

FEAMOE: Fair, Explainable and Adaptive Mixture of Experts

Shubham Sharma¹, Jette Henderson² and Joydeep Ghosh¹

¹The University of Texas at Austin

²TecnoTree

{shubham_sharma, jghosh}@utexas.edu, jette.henderson@gmail.com

Abstract

Three key properties that are desired of trustworthy machine learning models deployed in high-stakes environments are fairness, explainability, and an ability to account for various kinds of "drift". While drifts in model accuracy have been widely investigated, drifts in fairness metrics over time remain largely unexplored. In this paper, we propose FEAMOE, a novel "mixture-of-experts" inspired framework aimed at learning fairer, more interpretable models that can also rapidly adjust to drifts in both the accuracy and the fairness of a classifier. We illustrate our framework for three popular fairness measures and demonstrate how drift can be handled with respect to these fairness constraints. Experiments on multiple datasets show that our framework as applied to a mixture of linear experts is able to perform comparably to neural networks in terms of accuracy while producing fairer models. We then use the large-scale HMDA dataset and show that various models trained on HMDA demonstrate drift and FEAMOE can ably handle these drifts with respect to all the considered fairness measures and maintain model accuracy. We also prove that the proposed framework allows for producing fast Shapley value explanations, which makes computationally efficient feature attribution based explanations of model decisions readily available via FEAMOE.

1 Introduction

The field of responsible artificial intelligence has several desiderata that are motivated by regulations such as the General Data Protected Regulation [Butterworth, 2018]. These include: ensuring that an AI model is non-discriminatory and transparent; individuals subject to model decisions should have access to explanations that point a path towards recourse; and models should adapt to any changes in the characteristics of the data post-deployment so as to maintain their quality and trustworthiness.

Most approaches towards the mitigation of any form of bias assume a static classifier. A practitioner decides on some definition of fairness, trains a model that attempts to enforce

this notion of fairness and then deploys the model. Many of the fairness definitions are based on model outcomes or on error rates (the gap between true and/or false positive rates) that are associated with different subgroups specified by a protected attribute. The goal is to reduce the difference between these error rates across relevant subgroups. For example, average odds difference [Bellamy *et al.*, 2018] is a measure signifying equalized odds and is given by the sum of the differences in both true positive and false positive rates between two groups, scaled by a factor of 0.5. Equality of opportunity and demographic parity [Barocas *et al.*, 2019] are also popular definitions of fairness. Recently, fairness in terms of a gap of recourse has been proposed, where recourse is defined as the ability to obtain a positive outcome from the model [Sharma *et al.*, 2020a]. While the suitability of a fairness measure is application dependent [Mehrabi *et al.*, 2019; Barocas *et al.*, 2019], demographic parity and equalized odds remain the most popularly used, and the need for recourse gap-based fairness is being increasingly recognized [Karimi *et al.*, 2020].

However, static models can encounter drift once deployed, as the statistical properties of real data often change over time. This can lead to deteriorating performance. Model drift can occur when the properties of the target variable change (concept drift) or when the input data distribution changes, or both. The performance of models has largely been measured through accuracy-based metrics such as misclassification rates, F-score or AUC. [Stanley, 2003]. However, a model trained in the past and found to be fair at training time may act unfairly for data in the present. Addressing drift with respect to fairness in addition to accuracy has remained largely unexplored though it is an important aspect of trustworthy AI in practice.

Explainability of individual model outcomes is another principal concern for trustworthy ML. Among many methods of explanations in terms of feature attribution, [Bhatt *et al.*, 2020], the SHAP approach based on Shapley values is particularly popular as it enjoys several axiomatic guarantees [Lundberg and Lee, 2017a]. While computation of SHAP values is fast for linear and tree-based models, it can be very slow for neural networks and several other model types, especially when the data has a large numbers of features or when a large number of explanations are required [Molnar, 2019]. This poses a barrier to deployments that demand fast explanations in real-time, production settings.

In this paper, we address these fairness, data/model drift,

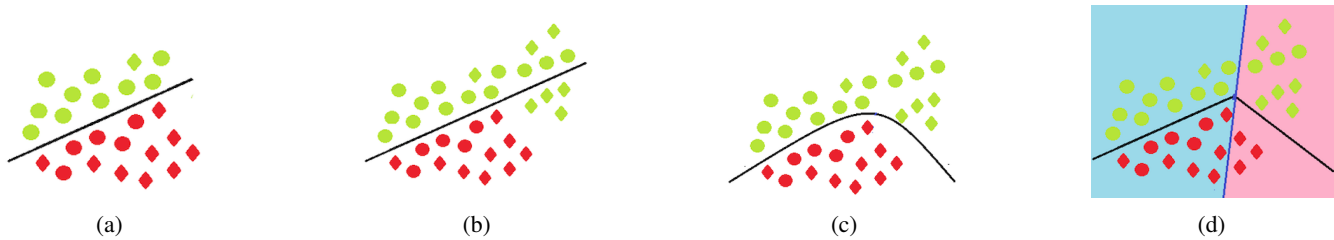


Figure 1: A toy example demonstrating the need and use of FEAMOE. The color of every datapoint corresponds to the original class label. Diamonds represent the underprivileged group and circles represent the privileged group. (a) Represents a perfectly accurate linear classifier, (b) represents the same classifier mis-classifying new data points and inducing bias (drift), (c) represents an alternate non-linear model that corrects for drift but has a complex decision boundary and (d) represents FEAMOE where the blue and pink regions show the regions of operation for each of the two experts, separated by the gating network

and explainability concerns by proposing FEAMOE: Fair, Explainable and Adaptive Mixture of Experts, an incrementally grown mixture of experts (MOE) with fairness constraints. In the standard mixture of experts setup, each expert is a machine learning model, and so is the gating network. The gating network learns to assign an input-dependent weight $g_u(\mathbf{x})$ to the u^{th} expert for input \mathbf{x} , and the final output of the model is a weighted combination of the outputs of each expert. Hence, each expert contributes differently for every data point towards the final outcome, which is a key difference from standard ensembles.

Many types of MOE’s exist in the literature [Yuksel *et al.*, 2012] - the architecture is not standard. For FEAMOE, we chose this family, with some novel modifications described later, for three main reasons: 1) Suitable regularization penalties that promote fairness can be readily incorporated into the loss function. 2) Online learning is possible, so changes in the data can be tracked. Crucially, since localized changes in data distribution post-deployment may impact only one or a few experts, the other experts may not need to be adjusted, making the experts localized and only loosely coupled. This allows for handling drift and avoiding catastrophic forgetting, which is a prime concern in widely used neural network models [Robins, 1995]. 3) Simpler models can be used to fit a more complex problem in the mixture of experts, as each model needs to fit well in only a limited part of the input space. In particular, even linear models, which provide very fast SHAP explanations, can be used. The overall mixture of experts, even with such simple base models (the "experts") often has predictive power that is comparable to a single complex model such as a neural network, as shown by our experiments as well as in many previous studies [Yuksel *et al.*, 2012].

A motivating toy example of why FEAMOE is needed and how it works is shown in Figure 1. Consider a linear binary classifier (1a) that has perfect accuracy. The colors represent the ground truth labels, and green is the positive (desired) class label. The circles are the privileged group and diamonds are the underprivileged group. As can be seen in the figure, more diamonds receive a negative outcome and more circles receive a positive outcome. Consider new data that arrives for predictions. This classifier (1b) not only misclassifies individuals but also gives more underprivileged individuals that were actually in the positive class a negative outcome, hence inducing bias with respect to equalized odds. There is

drift with respect to accuracy and fairness. A more complex model (1c) such as a neural network, if retrained, may handle some of these concerns but would be less explainable.

FEAMOE can address these imperative concerns, as shown in 1d. Trained in an online manner, a new linear model is added (i.e., an expert) once the new data arrives. The gating network dictates which region each expert operates in (shown by the blue and pink colors), and FEAMOE is able to adapt automatically with respect to accuracy and fairness. This dynamic framework enables the overall model to be fairer, adjust to drift, maintain accuracy, while also remaining explainable since the decision boundary is locally linear.

We show how three fairness constraints—demographic parity, equalized odds, and burden-based fairness—can be incorporated into the mixture of experts training procedure in order to encourage fitting fairer models (according to these measures). We use these three popular fairness measures as illustrative examples to demonstrate the effectiveness of FEAMOE, but our method can be adapted to incorporate other fairness constraints as well. We then describe a new algorithm for training to account for drift, where the drift in question can be due to accuracy or fairness. We show experimentally that by using a set of logistic regression experts, the accuracy of the mixture is comparable to using a complex model like a neural network. Additionally, we show we can efficiently compute Shapley value explanations when explanations for every individual expert can be computed quickly. To the best of our knowledge, this is the first work that addresses the problem of drift with respect to fairness in a large-scale real world dataset. We then introduce a framework that can flexibly adapt to drifts in both fairness and accuracy with the added benefit of delivering explanations quickly, while comparing to the less explainable neural network model class trained in online mode.

The key contributions of this work are: a mixture of experts framework that can incorporate multiple fairness constraints, a method to handle drift, where drift can be with respect to accuracy or fairness, empirical evidence of the presence of drift with respect to fairness in a real-world, large-scale dataset, a theoretical proof that FEAMOE leads to the generation of fast explanations given a suitable choice of experts, and extensive experimentation on three datasets to show that our method has predictive performance similar to neural networks while being fairer, handling different types of drift, and generating faster explanations.

2 Related Work

The mixture of experts (MOE) [Jacobs *et al.*, 1991; ?] represent a class of co-operative ensemble models; detailed surveys on their design and use can be found in [Yuksel *et al.*, 2012] and [Masoudnia and Ebrahimpour, 2014]. Very recently, the deep learning community has started recognizing and leveraging several advantageous properties that MOE’s have for efficient design of complex, multi-purpose learners [Riquelme *et al.*, 2021]. This paper contributes to this expanding literature by proposing a new algorithm to train this model class to account for both fairness and drift, and by also adding an explainability module.

Fairness in machine learning is a growing field of research [Hacker, 2018]. Mitigating biases in models can be done through pre-processing, in-processing, or post-processing techniques. A description of these techniques can be found in [Bellamy *et al.*, 2018]. In-processing techniques for fairness have been gaining traction [Zhang *et al.*, 2018; Mehrabi *et al.*, 2019; ?]. However, there is limited work on investigating the usefulness of ensemble models in dealing with biases. [Grgić-Hlača *et al.*, 2017] show that an ensemble of fair classifiers is guaranteed to be fair for several different measures of fairness, an ensemble of unfair classifiers can still achieve fair outcomes, and an ensemble of classifiers can achieve better accuracy-fairness trade-offs than a single classifier. However, they neither provide experimental evidence nor discuss specific methods to incorporate fairness into ensemble learning. [Madras *et al.*, 2017] develop a method to learn to defer in the case of unfair predictions. [Bhaskaruni *et al.*, 2019] use an AdaBoost framework to build a fairer model. [Nejdl,] use adaptive random forest classifiers to account for fairness in online learning.

Accounting for drift is a widely explored problem, and is now appearing in commercial products as well (e.g. model monitoring is a key part of MLOPs) as ML solutions get deployed in business environments. Details on many such approaches can be found in [Gama *et al.*, 2014; Lu *et al.*, 2018]. Among these approaches, the one that comes closest to ours is [Stanley, 2003] which uses a committee of decision trees to account for drift. However, ensuring fairness in the presence of drift remains an open problem. [Biswas and Mukherjee, 2020] is a very recent work on achieving a fairer model by building a set of classifiers in the presence of prior distribution shifts. The method is built for a shift between the training and test distributions, and not for online learning. [Zhang and Ntoutsis, 2019; Zhang *et al.*, 2021] provide online learning methods for tree-based models. We show experimentally (in the supplementary material ¹) that FEAMOE can work comparably or better to adapt for drift. Recently, some researchers have studied fairness in online learning [Ntoutsis *et al.*, 2020; Bechavod *et al.*, 2020]. Innovations include the notion of cumulative fairness monitoring to account for discriminatory outcomes from the beginning of the stream until a time point.

There are many ways to explain a machine learning model [Burkart and Huber, 2021; Molnar, 2019]. In this paper, we focus on Shapley values-based explanations, which are widely

used in practical applications [Bhatt *et al.*, 2020]. [Lundberg *et al.*, 2018] propose the computation of Shapley values for tree ensembles, which is a faster way to get Shapley values than through the more broadly applicable method, KernelShap [Lundberg and Lee, 2017a]. We show that in FEAMOE, the Shap values for the overall model are just a data-dependent linear combination of the values from individual experts. Thus the mixture approach does not add any significant complexity to the computation of feature attribution scores.

3 Theory

We first summarize the original mixture of experts framework and then describe the addition of fairness constraints. Then, we introduce the algorithm to detect and mitigate data drift when the data input is sequential (online learning). Thereafter, we show how using the proposed mixture of experts architecture leads to computing faster Shapley value explanations for the overall non-linear model.

Mixture of Experts (MoE) [Jacobs *et al.*, 1991] is a technique where multiple experts (learners) can be used to softly divide the problem space into regions. A gating network decides which expert to weigh most heavily for each input region. Learning thus consists of the following: 1) learning the parameters of individual learners and 2) learning the parameters of the gating network. Both the gating network and every expert have access to the input \mathbf{x} . The gating network has one output g_i for every expert i . The output vector is the weighted (by the gating network outputs) mean of the expert outputs: $y(\mathbf{x}) = \sum_{i=1}^m g_i(\mathbf{x})y_i(\mathbf{x})$. Consistent with [Jacobs *et al.*, 1991], the error associated with training the mixture of experts for case j for an accurate prediction is given by:

$$E_{acc}^j = -\log \sum_i g_i^j e^{-\frac{1}{2} \|\mathbf{d}^j - \mathbf{y}_i^j\|^2} \tag{1}$$

,where \mathbf{y}_i^j is the output vector of expert i on case j , g_i^j is the proportional contribution of expert i to the combined output vector, and \mathbf{d}^j is the desired output vector.

3.1 Fairness Constraints

In this paper, we incorporate three diverse fairness definitions into the mixture of experts framework: demographic parity only depends on the model outcome, equalized odds is conditioned on the ground-truth label, and burden-based fairness depends on the distance of the input to the boundary. These three popular definitions have been chosen as illustrative metrics; our approach can be readily extended to several other fairness metrics as well.

For simplicity, we consider a binary classification setting with a binary protected attribute (our approach readily extends to multi-class and multi-protected attribute problems, where a protected attribute is a feature such as race or gender). Let $y_i^j = 1$ be the positive outcome. Let $A = 0$ and $A = 1$ represent the underprivileged and privileged protected attribute groups, respectively. For a given dataset D , let D_{ad} represent all individuals that belong to the protected attribute group a and original class label d .

Statistical parity difference (SPD), which is a measure of demographic parity, measures the difference between the probability of getting a positive outcome between protected attribute

¹<https://drive.google.com/file/d/112qz50Flvj4VAEvrRr-H4Gdy3QAmCRnY/view?usp=sharing>

groups [Bellamy *et al.*, 2018; Sharma *et al.*, 2020b]. Let D_0 be the set of individuals in the underprivileged group and D_1 be the set of individuals in the privileged group. Inspired by [Slack *et al.*, 2020], the associated penalty for demographic parity for case j is:

$$E_{SPD}^j = 1[j \in D_0](1 - \sum_i g_i \mathbf{y}_i^j) + 1[j \in D_1](\sum_i g_i \mathbf{y}_i^j). \quad (2)$$

The idea behind this term is that individuals belonging to the underprivileged group predicted as getting a negative outcome are assigned a higher penalty. Similarly, individuals belonging to the privileged group predicted to have a positive outcome are assigned a higher penalty, thereby encouraging an SPD value closer to zero.

Average odds difference (AOD), which is a measure of equalized odds, measures the difference in true and false rates between protected attribute groups. Details on the measure can be found in [Bellamy *et al.*, 2018; Sharma *et al.*, 2020b]. The associated penalty for equalized odds is:

$$E_{AOD}^j = 1[j \in D_{01}](1 - \sum_i g_i \mathbf{y}_i^j) + 1[j \in D_{11}](1 - \sum_i g_i \mathbf{y}_i^j) + 1[j \in D_{10}](\sum_i g_i \mathbf{y}_i^j) + 1[j \in D_{00}](\sum_i g_i \mathbf{y}_i^j).$$

This term encourages the true and false positive rate gaps between groups to reduce by conditioning the indicator function on the ground truth label in addition to the protected attribute (as was in the demographic parity formulation).

Burden for a protected attribute group is a measure of the ability to obtain recourse for individuals in that group. As shown in [Sharma *et al.*, 2020a], burden-based fairness can be calculated as:

$$Burden = \left| \mathbb{E}_{\mathbf{x}|A=0} [d(\mathbf{x}, \mathcal{B})] - \mathbb{E}_{\mathbf{x}|A=1} [d(\mathbf{x}, \mathcal{B})] \right| \quad (3)$$

,where $d(\mathbf{x}, \mathcal{B})$ represents the distance to the boundary for a given \mathbf{x} that is classified as being in the negative class. Then, the associated penalty for burden based fairness is:

$$E_{Burden}^j = \left| \mathbb{E}_{\mathbf{x}|A=0} [d(\mathbf{x}, \mathcal{B})] - \mathbb{E}_{\mathbf{x}|A=1} [d(\mathbf{x}, \mathcal{B})] \right|. \quad (4)$$

The overall loss for case j is then given by:

$$E_{MOE}^j = E_{acc}^j + \lambda_1 E_{SPD}^j + \lambda_2 E_{AOD}^j + \lambda_3 E_{Burden}^j \quad (5)$$

Note that there may be a trade-off between accuracy and fairness measures when the true positive outcome rates differ among the different population segments [Sharma *et al.*, 2020a; Bellamy *et al.*, 2018], and the practitioner needs to decide the relative importance of the different constraints for the given application.

3.2 Data Drift and the FEAMOE Algorithm

Data Drift means that the statistical properties of the data, embodied in the underlying joint distribution of independent and

Algorithm 1 Learning FEAMOE

Inputs: data X , labels Y
 Hyperparameters: $k, \Delta\lambda_1, \Delta\lambda_2, \Delta\lambda_3$
 $s = 1$
 # Select k points $\{x\}_{l=1}^k \subset X$
 $X^{(s)} = X \setminus \{x\}_{l=1}^k$
 # Learn the first expert, m_s
 # Initialize w_s , weights of m_s
for j in $\{1, \dots, k\}$ **do**
 # Take gradient steps to minimize MoE loss:
 $w_s^j = w_s^{j-1} - \nu \nabla E_{acc}^j$
end for
 $\lambda_1 = 0, \lambda_2 = 0$, and $\lambda_3 = 0$
while $X^{(s)}$ is not empty **do**
 $s += 1$
 $\lambda_1 = \lambda_1 + \Delta\lambda_1, \lambda_2 = \lambda_2 + \Delta\lambda_2, \lambda_3 = \lambda_3 + \Delta\lambda_3$
 # Select k points $\{x\}_{l=1}^k \subset X^{(s)}$
 $X^{(s)} = X^{(s)} \setminus \{x\}_{l=1}^k$
 # Learn the subsequent expert, m_s
 Initialize w_s
 for j in $\{1, \dots, s * k\}$ **do**
 # Take gradient steps to minimize Equation 5
 for l in $\{1, \dots, s\}$ **do**
 $w_s^j = w_s^{j-1} - \nu \nabla E_{acc}^j - \lambda_1 \nabla E_{SPD} - \lambda_2 \nabla E_{AOD} - \lambda_3 \nabla E_{Burden}^j$
 end for
 end for
end while

dependent variables, can change over time, often in unforeseen ways. The change could be in the class priors, the class conditional distributions (concept drift), in the distribution of the independent variables etc. Drift can cause the model to become less accurate as time passes. However, drift can also cause other properties associated with the model to change, such as fairness. We develop an algorithm that can handle drift with respect to both accuracy and fairness.

Consider an online learning setup where input data points are observed sequentially. The algorithm to learn FEAMOE (Algorithm 1) is as follows: start with a single model. Begin to train with data points (using stochastic gradient descent) and train the current model for a certain number of data points k using only E_{acc} (equation 2). After k points, introduce a new logistic regression model and train the mixture of experts with a softmax gating function using the loss in Equation 5. Simultaneously, introduce the fairness penalties. Then, continue training for the next k points, and then add another expert. As more experts are added, gradually increase the hyperparameters (λ' values) associated with the three fairness losses. This process is continued until available data is seen.

The motivation behind this training scheme is two-fold: in beginning with the accuracy penalty for the first expert, we ensure that the fairness measures do not interfere with training an accurate classifier, since high weights on the fairness terms would result in a less accurate classifier (as shown in experiments). Then, we slowly increase the weights on the fairness penalties with the goal that for individuals that are

Type	None	SPD	AOD	Burden	All
MOE	MOE	FEAMOE1	FEAMOE2	FEAMOE3	FEAMOE
NN	NN	x	x	x	FairNN

Table 1: Names of the models compared in figure 2, based on the class of models (Mixture of Experts (MOE) or Neural Network (NN)) and type of fairness constraints (None, SPD, AOD, Burden, or all). Experiments on models marked x are in the supplementary material

classified unfairly with respect to these group fairness measures, another expert takes over in this data regime to train for these individuals over time. This is because the mixture of experts framework allows some or all of the experts to learn on different regions of the data. Secondly, the algorithm allows us to account for drift, both with respect to the accuracy of the classifier and the fairness, since our framework allows for fairness constraints. If there is a change in the statistical properties of the data that impacts any of the loss terms, the mixture of experts adapts to this change over time through the addition of experts.

3.3 Fast Shapley Value Explanations

A prominent class of feature attribution methods is based on Shapley values from cooperative game theory [Shapley, 1953]. Details about Shapley value explanations can be found in [Lundberg and Lee, 2017b], [Sundararajan *et al.*, 2017], and [Aas *et al.*, 2019]. While computing Shapley values for a linear model is fast, doing so for non-linear models with methods like KernelShap [Lundberg and Lee, 2017a] requires approximations and methods that cause the overall computation to become slow [Molnar, 2019; Aas *et al.*, 2019]. Another method, TreeShap, [Lundberg *et al.*, 2018] works only for tree models. Though the mixture of experts model proposed is non-linear, as the individual experts are linear, the theorem below shows how to compute them for the whole model quickly and efficiently.

Consider a mixture of experts model with m experts. Let $\phi_j(m(\mathbf{x}))$ be the Shapley value associated with expert m for feature j for an input instance \mathbf{x} .

Theorem 1. *For a mixture of experts model, the Shapley value for a given instance \mathbf{x} and feature j for the model prediction is given by:*

$$\phi_j(y(\mathbf{x})) = \sum_{i=1}^m g_i(\mathbf{x}) \phi_j(m(\mathbf{x})) \quad (6)$$

The proof is provided in the supplementary material. This result shows that the Shapley value for a given feature and input instance for the mixture of linear experts is a linear combination of the Shapley values of the feature and that input instance from every expert, weighted by the gating network’s assigned weights for that input. This means that so long as the Shap values for individual models can be quickly computed (as is the case for linear/logistic regression, decision trees, XGBoost), the FEAMOE system-level Shap computation is also very quick. In this paper, we illustrate FEAMOE using logistic regression experts, so this desirable property holds, even though the mixture model is able to construct non-linear models of arbitrary complexity by including as many logistic regression-based experts as needed.

4 Experiments and Results

Experiments are performed using the mixture of experts with logistic regression experts and a softmax gating function. We implement the logistic regression models using scikit learn with default parameters. We show that using logistic regression experts within the MOE produces accuracies similar to using appropriately sized neural networks while allowing for the generation of faster explanations. All neural networks are multilayer perceptrons. There are two sets of experiments, highlighting different aspects of FEAMOE:

(a) **Fairness Study.** We use two classification datasets that are very well studied in the fairness community: UCI Adult [Kohavi, 1996] and COMPAS [ProPublica, 2016]. Gender is considered as the protected attribute for UCI Adult. A two layer multilayer perceptron with 30 hidden units in each layer was trained for the UCI Adult dataset. Additional or larger hidden layers, or ensemble methods such as xgboost do not provide extra benefit for these two tabular datasets, and hence are omitted for comparison purposes. Experiments on the COMPAS dataset are in the supplementary material.

(b) **Drift Study.** The large HMDA (Home Mortgage Disclosure Act) dataset [Bureau, 2020] reflects data from multiple years, with the underlying data statistics varying considerably over the years. Gender is the protected attribute.

A five layer multilayer perceptron with 50 hidden units in each layer is trained for the HMDA dataset as the baseline neural network. Experiments on a synthetic streaming version of the UCI Adult dataset are in the supplementary material.

First, we use UCI Adult to demonstrate the effects of incorporating the proposed fairness constraints on the mixture of experts models. Similar results for the COMPAS dataset are provided in the supplementary material. We then show that the HMDA dataset demonstrates drift with respect to both fairness and accuracy, and that FEAMOE can adapt to such drifts. Comparisons are made to neural networks (both with and without fairness constraints), which is the state-of-the-art model class for accuracy-based performance across these datasets in all experiments. Experiments on faster Shapley value explanations are in the supplementary material.

4.1 Fairness Constraints

Experiments are performed on UCI Adult in seven different regimes based on model types and fairness constraints. Details of these regimes are in Table 1. Experts are added every 4000 data points for the UCI Adult dataset. Hyperparameters associated with the fairness constraints are incremented in levels of 0.02 per expert for the UCI Adult dataset. The parameters are found using grid search and vary based on dataset size and extent of prevalent bias (details in supplementary material). The results are averaged across five runs. We report the accuracy and the absolute value of the three

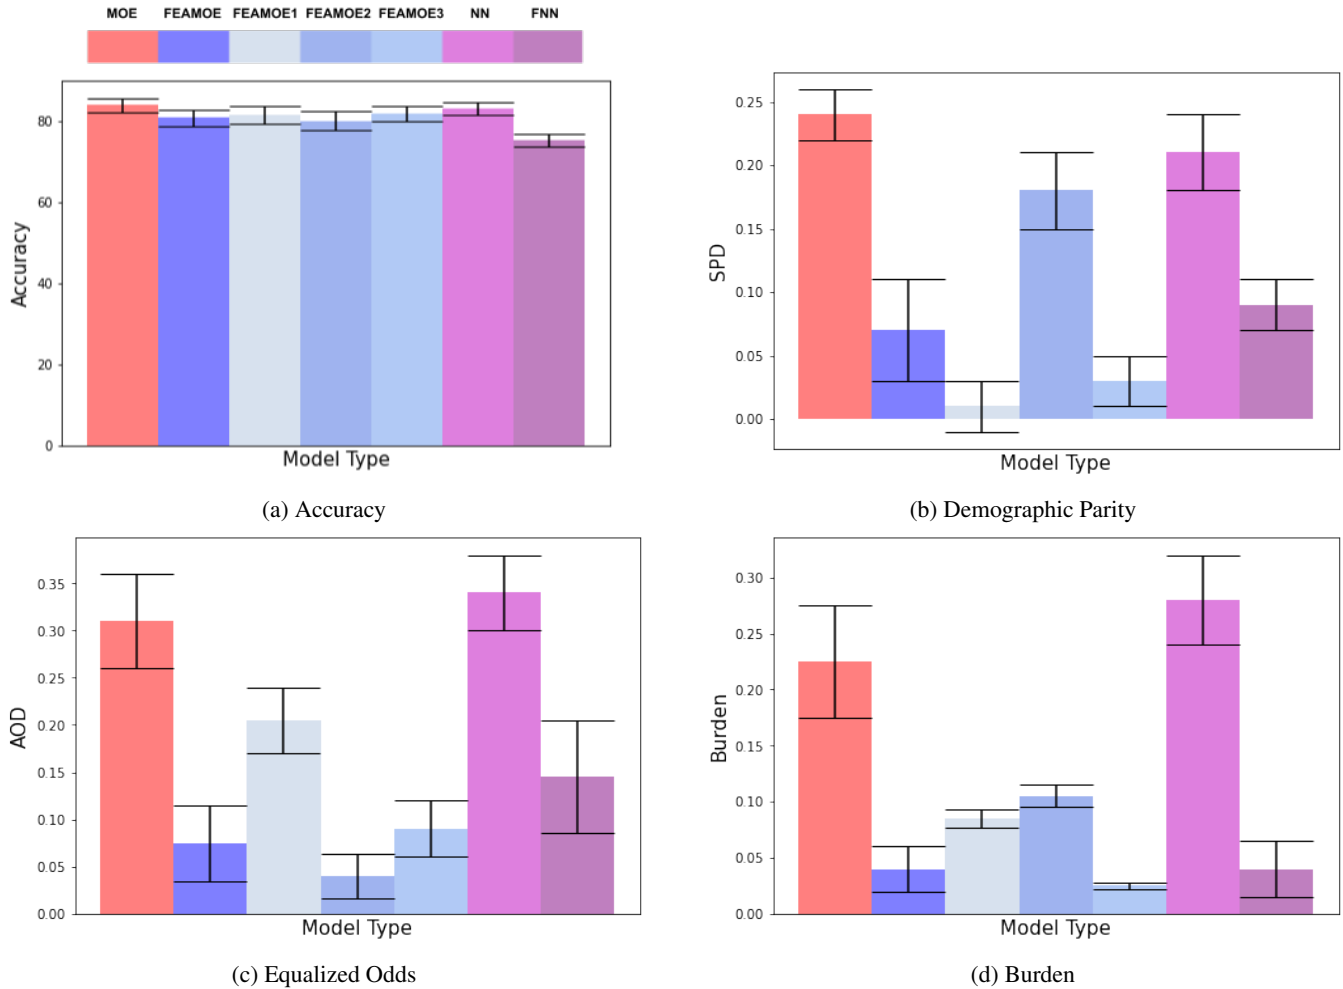


Figure 2: Results for the UCI Adult Dataset on different fairness constraints being incorporated. Details on model names are provided in Table 1. A higher accuracy is better and a lower bias is better

fairness measures. We provide comparisons to other methods for bias reduction [Calmon *et al.*, 2017; Sharma *et al.*, 2020a; Agarwal *et al.*, 2018] in the supplementary material.

The results are shown in Figure 2. The accuracy across different model types remains similar for the UCI Adult dataset, but using just a neural network with fairness constraints works poorly, as shown in Figure 2a. As seen in 2b,c,d, the fairness measures also work well even in isolation from each other. That is, in trying to improve based on just one measure, the other measures also improve. In this regard, the burden-based fairness measure (FEAMOE3) has the best effect; just using burden-based fairness alone helps significantly improve the other fairness measures while maintaining reasonable accuracy. The fair neural network (FairNN) performs worse for demographic parity and equalized odds compared to FEAMOE. We hypothesize that this happens because our learning process slowly induces fairness with every expert.

4.2 Real-World Drift: The HMDA Dataset

We first demonstrate that the HMDA dataset exhibits drift across years, and then show FEAMOE’s effectiveness in han-

dling it. The HMDA dataset has millions of records of individuals spanning several years. It contains consumer characteristics; the target variable indicates whether a consumer received a mortgage. While this dataset considered as a whole has been previously shown to exhibit bias, there is no investigation into how such bias varies across the years. First, to quantify drift in this dataset in both fairness and accuracy, we trained one neural network per year from 2007 to 2017, each on 100,000 random samples in that year, and tested each of these networks on data from the years 2016-2017 (Once trained, these models, which we call fixed neural networks, cannot be updated). The results are shown in Figure 3 by the blue points. In general, the farther the training data is away from the test year the more the accuracy and fairness measures degrade (i.e., accuracy decreases and fairness differences increase). Also training a single model on a dataset of the same size but sampled uniformly over all the previous years does not help either as the data is non-stationary.

We now study how fairness aware neural networks with on-line updates compare with FEAMOE in their ability to handle drift. Note that for our setting, we cannot use certain mod-

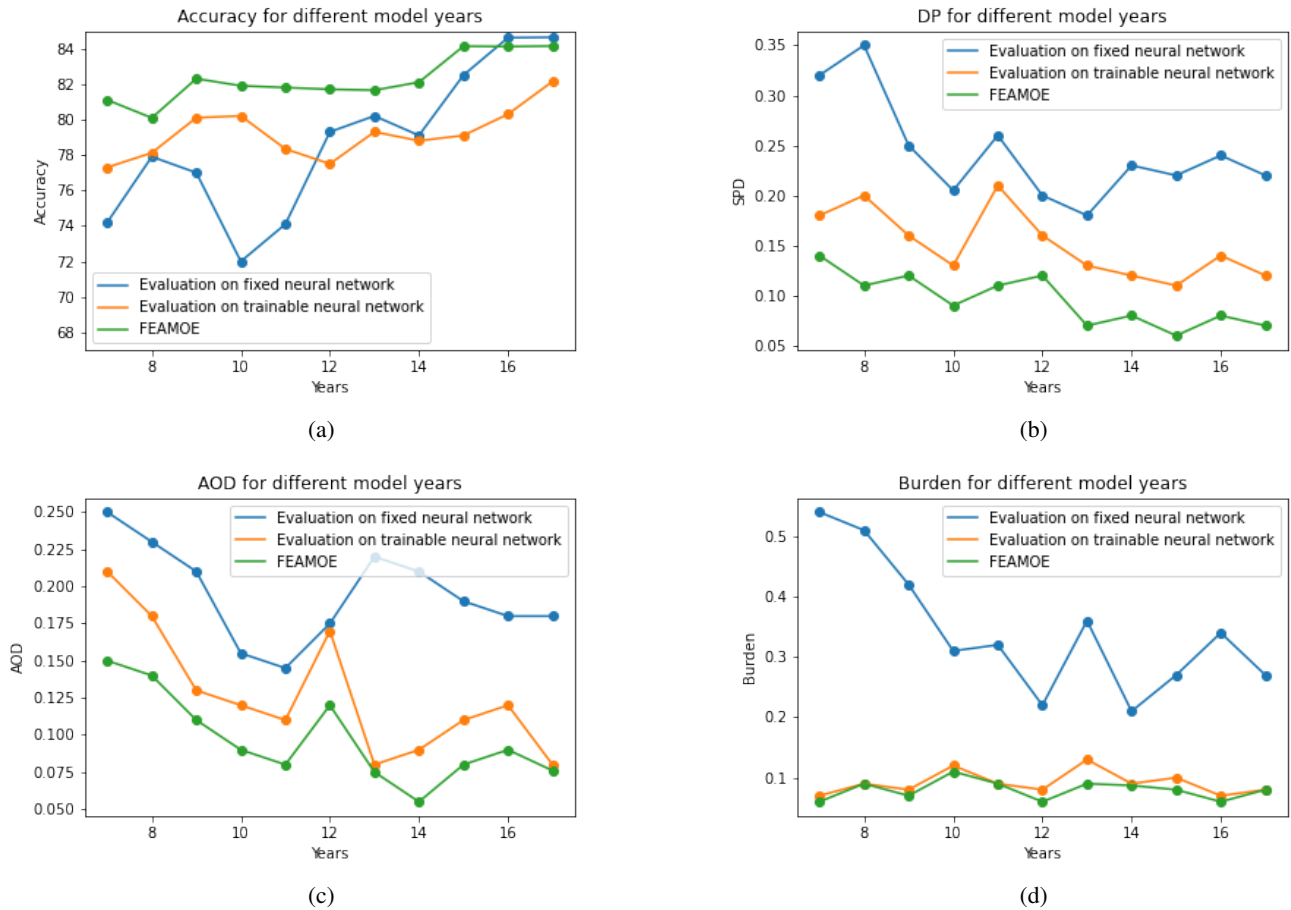


Figure 3: Drift Handling on the HMDA dataset. 1) Blue: baseline neural networks (fixed neural network) trained without fairness constraints on a previous year (20XX, indicated by x-axis; 2016 and 2017 are the "future" years) and not updated with new data; 2) Orange: (fair and trainable neural network: Neural networks with fairness constraints incorporated; also updated with streaming data from the "future" years and 3) Green: FEAMOE, also update with a single pass on streaming data from the "future" years).

els such as popular ensemble approaches (XGBoost etc), that need multiple passes through new data in batch mode after the initial model is built and deployed. For FEAMOE, we train models for each year separately and consider each of those as one expert. Then, we add on experts based on a single pass on new data (from the 2016-2017 years). We compare this to neural networks (with fairness constraints) that are also trained on each of the past years in batch mode, and then updated online with a single pass on the new data. The FEAMOE results are shown in Figure 3 by the green points, and the trainable neural networks results by orange points. FEAMOE is noticeably better at maintaining good accuracy and keeping lower bias across all bias metrics, irrespective of how old the original model was, even when compared to adaptive neural networks. The loosely coupled architecture and adaptive model complexity is key to FEAMOE’s success in handling drift. Also, using the FEAMOE architecture provides for much faster Shapley value explanations compared to the trainable neural network.

Additional experiments and findings are reported in the supplementary material. Some key information includes a com-

parison to other methods, further details on hyperparameters, experiments on fast Shapley value explanations, experiments on the COMPAS dataset for the fairness study, and on the synthetic streaming version of the UCI Adult dataset.

5 Conclusion and Future Work

We propose FEAMOE: a novel mixture of experts architecture and learning framework that can better maintain the fairness of a model in the face of data drift. We show how three fairness constraints can be incorporated into this framework. We prove that by using this mixture of experts, Shapley value explanations can be computed efficiently even though the overall model is non-linear. Experiments are performed on three datasets to demonstrate the various properties and effectiveness of FEAMOE. In particular, we identified a large-scale, real-world dataset that induces drift with respect to fairness over time in non-adaptive models, and show that our framework can adequately address this challenge. We would now like to extend it to incorporate other potential forms of drift, such as those that cause changes in adversarial robustness.

References

- [Aas *et al.*, 2019] Kjersti Aas, Martin Jullum, and Anders Løland. Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *arXiv preprint arXiv:1903.10464*, 2019.
- [Agarwal *et al.*, 2018] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *International Conference on Machine Learning*, pages 60–69. PMLR, 2018.
- [Barocas *et al.*, 2019] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness and machine learning. <http://www.fairmlbook.org>, 2019. Accessed: 2023-06-05.
- [Bechavod *et al.*, 2020] Yahav Bechavod, Christopher Jung, and Steven Z Wu. Metric-free individual fairness in on-line learning. *Advances in neural information processing systems*, 33:11214–11225, 2020.
- [Bellamy *et al.*, 2018] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, et al. Ai fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943*, 2018.
- [Bhaskaruni *et al.*, 2019] Dheeraj Bhaskaruni, Hui Hu, and Chao Lan. Improving prediction fairness via model ensemble. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1810–1814. IEEE, 2019.
- [Bhatt *et al.*, 2020] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José MF Moura, and Peter Eckersley. Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 648–657, 2020.
- [Biswas and Mukherjee, 2020] Arpita Biswas and Suvam Mukherjee. Ensuring fairness under prior probability shifts. *arXiv preprint arXiv:2005.03474*, 2020.
- [Bureau, 2020] Consumer Finance Protection Bureau. Hmda dataset. [Online]. Available from: <https://www.consumerfinance.gov/data-research/hmda/historic-data>, 2020. Accessed 2023-02-05.
- [Burkart and Huber, 2021] Nadia Burkart and Marco F Huber. A survey on the explainability of supervised machine learning. *Journal of Artificial Intelligence Research*, 70:245–317, 2021.
- [Butterworth, 2018] Michael Butterworth. The ico and artificial intelligence: The role of fairness in the gdpr framework. *Computer Law & Security Review*, 34(2):257–268, 2018.
- [Calmon *et al.*, 2017] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pages 3992–4001, 2017.
- [Gama *et al.*, 2014] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014.
- [Grgić-Hlača *et al.*, 2017] Nina Grgić-Hlača, Muhammad Bilal Zafar, Krishna P Gummadi, and Adrian Weller. On fairness, diversity and randomness in algorithmic decision making. *arXiv preprint arXiv:1706.10208*, 2017.
- [Hacker, 2018] Philipp Hacker. Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under eu law. *Common Market Law Review*, 55(4):1143–1185, 2018.
- [Jacobs *et al.*, 1991] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [Karimi *et al.*, 2020] Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. A survey of algorithmic recourse: definitions, formulations, solutions, and prospects. *arXiv preprint arXiv:2010.04050*, 2020.
- [Kohavi, 1996] Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *Kdd*, volume 96, pages 202–207. Citeseer, 1996.
- [Lu *et al.*, 2018] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12):2346–2363, 2018.
- [Lundberg and Lee, 2017a] Scott Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*, 2017.
- [Lundberg and Lee, 2017b] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pages 4765–4774. Curran Associates, Inc., 2017.
- [Lundberg *et al.*, 2018] Scott M Lundberg, Gabriel G Erion, and Su-In Lee. Consistent individualized feature attribution for tree ensembles. *arXiv preprint arXiv:1802.03888*, 2018.
- [Madras *et al.*, 2017] David Madras, Toniann Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. *arXiv preprint arXiv:1711.06664*, 2017.
- [Masoudnia and Ebrahimpour, 2014] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42(2):275–293, 2014.
- [Mehrabi *et al.*, 2019] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- [Molnar, 2019] Christoph Molnar. Interpretable machine learning. <https://christophm.github.io/interpretable-ml-book/>, 2019. Accessed: 2023-06-05.

- [Nejdl,] Wolfgang Nejdl. Farf: A fair and adaptive random forests classifier. In *Advances in Knowledge Discovery and Data Mining: 25th Pacific-Asia Conference, PAKDD 2021, Virtual Event, May 11–14, 2021, Proceedings, Part II*, page 245. Springer Nature.
- [Ntoutsis *et al.*, 2020] Eirini Ntoutsis, Pavlos Falalios, Ujwal Gadiraju, Vasileios Iosifidis, Wolfgang Nejdl, Maria-Esther Vidal, Salvatore Ruggieri, Franco Turini, Symeon Papadopoulos, Emmanouil Krasanakis, et al. Bias in data-driven artificial intelligence systems—an introductory survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(3):e1356, 2020.
- [ProPublica, 2016] ProPublica. Compas dataset. [Online]. Available from: <https://www.propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis>, 2016. Accessed: 2023-02-05.
- [Riquelme *et al.*, 2021] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *arXiv preprint arXiv:2106.05974*, 2021.
- [Robins, 1995] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995.
- [Shapley, 1953] Lloyd S Shapley. A value for n-person games. In *Contributions to the Theory of Games II*, pages 307–317. 1953.
- [Sharma *et al.*, 2020a] Shubham Sharma, Alan H Gee, David Paydarfar, and Joydeep Ghosh. Fair-n: Fair and robust neural networks for structured data. *arXiv preprint arXiv:2010.06113*, 2020.
- [Sharma *et al.*, 2020b] Shubham Sharma, Yunfeng Zhang, Jesús M Ríos Aliaga, Djallel Bouneffouf, Vinod Muthusamy, and Kush R Varshney. Data augmentation for discrimination prevention and bias disambiguation. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 358–364, 2020.
- [Slack *et al.*, 2020] Dylan Slack, Sorelle A Friedler, and Emile Givental. Fairness warnings and fair-maml: learning fairly with minimal data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 200–209, 2020.
- [Stanley, 2003] Kenneth O Stanley. Learning concept drift with a committee of decision trees. *Informe técnico: UT-AI-TR-03-302, Department of Computer Sciences, University of Texas at Austin, USA*, 2003.
- [Sundararajan *et al.*, 2017] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70 (ICML 2017)*, pages 3319–3328. Journal of Machine Learning Research, 2017.
- [Yuksel *et al.*, 2012] Seniha Esen Yuksel, Joseph N Wilson, and Paul D Gader. Twenty years of mixture of experts. *IEEE transactions on neural networks and learning systems*, 23(8):1177–1193, 2012.
- [Zhang and Ntoutsis, 2019] Wenbin Zhang and Eirini Ntoutsis. Faht: an adaptive fairness-aware decision tree classifier. *arXiv preprint arXiv:1907.07237*, 2019.
- [Zhang *et al.*, 2018] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340. ACM, 2018.
- [Zhang *et al.*, 2021] Wenbin Zhang, Albert Bifet, Xiangliang Zhang, Jeremy C Weiss, and Wolfgang Nejdl. Farf: A fair and adaptive random forests classifier. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 245–256. Springer, 2021.