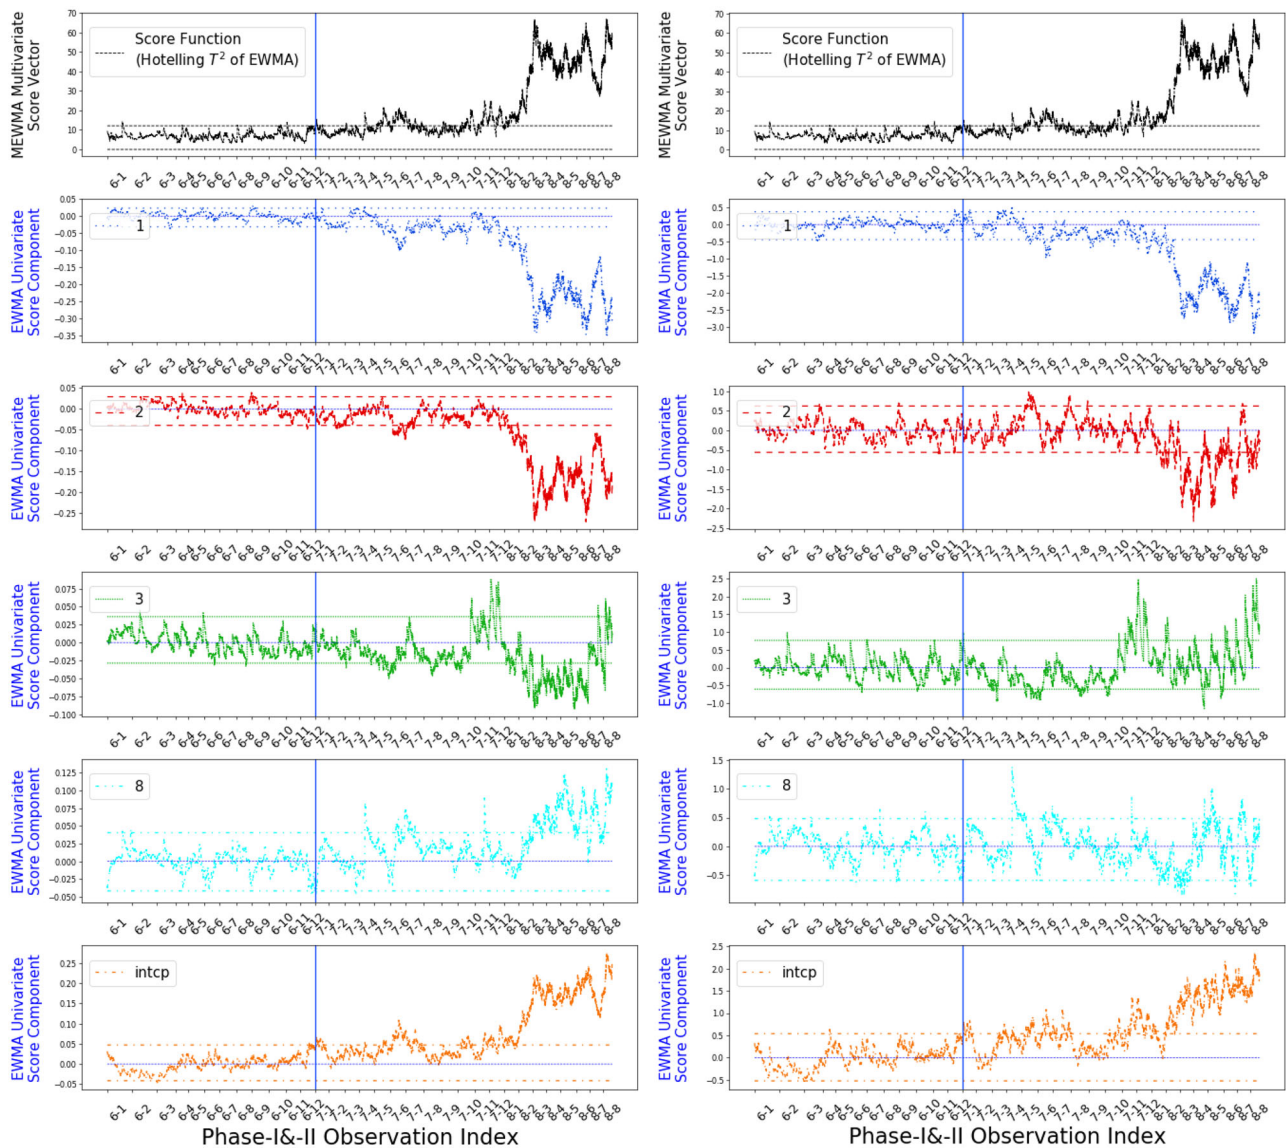


Figure 2. MEWMA (the top row) and various univariate component EWMA (the other rows, for θ_1 , θ_2 , θ_3 , θ_8 , and the intercept θ_0) control charts for the logistic regression model in the credit risk example over the Phase-I and Phase-II data. The left column are the original score component EWMA control charts (9), and the right column are the decoupled component EWMA control charts (12). The horizontal lines are the control limits and also a line indicating the zero value for the EWMA. The blue vertical line is the boundary between the Phase-I and Phase-II data. The decoupled component EWMA control charts show clear drift in a number of parameters, especially the intercept (the bottom right) and θ_1 (the second from the top, the right column).

concept drift) the performance of the models may degrade. We have developed a comprehensive, general, and powerful score-based framework for monitoring and diagnosing concept drift. The framework is general in that it applies to broad classes of parametric models, either regression or classification. It can be used retrospectively (to analyze whether $P(Y|X)$ was stable over a past dataset and aid in the model building procedure) and prospectively (to quickly detect changes in $P(Y|X)$ so that predictive models currently in use can be updated accordingly). We have provided theoretical arguments that, under reasonably general conditions, concept drift occurs if and only if the mean of the score vector changes. Consequently, our score-based procedure is based on monitoring and analyzing changes in the mean of the score vector. For this, we have adopted procedures (MEWMA and univariate EWMA) that were developed in the SPC literature for monitoring for changes in the mean of multivariate vectors in general. As part of the framework,

we have also developed a diagnostic approach that involves decoupling changes in the parameters from the change in the mean of the score vector.

In simulation and real data examples, we have demonstrated that our score-based monitoring procedure provides much more powerful detection of changes in $P(Y|X)$ than the state-of-the-art existing approach (error-based EWMA). In particular, there are examples (e.g., Figure B1 in the supplementary materials) in which the score-based approach quickly detects the change in $P(Y|X)$, but the error-based approach is completely unable to detect the change because they do not result in a change in the error rate.

We note that although multivariate EWMA monitoring charts (e.g., Figure 1 and the top row of Figure 2) can be used with complex models, the univariate EWMA diagnostic plots (e.g., the remaining rows of Figure 2) are intended more for models in which the parameters are meaningful, such as

for linear and logistic regression models. However, for complex models like neural networks with many layers, one could construct diagnostic plots for only the parameters in the last layer as described in Section D, supplementary materials, which represent the effects of features extracted in the previous layers.

One of the assumptions when fitting a supervised learning model of the type considered in this article is that the training data are stationary over time, in the sense that there is no drift in the predictive relationship between the response variable and the predictor variables. Our concept drift detection approach can be viewed as a way of determining whether this stationarity assumption is violated and, if not, at which point(s) in time the relationship drifted. Moreover, if the univariate diagnostic plots are used, the approach can help indicate which model parameters have drifted.

Supplementary Materials

In the online supplementary materials of this article, we provide a file containing derivations of Fisher Decoupling Equations (10) and (11), examples demonstrating effectiveness of our methods, simulation study results, and a discussion on potential extensions of our methods. Additionally, the zip file contains python code and data necessary to reproduce some results presented in this article.

Funding

This work was funded in part by the Air Force Office of Scientific Research Grant (#FA9550-18-1-0381), which we gratefully acknowledge. This work was also supported by a startup grant (#DMR200021) from the Extreme Science and Engineering Discovery Environment (XSEDE) (Towns et al. 2014), which is supported by the National Science Foundation (#ACI-1548562). We also acknowledge the many helpful comments from the Editor (Dr. Roshan Joseph) and the anonymous Associate Editor and Referees.

ORCID

Kungang Zhang  <http://orcid.org/0000-0001-7790-4903>

Anh T. Bui  <http://orcid.org/0000-0003-0397-1307>

References

- Baena-Garcia, M., del Campo-Ávila, J., Fidalgo, R., Bifet, A., Gavalda, R., and Morales-Bueno, R. (2006), "Early Drift Detection Method," in *Fourth International Workshop on Knowledge Discovery from Data Streams* (Vol. 6), pp. 77–86. [139]
- Barros, R., S. M., and Santos, S. G. T. C. (2018), "A Large-Scale Comparison of Concept Drift Detectors," *Information Sciences*, 451, 348–370. [138,139]
- Bickel, P. J., and Doksum, K. A. (2015), *Mathematical Statistics: Basic Ideas and Selected Topics*, Chapman & Hall/CRC Texts in Statistical Science Book 117 (Vol. 1), Boca Raton, FL: CRC Press. [140]
- Bifet, A., and Gavalda, R. (2007), "Learning from Time-Changing Data with Adaptive Windowing," in *Proceedings of the 2007 SIAM International Conference on Data Mining*, pp. 443–448, SIAM. [138]
- Bottou, L., Curtis, F. E., and Nocedal, J. (2018), "Optimization Methods for Large-Scale Machine Learning," *SIAM Review*, 60, 223–311. [141]
- Box, G. E. P. (1976), "Science and Statistics," *Journal of the American Statistical Association*, 71, 791–799. [141]
- Calandra, R., Raiko, T., Deisenroth, M. P., and Pouzols, F. M. (2012), "Learning Deep Belief Networks from Non-stationary Streams," in *International Conference on Artificial Neural Networks*, pp. 379–386, Springer. [138]
- Carmona-Cejudo, J. M., Baena-García, M., del Campo-Ávila, J., Morales-Bueno, R., and Bifet, A. (2010), "GNUmail: Open Framework for On-line Email Classification," in *Frontiers in Artificial Intelligence and Applications* 215, pp. 1141–1142, IOS Press. [138]
- Crook, J. N., Thomas, L. C., and Hamilton, R. (1992), "The Degradation of the Scorecard Over the Business Cycle," *IMA Journal of Management Mathematics*, 4, 111–123. [137]
- Donoho, S. (2004), "Early Detection of Insider Trading in Option Markets," in *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 420–429. [138]
- Fong, Y., Di, C., and Permar, S. (2015), "Change Point Testing in Logistic Regression Models with Interaction Term," *Statistics in Medicine*, 34, 1483–1494. [138]
- Frias-Blanco, I., del Campo-Ávila, J., Ramos-Jiménez, G., Morales-Bueno, R., Ortiz-Díaz, A., and Caballero-Mota, Y. (2015), "Online and Non-parametric Drift Detection Methods based on Hoeffding's Bounds," *IEEE Transactions on Knowledge and Data Engineering*, 27, 810–823. [138]
- Gama, J., Medas, P., Castillo, G., and Rodrigues, P. (2004), "Learning with Drift Detection," in *Brazilian Symposium on Artificial Intelligence*, pp. 286–295, Springer. [139]
- Gonçalves, Jr., P. M., de Carvalho Santos, S. G. T., Barros, R. S. M., and Vieira, D. C. L. (2014), "A Comparative Study on Concept Drift Detectors," *Expert Systems with Applications*, 41, 8144–8156. [138]
- Harel, M., Mannor, S., El-Yaniv, R., and Crammer, K. (2014), "Concept Drift Detection through Resampling," in *International Conference on Machine Learning*, pp. 1009–1017. [138]
- Hotelling, H. (1947), "Multivariate Quality Control Illustrated by Air Testing of Sample Bombsights," In *Techniques of Statistical Analysis*, eds. C. Eisenhart, M. W. Hastay, and W. A. Wallis, pp. 111–184, New York: McGraw-Hill. [142]
- Hulten, G., Spencer, L., and Domingos, P. (2001), "Mining Time-Changing Data Streams," in *Proceedings of the seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 97–106. [137]
- Im, J. K., Apley, D. W., Qi, C., and Shan, X. (2012), "A Time-Dependent Proportional Hazards Survival Model for Credit Risk Analysis," *Journal of the Operational Research Society*, 63, 306–321. [145]
- Kingma, D. P., and Ba, J. (2014), "Adam: A Method for Stochastic Optimization," arXiv preprint arXiv:1412.6980. [141]
- Kuan, C.-M., and Hornik, K. (1995), "The Generalized Fluctuation Test: A Unifying View," *Econometric Reviews*, 14, 135–161. [139]
- Kukar, M. (2003), "Drifting Concepts as Hidden Factors in Clinical Studies," in *Conference on Artificial Intelligence in Medicine in Europe*, 355–364, Springer. [138]
- Lowry, C. A., Woodall, W. H., Champ, C. W., and Rigdon, S. E. (1992), "A Multivariate Exponentially Weighted Moving Average Control Chart," *Technometrics*, 34, 46–53. [142]
- Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., and Zhang, G. (2018), "Learning under Concept Drift: A Review," *IEEE Transactions on Knowledge and Data Engineering*, 31, 2346–2363. [139]
- Mejri, D., Limam, M., and Weihs, C. (2018), "A New Dynamic Weighted Majority Control Chart for Data Streams," *Soft Computing*, 22, 511–522. [137,143]
- (2021), "A New Time Adjusting Control Limits Chart for Concept Drift Detection," *IFAC Journal of Systems and Control*, 17, 100170. [137,143]
- Montgomery, D. C. (2007), *Introduction to Statistical Quality Control*, New York: Wiley. [146]
- Moreno-Torres, J. G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N. V., and Herrera, F. (2012), "A Unifying View on Dataset Shift in Classification," *Pattern Recognition*, 45, 521–530. [137]
- Nelder, J. A., and Wedderburn, R. W. M. (1972), "Generalized Linear Models," *Journal of the Royal Statistical Society, Series A*, 135, 370–384. [140]
- Pechenizkiy, M., Tsymbal, A., Puuronen, S., Shifrin, M., and Alexandrova, I. (2005), "Knowledge Discovery from Microbiology Data: Many-Sided Analysis of Antibiotic Resistance in Nosocomial Infections," in *Biennial Conference on Professional Knowledge Management/Wissensmanagement*, pp. 360–372, Springer. [138]

- Raza, H., Prasad, G., and Li, Y. (2015), "EWMA Model based Shift-Detection Methods for Detecting Covariate Shifts in Non-stationary Environments," *Pattern Recognition*, 48, 659–669. [137,143]
- Roberts, S. W. (1959), "Control Chart Tests based on Geometric Moving Averages," *Technometrics*, 1, 239–250. [142]
- Ross, G. J., Adams, N. M., Tasoulis, D. K., and Hand, D. J. (2012), "Exponentially Weighted Moving Average Charts for Detecting Concept Drift," *Pattern Recognition Letters*, 33, 191–198. [138,139]
- Shannon, C. E. (1948), "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27, 379–423. [140]
- Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., Hazlewood, V., Lathrop, S., Lifka, D., Peterson, G. D., and Roskies, R. (2014), "XSEDE: Accelerating Scientific Discovery," *Computing in Science & Engineering*, 16, 62–74. [148]
- Tsymbol, A. (2004), "The Problem of Concept Drift: Definitions and Related Work," Computer Science Department, Trinity College Dublin, 106, 2. [138]
- Tsymbol, A., Pechenizkiy, M., Cunningham, P., and Puuronen, S. (2008), "Dynamic Integration of Classifiers for Handling Concept Drift," *Information Fusion*, 9, 56–68. [138]
- Wang, H., and Abraham, Z. (2015), "Concept Drift Detection for Streaming Data," in *2015 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–9. IEEE. [139]
- Wang, H., Fan, W., Yu, P. S., and Han, J. (2003), "Mining Concept-Drifting Data Streams using Ensemble Classifiers," in *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 226–235. ACM. [138]
- Widmer, G., and Kubat, M. (1996), "Learning in the Presence of Concept Drift and Hidden Contexts," *Machine Learning*, 23, 69–101. [138]
- Xia, Z., Guo, P., and Zhao, W. (2009), "Monitoring Structural Changes in Generalized Linear Models," *Communications in Statistics—Theory and Methods*, 38, 1927–1947. [139]
- Zeileis, A. (2005), "A Unified Approach to Structural Change Tests based on ML Scores, F Statistics, and OLS Residuals," *Econometric Reviews*, 24, 445–466. [139]
- Zeileis, A., and Hornik, K. (2007), "Generalized M-Fluctuation Tests for Parameter Instability," *Statistica Neerlandica*, 61, 488–508. [139]
- Žliobaitė, I., Bakker, J., and Pechenizkiy, M. (2012), "Beating the Baseline Prediction in Food Sales: How Intelligent an Intelligent Predictor is?," *Expert Systems with Applications*, 39, 806–815. [138]
- Žliobaitė, I., Pechenizkiy, M., and Gama, J. (2016), "An Overview of Concept Drift Applications," in *Big Data Analysis: New Algorithms for a New Society*, eds. N. Japkowicz and J. Stefanowski, pp. 91–114, Cham: Springer. [137,138]