

2021

## Underestimation Bias and Underfitting in Machine Learning

Padraig Cunningham

Sarah Jane Delany

Follow this and additional works at: <https://arrow.tudublin.ie/scschcomcon>



Part of the [Engineering Commons](#)

---

This Conference Paper is brought to you for free and open access by the School of Computer Sciences at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact [arrow.admin@tudublin.ie](mailto:arrow.admin@tudublin.ie), [aisling.coyne@tudublin.ie](mailto:aisling.coyne@tudublin.ie), [gerard.connolly@tudublin.ie](mailto:gerard.connolly@tudublin.ie).



This work is licensed under a [Creative Commons Attribution-NonCommercial-Share Alike 4.0 License](#)

# Underestimation Bias and Underfitting in Machine Learning

Pádraig Cunningham<sup>1</sup> and Sarah Jane Delany<sup>2</sup>

<sup>1</sup> University College Dublin

`padraig.cunningham@ucd.ie`

<sup>2</sup> Technological University Dublin

`sarahjane.delany@tudublin.ie`

**Abstract.** Often, what is termed algorithmic bias in machine learning will be due to historic bias in the training data. But sometimes the bias may be introduced (or at least exacerbated) by the algorithm itself. The ways in which algorithms can actually accentuate bias has not received a lot of attention with researchers focusing directly on methods to eliminate bias - no matter the source. In this paper we report on initial research to understand the factors that contribute to bias in classification algorithms. We believe this is important because underestimation bias is inextricably tied to regularization, i.e. measures to address overfitting can accentuate bias.

**Keywords:** Machine Learning · Algorithmic Bias

## 1 Introduction

Research on bias in Machine Learning (ML) has focused on two issues; how to measure bias and how to ensure fairness [12]. In this paper we examine the contribution of the classifier algorithm to bias. An algorithm would be biased if it were more inclined to award desirable outcomes to one side of a sensitive category. The desirable outcome could be a loan or a job, the category might be gender or race. It is clear that there are two main sources of bias in classification [10]:

- **Negative Legacy:** the bias is there in the training data, either due to poor sampling, incorrect labeling or discriminatory practices in the past.
- **Underestimation:** the classifier under-predicts an infrequent outcome for the minority group. This can occur when the classifier *underfits* the data, thereby focusing on strong signals in the data and missing more subtle phenomena.

In most cases the data (negative legacy) rather than the algorithm itself is the root of the problem. This question is neatly sidestepped in most fairness research by focusing on fair outcomes no matter what is the source of the problem.

We argue that it is useful to explore the extent to which algorithms accentuate bias because this issue is inextricably tied to regularisation, a central issue in ML.

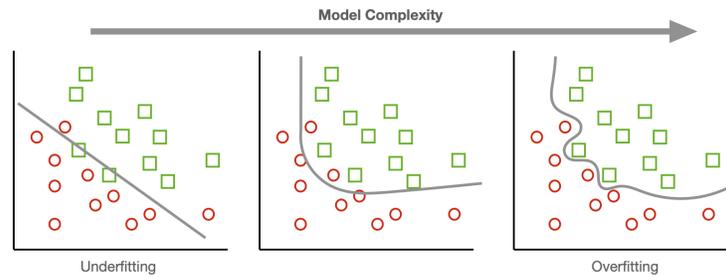


Fig. 1: The relationship between model fitting and model complexity. The fitted model is shown as a grey line that separates two classes in the data (red circles and green squares). Overfitted models can fit to noise in the training data while underfitted models miss a lot of the detail.

In developing ML models a key concern is to avoid overfitting. Overfitting occurs when the model fits to noise in the training data, thus reducing generalisation performance (see Figure 1). Regularisation controls the complexity of the model in order to reduce the propensity to overfit. The way regularisation is achieved depends on the model. The complexity of decision trees can be controlled by limiting the number of nodes in the tree. Overfitting in neural networks can be managed by limiting the magnitude of the weights. In this paper we discuss models at the other end of the complexity spectrum that underfit the training data. We show that underestimation bias can occur with underfitted models.

## 2 Background

Before proceeding to the evaluation it is worth providing some background on the key concepts discussed in this paper. In the next section we provide some definitions relating to regularisation. In section 2.2 we review some relevant research on bias and fairness in ML. Section 2.3 introduces the concept of Illusory Correlation and discusses its relevance to algorithmic bias.

### 2.1 Regularisation

Regularisation is a term from Statistics that refers to policies and strategies to prevent overfitting. In training a supervised ML model, the objective is to achieve good generalisation performance on unseen data. This can be assessed by holding back some test data from the training process and using it to get an assessment of generalisation performance. The relative performance on these training and test sets provides some insight on underfitting and overfitting [8]:

- **Underfitting** occurs when a model does not capture the detail of the phenomenon; these models will have some errors on the training data (see Figure 1).

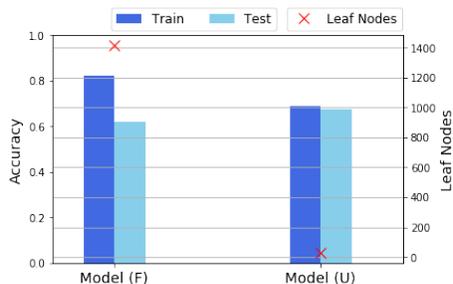


Fig. 2: A demonstration of underfitting Model(U) and overfitting Model(F) on the Recidivism task presented in section 3.2.

- **Overfitting** occurs when the model fits too closely to the training data picking up detail that does not generalise to unseen data. An overfitted model will have low error on the training data but higher error on test data.

Controlling overfitting is one of the central concerns in ML. It is difficult because it depends on the interplay between the volume of data available and the *capacity* of the model. If the training data provides comprehensive coverage of the phenomenon then overfitting is not a problem – it need only be an issue when the training data is not comprehensive. The capacity of the model refers to its expressive power; low capacity models will underfit while high capacity models can overfit. Figure 2 shows examples of under and overfitting on the Recidivism task described in section 3.2. In this case the model is a decision tree and the capacity is managed by controlling the number of leaf nodes. Limiting the number of leaf nodes to 30 results in an underfitted model whereas the overfitted model has been allowed to extend to 1413 leaf nodes. The overfitted model has good performance on the training data but does not generalise well. The underfitted model has better performance on the test (unseen) data.

## 2.2 Bias and Fairness

The connection between regularisation and bias arises from the context in which bias occurs. Consider a scenario where the desirable classification outcome is the minority class (e.g. job offer) and the sensitive feature represents groups in the population where minority groups may have low *base rates* for the desirable outcome [1]. So samples representing good outcomes for minority groups are scarce. Excessive regularisation causes the model to ignore or under-represent these data.

We are not aware of any research on bias in ML that explores the relationship between underestimation and regularisation. This issue has received some attention ([1,10]) but it is not specifically explored. Instead research has addressed algorithmic interventions that ensure fairness as an outcome[14,2,12,15].

Fundamental to all research on bias in ML are the measures to quantify bias. We define  $Y$  to be the outcome (class variable) and  $Y = 1$  is the ‘desirable’

outcome.  $\hat{Y} = 1$  indicates that the classifier has predicted the desirable outcome.  $S$  is a ‘sensitive’ feature and  $S \neq 1$  is the minority group (e.g. non-white). In this notation the Calders Verwer *discrimination score*[2] is:

$$CV \leftarrow P[\hat{Y} = 1|S = 1] - P[\hat{Y} = 1|S \neq 1] \quad (1)$$

If there is no discrimination, this score should be zero.

In contrast the *disparate impact* ( $DI_S$ ) definition of unfairness [7] is a ratio rather than a difference:

$$DI_S \leftarrow \frac{P[\hat{Y} = 1|S \neq 1]}{P[\hat{Y} = 1|S = 1]} < \tau \quad (2)$$

$\tau = 0.8$  is the 80% rule, i.e. outcomes for the minority should be within 80% of those for the majority.

For our evaluations we define an *underestimation score* ( $US_S$ ) in line with  $DI_S$  that quantifies the underestimation effect described above:

$$US_S \leftarrow \frac{P[\hat{Y} = 1|S \neq 1]}{P[Y = 1|S \neq 1]} \quad (3)$$

If  $US_S < 1$  the classifier is predicting fewer desirable outcomes than are present in the training data.

### 2.3 Illusory Correlation

Algorithmic bias due to underestimation is similar in some respects to the concept of Illusory Correlation in Psychology [3]. With Illusory Correlation people associate the frequent class with the majority group and the rare class with the minority, the infrequent class is overestimated for the minority group [5]. For instance, if the infrequent class is antisocial behaviour, the incidence will be over-associated with the minority. If the frequent class is a good credit rating, it will be over-associated with the majority. This over-estimation also happens for the frequent class in ML. However, the impact for the infrequent class and the minority feature is the opposite; the algorithm will likely accentuate the under-representation. In ML, the infrequent class outcome will be underestimated for the minority feature.

In Figure 3 we see the similarities and differences between Illusory Correlation and ML bias. For the majority feature and the frequent class the behaviour is the same, the association is accentuated. However, the impact on the infrequent side of the classification is different, the ML bias is always *towards* the frequent class. This is a problem if the infrequent outcome is the desirable one (e.g. loan approval, job offer). This difference is probably due to the different underlying mechanisms in humans and computers. Hamilton and Gifford [9] argue that Illusory Correlation is due to cognitive mechanisms that produce distortions in judgement and contribute to the formation of stereotypes. Whereas we argue here that underestimation results from a combination of poor data coverage and restricted model capacity.

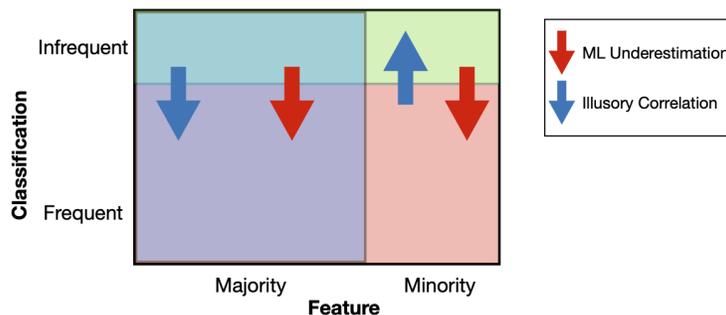


Fig. 3: The relationship between Illusory Correlation and ML Underestimation. Illusory Correlation accentuates the frequent class for the majority and the infrequent class for the minority. With ML Underestimation the bias is all towards the Frequent outcome.

### 3 Evaluation

We present preliminary results on three binary classification datasets from the UCI repository<sup>3</sup>. We also include an analysis on an anonymised and reduced version of the *ProPublica* Recidivism dataset [11,6]. Our analysis is based on the same set of seven features as used in the study by Dressel and Farid [6]. The UCI datasets are *Adult*, *Wholesale Customers* and *User Knowledge Modelling* (see Table 1).

We use the *Adult* and *Recidivism* datasets to demonstrate the relationship between underestimation and regularisation (section 3.2) and then we explore this in more detail in the other two datasets. In this second part of the evaluation, in order to control the propensity for bias, we introduce a sensitive binary feature  $S$  with  $S \neq 0$  representing the sensitive minority. To show bias due to underestimation we set  $P[S \neq 1|Y = 1] = 0.15$  and  $P[S \neq 1|Y \neq 1] = 0.3$ , i.e. the sensitive group is under-represented in the desirable class. We also report baseline results with the incidence of  $S \neq 1$  balanced with the sensitive group occurring 30% of the time in both classes. We are interested in uncovering situa-

Table 1: The datasets used in the evaluation. For consistency, the Positive outcome is the minority class in all datasets.

Dataset	Samples	Features	% Positive	Train : Test
Adult	48,842	14	25%	2:1
Recidivism	7,214	7	45%	2:1
Knowledge	403	5	26%	1:2
Wholesale	440	7	32%	1:1

<sup>3</sup> [archive.ics.uci.edu/ml/](http://archive.ics.uci.edu/ml/)

tions where this under-representation is accentuated by the classifier. In section 3.3 we look at the untuned performance of seven classifiers from `scikit-learn`. In section 3.4 we look at the impact of regularisation on neural net performance. In the next section we provide a more formal account of our bias measures.

### 3.1 Evaluation Measures

If a binary classifier is biased this will show up as a mismatch between the predicted positives  $P'$  and the actual positives  $P$  (see Table 2). This is likely to happen when the training data is significantly imbalanced resulting in predictions that underestimate the overall minority class [4]. From this perspective, bias is

Table 2: Confusion matrix for binary classification.

		Predicted		
		Pos	Neg	
Actual	Pos	TP	FN	P
	Neg	FP	TN	N
		$P'$	$N'$	

independent of accuracy so whether predictions are correct or not (True or False) is not relevant. So the minority class bias is effectively an underestimation at a class level, i.e.:

$$US \leftarrow \frac{TP+FP}{TP+FN} = \frac{P'}{P} \quad (4)$$

When the classifier is biased away from the positive class,  $P' < P$  this US score is less than 1. The  $US_S$  score defined in equation 3 measures the same fraction, but for samples where the sensitive feature  $S \neq 1$ .

$$US_S \leftarrow \frac{P'_{S \neq 1}}{P_{S \neq 1}} \quad (5)$$

While bias can be considered independently of accuracy, there is an important interplay between bias and accuracy. So the final evaluation measure we consider is the overall accuracy:

$$Acc \leftarrow \frac{TP+TN}{TP+FN+FP+FN} \quad (6)$$

### 3.2 Underestimation in Action

We use the `Adult` and `Recidivism` datasets to show the impact of underestimation. In Table 3 we see that in the `Adult` dataset Females with salaries greater than 50K account for just 4% of cases and in the `Recidivism` dataset Caucasians are relatively underrepresented among repeat offenders. We will see that this under-representation is accentuated by the classifiers.

Table 3: Summary statistics for the Adult and Recidivism datasets. In both cases a feature/class combination is significantly underrepresented in the data. Key feature-specific percentages are shown in red, for instance 11% of females are in the > 50K salary category.

Adult				Recidivism			
Salary	Gender			Recidivism	Caucasian		
	Male	Female	Totals		No	Yes	Totals
	< 50k	47%	29%		76%	No	34%
> 50k	20% (30%)	4% (11%)	24%	Yes	32%(48%)	13%(39%)	45%
Totals	67%	33%	100%	Totals	66%	34%	100%

**Adult Dataset:** This dataset is much studied in research on bias in ML because there is clear evidence of Negative Legacy [2]. At 33%, females are underrepresented in the dataset. This under-representation is worse in the > 50K salary category where only 1/6 are female.

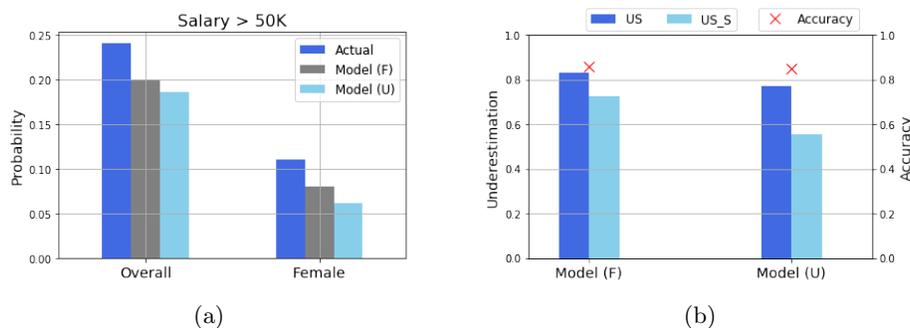


Fig. 4: A demonstration of model bias and underestimation on the **Adult** dataset. (a) 24% of all cases and 11% of Females have Salary > 50K. Both models underestimate these outcomes with the underfitted model Model (U) underestimating most.

(b) Both models have similar accuracies and both underestimate overall (US). The underestimation for Females (US\_S) is worse and is worst in the underfitted model Model (U).

To illustrate underestimation we use a gradient boosting classifier [13]. We build two classifiers, one with just 5 trees (**Model (U)**) and one with 50 (**Model (F)**). Both models have good accuracy, 85% and 86% respectively. Figure 4(a) shows the actual incidence of Salary > 50 overall and for females. It also shows the predicted incidence by the two models. We can see that both models underestimate the probability of the Salary > 50 class overall. On the right in Figure 4(a) we can see that this underestimation is exacerbated for females. This under-

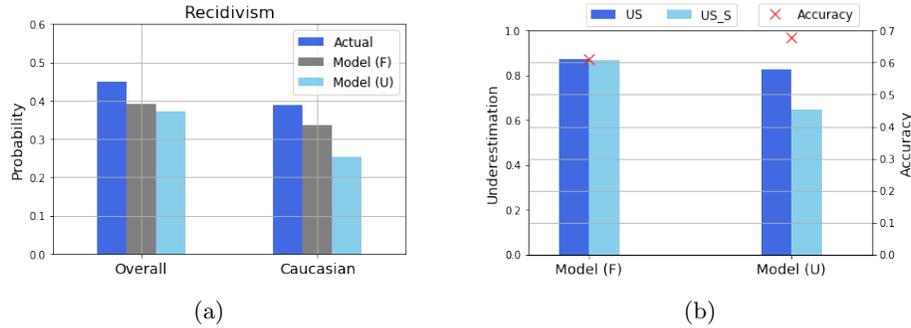


Fig. 5: A demonstration of bias and underestimation on the Recidivism dataset. (a) 46% of all cases and 38% of Caucasians are recidivists. Both models underestimate these outcomes with the underfitted model Model (U) underestimating most. (b) The underfitted model has better generalisation accuracy. Both models underestimate overall (US) with the underfitted model underestimating more for Caucasians (US.S).

estimation is worse in the underfitted model. The actual occurrence of salaries  $> 50K$  for females is 11% in the data, the underfitted model is predicting 6%. The extent of this underestimation is quantified in Figure 4(b).

So underestimation is influenced by three things, underfitting, minority class and minority features. The underfitted model does a poor job of modelling the minority feature (female) in the minority class ( $> 50K$ ). This is not that easy to fix because it is often desirable not to allow ML models overfit the training data.

**Recidivism Dataset:** We use the decision tree classifier from `scikit-learn`<sup>4</sup> to demonstrate underestimation on the Recidivism dataset. In this case we control overfitting by constraining the size of the tree. The underfitted model (Model (U)) has 30 leaf nodes and the other model (Model (F)) has 1349 leaves. The picture here is similar but in this case the underfitted model has better accuracy. This accuracy comes at the price of increased underestimation. The underestimation is particularly bad for the minority feature with the level of recidivism for Caucasians significantly underestimated. As reported in other analyses [11,6] the input features do not provide a strong signal for predicting recidivism. So the fitted model does not generalise well to unseen data. On the other hand the model that is forced to underfit generalises better but fails to capture the Caucasian recidivism pattern.

<sup>4</sup> `scikit-learn.org`

### 3.3 Baseline Performance of Classifiers

We move on now to look at the impact of a synthetic minority feature injected into the other two datasets. This synthetic feature is set up to be biased (negative legacy)  $P[S \neq 1|Y = 1] = 0.15$  and  $P[S \neq 1|Y \neq 1] = 0.3$ . We test to see if this bias is accentuated (i.e.  $US_S < 0$ ) for seven popular classifiers available in `scikit-learn`. For this assessment, the default parameters for the classifiers are used. There are two exceptions to this; the number of iterations for Logistic Regression and the Neural Network were increased to get rid of convergence warnings. For the results shown in Figure 6 the main findings are:

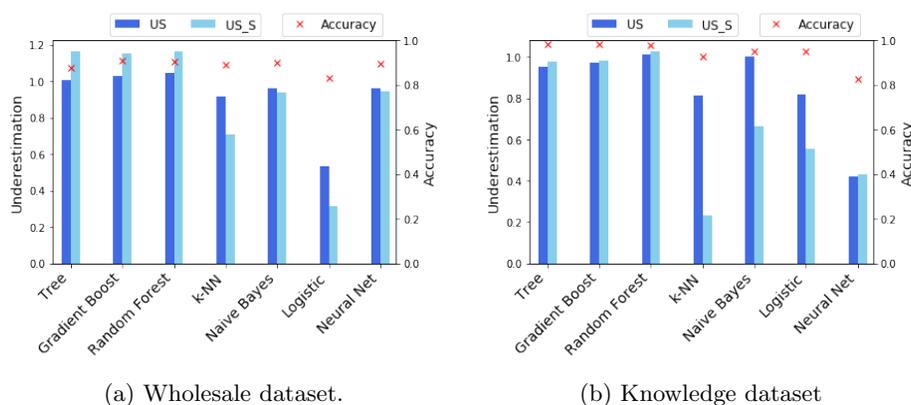


Fig. 6: The varying impact of underestimation across multiple classifiers. A sensitive feature  $S \neq 1$  has been added to both datasets (15% in the desirable class and 30% in the majority class). With the exception of the three tree-based classifiers, underestimation is exacerbated for this sensitive feature (US\_S).

- The tree-based classifiers (Decision Tree, Gradient Boost & Random Forest) all perform very well showing no bias (or a positive bias), both overall (US) and when we filter for the sensitive feature (US\_S).
- The other classifiers ( $k$ -Nearest Neighbour, Naive Bayes, Logistic Regression & Neural Network) all show bias (underestimation), overall and even more so in the context of the sensitive feature.
- The accentuation effect is evident for the four classifiers that show bias, i.e. they predict even less than 15% of the desirable class for the sensitive feature.

This base-line performance by the tree based classifiers is very impressive. However, it is important to emphasise that the performance of the other methods can be improved significantly by parameter tuning – as would be normal in configuring an ML system. Finally, it should not be inferred that tree-based methods are likely to be free of underestimation problems. In particular, the Decision Tree implementation in `scikit-learn` provides a number of mechanisms for regularisation that may introduce underestimation as a side effect.

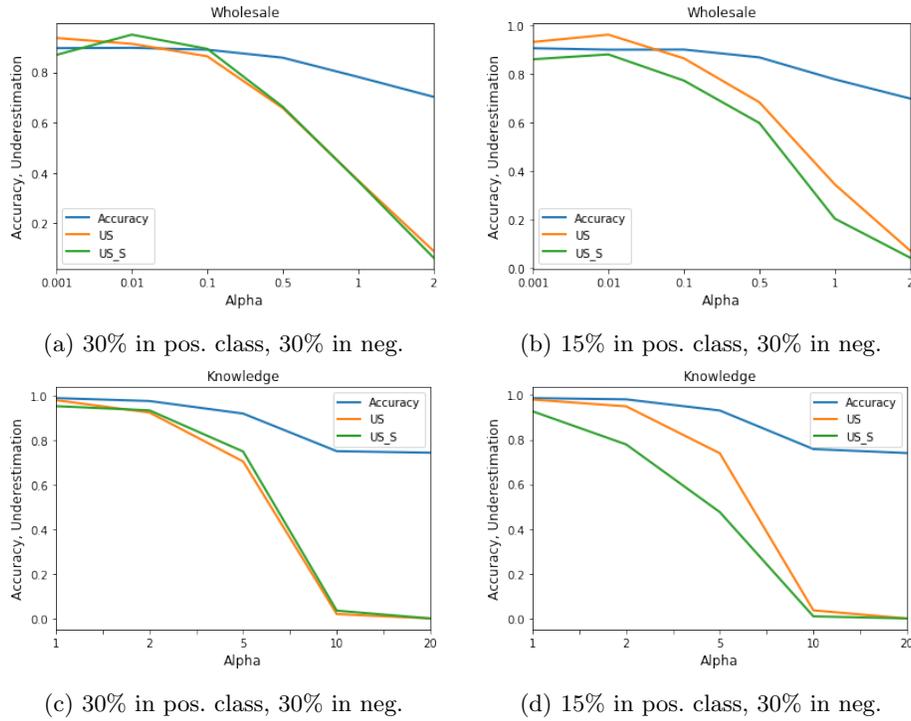


Fig. 7: The impact of underfitting on bias. A synthetic sensitive attribute has been added, it is underrepresented in the plots on the right (15% compared with 30%). Higher values of Alpha result in underfitting. Furthermore, when the sensitive attribute is underrepresented in the desirable class (b) & (d) the bias is exacerbated.

### 3.4 Impact of Underfitting

The standard practice in training a classifier is to ensure against overfitting in order to get good generalisation performance. Kamishima *et al.* [10] argue that bias due to underestimation arises when a classifier underfits the phenomenon being learned. This will happen when the data available is limited and samples covering the sensitive feature and the desirable class are scarce.

The `scikit-learn`<sup>5</sup> neural network implementation provides an  $\alpha$  parameter to control overfitting. It works by providing control over the size of the weights in the model. Constraining the weights reduces the potential for overfitting. The plots in Figure 7 show how underestimation varies with this  $\alpha$  – high values cause underfitting. These plots show Accuracy and Underestimation for different values of  $\alpha$ . For the plots on the left (Figure 7 (a)&(c)) the incidence of the sensitive

<sup>5</sup> [scikit-learn.org](https://scikit-learn.org)

feature is the same in both the positive and negative class (30%). For the plots on the right ((b)&(d)) the sensitive feature is underrepresented (15%) in the positive class.

In Figure 7 (a) and (c) we see that high values of  $\alpha$  (i.e. underfitting) result in significant bias. When the base rates for the minority group in the positive and negative classes are the same the US and  $US_S$  rates are more or less the same.

When the prevalence of the sensitive group in the desirable class is lower ((b)&(d)) the bias is exacerbated. It is important to emphasise that a good  $US_S$  score simply means that underestimation is not present. There may still be bias against the minority group (i.e. poor CV or  $DI_S$  scores).

## 4 Conclusions & Future Work

In contrast to what Illusory Correlation tells us about how humans perceive things, underestimation occurs in ML classifiers because they are inclined to over-predict common phenomena and under-predict things that are rare in the training data. We have shown, on two illustrative datasets, **Adult** and **Recidivism**, how the impact of under-representation in data leads to underestimation of the classifiers built on that data. We believe classifier bias due to underestimation is worthy of research because of its close interaction with regularisation. We have demonstrated this interaction on two datasets where we vary the levels of under-representation and regularisation showing the impact on underestimation. Underfitting data with an under-represented feature in the desirable class leads to increased underestimation or bias of the classifiers.

We are now exploring how sensitive underestimation is to distribution variations in the class label and the sensitive feature. Our next step is to develop strategies to mitigate underestimation.

## Acknowledgements

This work was funded by Science Foundation Ireland through the SFI Centre for Research Training in Machine Learning (Grant No. 18/CRT/6183)

## References

1. S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning*. fairml-book.org, 2019. <http://www.fairmlbook.org>.
2. T. Calders and S. Verwer. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.
3. L. J. Chapman and J. P. Chapman. Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of abnormal psychology*, 74(3):271, 1969.
4. N. V. Chawla, A. Lazarevic, L. O. Hall, and K. W. Bowyer. Smoteboost: Improving prediction of the minority class in boosting. In *European conference on principles of data mining and knowledge discovery*, pages 107–119. Springer, 2003.

5. F. Costello and P. Watts. The rationality of illusory correlation. *Psychological review*, 126(3):437, 2019.
6. J. Dressel and H. Farid. The accuracy, fairness, and limits of predicting recidivism. *Science advances*, 4(1):eaao5580, 2018.
7. M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 259–268, 2015.
8. I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
9. D. L. Hamilton and R. K. Gifford. Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12(4):392–407, 1976.
10. T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer, 2012.
11. J. Larson, S. Mattu, L. Kirchner, and J. Angwin. How we analyzed the compas recidivism algorithm. *ProPublica (5 2016)*, 9, 2016.
12. A. K. Menon and R. C. Williamson. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency*, pages 107–118, 2018.
13. L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin. Catboost: unbiased boosting with categorical features. In *Advances in neural information processing systems*, pages 6638–6648, 2018.
14. R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.
15. Y. Zhang and L. Zhou. Fairness assessment for artificial intelligence in financial industry. *arXiv preprint arXiv:1912.07211*, 2019.