

UNIVERSITY OF CALIFORNIA

Los Angeles

Incorporating minority groups for the improvement of AI performance in majority groups

A thesis submitted in partial satisfaction

of the requirements for the degree

Master of Science in Electrical and Computer Engineering

by

Shreeram Seshathri Athreya

2023

© Copyright by  
Shreeram Seshathri Athreya  
2023

## ABSTRACT OF THE THESIS

Incorporating minority groups for the improvement of AI performance in majority groups

by

Shreeram Seshathri Athreya

Master of Science in Electrical and Computer Engineering

University of California, Los Angeles, 2023

Professor Achuta Kadambi, Chair

Numerous studies have rightly emphasized the importance of incorporating minority groups into artificial intelligence (AI) training data to enhance test inference not only for minority groups but also for society as a whole. A comprehensive society includes both minority and majority stakeholders. A common misunderstanding is that the inclusion of minority groups does not lead to improved performance solely for majority groups. In this thesis, I make a remarkable discovery that incorporating data from minority groups can, in fact, result in a reduction in test errors for the majority group. Put simply, the integration of minority groups results in performance improvements for the majority group, an effect termed as Majority Group Enhancements through Minority Inclusion (MIME). To support this finding, I present a theoretical existence proof of the MIME effect, which demonstrates that this phenomenon is not only possible but also well-founded. The theoretical results align with experimental outcomes gathered from six diverse datasets, further validating the MIME effect's existence and relevance. By acknowledging and incorporating the MIME effect in AI development, we can create more robust and effective systems that cater to the needs of both minority and majority populations. This approach enables us to harness the full potential of artificial

intelligence by ensuring that it is more inclusive and efficient, ultimately benefiting society as a whole. An understanding of the MIME effect has substantial implications for the future of AI research and development, as it underscores the importance of fostering diversity in training data. By embracing a more inclusive approach to AI development, we can ensure that our models are better equipped to handle a wider range of scenarios, ultimately leading to more accurate and reliable outcomes. Additionally, the MIME effect has the potential to inspire new research directions, encouraging scholars to explore the benefits of diverse data sets in various AI applications and domains. It is crucial that the AI community recognizes the value of minority inclusion in achieving performance enhancements for all stakeholders, and actively works towards creating systems that are truly representative of the diverse society we live in. By doing so, we can not only maximize the potential of AI technologies, but also contribute to a more equitable and just society where artificial intelligence serves the needs of everyone, regardless of their background or status. This paradigm shift will not only revolutionize the field of AI, but also pave the way for a more inclusive and sustainable future for all.

The thesis of Shreeram Seshathri Athreya is approved.

Jonathan Kao

Vwani Roychowdhury

Achuta Kadambi, Committee Chair

University of California, Los Angeles

2023

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Contributions	3
1.2	Outline of Theoretical Scope	4
<b>2</b>	<b>Background</b>	<b>5</b>
2.1	Debiasing and fairness	5
2.2	Learning from multiple domains	5
2.3	Dataset diversity	6
<b>3</b>	<b>Statistical Origins of the MIME Effect</b>	<b>7</b>
3.1	Key Result 1: A minority sample can be more valuable for majority classifiers than another majority sample	9
3.2	Key Result 2: MIME holds under domain gap	17
3.3	Key Result 3: MIME holds for general distributions	21
<b>4</b>	<b>Verifying MIME Theory on Real Tasks</b>	<b>24</b>
4.1	The Assumptions	24
4.1.1	Verifying Gaussianity	24
4.1.2	Verifying minority/majority definitions	25
4.2	MIME Effect Across Six, Real Datasets	26
4.2.1	Implementation	26
4.2.2	MIME effect on gender classification	26
4.2.3	MIME effect on animal species identification	27

4.2.4	MIME effect on age classification . . . . .	27
4.2.5	MIME effect on X-ray diagnosis Classification . . . . .	27
<b>5</b>	<b>Discussion . . . . .</b>	<b>33</b>
5.1	Secondary validation and analysis . . . . .	33
5.2	Optimality of inclusion ratios . . . . .	33
5.3	Limitations . . . . .	34
<b>A</b>	<b>Generalizations and Implementation . . . . .</b>	<b>35</b>
A.1	MIME Existence Beyond 1D Settings . . . . .	35
A.2	Feature Space Analysis . . . . .	37
A.2.1	Constructing the Projected Feature Histograms . . . . .	37
A.2.2	Estimating the Overlap . . . . .	38
A.2.3	Estimating the Domain Gap . . . . .	38
A.2.4	Notes on the Estimated Measures . . . . .	39
A.2.5	Analysis of Feature Space Gaussian-like Behavior . . . . .	40
A.3	Implementation Details . . . . .	41
A.3.1	Analysis measures . . . . .	41
A.3.2	Network Architectures Used . . . . .	43
A.3.3	General Experiment Details . . . . .	44
A.3.4	Dataset Specific Information . . . . .	45
<b>B</b>	<b>Additional Analysis . . . . .</b>	<b>46</b>
B.1	Secondary Analysis of MIME . . . . .	46
B.1.1	MIME effect with debiasing methods: . . . . .	46

B.1.2 Reconciling MIME with existing equal representation (ER) datasets:	46
B.2 Hard Mining Baseline Implementation . . . . .	47
B.3 Negative impacts and mitigation: . . . . .	48
<b>References . . . . .</b>	<b>50</b>



LIST OF FIGURES

1.1 **This thesis proves that including minorities improves majority performance.** The provable guarantees are certifiable for fixed backbone binary classification - when one uses a head network with pretrained weights and fine-tunes a downstream layer for classification. . . . . 2

1.2 **Inclusion of minorities can improve performance for majorities.** We theoretically describe an effect called Minority Inclusion, Majority Enhancement (MIME). The figure depicts test classification of blue mimes, and an initial training stack, also of blue mimes. If allowed to add one more training sample, it can be better to push an orange mime onto the training stack rather than a blue mime. Test accuracy can increase by pushing orange, even though the test set consists of blue mimes alone. . . . . 3

3.1 **Visualization of Gaussian Mixture Model parameters.** We plot GMMs with different task complexities. The domain gap  $\delta$  is visualized as the difference in the ideal threshold locations. The overlap/task complexity metric can be visually seen. . . . . 9

4.1 **The use of Gaussian mixtures to represent minority and majority distributions is consistent with behaviors in modern neural networks, on real-world datasets.** (top row) The last layer of common neural architectures is a linear classifier on features. Histograms of the penultimate layer projections are generated for models with  $\beta = 0.5$ . (middle row) Minority histograms: note the greater difficulty due to less separation of data. (bottom row) Majority histograms: note smaller overlap and easier classification. Figure can be parsed on a per-dataset basis. Within each column, the reader can compare the domain gap and overlap in the two histograms. . . . . 29

4.2	<b>When domain gap is small, the MIME effect holds.</b> On four vision datasets, majority performance is maximized with some inclusion of minorities. All experiments are run for several trials and realizations (described in Section 4.2).	30
4.3	<b>When domain gap is small, the MIME effect holds.</b> On four vision datasets, majority performance is maximized with some inclusion of minorities. All experiments are run for several trials and realizations (described in Section 4.2).	31
4.4	<b>MIME effect is observed in non-vision datasets, and is absent in the case of large domain gap.</b> (a) The Adult Dataset uses Census data to predict an income label. (b) On dataset six, gender classification is re-scoped to occur in a high domain gap setting. Majority group is chickens and minority group is humans. . . . .	32
A.1	<b>The MIME effect holds in a multidimensional setting as well.</b> We show the support for the two finite distributions. Weight vector updates arising out of samples from regions R3, R4, R5 and R6 lead to an update with a large vertical (corrective) component (favorable update). Updates arising out of regions R1 and R2 result in an overall update in the horizontal direction (unfavorable update).	36
B.1	<b>The MIME effect is complementary to data debiasing methods and consistent with research aimed at equal representation (ER) datasets.</b> (a) Training configurations using data debiasing methods show the MIME effect. (b) While ER datasets are not optimal for the MIME effect (Figures 4.2, 4.3 and 4.4), optimal overall performance is observed close to ER. . . . .	49

## LIST OF TABLES

4.1	<b>Experimental measures of overlap and domain gap are consistent with the theory in Chapter 3.</b> Note that the majority group consistently has lower overlap. Domain gaps are found to be small. DS-1 is FairFace, DS-2 is Pet Images, DS-4 is Chest-Xray14 and DS-5 is Adult. DS-6 is the high domain gap gender classification experiment. DS-3 is excluded here since it deals with a 9 class classification problem. . . . .	25
5.1	<b>Additional evaluation metrics provide further evidence of MIME existence across all datasets.</b> The table highlights: (i) number of trials with MIME performance gain (i.e. majority accuracy at some $\beta > 0$ is greater than majority accuracy at $\beta = 0$ ), and (ii) the mean MIME performance gain across trials (in % points). . . . .	34
A.1	<b>Chi-Squared goodness of fit measures for all distributions.</b> Distributions with bolded values show the estimated statistics that are lower than the critical value, indicating that the null hypothesis (Gaussian distribution) cannot be rejected.	42
A.2	<b>Training configuration and parameters for all datasets and experiments.</b> Parameters for each dataset are chosen so as to maximize performance.	44
A.3	<b>Random seeds used for the trials.</b> Seeds were chosen at random for trials to generate average trends and error bounds. . . . .	45

## ACKNOWLEDGMENTS

My sincere appreciation goes first to my thesis advisor, Professor Achuta Kadambi, whose expert guidance and unwavering support were instrumental in shaping both the direction and the success of this thesis. His willingness to engage in extensive discussions and provide valuable feedback at all hours has been a cornerstone of this journey. Next, I extend my gratitude to Professor Jonathan Kao and Professor Vwani Roychowdhury, whose roles on my thesis committee were marked by invaluable insights and support throughout this academic endeavor.

A special acknowledgment is owed to Pradyumna Chari, whose dedication extended beyond the call of duty. Our collaborative efforts have not only enriched this manuscript but also facilitated its progress towards publication. Pradyumna's influence and the knowledge shared have been instrumental in navigating the complexities of this project. I truly wish to work with him in the future. My heartfelt thanks also go to all my co-authors — Yunhao Ba, Ellin Zhao, Alexander (Sasha) Vilesov, Dr. Laleh Jalilian and colleagues from VMG — for their essential contributions and guidance.

I must express my deep gratitude to my family, friends, and peers, whose support has been a constant throughout this process. Finally, I reserve the most profound acknowledgment for my parents. Their endless love, encouragement, and belief in my capabilities have been the unwavering pillars of my strength. To my mother Saroja, in particular, I owe a debt of gratitude for her unconditional support. Her faith in me has been a source of motivation, even in the most trying of times. Her sacrifices have not gone unnoticed, and it is with her spirit in mind that I have pursued and completed this work. Thank you for being my constant inspiration and for instilling in me the resilience and determination I so greatly needed.

## PREVIOUS PUBLICATIONS

This thesis revises the following publication:

P. Chari, Y. Ba, S. Athreya, A. Kadambi, “MIME: Minority Inclusion for Majority Group Enhancement of AI Performance,” European Conference on Computer Vision, pp. 326-343, Springer, Cham, 2022 [1].

# CHAPTER 1

## Introduction

Inclusion of minorities in a dataset impacts the performance of artificial intelligence (AI). Recent research has presented the value of inclusive datasets to improve AI performance on minorities and also for society-at-large [2, 3, 4, 5, 6, 7, 8, 9, 10]. A society-at-large consists of both majority and minority stakeholders. However, an objection (often silently posed) to minority inclusion efforts, is that the inclusion of minorities can diminish performance for the majority. This is based on a “rule of thumb” that AI performance is maximized when one trains and tests on the same distribution. A devil’s advocate position against minority inclusion might be presented as: “In a fictitious society where we are absolutely certain that only blue-skinned humans will exist in the test set, why include out of distribution orange-skinned humans in the training set?”.

In this thesis, we make the surprising finding that inclusion of minority samples improves AI performance not just for minorities, not just for society-at-large, but *even for majorities*. We refer to this effect as Minority Inclusion, Majority Enhancement (MIME), illustrated in Figure 1.2. Specifically, we note that including some minority samples in the train set improves majority group test performance. However, continued addition of minority samples leads to performance drop. The effect holds under statistical conditions that are represented in traditional computer vision datasets including FairFace [11], UTKFace [12], pets [13], medical imaging datasets [14] and even non-vision data [15]. Although deep learning is used for these problems, the flattening layer of a network can be empirically approximated to elementary distributions like Gaussian Mixture Models (GMMs). A GMM facilitates

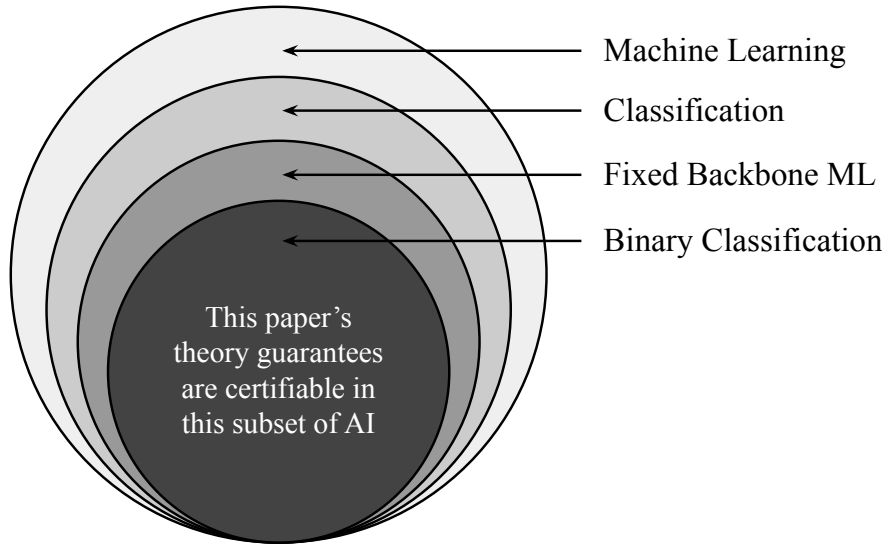


Figure 1.1: **This thesis proves that including minorities improves majority performance.** The provable guarantees are certifiable for fixed backbone binary classification - when one uses a head network with pretrained weights and fine-tunes a downstream layer for classification.

closed-form analysis to prove the existence of the MIME effect. Additionally, we show existence of MIME on general distributions. Classification experiments on neural networks validate using Gaussian mixtures: complex neural networks exhibit feature embeddings in flat layers, distributed with approximately Gaussian density, across six datasets, in and beyond computer vision, and across many realizations and configurations.

Fairness in machine learning is an exceedingly popular area, and our results benefit from several key papers published in recent years. Sample reweighting approaches recognize the need to preferentially weight difficult examples [16, 17, 18]. Active and online learning benefit from insights into sample “informativeness” (i.e. given a budget on the number of training samples, which would be the best sample to include [19, 20]). Domain randomization literature indicates that surprising perturbations to the training set can improve generalization performance [21, 22, 23]. We extend some of these theoretical insights to the sphere of analyzing benefits of minority inclusion on majority performance.




Figure 1.2: **Inclusion of minorities can improve performance for majorities.** We theoretically describe an effect called Minority Inclusion, Majority Enhancement (MIME). The figure depicts test classification of blue mimes, and an initial training stack, also of blue mimes. If allowed to add one more training sample, it can be better to push an orange mime onto the training stack rather than a blue mime. Test accuracy can increase by pushing orange, even though the test set consists of blue mimes alone.

## 1.1 Contributions

While some works [24, 4] have observed related phenomena for isolated tasks, to the best of our knowledge, characterizing benefits to majority groups by including minority data is largely unexplored theoretically. Our contributions are as follows:

- We introduce the Minority Inclusion Majority Enhancement (MIME) effect in a theoretical and empirical setting.
- Theoretically: we derive in closed form, the existence of the MIME effect both with and without domain gap (Key Results 1 and 2) and for general sample distributions (Key Result 3).
- Empirically: we test the MIME effect on six datasets, as varied as animals to medical images, and observe the existence of MIME consistent with theory.



## 1.2 Outline of Theoretical Scope

Figure 1.1 describes the theoretical scope. Through three key results (Theorem 1, Theorem 2 and Theorem 3), this thesis offers an existence proof of the MIME effect. An existence proof can leverage a tractable setting. As in Figure 1.2, training data is a stack of  $K - 1$  majority samples. Test data is all majority samples. We can push one additional training sample to increase the stack size to  $K$ . We are allowed the choice of having the  $K$ -th sample drawn from the minority or majority group. Theorem 1 proves that, under the assumptions in Chapter 3, pushing a minority sample is superior for majority group performance improvements. Theorem 2 generalizes this result to a more realistic scenario, with domain gap. Theorem 3 extends the existence proof to general sample distributions. Empirical results on real-world AI tasks offer validation for theoretical assumptions. From Figure 1.1, the provable guarantees are certifiable for fixed backbone binary classification. The fixed backbone ML is far from a toy scenario (it is considered SoTA by some authors [25]) and also enables provable certification - ordinarily it is hard to prove things for neural network settings.

# CHAPTER 2

## Background

### 2.1 Debiasing and fairness

It has been widely reported that biases in training data lead to biased algorithmic performance [26, 27, 3]. Work has been carried out in identifying and quantifying biases [28, 29, 30] and a range of methods exist to address them [9, 7]. Early approaches suggest oversampling strategies [31, 32]. Other methods propose resampling based on individual performance [6]. Some works utilize information bottlenecks to disentangle biased attributes [33]. Still other methods propose bias mitigation solutions based on adversarial learning [34] or include considerations like protected class-specific classifiers [35]. Generative models have also found use in creating synthetic datasets with debiased attributes [36]. Xu *et al.* [37] identify inherent bias amplification as a result of adversarial training and propose a framework to mitigate these biases. Our goals are different – while these aim to reduce test time performance bias across groups, we analyze influence of minority samples on majority group performance.

### 2.2 Learning from multiple domains

Domain adaptation literature explores learning from multiple sources [38]. It could therefore be one potential way to analyze our problem of training on combinations of majority and minority data. In our setting, data arising from distinct domains is seen as being drawn from different distributions with a domain gap [39]. Between these domains, [40] establishes error bounds for learning from combinations of domains. However, these error estimates

and bounds do not take into account the notion of majority and minority groups; therefore, describing the MIME effect is outside their scope.

## 2.3 Dataset diversity

An important push towards fairness is through analysis of dataset composition. Several works indicate the importance of diverse datasets [2, 8]. Ryu *et al.* [5] note that class imbalance in the training set leads to performance reduction. Wang *et al.* [30] highlight that perfectly balanced datasets may still not lead to balanced performance. For designing medical devices, [10] emphasizes the importance of diverse datasets. Through experiments on X-ray datasets, [4] observe that imbalanced training sets adversely affect performance on the disadvantaged group. They also observe that an unbiased training set shows the best overall accuracy. However, their inferences are related empirical observations on a few medical tasks and datasets. From an application perspective, the task of remote photoplethysmography enables analysis of the bias problem. Prior work notes that camera-based heart rate estimation exhibits skin tone bias [41], and [42, 43] propose synthetic augmentations to mitigate this. Additionally, [44, 45] establish that camera based heart rate estimation is fundamentally biased against dark skin tone subjects, establishing a notion of task complexity. While all these works recognize that data composition affects bias, none to our knowledge describe the effect of varying minority group proportions on majority group accuracy.

## CHAPTER 3

### Statistical Origins of the MIME Effect

For more concise exposition, we have made assumptions in the derivation outlined in this chapter of the thesis and defer extended generality to Appendix A. The assumptions are explained as follows:

- Assumption 1: one-dimensional data samples and binary labels,  $x \in \mathbb{R}$ ,  $y \in \{1, 2\}$ . This is relevant to modern classification problems since the final classification decision is based on a one dimensional projection of the feature representation of the sample with respect to the learnt hyperplane (illustrated in Figures 1.1, 4.1 and discussed in Chapter 4). Additionally, existence proof of MIME holds for more general vectorized notation, as discussed in Appendix A.
- Assumption 2: the binary classifier used is a perceptron: this assumption relates to real neural networks since the last layer is perceptron-like [46].

We now introduce some key definitions that follow from these assumptions.

**Definition 1:** (Task complexity): *For binary classification we define task complexity for a group of data  $\theta$  as a continuous variable in  $[0, 1]$ , such that,*

$$\theta = \arg \min_{h \in H} \epsilon(h), \tag{3.1}$$

*where  $\epsilon(h)$  is the classification error for hypothesis  $h$  (the classifier),  $H$  is the space of feasible hypotheses. It is noted later that this is empirically equivalent to distributional overlap. This*

definition is not new. Hard-sample mining [16] establishes the use of performance measures as an indicator of difficulty.

**Definition 2:** (Majority Group): Group class (i.e. group label  $g = \text{major}$ ) on which the task performs better. Quantified by training a network only with majority group data and evaluating test performance:  $\theta^{\text{major}} = \arg \min_{h \in H} \epsilon^{\text{major}}(h)$ .

**Definition 3:** (Minority Group): Group class (i.e. group label  $g = \text{minor}$ ) on which the task performs worse. Quantified by training a network only with minority class data and evaluating test performance:  $\theta^{\text{minor}} = \arg \min_{h \in H} \epsilon^{\text{minor}}(h)$ .

**Definition 4:** (Minority Training Ratio ( $\beta$ )): Ratio of minority to majority samples in the data under consideration (training set, in the context of this paper).

**Definition 5:** (MIME Domain Gap): Measure of how classification differs for minorities and majorities. Quantified as a difference between ideal hyperplanes. Note that this definition for domain gap could be different from other definitions. In this work, domain gap should be taken to mean MIME domain gap.

Empirical observations on cutting-edge machine learning tasks demonstrate the real-world applicability of the assumptions above. We now discuss three key results. For ease of understanding, we make two simplifying assumptions for Key Results 1 and 2: (i) simplified distributions that follow a symmetric Gaussian Mixture Model, and (ii) equally likely class labels, i.e.  $Pr(y = 1) = Pr(y = 2)$ . These assumptions are relaxed in Key Result 3.

### 3.1 Key Result 1: A minority sample can be more valuable for majority classifiers than another majority sample

Our first key result shows that it can benefit performance on the majority group more if one adds minority data (instead of majority data). Consider a binary classification setting with data samples  $x \in \mathbb{R}$  and labels  $y \in \{1, 2\}$ . Samples from the two classes are drawn from distributions with distinct means:

$$\begin{aligned} x|y = 1 &\sim p_1(x|\mu_1, \sigma_1) \\ x|y = 2 &\sim p_2(x|\mu_2, \sigma_2). \end{aligned} \tag{3.2}$$

#### Visualizing domain gap and overlap in minority and majority distributions




Figure 3.1: **Visualization of Gaussian Mixture Model parameters.** We plot GMMs with different task complexities. The domain gap  $\delta$  is visualized as the difference in the ideal threshold locations. The overlap/task complexity metric can be visually seen.

Maximum likelihood (ML) can be used to estimate the label as

$$\hat{y} = \arg \max_y \mathcal{L}(x|y). \quad (3.3)$$

An ideal hyperplane for ML  $\mathcal{H}_{\text{ideal}}$  is a set of data samples such that:

$$\mathcal{H}_{\text{ideal}} = \{x \mid \mathcal{L}(x|y = 1) = \mathcal{L}(x|y = 2)\}. \quad (3.4)$$

We consider the hyperplane’s geometry to be linear in this one dimensional setting. Therefore the hyperplane can be represented as a normal vector:  $\mathbf{h}_{\text{ideal}}$ . The normalized hyperplane is represented by a two dimensional vector,  $\mathbf{h} = [1 \ b]^T$ . Here,  $b$  is the offset/bias. In general, a hyperplane  $\mathbf{h}$  may not be ideal. The accuracy of a hyperplane is based on a performance measure  $\mathcal{P}\{\mathbf{h}\}$ , where the operator  $\mathcal{P}$  takes as input the hyperplane and outputs the closeness to the ideal hyperplane  $\mathbf{h}_{\text{ideal}}$ . A goal of a learning based classifier is to obtain:

$$\hat{\mathbf{h}} = \arg \min_{\mathbf{h}} \mathcal{P}\{\mathbf{h}\} = \arg \min_{\mathbf{h}} \|\mathbf{h} - \mathbf{h}_{\text{ideal}}\|, \quad (3.5)$$

where  $\hat{\mathbf{h}}$  is the best learnt estimate of the ideal hyperplane. The ideal hyperplane is the global minimizer of this objective. Now, assume we are provided a finite training set of labelled data  $\mathcal{D}_{K-1} = \{(x_i, y_i)\}_{i=1}^{K-1}$ . Let the estimated hyperplane be  $\mathbf{h}_{K-1}$ , denoting that  $K-1$  samples have been used to learn the hyperplane. If one additional data sample is made available, then the learnt hyperplane would be  $\mathbf{h}_K$ . From Equation 3.2, the  $k$ -th sample is drawn from one of two distributions:

$$\begin{aligned} x_k|y = 1 &\sim p_1(x|\mu_1, \sigma_1) \\ x_k|y = 2 &\sim p_2(x|\mu_2, \sigma_2). \end{aligned} \quad (3.6)$$

We now introduce the notion of majority and minority sampling.

**Introducing Majority/Minority Distributions:** Suppose that the  $k$ -th data sample could be drawn for the same classification task from a minority or majority group. Let  $g \in \{\text{major}, \text{minor}\}$  denote the group label (for the group class). Equation 3.2 can now be

conditioned on the group label, such that there are four possible distributions from which the  $k$ -th sample can be drawn:

$$\left. \begin{aligned} x_k | g = \text{major}, y = 1 &\sim p_1^{\text{major}}(x | \mu_1^{\text{major}}, \sigma_1^{\text{major}}) \\ x_k | g = \text{major}, y = 2 &\sim p_2^{\text{major}}(x | \mu_2^{\text{major}}, \sigma_2^{\text{major}}) \end{aligned} \right\} \begin{array}{l} \text{Majority} \\ \text{group} \end{array} \quad (3.7)$$

$$\left. \begin{aligned} x_k | g = \text{minor}, y = 1 &\sim p_1^{\text{minor}}(x | \mu_1^{\text{minor}}, \sigma_1^{\text{minor}}) \\ x_k | g = \text{minor}, y = 2 &\sim p_2^{\text{minor}}(x | \mu_2^{\text{minor}}, \sigma_2^{\text{minor}}) \end{aligned} \right\} \begin{array}{l} \text{Minority} \\ \text{group} \end{array}$$

**Overlap:** Let the ideal decision hyperplane be located at  $x = d_{\text{ideal}}$ . Then, given equal likelihood of the two labels for  $y$ , the overlap for the majority group is defined as the probability of erroneous sample classification:

$$O_{\text{major}} = 0.5 \int_{x=-\infty}^{d_{\text{ideal}}} p_2^{\text{major}}(x) dx + 0.5 \int_{x=d_{\text{ideal}}}^{\infty} p_1^{\text{major}}(x) dx. \quad (3.8)$$

The same definition holds true for the minority class as well. Therefore, by definition,  $O_{\text{major}} < O_{\text{minor}}$ . The task complexities  $\theta^{\text{major}}$  and  $\theta^{\text{minor}}$  are empirical estimates of the respective overlaps. Hereafter, we assume that all four marginal distributions are Gaussian and symmetric (this is relaxed later for Key Result 3). Figure 3.1 visually highlights relevant parameters.  $O_{\text{minor}} > O_{\text{major}}$  occurs through the interplay of component means and variances. The expectation over the class label yields majority and minority sampling:

$$\begin{aligned} x_k^{\text{major}} &\triangleq x_k | g = \text{major} \sim \mathbb{E}_y [x_k | g = \text{major}, y] \\ x_k^{\text{minor}} &\triangleq x_k | g = \text{minor} \sim \mathbb{E}_y [x_k | g = \text{minor}, y], \end{aligned} \quad (3.9)$$

where we have defined  $x_k^{\text{major}}$  or  $x_k^{\text{minor}}$  as having the  $k$ -th sample come from the majority or minority distributions.

Armed with an expression for the  $k$ -th sample, we can consider a scope similar to active/online learning [47, 48, 49, 50, 51, 52, 20, 53]. Suppose a dataset of  $K - 1$  samples



has been collected on majority samples, such that there exists a dataset stack  $\mathcal{D}_{K-1}^{\text{major}} = \left\{ (x_i^{\text{major}}, y_i^{\text{major}}) \right\}_{i=1}^{K-1}$ . A hyperplane  $\mathbf{h}_{K-1}$  is learnt on this dataset and can be improved by expanding the dataset size. Consider pushing sample index  $K$ , denoted as  $x_K$  onto the stack. Now we have a choice of pushing  $x_K^{\text{major}}$  or  $x_K^{\text{minor}}$ , to create one of two datasets:

$$\begin{aligned} \mathcal{D}_K^+ &= \{ \mathcal{D}_{K-1}^{\text{major}}, x_K^{\text{major}} \} \\ \mathcal{D}_K^- &= \{ \mathcal{D}_{K-1}^{\text{major}}, x_K^{\text{minor}} \}, \end{aligned} \tag{3.10}$$

where  $\mathcal{D}_K^-$  represents the interesting case where we choose to push a minority sample onto a dataset with all majority samples (e.g. adding a dark skinned sample to a light skinned dataset). Denote  $\mathbf{h}_K^+$  and  $\mathbf{h}_K^-$  as hyperplanes learnt on  $\mathcal{D}_K^+$  and  $\mathcal{D}_K^-$ . We now arrive at the following result.

**Theorem 1:** *Let  $\mathcal{P}^{\text{major}}\{ \cdot \}$  be the performance of a hyperplane on the majority group. Let  $\Delta = \mathcal{P}^{\text{major}}\{ \mathbf{h}_{K-1} \}$ . Assume that the minority group distribution has an overlap  $O_{\text{minor}}$  while the majority group has an overlap  $O_{\text{major}} < O_{\text{minor}}$ . Both have the same ideal hyperplane  $\mathbf{h}_{\text{ideal}}$ . Under the definitions of  $\mathbf{h}_K^-$  and  $\mathbf{h}_K^+$  as above, assuming  $\Delta$  is sufficiently small and the group class distribution variances are not very large,*

$$\mathbb{E}_{x_K^{\text{minor}}} \mathcal{P}^{\text{major}}\{ \mathbf{h}_K^- \} < \mathbb{E}_{x_K^{\text{major}}} \mathcal{P}^{\text{major}}\{ \mathbf{h}_K^+ \}, \tag{3.11}$$

*stating that, perhaps surprisingly, expected performance for majorities improves more by pushing a minority sample on the stack, rather than a majority sample.*

**Proof:** The general idea is to show that samples closer to  $\mathbf{h}_{\text{ideal}}$  are more beneficial, and minority distributions may sample these with higher likelihood. Without loss of generality, we assume that  $\mathbf{h}_{K-1}$  is located, non-ideally, closer to the task class  $y = 2$  (arbitrarily called the positive class) than  $\mathbf{h}_{\text{ideal}}$ . For our perceptron update rule, the improvement in the estimated hyperplane due to  $x_K$  is proportional to the difference between the false negative

rate (FNR) and the false positive rate (FPR) for  $\mathbf{h}_{K-1}$ , with respect to the distribution of  $x_K$ . For sufficiently small  $\Delta$ ,  $FNR - FPR$  can be approximated in terms of the likelihood  $l$  that  $x_K$  is on the ideal hyperplane. The likelihood  $l$  is directly proportional to  $FPR - FNR$ . We consider the one-dimensional linear classifier setting, trained using the Perceptron algorithm. Given any  $x \in \mathbb{R}$ , the classifier evaluates an output  $y$  given by,

$$y = wx + b, \quad (3.12)$$

where  $w, b \in \mathbb{R}$ . The decision threshold in this case is at  $y = 0$ . For simplification, we reduce the redundant parameter, as follows:

$$y' = x + b'. \quad (3.13)$$

Note that the decision threshold is unaffected by this conversion. For notational simplicity, we use  $y = y'$  and  $b = b'$  here onward. We consider the perceptron decision and update rule, modified for our case. That is, for any training sample  $(x_i, y_i)$ , the predicted output is given by,

$$\hat{y}_i = \frac{\text{sign}(x_i + b) + 3}{2}, \quad (3.14)$$

where  $\text{sign}(\cdot)$  is the sign function. Readers will notice the unconventional form of this decision rule. The additional terms map the conventional perceptron labels in  $\{-1, 1\}$  to our chosen labels  $\{1, 2\}$  respectively.

For an appropriately chosen learning rate  $\gamma$ , the parameter update rule for this setting is given by:

$$b \leftarrow \begin{cases} b + \gamma, & \text{if } \hat{y}_i \neq y_i \text{ and } y_i = 2 \\ b - \gamma, & \text{if } \hat{y}_i \neq y_i \text{ and } y_i = 1 \end{cases}. \quad (3.15)$$

Let  $\mathbf{h}_{\text{ideal}} \triangleq [1, b_{\text{ideal}}]^T$  denote the ideal decision hyperplane. Under the current assumption of no domain gap, it can be shown that this ideal hyperplane is located at  $x = d_{\text{ideal}}$  such

that,

$$\begin{aligned} p_1^{\text{minor}}(d_{\text{ideal}}) &= p_2^{\text{minor}}(d_{\text{ideal}}) \\ p_1^{\text{major}}(d_{\text{ideal}}) &= p_2^{\text{major}}(d_{\text{ideal}}). \end{aligned} \tag{3.16}$$

This also implies that  $b_{\text{ideal}} = -d_{\text{ideal}}$ . Now, consider an initial training set of  $K - 1$  samples from the majority group,  $\mathcal{D}_{K-1}^{\text{major}}$ . A decision hyperplane  $\mathbf{h}_{K-1}$  is learnt from these samples. Then, without loss of generality, we can assume that,

$$d_{K-1} = d_{\text{ideal}} + \Delta. \tag{3.17}$$

That is, the real hyperplane  $\mathbf{h}_{K-1}$  is non-ideally located closer to the positive class ( $y = 2$ ) than  $\mathbf{h}_{\text{ideal}}$ .  $\Delta$  is a small positive value representing the error in the learnt decision hyperplane. Consider that the  $K$ -th sample is drawn from the majority group  $x_K^{\text{major}}$ . Recall that parameter updates for the Perceptron algorithm take place only in the event of incorrect label estimation  $\hat{y}_K \neq y_K$ . If we denote the change in the parameter  $b$  due to this sample as  $\Delta b$ , then three cases exist:

1. Sample from class 2 is classified as belonging to class 1 such that

$$x_K^{\text{major}} \sim p_2^{\text{minor}}(x), \quad x_K^{\text{major}} < d_{\text{ideal}} - \Delta. \quad \text{Associated } \Delta b = +\gamma.$$

2. Sample from class 2 is classified as belonging to class 1 such that

$$x_K^{\text{major}} \sim p_2^{\text{minor}}(x), \quad d_{\text{ideal}} - \Delta \leq x_K^{\text{major}} < d_{\text{ideal}} + \Delta. \quad \text{Associated } \Delta b = +\gamma.$$

3. Sample from class 1 classified as belonging to class 2 such that

$$x_K^{\text{major}} \sim p_1^{\text{minor}}(x), \quad x_K^{\text{major}} \geq d_{\text{ideal}} + \Delta. \quad \text{Associated } \Delta b = -\gamma.$$

Let the expected change in  $b$  due to one majority group sample be denoted as  $\Delta b^{\text{major}}$ .  $\Delta d^{\text{major}}$  is similarly defined for the expected change in  $d$ . Then, the following holds true:

$$\Delta b^{\text{major}} = \mathbb{E}_{x_K^{\text{major}}} [\Delta b]. \tag{3.18}$$

Writing out the expectation over all three cases,

$$\Delta b^{\text{major}} = \gamma \int_{x=-\infty}^{d_{\text{ideal}}-\Delta} p_2^{\text{major}}(x) dx + \gamma \int_{x=d_{\text{ideal}}-\Delta}^{d_{\text{ideal}}+\Delta} p_2^{\text{major}}(x) dx - \gamma \int_{x=d_{\text{ideal}}+\Delta}^{+\infty} p_1^{\text{major}}(x) dx. \quad (3.19)$$

Similar expressions can be identified if the  $K$ -th sample is drawn from the minority group. Under the assumption that the mixture models under consideration are symmetric Gaussian mixture models,

$$\int_{x=-\infty}^{d_{\text{ideal}}-\Delta} p_2^{\text{major}}(x) dx = \int_{x=d_{\text{ideal}}+\Delta}^{+\infty} p_1^{\text{major}}(x) dx. \quad (3.20)$$

Then, using Equation 3.19 and Equation 3.20,

$$\Delta b^{\text{major}} = \gamma \int_{x=d_{\text{ideal}}-\Delta}^{d_{\text{ideal}}+\Delta} p_2^{\text{major}}(x) dx. \quad (3.21)$$

The region between  $x = d_{\text{ideal}} - \Delta$  and  $d_{\text{ideal}} + \Delta$  determines the expected change in the classification parameter. If  $\Delta$  is small enough,  $\Delta b^{\text{major}} \approx 2\gamma p_2^{\text{major}}(d_{\text{ideal}})\Delta$ . Similarly,  $\Delta b^{\text{minor}} \approx 2\gamma p_2^{\text{minor}}(d_{\text{ideal}})\Delta$ .

We now identify a sufficient condition where  $p_2^{\text{minor}}(x) > p_2^{\text{major}}(x)$  for  $-\Delta \leq x \leq \Delta$ , given that the overlaps satisfy the condition  $O_{\text{minor}} > O_{\text{major}}$ , as defined in the main text. Under the GMM assumption,

$$p_2^{\text{major}}(x) = \frac{1}{\sqrt{2\pi(\sigma_2^{\text{major}})^2}} \exp\left(-\frac{(x - \mu_2^{\text{major}})^2}{2(\sigma_2^{\text{major}})^2}\right). \quad (3.22)$$

A similar expression exists for the minority group distribution as well. We wish to find the intersection point for the majority and minority distributions, that is  $p_2^{\text{major}}(x) = p_1^{\text{major}}(x)$  for some  $x$ . This expression reduces to,

$$\frac{(x - \mu_2^{\text{major}})^2}{\sigma_{\text{major}}^2} - \frac{(x - \mu_2^{\text{minor}})^2}{\sigma_{\text{minor}}^2} = 2 \ln \left| \frac{\sigma_{\text{minor}}}{\sigma_{\text{major}}} \right|. \quad (3.23)$$

We want to ensure that this intersection point occurs for an  $x > d_{\text{ideal}}$ . This sets up a hyperbolic equation for the condition. For our purposes of proving existence, we qualitatively note that if the majority group variance is not very large (meaning the likelihood of sampling at the ideal hyperplane is low for the majority group), and the minority group variance is not very large (such that it does not tend close to a uniform distribution),  $p_2^{\text{minor}}(x) > p_2^{\text{major}}(x)$ . Then,

$$\Delta b^{\text{minor}} > \Delta b^{\text{major}}. \quad (3.24)$$

$$\Delta d^{\text{minor}} < \Delta d^{\text{major}} < 0. \quad (3.25)$$

Our final task is to relate the expected change in the decision hyperplane over a choice of training sets  $\mathcal{D}_K^+$  and  $\mathcal{D}_K^-$ , with associated learnt hyperplanes  $\mathbf{h}_K^+$  and  $\mathbf{h}_K^-$ . As a reminder,

$$\begin{aligned} \mathcal{D}_K^+ &= \{\mathcal{D}_{K-1}^{\text{major}}, x_K^{\text{major}}\} \\ \mathcal{D}_K^- &= \{\mathcal{D}_{K-1}^{\text{major}}, x_K^{\text{minor}}\}, \end{aligned} \quad (3.26)$$

Consider a general training setting, where we use minibatches of size  $M > 1$ , over multiple epochs. Then, any minibatch containing the  $K$ -th sample can be split into the  $K$ -th sample and a random subset of  $M-1$  samples from  $\mathcal{D}_{K-1}^{\text{major}}$ . Therefore, on average, the only difference to the sample updates would be due to the contributions of the  $K$ -th sample. This brings us to our final observations,

$$\begin{aligned} \mathbb{E}_{x_K^{\text{minor}}} [d_K^+] &= d_{K-1}^{\text{major}} + \Delta d^{\text{major}} \\ \mathbb{E}_{x_K^{\text{minor}}} [d_K^-] &= d_{K-1}^{\text{major}} + \Delta d^{\text{minor}}. \end{aligned} \quad (3.27)$$

From Equations 3.25 and 3.27,

$$\mathbb{E}_{x_K^{\text{minor}}} [d_K^-] < \mathbb{E}_{x_K^{\text{minor}}} [d_K^+], \text{ and} \quad (3.28)$$

$$\mathbb{E}_{x_K^{\text{minor}}} [|d_{\text{ideal}} - d_K^-|] < \mathbb{E}_{x_K^{\text{minor}}} [|d_{\text{ideal}} - d_K^+|]. \quad (3.29)$$

The above holds for small enough  $\gamma$ . Since we know the relationship between the decision hyperplane  $\mathbf{h}$  and the associated  $d$  in our setting, the following equations hold true:

$$\mathbb{E}_{x_K^{\text{minor}}} \|\mathbf{h}_{\text{ideal}} - \mathbf{h}_K^-\| < \mathbb{E}_{x_K^{\text{minor}}} \|\mathbf{h}_{\text{ideal}} - \mathbf{h}_K^+\|, \quad (3.30)$$

$$\mathbb{E}_{x_K^{\text{minor}}} \mathcal{P}^{\text{major}}\{\mathbf{h}_K^-\} < \mathbb{E}_{x_K^{\text{major}}} \mathcal{P}^{\text{major}}\{\mathbf{h}_K^+\}. \quad (3.31)$$

Under the assumptions of the theorem, a direct relation is established between the overlap and  $l$  for each of the group classes. Then, it is shown that an additional minority sample, with overlap  $O_{\text{minor}} > O_{\text{major}}$  leads to greater expected gains as compared to an additional majority sample, concluding the proof. ■

### 3.2 Key Result 2: MIME holds under domain gap

In the previous key result we described the MIME effect in a restrictive setting where a minority and majority group have the same target hyperplane. However, it is rarely the case that minorities and majorities have the same decision boundary. We now consider the case with non-zero domain gap, to show that MIME holds on a more realistic setting. Domain gap can be quantified in terms of ideal decision hyperplanes. If  $\mathbf{h}_{\text{ideal}}^{\text{major}}$  and  $\mathbf{h}_{\text{ideal}}^{\text{minor}}$  denote ideal hyperplanes for the majority and minority groups respectively, then domain gap  $\delta = \|\mathbf{h}_{\text{ideal}}^{\text{major}} - \mathbf{h}_{\text{ideal}}^{\text{minor}}\|$ .

A visual illustration of domain gap is provided in Figure 3.1. Next, we define relative hyperplane locations in terms of halfspaces (since all hyperplanes in the one dimensional setting are parallel). We say two hyperplanes  $\mathbf{h}_1$  and  $\mathbf{h}_2$  lie in the same halfspace of a reference hyperplane  $\mathbf{h}_0$  if their respective offsets/biases satisfy the condition  $(b_1 - b_0)(b_2 - b_0) > 0$ . For occupancy in different halfspaces, the condition is  $(b_1 - b_0)(b_2 - b_0) < 0$ . We now enter into the second key result.

**Theorem 2:** Let  $\delta \neq 0$  be the domain gap between the majority and minority groups. Assume that the minority group distribution has an ideal hyperplane  $\mathbf{h}_{ideal}^{minor}$ ; while the majority group has an ideal hyperplane  $\mathbf{h}_{ideal}^{major}$ . Then, if  $\delta < \Delta$ ,  $\delta + \Delta$  is small enough, and the group class distribution variances are not very large, it can be shown that if either of the following two cases:

1.  $\mathbf{h}_{K-1}$  and  $\mathbf{h}_{ideal}^{minor}$  lie in different halfspaces of  $\mathbf{h}_{ideal}^{major}$ ,

or

2.  $\mathbf{h}_{K-1}$  and  $\mathbf{h}_{ideal}^{minor}$  lie in the same halfspace of  $\mathbf{h}_{ideal}^{major}$ , and if

$$\frac{O_{major}}{O_{minor}} < \left(1 - \frac{\delta}{\Delta}\right)f, \quad (3.32)$$

are true, then:

$$\mathbb{E}_{x_K^{minor}} \mathcal{P}^{major}\{\mathbf{h}_K^-\} < \mathbb{E}_{x_K^{major}} \mathcal{P}^{major}\{\mathbf{h}_K^+\}, \quad (3.33)$$

where  $f$  is a non-negative constant that depends on the majority and minority means and standard deviations for all the individual GMM components.

**Proof:** We prove independently for both cases. When  $\mathbf{h}_{K-1}$  and  $\mathbf{h}_{ideal}^{minor}$  lie in different halfspaces of  $\mathbf{h}_{ideal}^{major}$ , it can be shown that the expected improvement in the hyperplane is higher for the minority group as compared to the majority group, using a similar argument as in Theorem 1. Let  $\mathbf{h}_{ideal}^{major} \triangleq [1, b_{ideal}^{major}]^T$  denote the ideal decision hyperplane for the majority group. Let  $\mathbf{h}_{ideal}^{minor} \triangleq [1, b_{ideal}^{minor}]^T$  denote the ideal decision hyperplane for the minority group. Then, the ideal hyperplanes are located at  $x = d_{ideal}^{major}$  and  $x = d_{ideal}^{minor}$  respectively such that,

$$\begin{aligned} p_1^{minor}(d_{ideal}^{minor}) &= p_2^{minor}(d_{ideal}^{minor}) \\ p_1^{major}(d_{ideal}^{major}) &= p_2^{major}(d_{ideal}^{major}). \end{aligned} \quad (3.34)$$

This implies that  $b_{ideal}^{major} = -d_{ideal}^{major}$  and  $b_{ideal}^{minor} = -d_{ideal}^{minor}$ . Consider an initial training set of  $K - 1$  samples from the majority group,  $\mathcal{D}_{K-1}^{major}$ . Then, without loss of generality, we can

assume that  $d_{K-1} = d_{\text{ideal}}^{\text{major}} + \Delta$ , where  $\Delta > 0$ . Additionally, we consider the existence of domain gap in this case, that is,  $d_{\text{ideal}}^{\text{minor}} = d_{\text{ideal}}^{\text{major}} + \delta$ .

Let  $\delta < \Delta$ . Similar to the setting in Theorem 1 (Equation 3.19), we can set up the equation for expected parameter change in the case of the majority and minority groups as follows:

$$\Delta b^{\text{major}} = \gamma \int_{x=-\infty}^{d_{\text{ideal}}^{\text{major}} - \Delta} p_2^{\text{major}}(x) dx + \gamma \int_{x=d_{\text{ideal}}^{\text{major}} - \Delta}^{d_{\text{ideal}}^{\text{major}} + \Delta} p_2^{\text{major}}(x) dx - \gamma \int_{x=d_{\text{ideal}}^{\text{major}} + \Delta}^{+\infty} p_1^{\text{major}}(x) dx. \quad (3.35)$$

$$\Delta b^{\text{minor}} = \gamma \int_{x=-\infty}^{d_{\text{ideal}}^{\text{minor}} - (\Delta - \delta)} p_2^{\text{minor}}(x) dx + \gamma \int_{x=d_{\text{ideal}}^{\text{minor}} - (\Delta - \delta)}^{d_{\text{ideal}}^{\text{minor}} + (\Delta - \delta)} p_2^{\text{minor}}(x) dx - \gamma \int_{x=d_{\text{ideal}}^{\text{minor}} + (\Delta - \delta)}^{+\infty} p_1^{\text{minor}}(x) dx. \quad (3.36)$$

Under the assumption that the mixture models under consideration are symmetric Gaussian mixture models,

$$\Delta b^{\text{major}} = \gamma \int_{x=d_{\text{ideal}}^{\text{major}} - \Delta}^{d_{\text{ideal}}^{\text{major}} + \Delta} p_2^{\text{major}}(x) dx, \quad (3.37)$$

$$\Delta b^{\text{minor}} = \gamma \int_{x=d_{\text{ideal}}^{\text{minor}} - (\Delta - \delta)}^{d_{\text{ideal}}^{\text{minor}} + (\Delta - \delta)} p_2^{\text{minor}}(x) dx. \quad (3.38)$$

If  $\Delta + |\delta|$  is small enough,

$$\Delta b^{\text{major}} \approx 2\gamma p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}}) \Delta, \quad (3.39)$$

$$\Delta b^{\text{minor}} \approx 2\gamma p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}}) (\Delta - \delta). \quad (3.40)$$

By establishing the same conditions on group class variances as Theorem 1, we know that  $p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}}) > p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})$ . We now identify conditions under which  $\Delta b^{\text{minor}} > \Delta b^{\text{major}}$ .



**Case 1 -  $\delta < 0$ :** Under the same conditions as Theorem 1,  $(\Delta - \delta) > \Delta$ , and  $p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}}) > p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})$ . Therefore,

$$\Delta b^{\text{minor}} > \Delta b^{\text{major}}. \quad (3.41)$$

This proves the theorem for Case 1. When  $\mathbf{h}_{K-1}$  and  $\mathbf{h}_{\text{ideal}}^{\text{minor}}$  lie in the same halfspace of  $\mathbf{h}_{\text{ideal}}^{\text{major}}$ , and assuming that  $\mathbf{h}_{K-1}$  is located closer to the positive class, we approximate the  $FNR - FPR$  value as function of  $\delta$ ,  $\Delta$  and the likelihood  $l$  as defined for Theorem 1. Then, through algebraic manipulation, constraints can be established in terms of the two likelihoods  $l_{\text{minor}}$  and  $l_{\text{major}}$ . Under the assumptions of the theorem, a relation can be established between the ratios  $\frac{l_{\text{minor}}}{l_{\text{major}}}$  and  $\frac{O_{\text{minor}}}{O_{\text{major}}}$ .

**Case 2 -  $\delta > 0$ :**

$$\Delta b^{\text{major}} \approx 2\gamma p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})\Delta, \quad (3.42)$$

$$\Delta b^{\text{minor}} \approx 2\gamma p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}})(\Delta - \delta). \quad (3.43)$$

For  $\Delta b^{\text{minor}} > \Delta b^{\text{major}}$ ,

$$p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}})(\Delta - \delta) > p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})\Delta. \quad (3.44)$$

Rearranging Equation 3.44,

$$\frac{p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})}{p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}})} < \left(1 - \frac{\delta}{\Delta}\right). \quad (3.45)$$

Given the definitions of the majority and minority groups,

$$p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}}) < p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}}), \quad (3.46)$$

$$O_{\text{major}} < O_{\text{minor}}. \quad (3.47)$$

Since all four of these terms depend only on the means and variances of the Gaussian components, we can write,

$$\frac{O_{\text{major}}}{O_{\text{minor}}} = \frac{p_2^{\text{major}}(d_{\text{ideal}}^{\text{major}})}{p_2^{\text{minor}}(d_{\text{ideal}}^{\text{minor}})} f, \quad (3.48)$$

where  $f$  is a positive scalar constant that depends only on the component means and variances. From Equations 3.45 and 3.48,

$$\frac{O_{\text{major}}}{O_{\text{minor}}} < \left(1 - \frac{\delta}{\Delta}\right) f. \quad (3.49)$$

This proves the conditions in the theorem. Theorem 1 can now be used to show the existence of the MIME effect in the presence of domain gap, for these conditions. This concludes the proof. ■

**A Note on the Theorems:** Theorems 1 and 2 are existence theorems. That is, they show that there exist certain conditions under which the MIME effect can be observed. The theorems make these arguments based on the ‘usefulness’ of points close to the ideal hyperplane. The direct metric of correlation is the likelihood for a particular distribution to sample at the ideal hyperplane. However, since this cannot be easily measured in practice, we set up our proofs in terms of a correlated metric: the overlap.

### 3.3 Key Result 3: MIME holds for general distributions

We now relax the symmetric Gaussian and equally likely labels requirements to arrive at a general condition for MIME existence. Let  $p_1^{\text{major}}$  and  $p_2^{\text{major}}$  be general distributions describing the majority group  $y = 1$  and  $y = 2$  classes. Additionally,  $Pr(y = 1) \neq Pr(y = 2)$ . Minority group distributions are described similarly. We define the signed tail weight for the majority group as follows:

$$T^{\text{major}}(x_d) = \pi^{\text{major}} \int_{x=-\infty}^{x_d} p_2^{\text{major}}(x) dx - (1 - \pi^{\text{major}}) \int_{x=x_d}^{\infty} p_1^{\text{major}}(x) dx, \quad (3.50)$$

where  $\pi^{\text{major}} = Pr(x = 2)$  for the majority group.  $T^{\text{minor}}(\cdot)$  is similarly defined. This leads us to our third key result.

**Theorem 3:** Consider majority and minority groups, with general sample distributions and unequal prior label distributions. If,

$$\begin{aligned} \min \{T^{\text{minor}}(d_{\text{ideal}} + \Delta), -T^{\text{minor}}(d_{\text{ideal}} - \Delta)\} > \\ \max \{T^{\text{major}}(d_{\text{ideal}} + \Delta), -T^{\text{major}}(d_{\text{ideal}} - \Delta)\}, \end{aligned} \quad (3.51)$$

then  $\mathbb{E}_{x_K^{\text{minor}}} \mathcal{P}^{\text{major}}\{\mathbf{h}_K^-\} < \mathbb{E}_{x_K^{\text{major}}} \mathcal{P}^{\text{major}}\{\mathbf{h}_K^+\}$ .

**Proof:** This Theorem considers distributions with general prior distributions. The perceptron algorithm update rule is proportional to  $FNR - FPR$  (if  $\mathbf{h}_{K-1}$  is located closer to the positive class) or the  $FPR - FNR$  (if  $\mathbf{h}_{K-1}$  is located closer to the negative class). The MIME effect exists in the scenario where the worst case update for the minority group is better than the best case update for the majority group (described in Equation 3.51). Therefore, for the majority group, let

$$\begin{aligned} p_2^{\text{major}'}(x) &= \pi^{\text{major}} p_2^{\text{major}}(x), \\ p_q^{\text{major}'}(x) &= (1 - \pi^{\text{major}}) p_1^{\text{major}}(x). \end{aligned} \quad (3.52)$$

Similar definitions are made for the minority group as well. Then, assuming  $d_{K-1} = d_{\text{ideal}}^{\text{major}} + \Delta$ ,  $\Delta > 0$  (similar to Theorem 2), and  $\delta = 0$  (for now), and drawing from Equation 3.19), we can set up the equation for expected parameter change in the case of the majority group as follows:

$$\begin{aligned} \Delta b^{\text{major}} &= \gamma \int_{x=-\infty}^{d_{\text{ideal}} + \Delta} p_2^{\text{major}'}(x) dx - \gamma \int_{x=d_{\text{ideal}} + \Delta}^{+\infty} p_1^{\text{major}'}(x) dx \\ &= T^{\text{major}}(d_{\text{ideal}} + \Delta). \end{aligned} \quad (3.53)$$

A similar expression holds true for the minority group. Then, if  $T^{\text{major}}(d_{\text{ideal}} + \Delta) < T^{\text{minor}}(d_{\text{ideal}} + \Delta)$ , the MIME effect will hold true.

Similarly, if  $d_{K-1} = d_{\text{ideal}}^{\text{major}} - \Delta$ ,  $\Delta > 0$ ,

$$\begin{aligned} \Delta b^{\text{major}} &= -\gamma \int_{x=-\infty}^{d_{\text{ideal}} - \Delta} p_2^{\text{major}'}(x) dx + \gamma \int_{x=d_{\text{ideal}} - \Delta}^{+\infty} p_1^{\text{major}'}(x) dx \\ &= -T^{\text{major}}(d_{\text{ideal}} - \Delta). \end{aligned} \quad (3.54)$$

Then, if  $-T^{\text{major}}(d_{\text{ideal}} - \Delta) < -T^{\text{minor}}(d_{\text{ideal}} - \Delta)$ , the MIME effect will hold true.

Combining the two expressions, for a sufficient existence condition, we get,

$$\min \{T^{\text{minor}}(d_{\text{ideal}} + \Delta), -T^{\text{minor}}(d_{\text{ideal}} - \Delta)\} > \max \{T^{\text{major}}(d_{\text{ideal}} + \Delta), -T^{\text{major}}(d_{\text{ideal}} - \Delta)\}. \quad (3.55)$$

This completes the proof. ■

Note that the existence proof for Theorem 3 ignores the effect of domain gap  $\delta$ , in the interest of readability and brevity. A very similar existence proof can be established with domain gap. We omit the derivation and provide the final condition below (under the constraints on  $\delta$  and  $\Delta$  as in Theorem 2, and using the same notation):

$$\min \{T^{\text{minor}}(d_{\text{ideal}}^{\text{major}} + \Delta), -T^{\text{minor}}(d_{\text{ideal}}^{\text{major}} - \Delta)\} > \max \{T^{\text{major}}(d_{\text{ideal}}^{\text{major}} + \Delta), -T^{\text{major}}(d_{\text{ideal}}^{\text{major}} - \Delta)\}. \quad (3.56)$$

Generalizations of Theorem 3 to include domain gap are discussed in Appendix A. Theorems 1 and 2 are special cases of the general Theorem 3, describing MIME existence for specific group distributions.

## CHAPTER 4

### Verifying MIME Theory on Real Tasks

In the previous Chapter, we provide existence conditions for the MIME phenomenon for general sample distributions. However, experimental validation of the phenomenon requires quantification in terms of measurable quantities such as overlap. Theorem 2 provides us these resources. Here, we verify that the assumptions in Theorem 2 are validated by experiments on real tasks.

#### 4.1 The Assumptions

##### 4.1.1 Verifying Gaussianity

Theorem 2 assumes that data  $x$  is drawn from a Gaussian Mixture Model. At first glance, this quantification may appear to be unrelated to complex neural networks. However, as illustrated at the top of Figure 4.1, a ConvNet is essentially a feature extractor that feeds a flattened layer into a simple perceptron or linear classifier. The flattened layer can be orthogonally projected onto the decision boundary to generate, in analogy, an  $x$  used for linear classification (Figure 1.1, fixed-backbone configuration). We use this as a first approximation to the end-to-end configuration used in our experiments.

Plotting empirical histograms of these flattened layers (Figure 4.1) shows Gaussian-like distribution. This is consistent with the Law of Large Numbers – linear combination of several random variables follows an approximate Gaussian distribution. Hence, Theorem 2 is approximately related in this setting. Details about implementation and comparison to

Table 4.1: **Experimental measures of overlap and domain gap are consistent with the theory in Chapter 3.** Note that the majority group consistently has lower overlap. Domain gaps are found to be small. DS-1 is FairFace, DS-2 is Pet Images, DS-4 is Chest-Xray14 and DS-5 is Adult. DS-6 is the high domain gap gender classification experiment. DS-3 is excluded here since it deals with a 9 class classification problem.

Dataset (Task)	DS-1 [11] (Gender)	DS-2 [13] (Species)	DS-4 [14] (Diagnosis)	DS-5 [15] (Income)	DS-6 [12, 54] (Gender)
Major. overlap	0.186	0.163	0.294	0.132	0.09
Minor. overlap	0.224	0.198	0.369	0.208	0.19
Domain gap	0.276	0.518	0.494	0.170	1.62

Gaussians are deferred to Appendix A.

#### 4.1.2 Verifying minority/majority definitions

The MIME proof linked minority and majority definitions to distributional overlap and domain gap. Given the histogram embeddings from above, it is seen that minority groups on all four vision tasks have greater overlap. There also exists a domain gap between majority and minority but this is small compared to distribution spread (except for the high domain gap experiment). This establishes applicability of small domain gap requirements. Quantification is provided in Table 4.1. Code (algorithmic implementations) is provided in Appendix A.

## 4.2 MIME Effect Across Six, Real Datasets

### 4.2.1 Implementation

Six multi-attribute datasets are used to assess the MIME effect (five are in computer vision). For a particular *experiment*, we identify a task category to evaluate accuracy over (e.g. gender), and a group category (e.g. race). The best test accuracy on the majority group across all epochs is recorded as our accuracy measure. Each experiment is run for a fixed number of *minority training ratios* ( $\beta$ ). For each minority training ratio, the total number of training samples remains constant. That is, the minority samples replace the majority samples, instead of being appended to the training set. Each experiment is also run for a finite number of *trials*. Different trials have different random train and test sets (except for the FairFace dataset [11] where we use the provided test split). Averaging is done across trials. Note that minority samples to be added are randomly chosen – the MIME effect is not specific to particular samples. For the vision datasets, we use a ResNet-34 architecture [55], with the output layer appropriately modified. For the non-visual dataset, a fully connected network is used. Average accuracy and trend error, across trials are used to evaluate performance. Specific implementation details are provided in Appendix A.

### 4.2.2 MIME effect on gender classification

The FairFace dataset [11] is used to perform gender classification ( $y = 1$  is male,  $y = 2$  is female). The majority and minority groups  $g = \{\text{major}, \text{minor}\}$  are light and dark skin, respectively. Results are averaged over five trials. Figures 4.2 and 4.3 describe qualitative accuracy. The accuracy trends indicate that adding 10% of minority samples to the training set leads to approximately a 1.5% gain in majority group (light skin) test accuracy.

### 4.2.3 MIME effect on animal species identification

We manually annotate light and dark cats and dogs from the Pets dataset [13]. We classify between cats ( $y = 1$ ) and dogs ( $y = 2$ ). The majority and minority groups are light and dark fur color respectively. Figures 4.2 and 4.3 show qualitative results. Over five trials, we see a majority group accuracy gain of about 2%, with a peak at  $\beta = 10\%$ .

### 4.2.4 MIME effect on age classification

We use a second human faces dataset, the UTKFace dataset [12], for the age classification task (9 classes of age-intervals). We pre-process the UTKFace age labels into class bins to match the FairFace dataset format. The majority and minority groups are male and female respectively. The proportion of task class labels is kept the same across group classes. Results are averaged over five trials. Figures 4.2 and 4.3 show trends. We observe a smaller average improvement for the 10% minority training ratio. However, since these are average trends, this indicates consistent gain. Results on this dataset also empirically highlight the existence of the MIME effect beyond two class settings.

### 4.2.5 MIME effect on X-ray diagnosis Classification

We use the NIH Chest-Xray14 dataset [14] to analyze trends on a medical imaging task. We perform binary classification of scans belonging to ‘Atelectasis’ ( $y = 1$ ) and ‘Pneumothorax’ ( $y = 2$ ) categories. The male and female genders are the majority and minority groups respectively. Results are averaged over seven trials (due to noisier trends). From Figures 4.2 and 4.3, we observe noisy trends - specifically we see a performance drop for  $\beta = 0.2$ , prior to an overall gain for  $\beta = 0.3$ . The error bounds also have considerably more noise. However, confidence in the peak and the MIME effect, as seen from the average trends and the error bounds, remains high.



**MIME effect on income classification:** For validation in a non-vision setting, we use the Adult (Census Income) dataset [15]. The data consists of census information with annual income labels (income less than or equal to \$50,000 is  $y = 1$ , income greater than \$50,000 is  $y = 2$ ). The majority and minority groups are female and male genders respectively. Results are averaged over five trials. Figure 4.4(a) highlights a prominent accuracy gain for  $\beta = 0.6$ .

**MIME effect and domain gap:** Theorem 2 (Chapter 3) suggests that large domain gap settings will not show the MIME effect. We set up an experiment to verify this (Figure 4.4(b)). Gender classification among chickens (majority group) [54] and humans (minority group) [12] has a high domain gap due to minimal common context (validated by the domain gap estimates, Table 4.1). With increasing  $\beta$ , the majority accuracy decreases. This (and Figure 4.1, Table 4.1 that show low domain gap for other datasets) validates Theorem 2. Note that while this result may not be unexpected, it further validates our proposed theory.

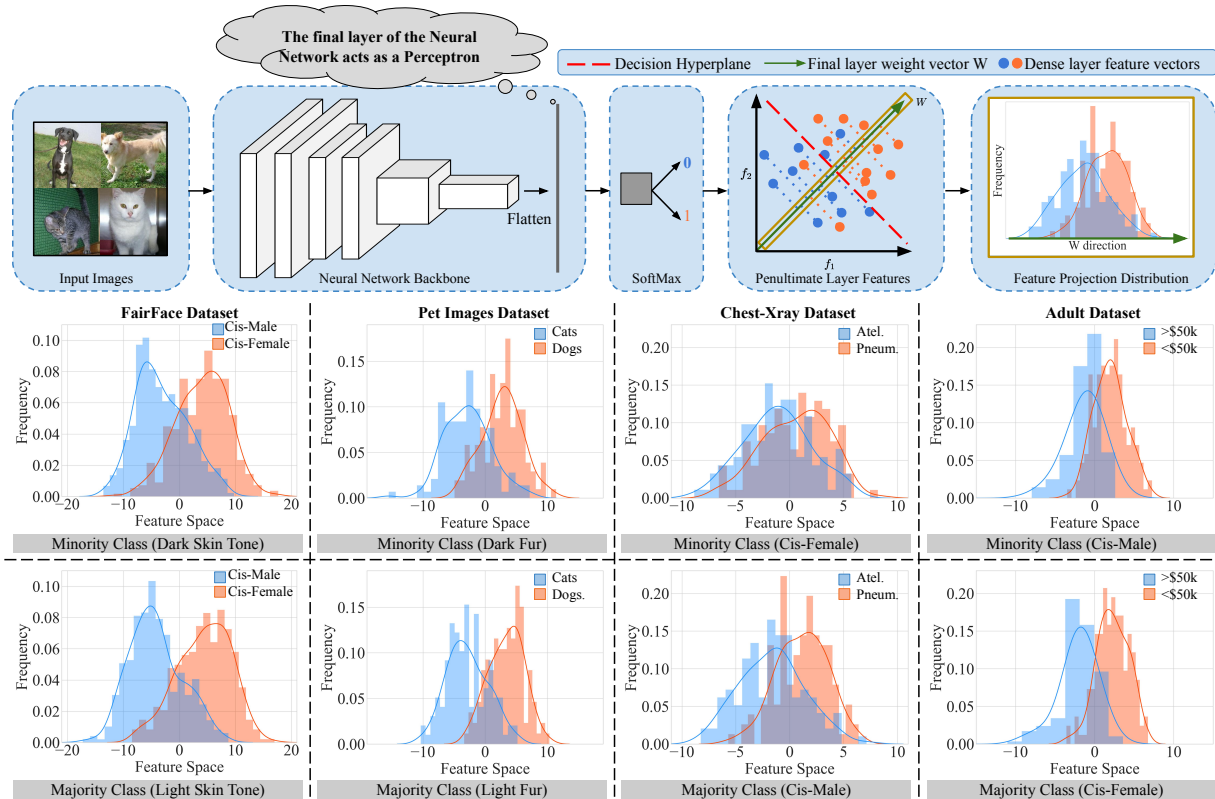


Figure 4.1: **The use of Gaussian mixtures to represent minority and majority distributions is consistent with behaviors in modern neural networks, on real-world datasets.** (top row) The last layer of common neural architectures is a linear classifier on features. Histograms of the penultimate layer projections are generated for models with  $\beta = 0.5$ . (middle row) Minority histograms: note the greater difficulty due to less separation of data. (bottom row) Majority histograms: note smaller overlap and easier classification. Figure can be parsed on a per-dataset basis. Within each column, the reader can compare the domain gap and overlap in the two histograms.

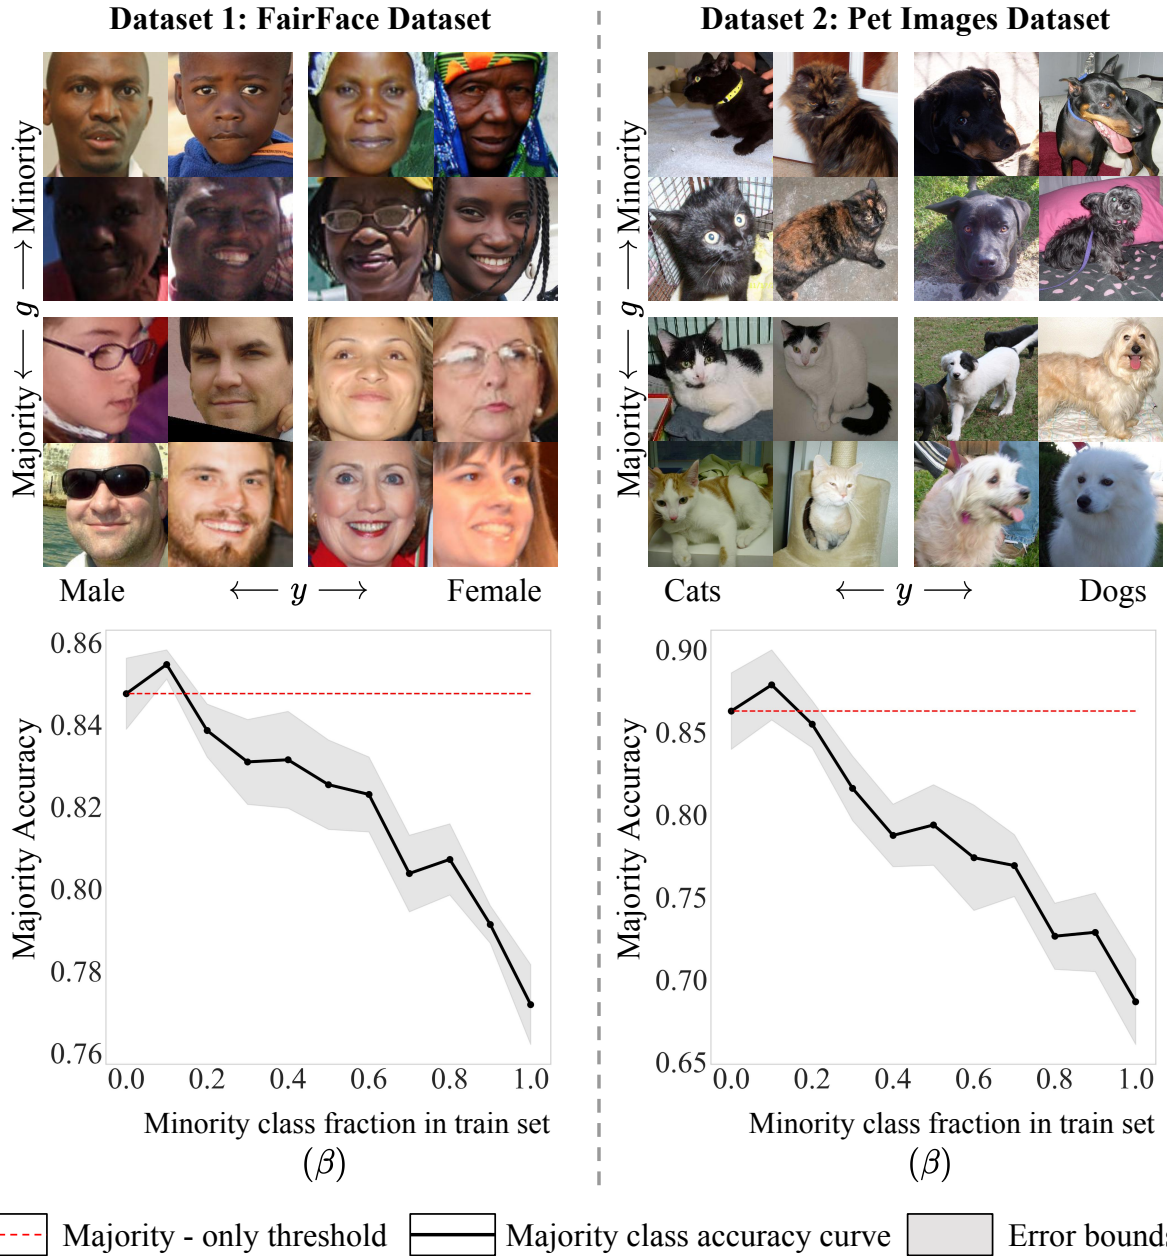


Figure 4.2: **When domain gap is small, the MIME effect holds.** On four vision datasets, majority performance is maximized with some inclusion of minorities. All experiments are run for several trials and realizations (described in Section 4.2).




Figure 4.3: **When domain gap is small, the MIME effect holds.** On four vision datasets, majority performance is maximized with some inclusion of minorities. All experiments are run for several trials and realizations (described in Section 4.2).

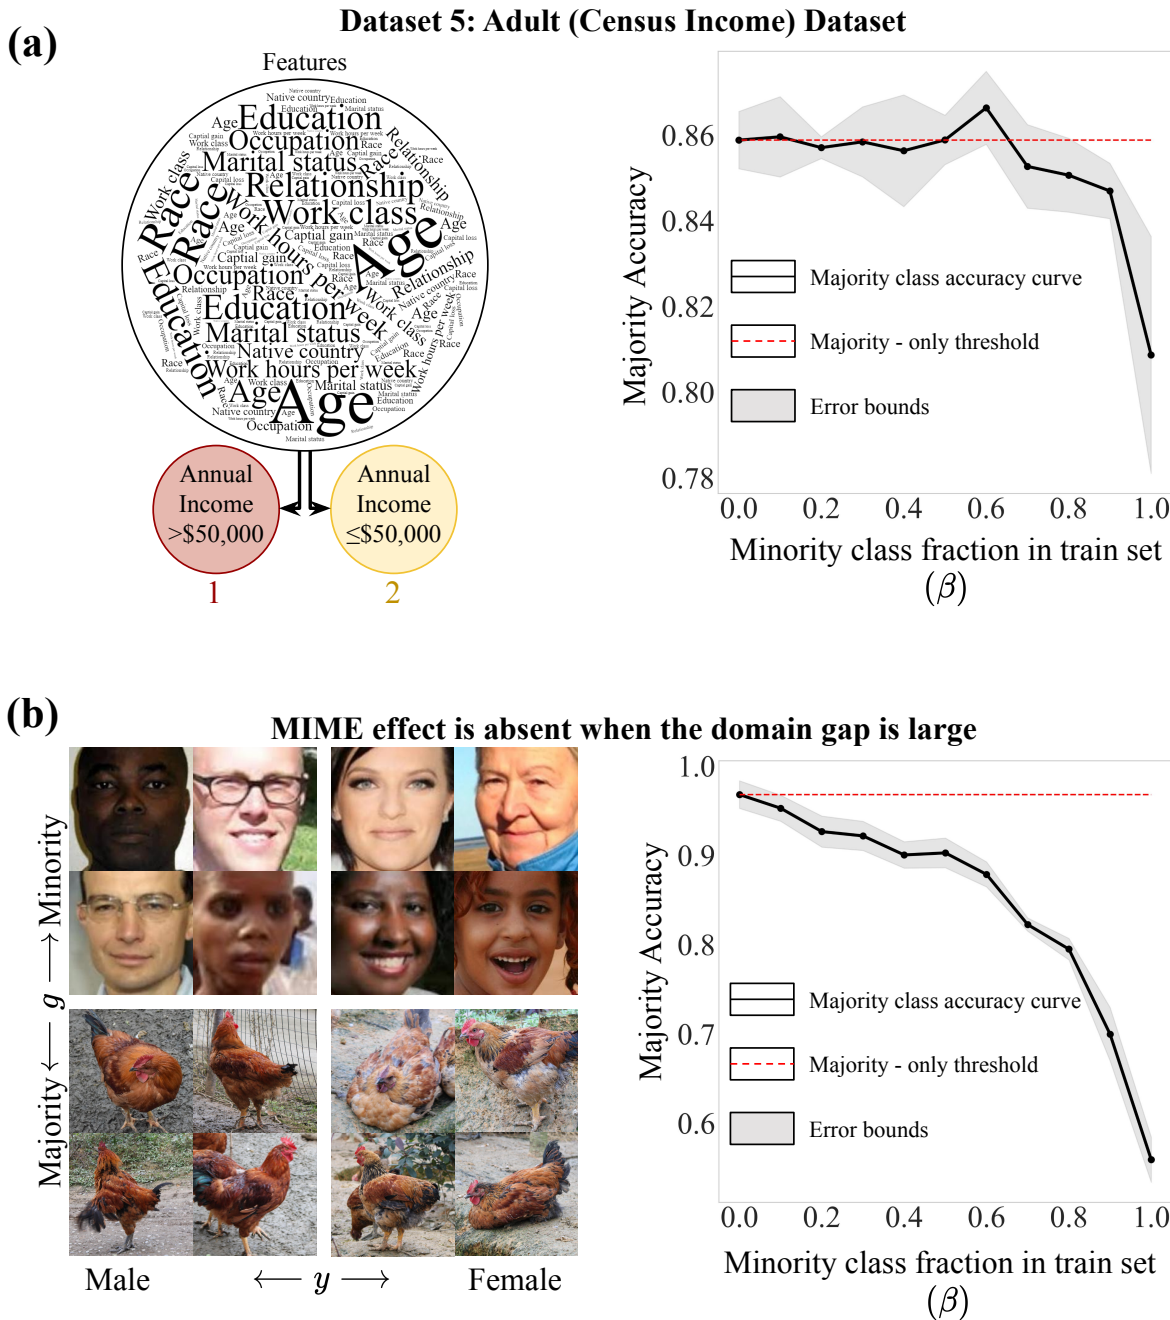


Figure 4.4: MIME effect is observed in non-vision datasets, and is absent in the case of large domain gap. (a) The Adult Dataset uses Census data to predict an income label. (b) On dataset six, gender classification is re-scoped to occur in a high domain gap setting. Majority group is chickens and minority group is humans.

# CHAPTER 5

## Discussion

### 5.1 Secondary validation and analysis

Table 5.1 provides additional metrics to analyze MIME. Across datasets, almost all trials show existence, with every dataset showing average MIME performance gain. Some readers may view the error bars in Figures 4.2, 4.3 and 4.4 as large, however they are comparable to other empirical ML works [56, 57]; they may appear larger due to scaling. Reasons for error bars include variations in train-test data and train set size (Tables A.2 and A.3, Appendix A). Further analysis, including interplay with debiasing methods (e.g. hard-sample mining [16]) and reconciliation with work on equal representation datasets [2, 3, 4, 5, 6, 7, 8, 9, 10] is deferred to Appendix B.

$$V_{imp} = \frac{1}{||W_{imp}||} \sum_{w \in W_{imp}} V_w \quad (5.1)$$

### 5.2 Optimality of inclusion ratios

Our experiments show that there can exist an optimal amount of minority inclusion to benefit the majority group the most. This appears true across all experiments in Figures 4.2, 4.3 and 4.4. However, beyond a certain amount, accuracy decreases consistently, with lowest accuracy on majority samples observed when no majorities are used in training. This optimal  $\beta$  depends on individual task complexities, among other factors. Since identifying it is outside our scope (Sections 1.1, 1.2), our experiments use 10% sampling resolution for  $\beta$ . Peaks at

Table 5.1: **Additional evaluation metrics provide further evidence of MIME existence across all datasets.** The table highlights: (i) number of trials with MIME performance gain (i.e. majority accuracy at some  $\beta > 0$  is greater than majority accuracy at  $\beta = 0$ ), and (ii) the mean MIME performance gain across trials (in % points).

Dataset	DS-1 [11]	DS-2 [13]	DS-3 [12]	DS-4 [14]	DS-5 [15]
#MIME trials/Total trials	4/5	4/5	5/5	6/7	4/5
Avg. MIME perf. gain	0.72%	1.84%	0.70%	1.89%	0.98%

$\beta = 10\%$  for some datasets are due to this lower resolution; optimal peak need not lie there for all datasets (e.g. X-ray [14] & Adult [15]). Future work can identify optimal ratios through finer analysis over  $\beta$ .

### 5.3 Limitations

The theoretical scope is certifiable within fixed-backbone binary classification, which is narrower than all of machine learning (Figure 1.1). Should this theory be accepted by the community, follow-up work can generalize theoretical claims. Another limitation is the definition-compatibility of majority and minority groups. Our theory is applicable to task-advantage definitions; some scholars in the community instead define majorities and minorities by proportion. Our theory is applicable to these authors as well, albeit with a slight redefinition of terminology. Additional considerations are included in Appendix A and B. In conclusion, majority performance benefits from a non-zero fraction of inclusion of minority data given a sufficiently small domain gap.

# APPENDIX A

## Generalizations and Implementation

### A.1 MIME Existence Beyond 1D Settings

Consider  $\mathbf{x} \in \mathbb{R}^n$ . The perceptron decisions are based on the metric  $y = \mathbf{w}^T \mathbf{x} + b$ , where  $\mathbf{w} \in \mathbb{R}^n$ , and  $y, b \in \mathbb{R}$ . Similar to Theorem 1, we consider the perceptron decision and update rule. That is, for any training sample  $(\mathbf{x}_i, y_i)$ , the predicted label is given by,

$$\hat{y}_i = \frac{\text{sign}(\mathbf{w}^T \mathbf{x}_i + b) + 3}{2}. \quad (\text{A.1})$$

We can rewrite this in terms of a single decision hyperplane by defining  $\tilde{\mathbf{w}} = [\mathbf{w}^T \ b]^T$  and  $\tilde{\mathbf{x}} = [\mathbf{x}^T \ 1]^T$ . For a small learning rate  $\gamma$ , the updated decision rule becomes,

$$\hat{\tilde{y}}_i = \frac{\text{sign}(\tilde{\mathbf{w}}^T \tilde{\mathbf{x}}_i) + 3}{2}. \quad (\text{A.2})$$

$$\tilde{\mathbf{w}} \leftarrow \begin{cases} \tilde{\mathbf{w}} + \gamma \tilde{\mathbf{x}}_i, & \text{if } \hat{\tilde{y}}_i \neq y_i \text{ and } y_i = 2 \\ \tilde{\mathbf{w}} - \gamma \tilde{\mathbf{x}}_i, & \text{if } \hat{\tilde{y}}_i \neq y_i \text{ and } y_i = 1 \end{cases}. \quad (\text{A.3})$$

We now refer to the hyperplane  $\tilde{\mathbf{w}}$  as the decision hyperplane. Let  $\mathbf{h}_{\text{ideal}}$  be the ideal decision hyperplane. In this setting, any domain gap  $\delta$  or error in real hyperplane estimation  $\Delta$  manifests as a direction/angle error in the hyperplane normal vector (since the bias term  $b$  is subsumed in the hyperplane). The updates change the normal vector of the hyperplane through a linear combination with the sample  $\tilde{\mathbf{x}}_i$ , scaled by the learning rate  $\gamma$ .

We now provide a qualitative description for the existence of the MIME effect, in terms of the likelihood of a favorable update to  $\tilde{\mathbf{w}}$ . We consider a simplified 2D case with symmetric






Figure A.1: **The MIME effect holds in a multidimensional setting as well.** We show the support for the two finite distributions. Weight vector updates arising out of samples from regions R3, R4, R5 and R6 lead to an update with a large vertical (corrective) component (favorable update). Updates arising out of regions R1 and R2 result in an overall update in the horizontal direction (unfavorable update).

distributions and  $\delta = 0$ . A finite support is assumed for the majority and minority groups, for ease of understanding. Consider that the bias term  $b$  is known, and only the hyperplane direction is to be refined. Again, we denote the hyperplane from our finite training set  $\mathcal{D}_{K-1}^{\text{major}}$  as  $\mathbf{h}_{K-1}$ . The error  $\Delta$  in this case is now the angular error between the normals for  $\mathbf{h}_{\text{ideal}}$  and  $\mathbf{h}_{K-1}$ . Figure A.1 indicates this setting. The learnt hyperplane  $\tilde{\mathbf{w}}_{K-1}$  is shown as a black solid line. The black dashed line represents the mirror image of the learnt hyperplane, defined for aid in simplification. Recall that updates to the weight vector take place on misclassification. On average, the updates due to samples in regions R1 for ( $y = 2$ ) and R2 (for  $y = 1$ ) lead to a net horizontal (leftward) weight update. This is an unfavorable update that increases  $\Delta$ . Therefore, the favorable updates on average are from regions R3 and R4 for  $y = 2$ , and R5 and R6 for  $y = 1$ . This is a net update with large vertical (upward) update. This is a favorable update that decreases  $\Delta$ . These regions are described based on the small angular deviation  $\Delta$ . Since the distributions have finite support along the direction parallel

to the ideal hyperplane (vertical direction in Figure A.1), the requirement again reduces to greater likelihood of sampling close to the ideal hyperplane (similar to Theorems 1 and 2), since  $\Delta$  is small. That is, distributions that sample close to the ideal hyperplane with greater probability have a greater expected likelihood of a favorable update. Under similar conditions as Theorem 1, MIME effect holds in this case.

The extension to include the bias term  $b$  is straightforward. We follow the setting in Equation A.2 and subsume the bias as part of the weights. In this case,  $\Delta$  includes the error in both the hyperplane normal direction as well as the bias. Extensions to greater number of dimensions can be done using the same arguments. Additionally, domain gap can also be introduced. We omit explicit mathematical expressions in the interest of brevity, and since our goal here is to establish existence.

## A.2 Feature Space Analysis

### A.2.1 Constructing the Projected Feature Histograms

Let  $f$  denote a feature vector, in the penultimate layer of a classification neural network. For example, in the case of ResNet-34 [55],  $\mathbf{f} \in \mathbb{R}^{512}$ . Similarly, let  $\mathbf{w}$  be the final layer weights. In the case of multiple final layer hyperplanes, we choose any one of the hyperplanes (since for the two class classification task, the two projected variables are correlated when trained against the cross entropy loss for 2 classes). Then, we define  $x \in \mathbb{R}$  as,

$$x = \mathbf{w}^T \mathbf{f}. \tag{A.4}$$

Classification decisions are made solely on the basis of the projected variable  $x$ . Therefore, we analyze the histogram distributions for  $x$ . Practically, for each dataset, we use the best performing (in terms of majority group performance) model trained using a minority training fraction ( $\beta$ ) of 0.5. This is chosen in order to obtain histograms of  $x$  for all four distributions – the two task classes for both the majority and minority groups. The histograms are created

using the test set samples.

## A.2.2 Estimating the Overlap

The overlap is estimated from the histograms, using the following Python code snippet:

```
def histogram_intersection(h1, h2, bins):  
    #INPUTS:  
    #h1, h2: normalized histograms  
    #bins: number of bins in the histograms (should be equal for the  
        two histograms)  
  
    #OUTPUTS:  
    #sm: overlap fraction  
    sm = 0  
    for i in range(bins):  
        sm += min(h1[i], h2[i])  
    return sm
```

## A.2.3 Estimating the Domain Gap

We follow a two step process to estimate the domain gap  $\delta$ . First, the ideal decision hyperplanes for the majority and minority groups are estimated, using Equation 3.34. We fit a fifth order polynomial to the two histograms. The central intersection point of the histograms (i.e the intersection point that lies between the means of the two classes) is then the location of the ideal decision threshold. The following Python code snippet describes this:

```
import numpy as np  
  
def ideal_hyperplane(h1, h2, z, ref=5):  
    #INPUTS:  
    #h1, h2: the two histograms, of equal length and identical bins  
    #z: a list of the histogram bin centers  
    #ref: Search space for the intersection of the two histograms-
```

```

                                                    default is from -5 to 5

#OUTPUT:
#z_dec: Ideal decision threshold between the two histogram
                                               distributions

z_dash = np.polyfit(z, h1, 5)
f1 = np.poly1d(z_dash)
# calculate polynomial
z_dash = np.polyfit(z, f2, 5)
f2 = np.poly1d(z_dash)
new_z = np.linspace(-ref, ref, 5000)
new_f1 = f1(new_z)
new_f2 = f2(new_z)
id_dec = np.argmin(np.abs(new_f1 - new_f2))
z_dec = new_z[id_dec]
return z_dec

```

The domain gap is the absolute difference between two ideal decision thresholds, for each of the two group classes. Figure 3.1 from Chapter 3 may be referred to for a graphical visualization.

#### A.2.4 Notes on the Estimated Measures

The latent feature space analysis is not perfect. This is because the feature extraction part of the network is jointly learnt along with decision hyperplane. Histograms are plotted on the 50% minority training ratio so as to enable a fair domain gap and overlap comparison between the two group classes. Specifically, note that we define task complexity in Chapters 3 and 4 in terms of the minority only and majority only train sets which deviates from the setting here. The estimates for overlap and domain gap are therefore approximate correlated estimates and not exact measures.

## A.2.5 Analysis of Feature Space Gaussian-like Behavior

We set up the Chi-Squared goodness of fit test on all 20 distributions under consideration (i.e. across 5 datasets and 4 distributions each per dataset). These statistics correspond to the distributions in Table 4.1 and Figure 4.1 from Chapter 4. Python code for testing the hypotheses is given below. The number of bins are chosen so as to ensure  $\geq 5$  samples per bin on average.

```
from scipy.stats import chisquare
from scipy.stats import norm
from scipy import stats
import pandas as pd

def chi_square_stats(vals, no_bins)
    #INPUTS:
    #vals: a list of samples whose Gaussianity is to be tested
    #no_bins: number of bins (thumb rule: no_bins < len(vals)/5)

    tot_vals = len(vals)
    # mean and standard deviation of given data
    mean = np.mean(vals)
    std = np.std(vals)

    interval = []
    for i in range(1, no_bins+1):
        val = stats.norm.ppf(i/no_bins, mean, std)
        interval.append(val)
    interval.insert(0, -np.inf)

    lower = interval[:-1]
    upper = interval[1:]

    df = pd.DataFrame({'lower_limit':lower, 'upper_limit':upper})
```

```

sorted_vals = list(sorted(vals))
df['obs_freq'] = df.apply(lambda x:sum([i>x['lower_limit'] and i<=
                                     x['upper_limit'] for i in
                                     sorted_vals]), axis=1)
df['exp_freq'] = tot_vals/no_bins

statistic = stats.chisquare(df['obs_freq'], df['exp_freq'])

p = 2      # number of parameters for 1D Gaussian
DOF = len(df['obs_freq']) - p - 1
thresh = stats.chi2.ppf(0.95, DOF)

return statistic, thresh

```

Table A.1 highlights the evaluated chi-square statistics, as well as related parameters. Note that a lower value of the statistic is better, and the null hypothesis is not rejected when the value of the statistic is lower than the critical value. We establish the null hypothesis at a 5% level of significance for each distribution to be that the samples are drawn from a Gaussian distribution. Distributions that are unable to reject the null hypothesis are indicated in bold. It can be seen that a large majority of the distributions indicate that the projected latent features follow a Gaussian-like distribution.

## A.3 Implementation Details

### A.3.1 Analysis measures

For each task, we estimate the test accuracy  $a_p^i(\beta)$  as a function of minority group fraction in the train set  $\beta \in [0, 1]$ , for a trial  $i \in \{1, \dots, N\}$ , for a group class  $g$  (e.g. dark skin tones).  $N$

Table A.1: **Chi-Squared goodness of fit measures for all distributions.** Distributions with bolded values show the estimated statistics that are lower than the critical value, indicating that the null hypothesis (Gaussian distribution) cannot be rejected.

Dataset	No. of samples per group per class	No. of Bins	Critical Value	Majority Group		Minority Group	
				$y = 1$	$y = 2$	$y = 1$	$y = 2$
DS-1 [11]	379	15	21.03	<b>13.65</b>	28.69	<b>10.25</b>	<b>15.39</b>
DS-2 [13]	126	15	21.03	<b>7.81</b>	<b>12.10</b>	<b>11.62</b>	<b>9.24</b>
DS-4 [14]	126	15	21.03	<b>10.43</b>	<b>17.57</b>	<b>17.10</b>	<b>4.48</b>
DS-5 [15]	159	15	21.03	<b>11.09</b>	24.05	<b>5.74</b>	<b>14.40</b>
DS-6 [54, 12]	43	5	5.99	25.48	<b>5.02</b>	<b>5.72</b>	<b>4.79</b>

is the total number of trials. Practically, we evaluate performance for a finite set of  $\beta$  values, represented by the set  $B = \{0, 0.1, 0.2, \dots, 1.0\}$ . We now define the following measures:

**Average accuracy** : For a given minority training ratio  $\beta_0$ , and for a given group class  $g$ , we define the average accuracy,

$$\bar{a}_g(\beta_0) = \frac{1}{N} \sum_{i=1}^N a_g^i(\beta_0). \quad (\text{A.5})$$

**Error bounds** : We also evaluate the *trend variation* among  $a_g^i(\beta)$  for various  $i$ . That is, we want to evaluate if across all the trials (for a particular task-dataset combination), the relative trend (of majority group performance gain) holds true. One candidate measure for this is  $std_i(a_g^i(\beta))$  for each  $\beta$ , where  $std_i(\cdot)$  is the standard deviation operator, over  $i$ . However, this measure will include average changes in accuracy for all splits, for a particular trial (arising out of unrelated effects such as different train or test set samples). This is unnecessary in our case. Therefore, we define our error measure  $\hat{\zeta}(\beta)$  as the  $\beta$ -mean subtracted standard deviation. That is,

$$\begin{aligned}\hat{\zeta}(\beta) &= \text{std}_i(a_g^i(\beta) - \bar{a}_g^i), \\ \bar{a}_g^i &= \frac{1}{|B|} \sum_{\beta \in B} a_g^i(\beta),\end{aligned}\tag{A.6}$$

where  $|\cdot|$  is the cardinality operator representing the size of a set. In our graphs, we plot the average accuracy  $\bar{a}_g(\beta)$  as well as the error bounds, from  $\bar{a}_g(\beta) - \hat{\zeta}(\beta)$  to  $\bar{a}_g(\beta) + \hat{\zeta}(\beta)$ ,  $\forall \beta \in B$ .

### A.3.2 Network Architectures Used

For all the vision-related experiments, we use the ResNet-34 architecture [55]. We only modify the output layer of the network so as to match the number of task classes (9 for Dataset 3, and 2 for all other tasks). For the Adult (Census) Dataset [15], we use a fully connected network with sigmoid outputs. The PyTorch [58] implementation for the model is included below.

```
#Model

def act(x):
    return F.relu(x)

class Network(nn.Module):
    def __init__(self,):
        super().__init__()
        self.fc1 = nn.Linear(101, 50)
        self.fc2 = nn.Linear(50, 50)
        self.fc3 = nn.Linear(50, 50)
        self.fcLast = nn.Linear(50,2)

    def forward(self, x):

        x = act(self.fc1(x))
```



```

# x = self.b1(x)
x = act(self.fc2(x))
x = act(self.fc3(x))
x = torch.sigmoid(self.fcLast(x))

return x

```

### A.3.3 General Experiment Details

Table A.2: **Training configuration and parameters for all datasets and experiments.** Parameters for each dataset are chosen so as to maximize performance.

Dataset (Task)	DS-1 [11] (Gender)	DS-2 [13] (Species)	DS-3 [12] (Age)	DS-4 [14] (Diagnosis)	DS-5 [15] (Income)	DS-6 [54, 12] (Gender)
Group class	Race	Skin tone	Gender	Gender	Gender	Species
Train set size	10900	1500	7700	1500	2600	750
Test set size (per group)	760	250	970	250	300	90
No. of trials	5	5	5	7	5	5
No. of epochs	35	60	65	40	250	20
Learning rate	0.0005	0.0006	0.0006	0.0006	0.0005	0.0005
Weight Decay	0.08	0.05	0.05	0.05	0.08	0.08
Input Shape/Config.	3x100x100	3x256x256	3x100x100	3x256x256	101x1	3x100x100

All experiments were carried out using PyTorch [58]. Table A.2 highlights the training parameters used for each dataset. We use different parameters for each of the datasets. These are experimentally chosen to maximize accuracy. All the models are trained using the AdamW optimizer [59] and the cross entropy loss. The train and test set sizes vary slightly across trials, due to different data splits. However, the train set size remains the same for all minority training ratios of a particular trial. A validation set is held out but given the small sample size of several datasets, we measure trends based on best test performance. This is to minimize the effect of sample specific performance gap in small datasets. Averaging of trends over multiple trials, and hence multiple train-test splits ensures that the trends do not

overfit to a particular configuration. Each trial is run using a unique random seed. Table A.3 highlights the random seeds used for our experiments, which were randomly chosen. Input images are resized to the chosen input size for each dataset. For the Adult dataset [15], we use a one-hot encoding scheme for the input. The group class information is dropped from the input before passing to the network. For all the datasets, across all minority training ratios for a particular trial, we use a fixed model initialization to ensure that the changes in accuracy are completely attributable to the train data configuration.

Table A.3: **Random seeds used for the trials.** Seeds were chosen at random for trials to generate average trends and error bounds.

Dataset (Task)	DS-1 [11] (Gender)	DS-2 [13] (Species)	DS-3 [12] (Age)	DS-4 [14] (Diagnosis)	DS-5 [15] (Income)	DS-6 [54, 12] (Gender)
Random Seeds	0, 1, 3, 5, 7	21, 42, 35, 28, 31	0, 55, 2, 15, 6	33, 42, 24, 36 54, 21, 28	13, 15, 17, 19, 21	0, 1, 3, 5, 9

#### A.3.4 Dataset Specific Information

To perform experiments on the **Pet Images Dataset**, we manually annotate light and dark fur cats and dogs from the larger dataset used in [13]. For the age classification task on the **UTKFace Dataset** [12], we pre-process the age labels to match the annotation format for the FairFace dataset [11]. For the large domain gap gender classification task using the **UTKFace and Chicken Images Datasets** [12, 54], we perform gender classification over human and chicken groups. Therefore, this experiment is over a new, composite dataset.

# APPENDIX B

## Additional Analysis

### B.1 Secondary Analysis of MIME

#### B.1.1 MIME effect with debiasing methods:

We now analyze the interaction of the MIME effect with existing debiasing methods. Specifically, while applying hard-sample mining [16] (as an exemplary case) across the task classes ( $y = 1, 2$ ), we sweep across various minority training ratios. Figure B.1(a) shows results on two datasets (implementation details may be found in the following section). The MIME effect continues to be observed. Debiasing methods act on the task classes ( $y = 1, 2$ ) in an effort to improve performance while MIME acts on majority and minority groups, regardless of the task class. Therefore, MIME is complementary to debiasing methods, rather than a competitor. In our experiments, hard-sample mining does not lead to significant performance gains since the task classes are balanced by experimental design. In other scenarios where this might not be the case, MIME and hard sample mining might together improve performance.

#### B.1.2 Reconciling MIME with existing equal representation (ER) datasets:

In this thesis, we focus only on majority group performance, for which ER training datasets are not optimal in general. In contrast, existing efforts [2, 3, 4, 5, 6, 7, 8, 9, 10] focus on ER datasets to maximize overall (majority+minority) performance. This need not be optimal but is a good thumb rule. This is because while majority group performance eventually

reduces with minority training ratio, minority group performance increases (Figure B.1(b) highlights this).

## B.2 Hard Mining Baseline Implementation

We implement a version of the method proposed in [16]. From a batch of 30 samples, 12 samples (6 of each task class) are retained and used in the training step. These are the samples with least confidence, with respect to ground truth targets. Code is shown below. Trial random seeds are the same as shown in Table A.3.

```
class compute_crossentropyloss_hardMine:
    """
    y0 is the vector with shape (batch_size,C)
    x shape is the same (batch_size), whose entries
    are integers from 0 to C-1. In our case, C=2.
    """
    def __init__(self, ignore_index=-100) -> None:
        self.ignore_index=ignore_index

    def __call__(self, y0, x):
        loss = 0.
        eps = 1e-5
        K = 6
        n_batch, n_class = y0.shape
        pos_score = torch.ones(n_batch).to(device)
        neg_score = torch.ones(n_batch).to(device)
        ix_pos = 0
        ix_neg = 0
        for y1, x1 in zip(y0, x):
            class_index = int(x1.item())
            score = torch.exp(y1[class_index])/(torch.exp(y1).sum()+eps)
            if class_index == 0:
```

```

        neg_score[ix_neg] = score
        ix_neg+=1
    else:
        pos_score[ix_neg] = score
        ix_pos+=1

pos_score,_ = torch.sort(pos_score,dim=0)
neg_score,_ = torch.sort(neg_score,dim=0)
pos_els = np.minimum(K,ix_pos)
neg_els = np.minimum(K,ix_neg)
for ix in np.arange(pos_els):
    loss = loss -torch.log(pos_score[ix])

for ix in np.arange(neg_els):
    loss = loss -torch.log(neg_score[ix])

loss = loss/(pos_els+neg_els)
torch.cuda.empty_cache()
return loss

```

### B.3 Negative impacts and mitigation:

This thesis focuses on highlighting the existence of the MIME effect, and not optimal configurations for performance gain. Nevertheless, potential negative outcomes may occur if the results are misinterpreted as guidance on dataset construction with respect to certain stakeholder groups. The rigor of our theoretical results emphasizes this nuance to computer scientists, and future work in diverse venues can extend the notion of minority inclusion for majority group performance gains to broader audiences.




Figure B.1: **The MIME effect is complementary to data debiasing methods and consistent with research aimed at equal representation (ER) datasets.** (a) Training configurations using data debiasing methods show the MIME effect. (b) While ER datasets are not optimal for the MIME effect (Figures 4.2, 4.3 and 4.4), optimal overall performance is observed close to ER.

## REFERENCES

- [1] Pradyumna Chari, Yunhao Ba, Shreeram Athreya, and Achuta Kadambi. Mime: Minority inclusion for majority group enhancement of ai performance. In *European Conference on Computer Vision*, pages 326–343. Springer, 2022.
- [2] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- [3] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.
- [4] Agostina J Larrazabal, Nicolás Nieto, Victoria Peterson, Diego H Milone, and Enzo Ferrante. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, 117(23):12592–12594, 2020.
- [5] Hee Jung Ryu, Hartwig Adam, and Margaret Mitchell. Inclusivefacenet: Improving face attribute detection with race and gender diversity. *arXiv preprint arXiv:1712.00193*, 2017.
- [6] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9572–9581, 2019.
- [7] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- [8] Eun Seo Jo and Timnit Gebru. Lessons from archives: Strategies for collecting socio-cultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 306–316, 2020.
- [9] Zhiqiang Gong, Ping Zhong, and Weidong Hu. Diversity in machine learning. *IEEE Access*, 7:64323–64350, 2019.
- [10] Achuta Kadambi. Achieving fairness in medical devices. *Science*, 372(6537):30–31, 2021.
- [11] Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1548–1558, 2021.

- [12] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5810–5818, 2017.
- [13] Philippe Golle. Machine learning attacks against the asirra captcha. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 535–542, 2008.
- [14] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.
- [15] Catherine L Blake and Christopher J Merz. Uci repository of machine learning databases, 1998, 1998.
- [16] Qi Dong, Shaogang Gong, and Xiatian Zhu. Class rectification hard mining for imbalanced deep learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1851–1860, 2017.
- [17] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343. PMLR, 2018.
- [18] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [19] Jiwoong Choi, Ismail Elezi, Hyuk-Jae Lee, Clement Farabet, and Jose M Alvarez. Active learning for deep object detection via probabilistic modeling. *arXiv preprint arXiv:2103.16130*, 2021.
- [20] Sanjoy Dasgupta. Two faces of active learning. *Theoretical computer science*, 412(19):1767–1781, 2011.
- [21] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 969–977, 2018.
- [22] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2100–2110, 2019.



- [23] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. FsdR: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021.
- [24] Matthew Gwilliam, Srinidhi Hegde, Lade Tinubu, and Alex Hanson. Rethinking common assumptions to mitigate racial bias in face recognition datasets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4123–4132, 2021.
- [25] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019.
- [26] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29:4349–4357, 2016.
- [27] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 771–787, 2018.
- [28] Guha Balakrishnan, Yuanjun Xiong, Wei Xia, and Pietro Perona. Towards causal benchmarking of bias in face analysis algorithms. In *Deep Learning-Based Face Analytics*, pages 327–359. Springer, 2021.
- [29] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, et al. Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5):4–1, 2019.
- [30] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5310–5319, 2019.
- [31] Charles Elkan. The foundations of cost-sensitive learning. In *International joint conference on artificial intelligence*, volume 17, pages 973–978. Lawrence Erlbaum Associates Ltd, 2001.
- [32] Steffen Bickel, Michael Brückner, and Tobias Scheffer. Discriminative learning under covariate shift. *Journal of Machine Learning Research*, 10(9), 2009.
- [33] Enzo Tartaglione, Carlo Alberto Barbano, and Marco Grangetto. End: Entangling and disentangling deep representations for bias correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13508–13517, 2021.

- [34] 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, 2018.
- [35] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8919–8928, 2020.
- [36] Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9301–9310, 2021.
- [37] Han Xu, Xiaorui Liu, Yaxin Li, Anil Jain, and Jiliang Tang. To be robust or to be fair: Towards fairness in adversarial training. In *International Conference on Machine Learning*, pages 11492–11501. PMLR, 2021.
- [38] Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younès Bennani. A survey on domain adaptation theory: learning bounds and theoretical guarantees. *arXiv preprint arXiv:2004.11829*, 2020.
- [39] Shai Ben-David, John Blitzer, Koby Crammer, Fernando Pereira, et al. Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems*, 19:137, 2007.
- [40] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- [41] Ewa M Nowara, Daniel McDuff, and Ashok Veeraraghavan. A meta-analysis of the impact of skin tone and gender on non-contact photoplethysmography measurements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 284–285, 2020.
- [42] Yunhao Ba, Zhen Wang, Kerim Doruk Karinca, Oyku Deniz Bozkurt, and Achuta Kadambi. Overcoming difficulty in obtaining dark-skinned subjects for remote-ppg by synthetic augmentation. *arXiv preprint arXiv:2106.06007*, 2021.
- [43] Zhen Wang, Yunhao Ba, Pradyumna Chari, Oyku Deniz Bozkurt, Gianna Brown, Parth Patwa, Niranjana Vaddi, Laleh Jalilian, and Achuta Kadambi. Synthetic generation of face videos with plethysmograph physiology. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20587–20596, 2022.
- [44] Pradyumna Chari, Krish Kabra, Doruk Karinca, Soumyarup Lahiri, Diplav Srivastava, Kimaya Kulkarni, Tianyuan Chen, Maxime Cannesson, Laleh Jalilian, and Achuta

- Kadambi. Diverse r-ppg: Camera-based heart rate estimation for diverse subject skin-tones and scenes. *arXiv preprint arXiv:2010.12769*, 2020.
- [45] Alexander Vilesov, Pradyumna Chari, Adnan Armouti, Anirudh B Harish, Kimaya Kulkarni, Ananya Deoghare, Laleh Jalilian, and Achuta Kadambi. Blending camera and 77 ghz radar sensing for equitable, robust plethysmography. In *ACM Trans. Graph. (SIGGRAPH)*, 2022.
- [46] Mehryar Mohri and Afshin Rostamizadeh. Perceptron mistake bounds. *arXiv preprint arXiv:1305.0208*, 2013.
- [47] Seyda Ertekin, Jian Huang, Leon Bottou, and Lee Giles. Learning on the border: active learning in imbalanced data classification. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 127–136, 2007.
- [48] Sheng-Jun Huang, Rong Jin, and Zhi-Hua Zhou. Active learning by querying informative and representative examples. *Advances in Neural Information Processing Systems*, 23:892–900, 2010.
- [49] Burr Settles. Active learning literature survey. *Computer Sciences Technical Report*, 2009.
- [50] Jan Kremer, Kim Steenstrup Pedersen, and Christian Igel. Active learning with support vector machines. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 4(4):313–326, 2014.
- [51] Sanjoy Dasgupta, Daniel J Hsu, and Claire Monteleoni. *A general agnostic active learning algorithm*. Citeseer, 2007.
- [52] Alina Beygelzimer, Daniel J Hsu, John Langford, and Tong Zhang. Agnostic active learning without constraints. *Advances in Neural Information Processing Systems*, 23:199–207, 2010.
- [53] Maria-Florina Balcan, Andrei Broder, and Tong Zhang. Margin based active learning. In *International Conference on Computational Learning Theory*, pages 35–50. Springer, 2007.
- [54] Yuanzhou Yao, Haoyang Yu, Jiong Mu, Jun Li, and Haibo Pu. Estimation of the gender ratio of chickens based on computer vision: Dataset and exploration. *Entropy*, 22(7):719, 2020.
- [55] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

- [56] Terrance Liu, Giuseppe Vietri, and Steven Z Wu. Iterative methods for private synthetic data: Unifying framework and new methods. *Advances in Neural Information Processing Systems*, 34:690–702, 2021.
- [57] Stéphane d’Ascoli, Marylou Gabrié, Levent Sagun, and Giulio Biroli. On the interplay between data structure and loss function in classification problems. *Advances in Neural Information Processing Systems*, 34:8506–8517, 2021.
- [58] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [59] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.