

Master Thesis



Czech  
Technical  
University  
in Prague

**F3**

Faculty of Electrical Engineering  
Department of Cybernetics

## Adaptation of CNN Classifiers to Prior Shift

**Tomáš Šipka**

Supervisor: Ing. et Ing. Milan Šulc, Ph.D.

Field of study: Open informatics

Subfield: Computer Vision and Image Processing

August 2021



## I. Personal and study details

Student's name: **Šipka Tomáš** Personal ID number: **457428**  
Faculty / Institute: **Faculty of Electrical Engineering**  
Department / Institute: **Department of Cybernetics**  
Study program: **Open Informatics**  
Specialisation: **Computer Vision and Image Processing**

## II. Master's thesis details

Master's thesis title in English:

**Adaptation of CNN Classifiers to Prior Shift**

Master's thesis title in Czech:

**Adaptace CNN klasifikátor na zm ny apriorních pravd podobnosti**

Guidelines:

1. Familiarize yourself with Convolutional Neural Network (CNN) classifiers and the standard training procedure of empirical risk minimization with cross entropy loss [1].
2. Familiarize yourself with common domain adaptation scenarios, with focus on prior shift (also known as label shift or target shift).
3. Review the state-of-the-art methods for prior shift adaptation of CNN and other probabilistic classifiers.
4. Review the state-of-the-art methods for calibration of CNN (and other probabilistic) classifiers.
5. Propose a new method (or methods) for classifier adaptation to a change in categorical priors (prior shift, label shift) between training and test data. Explore the problem of inconsistent estimates in the confusion matrix-based methods [7,9] and the possibility of regularization [3].
6. Experimentally evaluate the existing and proposed methods on different prior-shift scenarios and datasets, with and without classifier calibration.

Bibliography / sources:

- [1] Bengio, Yoshua, Ian Goodfellow, and Aaron Courville. Deep learning. Vol. 1. Massachusetts, USA: MIT press, 2017.
- [2] Amr Alexandari, Anshul Kundaje, and Avanti Shrikumar. Maximum likelihood with bias-corrected calibration is hard-to-beat at label shift adaptation. ArXiv, 1901.06852v5, 2019. 1
- [3] Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, and Animashree Anandkumar. Regularized learning for domain adaptation under label shifts. arXiv preprint arXiv:1903.09734,2019. 1, 2
- [4] Marthinus Christoffel du Plessis and Masashi Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. CoRR, abs/1206.4677, 2012. 1, 2, 4
- [5] Zachary C. Lipton, Yu-Xiang Wang, and Alex Smola. Detecting and correcting for label shift with black box predictors, 2018. 1, 2, 4
- [6] Geoffrey J. McLachlan. Discriminant Analysis and Statistical Pattern Recognition. John Wiley & Sons, Inc, Hoboken, NJ, USA, 1992-03-27. 2
- [7] Marco Saerens, Patrice Latinne, and Christine Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: A simple procedure. Neural Comput., 14(1):21–41, Jan.2002. 1, 2, 4
- [8] Milan Sulc and Jiri Matas. Improving CNNclassifiers by estimating test-time priors. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops, Oct 2019. 1, 2, 4
- [9] Slobodan Vucetic and Zoran Obradovic. Classification on data with biased class distribution. In European Conference on Machine Learning, pages 527–538. Springer, 2001. 1

Name and workplace of master's thesis supervisor:

**Ing. Milan Šulc, Ph.D., Visual Recognition Group, FEE**

Name and workplace of second master's thesis supervisor or consultant:

Date of master's thesis assignment: **28.01.2021** Deadline for master's thesis submission: **13.08.2021**

Assignment valid until: **30.09.2022**

\_\_\_\_\_  
Ing. Milan Šulc, Ph.D.  
Supervisor's signature

\_\_\_\_\_  
prof. Ing. Tomáš Svoboda, Ph.D.  
Head of department's signature

\_\_\_\_\_  
prof. Mgr. Petr Páta, Ph.D.  
Dean's signature

### III. Assignment receipt

The student acknowledges that the master's thesis is an individual work. The student must produce his thesis without the assistance of others, with the exception of provided consultations. Within the master's thesis, the author must state the names of consultants and include a list of references.

\_\_\_\_\_  
Date of assignment receipt

\_\_\_\_\_  
Student's signature

## Acknowledgements

I would like to thank my supervisor Ing. et Ing. Milan Šulc, Ph.D. for his valuable advice, relentless support and guidance not only during the time I was working on this thesis but also during my time at CMP, where the foundation stones of this work were laid. I also wish to thank prof. Ing. Jiří Matas, Ph.D. for the opportunity to be part of CMP and for his helpful advice and constructive criticism. The discussions with both of them were very helpful, inspiring and especially great lessons. Thanks should also go to my family for their support during my studies.

## Declaration

I declare that the presented work was developed independently and that I have listed all sources of information used within it in accordance with the methodical instructions for observing the ethical principles in the preparation of university theses.

Prague, 13. August 2021

## Abstract

In many classification tasks, the test set's relative class frequencies (class priors probabilities) differ from the relative class frequencies at training time. Such a phenomenon, called *label shift* or *prior shift*, can negatively affect the classifier's performance. Considering a probabilistic classifier approximating posterior probabilities, the predictions can be adapted to the label shift by re-weighting with a ratio of the test set and training set priors. Labels in the test set are usually unknown, therefore the prior ratio has to be estimated in an unsupervised manner. This thesis reviews existing methods for adapting probabilistic classifiers to label shift and for estimating test priors in an unlabeled test set. Moreover, we propose novel algorithms to address the problems of estimating new priors and prior ratio. The methods are designed to handle a known issue in confusion matrix-based methods, where inconsistent estimates of decision probabilities and confusion matrices lead to negative values in estimated priors. Experimental evaluation shows that our method improves the stability of prior estimation and the adapted classifier's accuracy compared to the baseline confusion matrix-based methods and achieves state-of-the-art performance in prior shift adaptation.

**Keywords:** Machine Learning, Computer Vision, Domain Adaptation, Transfer Learning, Classification, Label Shift, Prior Shift, Classifier Calibration, Convolutional Neural Networks, Confusion Matrix

**Supervisor:** Ing. et Ing. Milan Šulc, Ph.D.  
Visual Recognition Group,  
Department of Cybernetics,  
FEE CTU in Prague

## Abstrakt

V mnoha klasifikačních úlohách se relativní četnosti tříd (apriorní pravděpodobnosti tříd) na testovací sadě liší od relativních četností během trénování klasifikátoru. Tento jev, také nazývaný *label shift* nebo *prior shift*, může negativně ovlivnit přesnost klasifikace. Budeme-li uvažovat pravděpodobnostní klasifikátor aproximující aposteriorní pravděpodobnosti, mohou být jeho predikce adaptovány na label shift pomocí poměrem testovacích a trénovacích apriorních pravděpodobností. Jelikož jsou anotace v testovací sadě obvykle neznámé, musí být poměry apriorních pravděpodobností odhadnuty metodou učení bez učitele. Tato práce rekapituluje existující řešení pro adaptaci pravděpodobnostních klasifikátorů na label shift a odhad apriorních pravděpodobností na testovací sadě bez anotací. Navíc, navrhuje nové algoritmy pro odhad apriorních pravděpodobností na testovací sadě a pro odhad poměrů testovacích a trénovacích apriorních pravděpodobností. Tyto metody jsou navrženy tak, aby řešily známý problém v metodách založených na matici záměru, kde nekonzistentní odhad pravděpodobnosti rozhodnutí klasifikátoru a jeho chybové matice může vést k záporným hodnotám v odhadnutých četnostech. Experimentální vyhodnocení ukazuje, že naše metoda zlepšuje stabilitu odhadu apriorních pravděpodobností a přesnost adaptovaného klasifikátoru v porovnání s jinými metodami založenými na matici záměru a současně dosahuje nejlepších výsledků mezi metodami pro adaptaci klasifikátorů na prior shift.

**Klíčová slova:** Strojové učení, Pořítakové Vídění, Domain Adaptation, Transfer Learning, Klasifikace, Label Shift, Prior Shift, Kalibrace Klasifikátor, Konvoluční Neuronové Síť, Matice Záměru

**Překlad názvu:** Adaptace CNN klasifikátorů na změny apriorních pravděpodobností

# Contents

<b>1 Introduction</b>	<b>1</b>	4.2 Label Shift Adaptation . . . . .	37
<b>2 Related Work</b>	<b>5</b>	4.2.1 Re-weight or Retrain? . . . . .	37
2.1 Classification . . . . .	5	4.2.2 Improving Estimates from Confusion Matrix . . . . .	38
2.1.1 Training Classifier . . . . .	6	4.2.3 Comparing Regularization . .	39
2.2 Convolutional Neural Networks . .	8	4.2.4 Comparing Ratio Estimation	41
2.3 Unsupervised Domain Adaptation	8	4.2.5 Methods for MLE and MAP Prior Estimation . . . . .	41
2.4 Classifier Adaptation to Label Shift . . . . .	11	4.2.6 Dependence on the Number of Test Samples . . . . .	43
2.4.1 Confusion Matrix . . . . .	12	4.3 Confusion Matrices Illustrated on an Artificial Dataset . . . . .	43
2.4.2 Adaptation Based on Maximum Likelihood . . . . .	14	4.4 Convergence Speed . . . . .	47
2.4.3 Adaptation Based on Maximum Aposteriori . . . . .	16	<b>5 Conclusion</b>	<b>49</b>
2.4.4 BBSE . . . . .	16	<b>Bibliography</b>	<b>53</b>
2.4.5 RLLS . . . . .	18	<b>A Code</b>	<b>59</b>
2.5 Classifier Calibration . . . . .	19	A.1 Sample Code . . . . .	59
2.5.1 Evaluating Calibration . . . . .	21	A.2 Full Reproducibility Code . . . . .	59
2.5.2 Calibration Methods . . . . .	24	<b>B List of Attachments</b>	<b>61</b>
<b>3 The Proposed Method</b>	<b>27</b>		
3.1 Maximum Likelihood Approach .	28		
3.2 Maximum Aposteriori Approach	30		
3.3 Applying Regularization . . . . .	30		
3.4 Ratio Estimation . . . . .	31		
3.4.1 Projection onto the Priors Ratio Simplex . . . . .	32		
<b>4 Experiments</b>	<b>35</b>		
4.1 Implementation Details . . . . .	36		
4.1.1 Settings for Classifier Training	36		
4.1.2 Implementation of Label Shift Adaptation . . . . .	36		





# Chapter 1

## Introduction

*"It's going to be interesting to see how society deals with artificial intelligence, but it will definitely be cool."*

— Colin Angle

Due to its wide area of applications and great potential to reduce costs, human resources and managing risks, increase growth and productivity as well as the potential to overcome human limitations, *artificial intelligence* (AI) is a recently rapidly growing area of active research. There are two main approaches to define artificial intelligence [1]. The first one defines AI as the ability of machines or computers to act like human beings. From this point of view, the goal of AI is to design a machine simulating human-like behaviour. Based on this human-centred approach, in 1950 Alan Turing proposed the so-called Turing Test, designed to determine if machine intelligence is close to that of a human. The test was composed of several questions asked by the examiner. A computer passed the test if the examiner could not distinguish between the computer's and human's answers. The second approach to AI avoids humans in its definition and instead of comparing computers to humans, it compares their level of rationality. According to this rationalist approach, AI is the ability of an agent to learn from experience and solve problems to achieve its goals.

An important sub-field of artificial intelligence is *machine learning* [2, 3]. Machine learning studies algorithms, which are able to learn and adapt themselves from experience. The term "experience" is usually represented by data in electronic form [2], which were either collected in the past (offline learning) or directly at the training time (online learning). These data sets,

## 1. Introduction

also called "training data", are being provided to the algorithm. An example of such data are digital images, time series of stock prices, disease symptoms etc. The data serves as an input into the algorithm and based on these inputs, we aim to receive an output corresponding to a solution or prediction for some task. E.g. for image, the desired output could be a description of the captured scene, from stock prices in the past we may want to predict prices in the future, from symptoms we could predict disease etc. If the training data also contains the desired outputs, then the learning procedure is called *supervised*. Otherwise, we talk about *unsupervised* learning. A recently popular area of machine learning is deep learning [4]. Deep learning is specific by making use of artificial neurons, inspired by real neurons in the biological brain, and combining them into more complex structures, called artificial neural networks (ANNs). ANNs possess a great ability to extract complex information from the training data and learn from them. Different designs of neural networks were proposed for different data structures e.g. Convolutional Neural Network (CNN) [5] for image processing, Recurrent Neural Networks (RNN) [6] for time series data or Graph Neural Networks (GNN) [7] for graph-structured data. Machine learning and deep learning can find use in recommendation systems, medical diagnosis, speech recognition, computer vision and many other areas.

One of the fundamental tasks in machine learning is classification. Classification aims to assign a label (category, class) to a given observation. Usually, the number of categories is small, in a range from tens to thousands, but can also be unbounded [2]. The observation can be anything. For example, an image of a vehicle can be labelled as a car, truck or bus; animals can be assigned label corresponding to their species name and text can be labelled with a sentiment (positive, negative, or neutral). The algorithm assigning categories to observations is called a classifier. The input of the classifier is a feature vector, i.e. a vector of values (features) representing a particular observation. For instance, an animal can be represented by its colour, weight, height and width, yielding a 4-dimensional feature vector, but can also be represented by its digital grey-scale image with resolution  $1000 \times 1000$  pixels, yielding a 1 million-dimensional feature vector. The set of all categories, that can be assigned to the observations is referred to as label space, denoted as  $Y$ .  $D$ -dimensional space of feature vectors is called feature space and is denoted as  $X$ . It is common to train classifiers to output probabilities over label space, representing the classifier's confidence that a given datum belongs to a particular class. In other words, model's output  $f(X)$  is learned to approximate posterior probabilities  $p(Y/X)$ , i.e.  $f(X) \approx p(Y/X)$ , where  $X$

and  $Y$  are random variables over feature space and label space respectively. An example of such a classifier is a Convolutional Neural Network (CNN) trained with cross-entropy loss.

When deploying a classifier into a real-world application, it may happen that the real-world data, also referred to as "test data", are different from those during training. In the most general case, even the label space and feature space can differ. For example, if we do not have enough data to learn classifying dog's breeds, we can first train our classifier on an inter-species data set with a sufficient amount of samples and then reuse gained knowledge to classify dogs. The field dealing with knowledge transfer from one task (source domain) to another (target domain) is called transfer learning [8]. However, even though we have enough data for our specific task defined over some feature space and label space, there still can be a difference between training and test data. For instance, when the inputs into the classifier are images, they can be taken under different lighting conditions or disease symptoms can be slightly changed over time. In other words, the probability density generating training data is not equal to the density generating test data, i.e.  $p_T(Y; X) \neq p_E(Y; X)$ . The sub-field of transfer learning associated with adapting algorithms to a new data distribution, while feature space and label space are fixed, is called domain adaptation [9].

It is reasonable to simplify a general difference between the train density  $p_T(Y; X)$  and the test density  $p_E(Y; X)$  by decomposition  $p(Y; X) = p(Y|X)p(X) = p(X|Y)p(Y)$  and then make assumptions about its particular terms. Transferring knowledge without any assumptions about source and target domains would not be possible. For example, one could assume that density over feature space can change, i.e.  $p_T(X) \neq p_E(X)$ , but posterior probability remains the same, i.e.  $p_T(Y|X) = p_E(Y|X)$ . This scenario is called covariate shift [9, 10, 11, 12, 13]. As example of covariate shift consider a diagnosis of breast cancer. Suppose that the data for the training set were collected from mostly elderly women. Indeed such data set is not a representative sample of the population with respect to the age, resulting in a bias within the marginal distribution  $p(X)$ . Another data set shift is conditional shift [14] where the relative class frequencies are the same, but conditional densities  $p(X|Y)$  may differ. Last but not least is a situation where the class appearance over data sets does not change, i.e.  $p_T(X|Y) = p_E(X|Y)$ , and the only difference is in relative class frequencies (class priors), i.e.  $p_T(Y) \neq p_E(Y)$ . This phenomenon is called prior shift or label shift [9, 15, 16, 17, 18, 19, 20]. To have a better notion about the label shift, consider an image-based recognition system for some animal species

## 1. Introduction

in the wild to be applied in a real-world application. It is obvious that the distribution of species will vary from place to place in the world as well as from season to season during the year. However, we have to assume that the appearance of species we want to recognize does not change with place and season. Indeed, this is one of the drawbacks of this assumption, because in many cases this is not true (e.g. trees may look different in summer and winter). As another example consider a diagnosis of a disease. The number of infected people will depend on the season during the year or even on the location around the world. Yet another example, consider sentiment analysis, where emotions are assigned to text messages. The priors over emotions will differ depending on the environment where the classifier is deployed. For instance, emotions in messages exchanged between new partners in love versus messages between partners after a long marriage. In all those situations the class priors at test time have changed. To maintain close to the optimal performance we should adapt the classifier's predictions  $f(X) = p_T(Y|X)$  to reflect the change in the class priors.

This thesis aims to evaluate and extend the methodology for adapting classifier to prior shift and compare it to the existing methods. The algorithm proposed in this thesis is estimating test set priors, allowing to adapt classifier's predictions to label shift by re-weighting them in post-processing step. The test set prior estimation is based on classifier's confusion matrix and constrained likelihood maximization of probability of classifier's decision. By constraining the optimization, we handle the known problem [21, 16, 18, 20] of negative values in the test set prior estimated by confusion matrix. The evaluation is primarily performed in the domain of computer vision with CNN classifiers.

The thesis has the following structure: Chapter 2 gives a brief description of classification task and general approaches to domain adaptation. The main focus of Chapter 2 is on unsupervised methods for classifier adaptation to label shift and on classifier calibration. Chapter 3 proposes new algorithms based on confusion matrix to approach the label shift adaptation. Chapter 4 contains an evaluation of algorithms proposed in Chapter 3 and their comparison to the existing methods on standard long-tailed data sets. The final conclusion can be found in Chapter 5.

## Chapter 2

### Related Work

So many books, so little time.

Frank Zappa

### 2.1 Classification

Let  $X$  be a feature space and  $Y$  a finite set of  $K$  classes. The goal of classification task is to learn mapping (classifier)  $h : X \rightarrow Y$  such that for every sample  $x \in X$  with a ground truth label  $y \in Y$  the assignment  $h(x)$  is equal to  $y$ . It is common practice to mark the labels (in general label can be anything e.g. dog, cat,...) with integer values  $1; \dots; K$  and encode them as  $K$ -dimensional one-hot vectors. Meaning that for some label  $y \in \{1; \dots; K\}$ , the vector will have 1 at  $y$ -th position and 0 at all other positions. For example, set of four classes  $\{1; 2; 3; 4\}$  can be encoded as  $f(1; 0; 0; 0); (0; 1; 0; 0); (0; 0; 1; 0); (0; 0; 0; 1)$ . Such encoding allows us to train the classifier to output probability distribution over label space  $Y$ , representing classifier's confidence that sample  $x$  belongs to the particular class. From this point of view, the (probabilistic) classifier  $f$  is a mapping from feature space  $X$  to the probability simplex  $\Delta_{K-1}$ , i.e.  $f : X \rightarrow \Delta_{K-1}$ . Based on this probabilistic output, the input sample  $x$  is assigned label  $y \in Y$ , by taking class with the highest probability

$$d = h(x) := \arg \max_{k \in \{1; \dots; K\}} f(x) = \arg \max_{k \in \{1; \dots; K\}} (f^{(1)}; f^{(2)}; \dots; f^{(K)})(x); \quad (2.1)$$

## 2. Related Work

where  $\sum_{i=1}^K f^{(i)} = 1$  and  $f^{(i)} \geq 0$ . It is often assumed that the classifier  $f$  approximates posterior probability

$$f(x) = p(Y|X); \quad (2.2)$$

### 2.1.1 Training Classifier

To train a classifier, we need data from which our model will be able to extract some useful information to solve the classification task. In supervised learning the training data  $T = \{x_i; y_i\}_{i=1}^M$  are a set of features  $x_i$  and corresponding labels  $y_i$  (see Figure 2.1 for example). Once the model is trained, it is useful to know how well the classifier performs. For this purpose, we need to define some performance measure. The simplest choice is 0-1 loss, which is 1 for  $h(x) \neq y$  and 0 for  $h(x) = y$ . The expected value of 0-1 loss is called error rate. The opposite value to error rate is accuracy. The relation between accuracy ( $acc$ ) and error rate ( $err$ ) is as follow

$$acc = 1 - err; \quad (2.3)$$

Usually, we do not have access to all examples from feature space and therefore the error rate and accuracy have to be estimated from a finite number of samples. One could estimate the performance on the training set, but due to over-fitting, the estimated values do not reflect how well the classifier generalizes to new samples. Consequently, it is a common practice to estimate the values on a set of samples  $V = \{x_i; y_i\}_{i=1}^{M^0}$ , called validation set, unseen during the training stage. The accuracy and error rate can be estimated as

$$\begin{aligned} acc &= \frac{1}{M^0} \sum_{(x_i; y_i) \in V} \mathbb{1}_{h(x_i) = y_i} \\ err &= \frac{1}{M^0} \sum_{(x_i; y_i) \in V} \mathbb{1}_{h(x_i) \neq y_i} \end{aligned} \quad (2.4)$$

Minimizing the error rate directly is typically intractable, because it is a discontinuous and non-convex function, but can be minimized indirectly using some smooth, more tractable function under the Empirical Risk Minimization (ERM) [2] framework.

In general we would like to minimize true error

$$R(h) = E_{(X; Y) \sim P} [l(y; h(x))]; \quad (2.5)$$

Figure 2.1: An example of training data from CIFAR10 dataset. The dataset contains RGB images of size  $32 \times 32$  pixels. Each image is assigned one of ten classes.

where  $E$  denotes expected value and  $l: Y \times Y \rightarrow \mathbb{R}$  is a loss function. Since we do not have an access to  $\mathbb{P}(X; Y)$ , the problem is approached within the ERM framework by minimizing the empirical risk on a training set

$$R(h) = \frac{1}{M} \sum_{(x_i; y_i) \in \mathcal{T}} l(y_i; h(x_i)); \quad (2.6)$$

For a probabilistic classifier, one of the most common choices for  $l$  is the cross-entropy loss

$$l(y_i; f(x_i)) = - \sum_{j=1}^K y_i^{(j)} \log f^{(j)}(x_i); \quad (2.7)$$

where  $y_i^{(j)}$  denotes the  $j$ -th position of one-hot encoded ground truth label

## 2. Related Work

and  $f^{(j)}$  is the  $j$ -th output of probabilistic classifier. It can be observed that unlike error rate in Equation (2.4), the cross-entropy loss is smooth and convex and we can easily use this function to train our classifier with gradient descent.

### 2.2 Convolutional Neural Networks

Recently, Convolutional Neural Networks [4] are very popular and widely used in many computer vision tasks like classification, image segmentation, image completion etc.

The main building unit of CNN is the convolution operator. Depending on the form of data, the convolution can be defined over any number of dimensions. In the case of image data, CNNs are based on two-dimensional convolution

$$(I * F)(i; j) = \sum_m \sum_n I(m; n) F(i - m; j - n); \quad (2.8)$$

where  $I$  is an input image,  $F$  is a filter applied to the image and values  $F(k; l)$  are learnable parameters. Convolution is very similar to cross-correlation and in some literature, the operator in CNNs is called cross-correlation instead of convolution. The difference between them in a two-dimensional case is that convolution rotates the filter by 180 degrees. Since in deep learning the parameters in the filter are learnable, both operations will lead to the same result. In fact, in many frameworks cross-correlation is used for CNN implementation.

Neural networks are composed from several layers. A single CNN layer is made out of many filters, followed by a non-linear function. Their arrangement in the layer affects, together with the layer stacking, the network architecture. One of the first architectures was LeNet [5] (schema of LeNet-5 architecture is shown on image(2.2)), designed to recognize handwritten digits. Since then, many new architectures have been proposed: VGG2[2], GoogLeNet [23], ResNet [24], ResNeXt [25], ResNeSt [26] and Efficientnet [27] to name a few.

### 2.3 Unsupervised Domain Adaptation

At test time it is often assumed that the training set is a good reflection of test distribution  $p_E(X; Y)$ . When this assumption is not met, the classifier



Figure 2.2: Schema of a LeNet-5 CNN architecture. Six  $5 \times 5$  convolutional filters are applied on input,  $32 \times 32$  grayscale image, resulting in  $6 \times 28 \times 28$  feature map. Dimension of this feature map is reduced by subsampling layer and then another six convolutional filters are applied, again followed by subsampling. At the end, there are 3 fully connected layers. Image taken from [28].

fails to generalize to samples drawn from  $p_E(X; Y)$  and the performance is deteriorated. The aim of the domain adaptation is to train a model on the training distribution  $p_T(X; Y)$  capable to generalize well to a new samples from distribution  $p_E(X; Y)$ , when  $p_T(X; Y) \neq p_E(X; Y)$ . Domain adaptation is unsupervised when we have access to labeled source data  $\mathcal{D}_T = \{x_i; y_i\}_{i=1}^M$ , but unlabeled target data  $\mathcal{D}_E = \{x_i\}_{i=1}^N$ .

One of the approaches to domain adaptation is the importance re-weighting [14, 12, 13, 11, 29], commonly used in the label and covariate shift scenarios. This technique is based on minimization of an expected loss on the target domain by re-weighting the loss function on samples from the source domain:

$$\begin{aligned} R_E(h) &= E_{(X;Y) \sim p_E(X;Y)} [l(y; h(x))] = \int p_T(X; Y) \frac{p_E(X; Y)}{p_T(X; Y)} l(y; h(x)) dx dy = \\ &= E_{(X;Y) \sim p_T(X;Y)} \frac{p_E(X; Y)}{p_T(X; Y)} l(y; h(x)) : \end{aligned} \tag{2.9}$$

By factorizing the joint distribution  $p(X; Y)$  to  $p(X)p(Y|X)$  or  $p(Y)p(X|Y)$ , we get following ratios  $p_E(X) = p_T(X)$ ,  $p_E(Y|X) = p_T(Y|X)$  and  $p_E(Y) = p_T(Y)$ ,  $p_E(X|Y) = p_T(X|Y)$ , respectively, in the Equation (2.9). Estimating one of the two ratio pairs allows us to re-weight a loss function at the training stage and therefore to adapt the classifier to the target domain. Note, that estimating ratios containing conditional densities  $p(X|Y)$  and  $p(Y|X)$  is very difficult, as they are hard to model, unless we have some prior information about how the densities change. Consequently, the importance re-weighting is common

## 2. Related Work

Figure 2.3: Schema of an adversarial neural network for domain adaptation. Feature extractor (green) and label predictor (blue) together form a standard neural network for classification. Domain-invariant mapping is learnt in an adversarial manner with domain classifier (red) by applying reversed gradient in feature extractor with respect to the domain loss  $L_d$ . The image is taken from [35].

to use with the assumption that one of the conditional densities does not change.

Another technique is based on the assumption that some learn-able transformation [9, 30, 31, 32] exists between the domains. Unfortunately, this method tends to overfit when working with high-dimensional data and complicated domain transformation. This is common for example in computer vision, where feature space can have thousands or even millions of dimensions.

Instead of transforming one domain into another, an alternative approach attempts to find domain-invariant representation [9, 33, 34]. This technique learns a transformation from the feature space  $X$  into a possibly lower-dimensional feature space  $Z$ , which is invariant to the source and target domain. Attempting to reach this goal, promising results were recently achieved by methods [35, 36, 37, 38, 39] leveraging adversarial learning (AL) [40]. Domain adaptation based on AL is using a feature extractor that maps original features to domain-invariant space. Two classifiers are learned on top of this intermediate feature space. The first classifier is predicting class labels, while the second one predicts domain in an adversarial manner (see Figure 2.3).

## 2.4 Classifier Adaptation to Label Shift

Recall that label shift makes the following assumptions about the data generating distributions:

1.  $p_T(Y) \neq p_E(Y)$
2.  $p_T(X|Y) = p_E(X|Y)$

Let us assume probabilistic classifier  $p_T(x) = p_T(Y|x)$  trained on the training set  $T$  with relative class frequencies  $p_T(Y)$ . The formula for adapting the classifier's predictions to the new priors  $p_E(Y)$  in a test set  $E = \{x_i\}_{i=1}^N$  can be simply derived by applying Bayes theorem to the assumption about equality of class-conditional distributions

$$p_T(x|Y) = p_E(x|Y) = \frac{p_T(Y|x)p_T(x)}{p_T(Y)} = \frac{p_E(Y|x)p_E(x)}{p_E(Y)} \quad (2.10)$$

By expressing adapted posteriors, we get

$$p_E(Y|x) = p_T(Y|x) \frac{p_E(Y)p_T(x)}{p_T(Y)p_E(x)} \quad (2.11)$$

Note that the sum over posteriors is equal to one

$$\sum_{i=1}^K p_E(Y|x) = \sum_{i=1}^K p_T(Y|x) \frac{p_E(Y)p_T(x)}{p_T(Y)p_E(x)} = 1 \quad (2.12)$$

and the ratio  $p_T(x) = p_E(x)$  can be factored out before the sum, thus

$$\frac{p_T(x)}{p_E(x)} = \sum_{i=1}^K p_T(Y|x) \frac{p_E(Y)}{p_T(Y)} \quad (2.13)$$

is normalization constant and Equation (2.11) can be expressed as

$$p_E(Y|x) / p_T(Y|x) \frac{p_E(Y)}{p_T(Y)} \quad (2.14)$$

which gives us the recipe for adaption of classifier's predictions under the label shift.

The remaining question is how to estimate test and training priors, which are used to re-weight the predictions. If we have access to the training set

## 2. Related Work .....

T, we can estimate training priors  $p_T(Y)$  by simply counting ground-truth labels in the data set (although other options are possible, see experiments in the Section 4.2.1). Estimating test set priors  $p_E(Y)$  is typically more difficult because usually, we do not have access to the ground truth labels. This thesis is focused on answering the question "how to estimate test priors". An alternative approach is to estimate the priors ratio  $w(Y) = p_E(Y)/p_T(Y)$  directly. Note that once we have estimated priors  $p_T(Y)$  and  $p_E(Y)$ , instead of re-weighting the predictions on test set according to the Equation(2.14), we can use the ratio to train a new classifier, adapted to the target domain, by minimizing the expected loss in the Equation(2.9). This can be done by directly re-weighting the loss function during training or by sampling training examples according to the estimated importance weights.

### ■ 2.4.1 Confusion Matrix

One of the first procedures [41, 16] for prior estimation is based on a  $K \times K$  left stochastic matrix  $C_{d|y}$  called confusion matrix (CM). The value in the  $k$ -th column and  $i$ -th row is interpreted as probability  $p(D = i | Y = k)$  of the classifier  $h$  deciding for class  $i$  when the true class is  $k$ . Both  $D$  and  $Y$  are random variables over the label space, but  $D$  represents the classifier's decision, while  $Y$  represents the ground-truth label.

Useful property of confusion matrix is that it does not change with prior shift [20]:

Lemma 2.1. Let  $Y$  be a random variable representing the ground truth label of a sample  $x$  and  $D$  a random variable representing classifier's decision  $h(x)$ . If assumption  $p_T(x|Y) = p_E(x|Y)$  holds, then

$$p(D|Y) := p_T(D|Y) = p_E(D|Y);$$

$$C_{d|y} := C_{d|y}^T = C_{d|y}^E;$$

Proof. From the law of total probability

$$p_T(D|Y) = \sum_{x \in X} p_T(D; x|Y) = \sum_{x \in X} p_T(D|x; Y) p_T(x|Y):$$

Note that the classifier's decision depends only on  $x$ . From the property of conditional independence, equality  $p_T(D|x; Y) = p_T(D|x)$  holds. And since  $p_T(D|x)$  depends only on a fixed classifier  $h$ , it is equal to  $p_E(D|x)$ . Following

this observation and applying the label shift assumption, the equation above can be rewritten as

$$\prod_{x \in X} p_E(D|x; Y) p_E(x|Y) = \prod_{x \in X} p_E(D; x|Y) = p_E(D|Y):$$

□

Marginalizing over the joint density  $p(D; Y)$

$$p(D = i) = \sum_{k=1}^K p(D = i|Y = k) p(Y = k); \quad (2.15)$$

$$p(D) = C_{d|y} p(Y):$$

McLachnan [41] and Saerens et al. [6] simply compute the new priors  $p_E(Y)$  from the Equation (2.15):

$$\hat{p}_E(Y) = \hat{C}_{d|y}^{-1} \hat{p}_E(D); \quad (2.16)$$

using an estimate of  $C_{d|y}$  computed on a validation set  $V$

$$\hat{p}(D = i|Y = k) = \frac{|\{f(x_j; y_j) \in V : d(x_j) = i \wedge y_j = k\}|}{|\{f(x_j; y_j) \in V : y_j = k\}|}, \quad (2.17)$$

where  $j$  is the set cardinality, and an estimate of  $p(D)$  computed by counting the classifier's decisions on the test set  $E$

$$\hat{p}_E(D = i) = \frac{|\{f(x_j; y_j) \in E : d(x_j) = i\}|}{|E|}. \quad (2.18)$$

Let us also consider a soft confusion matrix<sup>1</sup> (SCM)  $C_{d|y}^{\text{soft}}$  estimated from the classifier's soft predictions  $f$  as

$$C_{:jk}^{\text{soft}} = \frac{1}{N_k} \sum_{x_i: y_i = k} f(x_i); \quad (2.19)$$

where  $C_{:jk}^{\text{soft}}$  denotes the  $k$ -th column of the SCM. The probability  $p_E^{\text{soft}}(D)$  can be estimated by averaging predictions  $f(x)$  over the test set. The new priors are then computed similarly to the Equation (2.16).

<sup>1</sup>Following the terminology of Lipton et al. [20].

### 2.4.2 Adaptation Based on Maximum Likelihood

Saerens et al. [6] suggested to estimate the test time priors  $p_E(Y)$  by maximizing the likelihood

$$L(E) = \prod_{i=1}^M p_E(x_i) = \prod_{i=1}^M \sum_{k=1}^K p_E(x_i; Y = k) = \prod_{i=1}^M \sum_{k=1}^K p_E(x_i | Y = k) p_E(Y = k): \quad (2.20)$$

The proposed EM algorithm is working iteratively in two steps:

$$\hat{p}_E^{(s)}(Y = k | x_i) := \frac{\hat{p}_T(Y = k | x_i) \frac{\hat{p}_E^{(s)}(Y = k)}{\hat{p}_T(Y = k)}}{\sum_{j=1}^K \hat{p}_T(Y = j | x_i) \frac{\hat{p}_E^{(s)}(Y = j)}{\hat{p}_T(Y = j)}}; \quad (2.21)$$

$$\hat{p}_E^{(s+1)}(Y = k) := \frac{1}{M} \sum_{i=1}^M \hat{p}_E^{(s)}(Y = k | x_i); \quad (2.22)$$

In the first step, it recomputes the posterior predictions (Equation (2.21)). The classifier outputs are denoted as  $\hat{p}_T(Y = k | x_i)$ . In the second step, it estimates the test priors by averaging all updated predictions (Equation (2.22)). The algorithm requires initialization of test priors  $\hat{p}_E^{(0)}(Y = k)$ . This can be done for example by the relative class frequencies in the training set, i.e.  $\hat{p}_T(Y = j) = M_j / M$ .

Du Plessis and Sugiyama [5] showed that the EM algorithm, proposed by Saerens et al. [6], can be derived from minimization of KL divergence between  $p_E(x)$  and its approximation  $p_E^0(x)$  modeled as

$$p_E^0(x) = \sum_{k=1}^K \hat{P}_k \hat{p}_T(x | Y = k); \quad (2.23)$$

where  $\hat{P}_k := \hat{p}_E(Y = k)$ . The KL divergence is expressed as

$$\begin{aligned} \text{KL}(p_E || p_E^0) &= \int p_E(x) \log \frac{p_E(x)}{p_E^0(x)} dx = \\ &= \int p_E(x) \log p_E(x) dx + \int p_E(x) \log \sum_{k=1}^K \hat{P}_k \hat{p}_T(x | Y = k) dx: \end{aligned} \quad (2.24)$$

The first integral depends only on  $x$ . Consequently it is a constant for a fixed input and it can be ignored during minimization. The second integral

can be viewed as an expectation over  $\log p_E^0(x)$  and can be approximated by empirical average

$$\int_{y \in Y} p_E(x) \log \sum_{k=1}^K \hat{p}_k p_T(x|Y=k) dx \approx \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \hat{p}_k p_T(x_i|Y=k): \quad (2.25)$$

These observations give the following optimization problem:

$$\begin{aligned} \hat{p} &= \arg \max_{\hat{p}} \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \hat{p}_k p_T(x_i|Y=k) = \\ &= \arg \max_{\hat{p}} \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \hat{p}_k \frac{p_T(Y=k|x_i)p_T(x_i)}{p_T(Y=k)} = \\ &= \arg \max_{\hat{p}} \frac{1}{N} \sum_{i=1}^N \log p_T(x_i) + \sum_{k=1}^K \hat{p}_k \frac{p_T(Y=k|x_i)}{p_T(Y=k)} = \\ &= \arg \max_{\hat{p}} \frac{1}{N} \sum_{i=1}^N \log p_T(x_i) + \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \hat{p}_k \frac{p_T(Y=k|x_i)}{p_T(Y=k)} = \\ &= \arg \max_{\hat{p}} \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \hat{p}_k \frac{p_T(Y=k|x_i)}{p_T(Y=k)} \\ &\text{s.t. } \sum_{k=1}^K \hat{p}_k = 1; \quad \hat{p}_k \geq 0 \end{aligned} \quad (2.26)$$

Indeed, this is equivalent to maximizing the logarithm of likelihood function from Equation (2.20). In other words, it is maximization of log-likelihood  $l(E) = \log L(E)$ .

Sulc and Matas [17] experimented with maximizing the log-likelihood function by projected gradient ascent

$$\hat{p}_k^{s+1} = \mathcal{P}_{\Delta_k} \left( \hat{p}_k^s + \frac{\partial l(E)}{\partial \hat{p}_k} \right); \quad (2.27)$$

where  $\mathcal{P}_{\Delta_k}(\cdot)$  denotes projection onto the probability simplex and

$$\frac{\partial l(E)}{\partial \hat{p}_k} = \sum_{i=1}^N \frac{p_T(Y=k|x_i)}{p_T(Y=k)} \left( \frac{p_T(Y=k|x_i)}{p_T(Y=k)} - \sum_{j=1}^K \hat{p}_j \frac{p_T(Y=j|x_i)}{p_T(Y=j)} \right); \quad (2.28)$$

The experimental results showed that the EM algorithm converged faster than the gradient ascent, while achieving similar results.

### 2.4.3 Adaptation Based on Maximum A Posteriori

As an extension to the projected gradient ascent algorithm given in the Equation (2.27), Sulc and Matas [17] proposed maximum a posteriori estimate

$$\begin{aligned} \hat{P} &= \arg \max_{\hat{P}} p(\hat{P}|E) = \arg \max_{\hat{P}} p(\hat{P})p(E|\hat{P}) = \\ &= \arg \max_{\hat{P}} \log p(\hat{P}) + \log p(E|\hat{P}); \end{aligned} \quad (2.29)$$

adding a hyper-prior  $p(\hat{P})$  to the log-likelihood function  $I(E)$ . The hyper-prior represents a prior knowledge about the categorical distribution  $p_E(Y)$  and was modeled by symmetric Dirichlet distribution

$$p(\hat{P}) = \text{Dir}(\alpha) = \frac{1}{B(\alpha)} \prod_{k=1}^K \hat{P}_k^{\alpha-1}; \quad (2.30)$$

where  $\alpha = (\alpha_1; \dots; \alpha_K)$ ,  $\alpha_1 = \dots = \alpha_K = \alpha$  and  $B(\alpha)$  is the multivariate beta function. The intuition behind the Dirichlet distribution is that for  $\alpha \gg 1$  the probability vector  $\hat{P}$  is forced to be a dense distribution, while for  $0 < \alpha < 1$  being a sparse one. See Figure 2.4 for visualization.

The solution to the maximum a-posteriori can be found by projected gradient ascent, adding the derivative of  $\log \text{Dir}(\alpha)$  into the  $I(E)$  function in the Equation (2.27)

$$\frac{\partial \log p(P)}{\partial P_k} = \frac{\partial \log \text{Dir}(\alpha)}{\partial P_k} = -\frac{1}{P_k}; \quad (2.31)$$

### 2.4.4 BBSE

So far, all of the mentioned algorithms were approaching the problem of prior shift adaptation through estimating the test time prior  $p_E(Y)$  and the training prior  $p_T(Y)$  separately. But, as proposed by Lipton et al. [20], it is also possible to estimate directly the prior ratio  $w(Y) = p_E(Y)/p_T(Y)$  using the confusion matrix  $C_{d,y}$  with joint probability  $p(D = i; Y = k)$ , unlike the conditional probability used in the Section 2.4.1.



Figure 2.4: Visualization of Dirichlet distribution on 3-dimensional probability simplex. Red encodes higher probabilities, blue smaller probabilities.

Let us again consider the Equation (2.15) and multiply it by  $\frac{p_T(Y=k)}{p_T(Y=k)}$

$$\begin{aligned}
 p_E(D = i) &= \sum_{k=1}^K p_T(D = i | Y = k) p_E(Y = k) = \\
 &= \sum_{k=1}^K p_T(D = i | Y = k) p_E(Y = k) \frac{p_T(Y = k)}{p_T(Y = k)} = \\
 &= \sum_{k=1}^K p_T(D = i; Y = k) \frac{p_E(Y = k)}{\underbrace{\frac{p_T(Y = k)}{w(Y = k)}}_{\{z = k\}}} = \quad (2.32) \\
 &= \sum_{k=1}^K p(D = i; Y = k) w(Y = k); \\
 p_E(D) &= C_{d,y} w(Y):
 \end{aligned}$$

The result is similar to the Equation (2.15), therefore it is not surprising that the ratio  $w(Y)$  is computed in the very same way

$$w(Y) = \hat{C}_{d,y}^{-1} p_E(D): \quad (2.33)$$

The only difference is in the normalization of the confusion matrix: while in the conditional confusion matrix  $C_{d,y}$  columns sum up to one, in the joint confusion matrix  $C_{d,y}$  all elements sum up to one. The CMC<sub>d,y</sub> with the joint probability is estimated as

$$p(D = i; Y = k) = \frac{\sum_{j \in V} \mathbb{1}(x_j = i \wedge y_j = k)}{|V|}: \quad (2.34)$$

## 2. Related Work

It is also possible to use the soft confusion matrix  $C_{d;y}^{\text{soft}}$  computed from the classifier's predictions

$$C_{:,k}^{\text{soft}} = \frac{1}{N} \sum_{x_i: y_i = k} f(x_i); \quad (2.35)$$

where  $C_{:,k}^{\text{soft}}$  denotes the  $k$ -th column of SCM. The probability  $p_E^{\text{soft}}(D)$  can be estimated by averaging predictions  $f(x)$  over the test set. The ratio is then computed in the same way as in the Equation (2.33).

This estimation is called the Black Box Shift Estimation (BBSE). A variant using a soft confusion matrix  $C_{d;y}^{\text{soft}}$  is denoted BBSE-S.

Note that the validation set used for the estimation of  $C_{d;y}$ , should have the same label distribution as the training set. Otherwise we have to re-weight rows of  $C_{d;y}$  by the ratio of training and validation set priors  $p_T(Y) = p_V(Y)$ .

The fact that  $C_{d;y}$  estimated on a validation set with  $p_T(Y) \neq p_V(Y)$  cannot be simply used for estimation of  $w(Y)$  is apparent from the following inequality:

$$p_V(D; Y) = p_V(D|Y)p_V(Y) = p_T(D|Y)p_V(Y) \neq p_T(D|Y)p_T(Y) = p_T(D; Y); \quad (2.36)$$

In order to use the validation set, we need to multiply  $p_V(D; Y)$  by  $p_T(Y) = p_V(Y)$ , to make it equal to  $p_T(D; Y)$ . This observation means that the BBSE also requires access to the training set.

### 2.4.5 RLLS

Azizzadenesheli et al. [9] proposed to estimate the prior ratio  $w(Y)$  in two steps. First, because in case of no label shift prior ratio is equal to the vector of ones, i.e.  $w(Y) = \mathbf{1}$ , they define amount of weight shift

$$\delta := w(Y) - \mathbf{1} \quad (2.37)$$

and propose to find  $\delta$  instead of  $w(Y)$ . This way it is possible to use  $l_2$ -regularization in the optimization procedure to penalize large shift. Plugging into the Equation (2.32) we get

$$\begin{aligned} p_E(D) &= C_{d;y}(\delta + \mathbf{1}) \\ p_E(D) - C_{d;y}\mathbf{1} &= C_{d;y}\delta \\ \mathbf{b} &= C_{d;y}\delta; \end{aligned} \quad (2.38)$$

where  $b := p_E(D) - C_{d,y}1$ . Azizzadenesheli et al. choose to find in the Equation (2.38) by minimizing  $l_2$ -norm with a regularization term to prevent large shifts

$$\hat{b} = \arg \min_{\hat{b}} \|\hat{C}_{d,y} \hat{b}\|_2 + c \|\hat{b}\|_2; \quad (2.39)$$

where  $c$  is a hyper-parameter, which can be found by cross validation and  $\hat{C}_{d,y}$  and  $p_E(D)$  are estimated in the same way as in the Section 2.4.4. The solution to the problem given by the Equation (2.39) is found by a third party solver.

In the second step the prior ratio is computed as

$$w(Y) = 1 + \hat{b}; \quad (2.40)$$

where

$$\hat{b} = \begin{cases} 8 & \\ < 1; & \text{if } N < \frac{1}{\frac{1}{2} \max(\min, \frac{1}{M})^2} \\ 0; & \text{otherwise} \end{cases} \quad (2.41)$$

$M$  is the number of samples in the training set,  $N$  is the number of samples in the test set,  $\min$  is the smallest singular value corresponding to the confusion matrix  $\hat{C}_{d,y}$  and  $\max$  represents our belief about the size of the shift.

Determination of  $\hat{b}$  in the Equation (2.41) serves as an indicator of how bad our estimate is. If the estimation of  $\hat{b}$  is based on an insufficient number of samples or the confusion matrix is close to the singular matrix (its smallest singular value is close to zero), then it is better to decide for no shift and set  $\hat{b}$  to zero.

## 2.5 Classifier Calibration

It is important to know how much we can trust the decision of our classifier. Therefore, having a probabilistic classifier  $f(x)$  trained to approximate posterior probabilities  $p(Y|X)$ , we would like the outputs to reflect the true probabilities of correct and incorrect classification. For example, from all the samples being classified with a confidence of 0.9, ninety percent of them should be classified correctly and ten percent of them incorrectly. Such property is desirable also in algorithms for adapting classifiers to the label shift because most of the methods work with the assumption that the classifiers approximate posterior probability. Unfortunately, common probabilistic classifiers, such as Convolution Neural Networks, tend to have overconfident predictions due to overfitting to the training set.

## 2. Related Work .....

When the classifier's outputs reflect its predictive uncertainty, we call the classifier calibrated. In the literature there are several definitions for calibrated classifier [42, 43]:

Definition 2.2. A probabilistic classifier  $f$  is multiclass-calibrated, or simply calibrated, if for set of all samples with prediction vector  $f(x)$ , the proportions of classes in this set are equal to  $\phi(x)$ :

$$p(Y|f(x)) = \phi(x): \tag{2.42}$$

Another weaker definition was proposed by [44]. The definition is weaker in the sense that it does not consider whole prediction vectors, but every class in the vector separately.

Definition 2.3. A probabilistic classifier  $f$  is classwise-calibrated, if for any class  $i$  and any predicted probability  $q_i$  for this class:

$$p(Y = i | f^{(i)}(x) = q_i) = q_i: \tag{2.43}$$

The definition basically says that among all samples predicted to belong into class  $i$  with probability  $q_i$ , the proportion of samples with ground truth class  $i$  is equal to  $q_i$ . Yet, one more even weaker definition exists [45]:

Definition 2.4. A probabilistic classifier  $f$  is confidence-calibrated if

$$p(Y = \arg \max f(x) | \max f(x)) = \max f(x): \tag{2.44}$$

In order to better understand the definitions 2.2, 2.3 and 2.4 let us demonstrate them on several examples:

Example 1. Suppose classification task into a three classes  $\mathcal{Y} = \{1, 2, 3\}$ , with label distribution  $p(Y) = (0.5; 0.1; 0.4)$ , and a classifier with a constant output  $f(x) = (0.5; 0.1; 0.4)$ . According to Definition 2.2 classifier  $f(x)$  is multiclass-calibrated. Note that the classifier is also classwise-calibrated and confidence-calibrated according to Definition 2.3 and 2.4 respectively.

Example 2. Let us consider classification task into a three classes  $\mathcal{Y} = \{1, 2, 3\}$ . Classifier  $f(x)$  outputs four predictions, displayed in the first column of Table 2.1. The true proportions of classes among samples having particular output vector  $f(x)$  are in the second column of Table 2.1. The classifier is classwise-calibrated and confidence calibrated, but not multiclass-calibrated.

$f(X)$	$p(Y f(X))$
(0.2, 0.5, 0.3)	(0.3, 0.4, 0.3)
(0.2, 0.3, 0.5)	(0.1, 0.3, 0.6)
(0.3, 0.5, 0.2)	(0.2, 0.6, 0.2)
(0.3, 0.2, 0.5)	(0.4, 0.2, 0.4)

Table 2.1: Classifier  $f(X)$  is given by its outputs in the first column. The true distribution of samples having prediction  $f(X)$  are in second column.

Example 3. Let us consider classification task into a three classes  $\mathcal{Y} = \{1, 2, 3\}$ . Table 2.2 gives outputs of classifier  $f(x)$  in the first column and corresponding ground truth relative frequencies of classes in the second column. The classifier is only confidence-calibrated according to Definition 2.4.

$f(X)$	$p(Y f(X))$
(0.2, 0.7, 0.1)	(0.0, 0.7, 0.3)
(0.2, 0.2, 0.6)	(0.3, 0.1, 0.6)

Table 2.2: Classifier  $f(X)$  is given by its outputs in the first column. The true distribution of samples having prediction  $f(X)$  are in second column.

Vaicenavicius et al. [43] formulated an evaluation framework for classifier calibration, under which all weaker definitions are only special cases of the Definition 2.2. In the framework, the authors introduced so-called calibration lenses which allow us to evaluate calibration on simpler induced problems. It is particularly useful in multiclass problems, where we usually do not have enough samples to evaluate general calibration. An example of such induced problem is transforming the  $K$ -classification task into binary classification, where a subset of  $K$  classes is merged into one class and the rest into the second class. Indeed, evaluating a classifier on an induced problem has one drawback: Although the calibrated classifier is also calibrated on the induced problem, it does not apply contrariwise [43].

### 2.5.1 Evaluating Calibration

Let us first define canonical calibration function [43]:

Definition 2.5. Canonical calibration function  $\mathcal{C} : \mathcal{K} \rightarrow \mathcal{K}$  of classifier  $f : X \rightarrow \mathcal{K}$  is

$$\mathcal{C}(q) := p(Y|f(x) = q) \quad \forall q \in \mathcal{K} \quad (2.45)$$

## 2. Related Work

Canonical calibration function is a mapping, which transforms the classifier's outputs into calibrated predictions. Therefore, the classifier is multiclass-calibrated if its calibration function is identity. Now we can define expected miscalibration [43]:

Definition 2.6. Expected miscalibration on  $A \subseteq \Delta_{K-1}$  with respect to some distance is

$$\hat{d} := E[d(\hat{f}(X)); f(X)] | f(X) \in A; \quad (2.46)$$

Here  $A$  only denotes a subset of probability simplex on which we want to evaluate the calibration of our classifier.

Let  $\mathcal{C} = \{C_i\}_{i=1}^L$  be a partition of  $A \subseteq \Delta_{K-1}$ , then the estimator of calibration function is defined as

$$\hat{c}^{(j)}(q) := \frac{|\{i : f(x_i) \in [q] \wedge y_i = j\}|}{|\{i : f(x_i) \in [q]\}|} \quad \text{for } j = 1 :: K; \quad (2.47)$$

where  $[q]$  denotes the partition for which  $q \in \Delta_k$  holds. To summarize it, we split subset  $A$  into  $L$  partitions (bins) and in every bin we compute relative frequencies for every class (from samples falling into the bin based on the classifier's predictions). These relative frequencies then determine the estimate of the calibration function in a particular bin. The estimator of the canonical function is constant on every partition, therefore we can define  $\hat{c}_i$  as a value of the estimator in the partition  $C_i$ .

Next, we can define an estimator of expected miscalibration as

$$\hat{d} := \sum_{i=1}^L \hat{p}_i d(\hat{c}_i; \hat{f}_i); \quad (2.48)$$

where  $\hat{p}_i$  is the proportion of samples from the data set falling into the bin  $C_i$

$$\hat{p}_i = \frac{|\{j : f(x_j) \in C_i\}|}{|\{j : f(x_j) \in A\}|} \quad (2.49)$$

and  $\hat{f}_i$  is the average of all predictions in the bin  $C_i$

$$\hat{f}_i = \frac{\sum_{j : f(x_j) \in C_i} f(x_j)}{|\{j : f(x_j) \in C_i\}|} \quad (2.50)$$

Note that whenever the denominator in Equations (2.47), (2.49), (2.50) is zero, the value of the particular estimator is set to 0. The estimator of expected

Figure 2.5: An example of reliability diagram, taken from [45]. The error denotes Expected calibration error.

miscalibration is nothing else than weighted distance between the histogram based estimation of calibration function and the classifier's predictions.

Vaicenavicius et al. [43] showed that as the number of partitions and the number of data samples grow to infinity, the estimator of expected miscalibration  $\hat{\epsilon}_d$  is equal to  $\epsilon_d$ . However, if the number of partitions is fixed and only the number of data samples grow to infinity, then the estimator is only a lower bound  $\hat{\epsilon}_d \leq \epsilon_d$  and  $\hat{\epsilon}_d$  is an inconsistent estimate. The intuition behind this fact is that when the predictions in the bin are averaged, the overconfident and underconfident estimates can cancel out.

Guo et al. [45] used expected miscalibration for induced binary classifier with total variation distance  $d(x; y) = \|x - y\|_1$ . This special case of expected miscalibration is called Expected calibration error (ECE).

The visualization of model calibration is done with reliability diagrams [43, 45, 46]. They are mostly used as one-dimensional graphs, plotting accuracy vs. confidence (see Figure 2.5), but multidimensional diagrams were proposed as well [43]. Note that the multidimensional diagrams suffer from a lack of data due to the curse of dimensionality. In general, reliability diagrams plot information about  $\hat{p}$  and  $f^{\wedge}$  for all partitions  $\mathcal{I}_i$ .

■ 2.5.2 Calibration Methods

Guo et al. [45] evaluated several calibration methods in the context of neural network classifiers and observed that temperature scaling performs best in most cases.

Temperature scaling (TS) has a single real valued parameter  $T > 0$ , called temperature. The neural network's logits are divided by  $T$  right before the softmax layer

$$f_{TS}^{(i)}(x) = P \frac{\exp(z_i(x)=T)}{\sum_{j=1}^K \exp(z_j(x)=T)}; \quad (2.51)$$

TS does not affect the accuracy of the model, because the class with maximum probability remains the same, only the confidence for the decision changes. For  $T > 1$  the resulting prediction is closer to the uniform distribution (confidence is decreasing)

$$\begin{aligned} \lim_{T \rightarrow 1} P \frac{\exp(z_i(x)=T)}{\sum_{j=1}^K \exp(z_j(x)=T)} &= \lim_{T \rightarrow 1} P \frac{1}{\sum_{j=1}^K \frac{\exp(z_j(x)=T)}{\exp(z_i(x)=T)}} = \\ &= P \frac{1}{\sum_{j=1}^K \lim_{T \rightarrow 1} \frac{\exp(z_j(x)=T)}{\exp(z_i(x)=T)}} = P \frac{1}{\sum_{j=1}^K \lim_{T \rightarrow 1} e^{(z_j(x) - z_i(x))/T}} = \\ &= P \frac{1}{\sum_{j=1}^K 1} = \frac{1}{K}; \end{aligned} \quad (2.52)$$

while for  $T < 1$  the prediction for the most likely class is getting closer to one and the prediction for the rest of the classes are getting closer to zero (the confidence is increasing). This can be seen by computing the limit

$$\begin{aligned} \lim_{T \rightarrow 0^+} P \frac{\exp(z_i(x)=T)}{\sum_{j=1}^K \exp(z_j(x)=T)} &= \lim_{T \rightarrow 0^+} P \frac{1}{\sum_{j=1}^K \frac{\exp(z_j(x)=T)}{\exp(z_i(x)=T)}} = \\ &= P \frac{1}{\sum_{j=1}^K \lim_{T \rightarrow 0^+} \frac{\exp(z_j(x)=T)}{\exp(z_i(x)=T)}} = P \frac{1}{\sum_{j=1}^K \lim_{T \rightarrow 0^+} e^{(z_j(x) - z_i(x))/T}} = \\ &= \frac{P}{1 + \sum_{j: j \neq i} \lim_{T \rightarrow 0^+} e^{(z_j(x) - z_i(x))/T}}; \end{aligned} \quad (2.53)$$

If the output for class  $i$  is not the maximum, then there is at least one  $j$  for which  $z_j > z_i$ , therefore making limit  $\lim_{T \rightarrow 0^+} e^{(z_j(x) - z_i(x))/T} = 1$  and thus the limit in the Equation (2.53) is equal to 0. If the output for class  $i$  is the maximum, then  $z_j \leq z_i \forall j \in \{1, \dots, K\}$  and  $\lim_{T \rightarrow 0^+} e^{(z_j(x) - z_i(x))/T} = 0 \forall j : j \neq i$ . Therefore the limit in Equation (2.53) is equal to 1.



Besides temperature scaling Guo et al. [5] experimented also with matrix scaling (MS) and vector scaling (VS). Matrix scaling transforms the logit vector  $z(x)$  by linear transformation and then applies the softmax

$$f_{MS}(x) = \text{softmax}(Wz(x) + b): \quad (2.54)$$

where  $\text{softmax}$  denotes softmax function. In case of vector scaling, the matrix  $W$  is restricted to be diagonal matrix. It can be rewritten in the form, where every logit is multiplied by a different scalar weight

$$f_{VS}^{(i)}(x) = \frac{\exp(W_i z_i(x) + b_i)}{\sum_{j=1}^K \exp(W_j z_j(x) + b_j)}: \quad (2.55)$$

While TS does not change the argument of maxima, in the case of MS and VS, there is no such guarantee.

Alexandari et al. [47] proposed the Bias-Corrected Temperature Scaling (BCTS) as an intermediary between vector scaling and temperature scaling

$$f_{BCTS}^{(i)}(x) = \frac{\exp(z_i(x) = T + b_i)}{\sum_j \exp(z_j(x) = T + b_j)}: \quad (2.56)$$

They showed that although the TS can effectively reduce the ECE error, it is not suitable in combination with prior shift adaptation, because of systematic bias and that BCTS in combination with EM algorithm [16] from the Section 2.4.2, outperforms both BBSL [20] and RLLS [19].

The parameters in all of the aforementioned methods are learned via negative log-likelihood (NLL), also called cross-entropy (given by Equation (2.7), on the validation set. The parameters are learned separately from the training procedure, otherwise, they would be learned along with other classifier's parameters to output overconfident predictions.



# Chapter 3

## The Proposed Method

When I find myself in the company of scientists, I feel like a shabby curate who has strayed by mistake into a room full of dukes.

W. H. Auden

As observed in the literature [21, 41, 18], in some cases the Equation (2.16) can result in a vector outside of the  $K - 1$  simplex, i.e. the estimate can contain negative values.

Following the Equation (2.15), the probability of the classifier's decisions  $p(D)$  is a convex combination of columns in  $C_{djy}$  as  $p(Y) \in \Delta_{K-1}$ . Since the columns of the confusion matrix are probability vectors, they define a convex set  $C$  of feasible values  $p(D)$  within the probability simplex  $\Delta_{K-1}$ . In other words, a classifier with the confusion matrix  $C$  will result in decisions from  $p(D) \in C$ . The class distribution  $p(Y)$  determines the value of  $p(D)$  within  $C$ . See Figure 3.1 for illustration. For the true distribution of decisions  $p(D)$  and confusion matrix  $C_{djy}$ , the Equation (2.15) holds. The problem occurs when we work with estimates of the distribution  $\hat{p}(D)$  and confusion matrix  $\hat{C}_{djy}$ . If the estimates computed from a limited set of samples are not consistent, meaning that the classifier with confusion matrix  $\hat{C}_{djy}$ , on data set with priors  $\hat{p}(Y)$ , would never produce decisions with distribution  $\hat{p}(D)$ , there may be no prior probability  $\hat{p}(Y)$  satisfying the Equation (2.15). For example, having

$$\hat{C}_{djy} = \begin{matrix} & \# \\ \begin{matrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{matrix} & ; \hat{p}(D) = \begin{matrix} 1 \\ 0 \end{matrix} \end{matrix} ;$$

the unique solution to the Equation (2.15) is  $\hat{p}(Y) = \begin{matrix} \# \\ 4 \\ 3, 1 \\ 3 \end{matrix}$ .

### 3. The Proposed Method

Figure 3.1: The convex set  $\mathcal{C}_{k-1}$  of all possible values of  $p(D)$  for a classifier with confusion matrix  $C$ .

The prior ratio estimation also suffers from the problem of inconsistent estimates. Let us consider the following example, resulting in a negative value in the estimated prior ratio, which as a ratio of two probabilities should be non-negative:

$$\hat{C}_{d;y} = \begin{matrix} & \text{"} & \text{\#} \\ \begin{matrix} 0.4 & 0.1 \\ 0.1 & 0.4 \end{matrix} & ; & \hat{p}(D) = \begin{matrix} \text{"} & \text{\#} \\ 1 & 0 \end{matrix} & ! & \hat{w}(Y) = \begin{matrix} \text{"} & \text{\#} \\ 0.8 & 0.2 \\ 0.3 & 0.7 \end{matrix} ; \end{matrix} \quad (3.1)$$

#### 3.1 Maximum Likelihood Approach

This thesis proposes a novel procedure for the prior estimation based on maximizing the likelihood of  $p_E(D)$ , which handles inconsistent estimates of  $\hat{p}(D)$  and the confusion matrix  $\hat{C}_{d;y}$  and can even work with singular matrices  $\hat{C}_{d;y}$ , as it does not use matrix inversion.

Let  $n = (n_1; \dots; n_K)$  be the numbers of the classifier's decisions for class  $1; \dots; K$  on test set  $E$  and let us denote  $Q = (q_1; \dots; q_K) := p_E(D)$  the probabilities of the classifier decisions on the test distribution  $p_E(X; Y)$ . Assuming the independence of classifier decisions on the test set  $E$ , the

### 3.1. Maximum Likelihood Approach

likelihood of  $Q$  follows by the multinomial distribution

$$L(Q) = p(n|Q) = \frac{(n_1 + \dots + n_K)!}{n_1! \dots n_K!} q_1^{n_1} \dots q_K^{n_K}; \quad (3.2)$$

Substituting the Equation (2.15) into the likelihood function  $L(Q)$ , we can express the likelihood function of class priors  $P$

$$L(P) = p(n|P) = \frac{(n_1 + \dots + n_K)!}{n_1! \dots n_K!} \prod_{k=1}^K (c_{kj} : P)^{n_k}; \quad (3.3)$$

where  $c_{kj}$  is the  $k$ -th row of  $C_{djy}$ .

The log-likelihood is

$$\ell(P) = \log p(n|P) = \sum_{k=1}^K n_k \log(c_{kj} : P) + \text{const}; \quad (3.4)$$

where  $\text{const}$  is constant for a fixed  $n$ .

We estimate the new class priors by maximizing the log-likelihood from Equation (3.4):

$$\hat{P} = \arg \max_P \ell(P) = \arg \max_P \sum_{k=1}^K n_k \log c_{kj} : P \quad (3.5a)$$

$$\text{s.t.} : \sum_{k=1}^K P_k = 1; \quad 0 \leq P_k \leq 1 \quad (3.5b)$$

By this formulation we are avoiding estimation of  $p_E(D)$  as a potential source of inconsistency. Note that  $p_E(D)$  is estimated implicitly with  $P$ . The convex objective can be iteratively maximized using projected gradient ascent

$$\hat{P}^{s+1} = \Pi(\hat{P}^s + r \ell'(P^s)); \quad (3.6)$$

where  $\Pi(\cdot)$  denotes projection onto the probability simplex and the gradient is computed as

$$r \ell'(P) = \sum_{k=1}^K \frac{n_k}{c_{kj} : P} c_{kj}; \quad (3.7)$$

We denote this method using  $C_{djy}$  as  $CM^L$  and the one using soft confusion matrix  $C_{djy}^{\text{soft}}$  as  $SCM^L$ .

### 3.2 Maximum Aposteriori Approach

Additional knowledge about the distribution  $P$  can be formulated as a hyper-prior  $p(P)$ . We can then extend the proposed procedure from Section 3.1 to formulate maximum a-posteriori (MAP) estimation:

$$\begin{aligned} \hat{P}_{\text{MAP}} &= \arg \max_P p(P|n) = \arg \max_P p(P)p(n|P) \\ &= \arg \max_P \log p(P) + \arg \max_P \log p(n|P) \\ \text{s.t.: } & \forall k : P_k \geq 0; \quad \sum_{k=1}^K P_k = 1; \end{aligned} \quad (3.8)$$

where  $p(P)$  denotes a hyper-prior on  $P$  and  $\log p(n|P)$  is log-likelihood given by Equation (3.4).

Following [17] we use a symmetric Dirichlet hyper-prior  $\text{Dir}(\alpha)$ , given by Equation (2.30), favoring dense distributions  $P$  with  $\alpha > 1$ , and a sparse distribution for  $0 < \alpha < 1$ .

The solution to maximum a-posteriori can be found by projected gradient ascent, adding the hyper-prior derivative from Equation (2.31) into the  $\mathcal{L}(P)$  function in Equation (3.6).

We denote this method using  $C_{\text{dij}}$  as  $\text{CM}^M$  and the one using a soft confusion matrix  $C_{\text{dij}}^{\text{soft}}$  as  $\text{SCM}^M$ .

### 3.3 Applying Regularization

As extension to the method described in the Section 3.1, we propose to add a regularization term in the Equation (3.5):

$$\begin{aligned} \hat{P} &= \arg \max_P \mathcal{L}(P) - \lambda \|P - P_T\|_F^2 = \\ &= \arg \max_P \sum_{k=1}^K n_k \log c_{kj} : P - \lambda \|P - P_T\|_F^2 \end{aligned} \quad (3.9a)$$

$$\text{s.t.: } \sum_{k=1}^K P_k = 1; \quad \forall k : P_k \geq 0; \quad (3.9b)$$

where  $\lambda$  is a hyperparameter. The intuition behind the regularization term is to penalize too large shifts from the training priors  $p_T(Y)$ . Note the minus sign before regularization term to ensure minimization of the  $l_2$ -norm.

The task given by Equation (3.9) can be solved by projected gradient ascent

$$\hat{P}^{s+1} = \Pi(P^s + r^{\lambda}(P^s)) \quad \text{rjj } P^s \quad p_T(Y) \text{jj}_2^2 ; \quad (3.10)$$

where  $r^{\lambda}(P)$  is given by Equation (3.7) and

$$\text{rjj } P \quad p_T(Y) \text{jj}_2^2 = 2(P \quad p_T(Y)) : \quad (3.11)$$

We denote this method using  $C_{d;y}$  as  $CM^{LR}$  and the one using a soft confusion matrix  $C_{d;y}^{soft}$  as  $SCM^{LR}$ .

### 3.4 Ratio Estimation

Following the derivation of  $CM^L$  and  $SCM^L$  methods in Section 3.1, we can also estimate priors ratio  $w(Y) = p_E(Y) = p_T(Y)$ , by using joint confusion matrix  $C_{d;y}$  instead of conditional one  $C_{d;y}$

$$w = \arg \max_w l(w) = \arg \max_w \sum_{k=1}^K n_k \log c_{k;:} w \quad (3.12a)$$

$$\text{s.t.: } w_i \geq 0 \quad \sum_{i=1}^K w_i = 1; \quad w^T p_T(Y) = 1 : \quad (3.12b)$$

The solution to the problem given by Equation (3.12) can be found by projected gradient ascent

$$w^{s+1} = \Pi(w^s + r^{\lambda}(w^s)) ; \quad (3.13)$$

where  $\Pi(\cdot)$  is projection onto the priors ratio simplex (see Section 3.4.1) and

$$r^{\lambda}(w) = \sum_{k=1}^K \frac{n_k}{c_{k;:} w} c_{k;:} \quad (3.14)$$

We denote this method using  $C_{d;y}$  as  $CM^W$  and the one using a soft confusion matrix  $C_{d;y}^{soft}$  as  $SCM^W$ .

### 3. The Proposed Method

#### 3.4.1 Projection onto the Priors Ratio Simplex

The projection of vector  $y$  to the probability ratio  $w$  (in this section bold  $w$  will be used instead of  $w$  to emphasize that it is a vector) is a quadratic problem:

$$w = \arg \min_w \frac{1}{2} \|w - y\|_2^2 \quad (3.15a)$$

$$\text{s.t.}: w \geq 0; \quad w^T z = 1; \quad (3.15b)$$

where  $z_i > 0 \forall i \in \{1, \dots, K\}$ . Note that in our task for the ratio estimation (3.12) the vector  $z = \hat{p}_T(Y)$  and thus sums up to 1. However, the algorithm proposed to solve the task given by the Equation (3.15) only requires the vector  $z$  to have positive values, therefore we do not mention the condition here as it is only a special case. We denote the set of possible solutions, given by the constraints in Equation (3.15b), as priors ratio simplex. The solution to the task can be found by the Algorithm 1, which is a generalization of the algorithm for projection onto the probability simplex, described in [48]:

---

#### Algorithm 1: Projection onto the priors ratio simplex

---

Input :  $y; z \in \mathbb{R}^D$ , where  $z_i > 0 \forall i \in \{1, \dots, D\}$

Compute  $y_i^0 = \frac{y_i}{z_i} \forall i \in \{1, \dots, D\}$

Sort  $y_i^0$  into  $u : u_1 \leq u_2 \leq \dots \leq u_D$  Find  $j = \max \{1 \leq j \leq D : u_j + \frac{1}{\sum_{i=1}^j z_i^2} \leq 1\}$

Define  $w_i = \frac{1}{\sum_{i=1}^j z_i^2} \sum_{i=1}^j u_i z_i^2$

Output:  $w$ , where  $w_i = \max\{y_i + z_i; 0\}$

---

We proof correctness of the Algorithm 1 by modifying the derivation of algorithm for projection onto the probability simplex from [48]:

Proof. Lagrangian of the problem, given by the Equation 3.15 is

$$L(x; \lambda, \mu) = \frac{1}{2} \|w - y\|_2^2 - \lambda (w^T z - 1) - \mu^T w; \quad (3.16)$$

where  $\lambda$  is a lagrange multiplier for equality constraint and  $\mu$  is a lagrange multiplier for inequality constraint. Let us write KKT conditions for optimal



solution w

$$w_i = y_i + z_i \quad i = 0; \dots; K-1 \tag{3.17a}$$

$$w_i = 0; \quad i = 0; \dots; K-1 \tag{3.17b}$$

$$z_i = 0; \quad i = 0; \dots; K-1 \tag{3.17c}$$

$$w_i = 0; \quad i = 0; \dots; K-1 \tag{3.17d}$$

$$\sum_{i=1}^D w_i z_i = 1; \quad i = 1; \dots; D \tag{3.17e}$$

Next, let us define  $w_i^0 = w_i + z_i$ ,  $y_i^0 = y_i + z_i$  and  $z_i^0 = z_i$ . Note that  $z_i > 0$ , hence  $w_i^0$  and  $y_i^0$  have the same sign as  $w_i$  and  $y_i$  respectively. From condition (3.17d) we observe that if  $w_i > 0$ , then  $z_i = 0$  and  $w_i = y_i + z_i \Rightarrow w_i^0 = y_i^0$ . If  $w_i = 0$ , then  $z_i = 0$  and  $y_i + z_i = y_i < 0 \Rightarrow y_i^0 = y_i < 0$ . This observation means that smaller components in  $w^0$  correspond to smaller components in  $y^0$  and when  $w_i^0 = 0$ , then  $y_i^0 < 0$ . Without the loss of generality we can assume the components in  $w^0$  and  $y^0$  to be sorted

$$\begin{aligned} w_1^0 & \geq \dots \geq w_{\delta}^0 > w_{\delta+1}^0 > \dots > w_D^0; \\ y_1^0 & \geq \dots \geq y_{\delta}^0 > y_{\delta+1}^0 > \dots > y_D^0; \end{aligned}$$

Here  $\delta$  denotes the last index where  $w_i^0$  is greater than zero, i.e.  $w_{\delta}^0 > 0$  and  $w_{\delta+1}^0 = 0$ . From the last condition in Equation (3.17e), we get

$$1 = \sum_{i=1}^D w_i z_i = \sum_{i=1}^{\delta} w_i z_i^2 = \sum_{i=1}^{\delta} (y_i^0 + z_i^0) z_i^2$$

Now we can express as

$$= \sum_{i=1}^{\delta} \frac{1}{z_i^2} \left( 1 + \sum_{i=1}^{\delta} y_i^0 z_i^2 \right) \tag{3.18}$$

Therefore to obtain final solution w we need to know only  $z_i$ , which will allow us to compute  $w_i = y_i + z_i$ , for  $i = 1; \dots; \delta$ ; and  $w_i = 0$ , for  $i = \delta + 1; \dots; D$ . The  $z_i$  can be found by computing  $y_j^0 + \frac{1}{z_j^2} \left( 1 + \sum_{i=1}^{\delta} y_i^0 z_i^2 \right)$  for every  $j = 1; \dots; D$  and then selecting the greatest  $z_j$  for which the quantity is greater than 0. Whence in the last step of the proof we need to show that  $y_j^0 + \frac{1}{z_j^2} \left( 1 + \sum_{i=1}^{\delta} y_i^0 z_i^2 \right) > 0; \delta \leq j$  and  $y_j^0 + \frac{1}{z_j^2} \left( 1 + \sum_{i=1}^{\delta} y_i^0 z_i^2 \right) < 0; \delta > j$ :

### 3. The Proposed Method

- For  $j = 0$  :

$$y^0 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 = y^0 + \sum_{i=1}^j y_i^0 > 0$$

- For  $j < 0$  :

$$\begin{aligned} y_j^0 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A \\ &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A \\ &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A > 0 \end{aligned}$$


Note that for  $j = 0$  is  $y_j^0 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 > 0$ ;  $\delta_j < 0$ .

- For  $j > 0$  :

$$\begin{aligned} y_j^0 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A \\ &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A \\ &= \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 A < 0 \end{aligned}$$

Note that because of ordering  $y_j^0 \geq y_i^0 \geq y_i^0 \geq 0$  holds, hence  $y_j^0 + \prod_{i=1}^j \frac{1}{z_i^2} \sum_{i=1}^j y_i^0 z_i^2 < 0$ ;  $\delta_j > 0$ .

□



## Chapter 4

### Experiments

No amount of experimentation can ever prove me right; a single experiment can prove me wrong.

Albert Einstein

In this section, we compare the existing and proposed methods for prior shift adaptation on existing long-tailed versions of standard image classification datasets: the CIFAR100-LT [49], Places365-LT [50] and ImageNet-LT [50] (see Figure 4.1 for corresponding class distributions on the training sets). Unlike Cao et al. [49], our experiments require a validation set. Therefore, our training set, denoted as CIFAR100-LT, is smaller than the original CIFAR100-LT, keeping 50 samples from each class for the validation set. Using the same script<sup>1</sup> as Cao et al. [49] to sample the training set, the resulting imbalance ratio of 112.5 slightly differs from the original ratio of 100. For Places365-LT and ImageNet-LT, we use the same training and validation splits as Liu et al. [50]. Networks trained on these long-tailed datasets are then evaluated on uniformly distributed test sets (UNI). We also provide experiments in the other direction, denoted as UNI-LT, where the networks trained on the full CIFAR100 and Places365 datasets are evaluated on test sets subsampled from the full test sets following the prior distributions of CIFAR100-LT and Places365-LT.

The classifiers are evaluated on each dataset twice, without calibration and with calibration computed on the validation set. We choose BCTS [7] as the calibration method since it is to the best of our knowledge the best performing method in combination with label shift adaptation.

---

<sup>1</sup><https://github.com/kaidic/LDAM-DRW>

## 4. Experiments

Figure 4.1: Long-tailed class distributions used in the CIFAR100-LT [49], Places365-LT [50] and ImageNet-LT [50] datasets. Note that our CIFAR-100-LT slightly differs from the original CIFAR-100-LT [49], which did not have a validation set.

### 4.1 Implementation Details

#### 4.1.1 Settings for Classifier Training

In the experiments conducted on ImageNet and Places365, we used the ResNet-18 [51] classifier architecture. On ImageNet, the networks were trained from scratch for 90 epochs with the Stochastic Gradient Descent (SGD), initial learning rate set to 0.1 and decaying every 30 epochs by a factor of 10. On Places365, the network was pre-trained from an ImageNet pre-trained checkpoint for 30 epochs with the initial learning rate set to 0.01 and decaying by a factor of 10 every 10 epochs. For experiments on CIFAR100, we trained a ResNet-32 [52] adjusted to image input of 32x32. The network was trained from scratch for 200 epochs, with a learning rate of 0.1 and decaying every 80 epochs by a factor of 10. In all experiments, the batch size was set to 256. Momentum and weight decay were set to 0.9 and 0.0001 respectively.

#### 4.1.2 Implementation of Label Shift Adaptation

This section describes the setting of the hyperparameters in methods for label shift adaptation. The proposed methods (S)CM<sup>L</sup> and (S)CM<sup>W</sup> take a single hyperparameter - learning rate in gradient ascent. Since both methods have well-defined objectives, we can find the learning rate by solving the task for several different rates and then select the one maximizing objective given by the Equation (3.5) and (3.12) for (S)CM<sup>L</sup> and (S)CM<sup>W</sup> respectively. In the same way, we select the learning rate for (S)CM<sup>R</sup>, but the objective is given by the Equation (3.9). The regularization constant is selected manually as

$\epsilon = 0.001$  and  $\alpha = 1$ . Note that  $\alpha$  cannot be determined by cross-validation because in general, the validation set has different label distribution from the test set. (S)CM<sup>M</sup> also has two parameters: learning rate and  $\alpha$ -parameter of symmetric Dirichlet distribution. We search only for the learning rate by maximizing objective given by the Equation (3.8) and set  $\alpha = 3$ , following [17]. In all proposed methods the gradient ascent runs for 1000 iterations. The results for RLLS are based on the authors' code<sup>2</sup> using the cvxpy optimizer. As termination condition for the EM and MAP algorithms, we set a threshold to  $l_2$ -distance between two consecutive solutions. The threshold is set to 0.001.

## 4.2 Label Shift Adaptation

### 4.2.1 Re-weight or Retrain?

Recall from Section 2.4 that we have basically two options how to deal with the label shift if we know the test priors  $p_E(Y)$  or priors ratio  $w(Y)$ :

1. Adapting the predictions of a previously trained classifier  $f_T(x)$   $p_T(Y|x)$ , according to Equation 2.14.
2. Training a classifier  $f_E(x)$  with a sampler following the known new class priors  $p_E$ .

When we choose a method estimating the test priors  $p_E(Y)$  rather than the priors ratio, the training priors can be determined either as a proportion of class labels in the training set, i.e.  $p_T^M(Y = k) = \frac{M_k}{M}$ , or as an average of predictions  $f(x)$  on the training set,  $p_T^f(Y) = \frac{1}{M} \sum_{i=1}^M f(x_i)$ . The intuition behind the latter one is that the classifier has learned some inner trained prior, possibly different from the training set prior, due to some unknown phenomenon at the training stage.

The retraining and re-weighting strategies are experimentally compared in the Table 4.1 on datasets CIFAR100-LT, Places365-LT and ImageNet-LT. The results show that re-weighting the classifier by the ratio of test and training priors outperforms retraining in all of the cases. In most of the cases, re-weighting with the training priors estimated by averaging predictions on

<sup>2</sup><https://github.com/Angie-Liu/labelshift/tree/5bbe517938f4e3f5bd14c2c105de973dcc2e0917>

#### 4. Experiments

D	C	Standard training sampler			Sampler follows $p_E$
		NA	$\frac{p_E}{\hat{p}_T^M}$	$\frac{p_E}{\hat{p}_T^f}$	NA
CIFAR100*					
LT! UNI	7	31.66 <sup>1.27</sup>	33.99 <sup>1.41</sup>	34.06 <sup>1.35</sup>	22.44
LT! UNI	3	31.71 <sup>1.29</sup>	30.41 <sup>1.57</sup>	34.54 <sup>1.32</sup>	22.44
UNI! LT	7	63.83 <sup>0.82</sup>	69.14 <sup>0.61</sup>	69.13 <sup>0.58</sup>	67.34
UNI! LT	3	63.83 <sup>0.82</sup>	70.63 <sup>0.75</sup>	70.65 <sup>0.75</sup>	67.34
Places365					
LT! UNI	7	25.14 <sup>0.14</sup>	28.03 <sup>6.09</sup>	32.99 <sup>0.46</sup>	24.98
LT! UNI	3	25.16 <sup>0.14</sup>	27.36 <sup>6.18</sup>	33.51 <sup>0.25</sup>	25.00
ImageNet					
LT! UNI	7	34.30 <sup>0.19</sup>	37.36 <sup>0.07</sup>	37.34 <sup>0.15</sup>	30.01
LT! UNI	3	34.31 <sup>0.19</sup>	36.07 <sup>0.30</sup>	37.45 <sup>0.19</sup>	30.01

Table 4.1: Adapt or Re-Train? Accuracy ( std. dev.) of classifiers adapted to new known priors  $p_E(Y)$  with different estimates of trained priors ( $\hat{p}_T^M, \hat{p}_T^f$ ), compared to training a classifier with a sampling strategy following  $p_E(Y)$ . NA denotes no adaptation of predictions. Results of classifier adaptation on CIFAR are averaged from 10 experiments, on Places365 and ImageNet from 5 experiments respectively. Re-training the classifier with a sampler following  $p_E(Y)$  was only experimented once for each dataset. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.

a training set performs better than label proportions. Hence, we will use  $\hat{p}_T(Y) = \hat{p}_T^f(Y)$  in all of our following experiments.

#### 4.2.2 Improving Estimates from Confusion Matrix

Table 4.2 compares accuracy after adaptation with new prior estimate based on confusion matrix (CM) inversion [16] and our proposed method from Section 3.1 ( $CM^L$ ). The proposed method handles inconsistent estimates  $\hat{p}(D)$  and  $\hat{C}_{dij}$  and consistently improves the results both using the confusion matrix ( $CM^L$ ) and the soft confusion matrix ( $SCM^L$ ). In all cases, the proposed  $SCM^L$  method using soft confusion matrix achieves the best results.

## 4.2. Label Shift Adaptation

D	C	NA	CM	CM <sup>L</sup>	SCM	SCM <sup>L</sup>	Oracle
CIFAR100*							
LT! UNI	<b>7</b>	31.66 <sup>1.27</sup>	21.37 <sup>2.68</sup>	33.00 <sup>1.56</sup>	<b>26.64</b> <sup>3.85</sup>	33.47 <sup>1.33</sup>	34.06 <sup>1.35</sup>
LT! UNI	<b>3</b>	31.67 <sup>1.27</sup>	19.11 <sup>2.98</sup>	32.41 <sup>1.50</sup>	<b>26.98</b> <sup>3.76</sup>	33.42 <sup>1.51</sup>	34.40 <sup>1.38</sup>
UNI! LT	<b>7</b>	63.83 <sup>0.82</sup>	68.06 <sup>0.92</sup>	68.08 <sup>0.75</sup>	<b>68.10</b> <sup>0.81</sup>	68.24 <sup>0.75</sup>	69.13 <sup>0.58</sup>
UNI! LT	<b>3</b>	63.83 <sup>0.82</sup>	69.08 <sup>0.94</sup>	69.10 <sup>0.98</sup>	<b>69.31</b> <sup>0.94</sup>	69.40 <sup>0.77</sup>	70.65 <sup>0.75</sup>
Places365							
LT! UNI	<b>7</b>	25.14 <sup>0.14</sup>	17.45 <sup>0.30</sup>	27.77 <sup>0.45</sup>	<b>19.78</b> <sup>2.21</sup>	28.47 <sup>0.14</sup>	32.99 <sup>0.46</sup>
LT! UNI	<b>3</b>	25.14 <sup>0.14</sup>	16.24 <sup>1.39</sup>	27.69 <sup>0.51</sup>	<b>18.88</b> <sup>1.61</sup>	27.83 <sup>0.27</sup>	33.38 <sup>0.31</sup>
UNI! LT	<b>7</b>	58.17 <sup>1.01</sup>	81.16 <sup>0.61</sup>	81.64 <sup>0.63</sup>	<b>82.04</b> <sup>0.15</sup>	82.04 <sup>0.63</sup>	88.14 <sup>0.27</sup>
UNI! LT	<b>3</b>	58.17 <sup>1.01</sup>	81.20 <sup>0.61</sup>	81.65 <sup>0.61</sup>	<b>82.04</b> <sup>0.15</sup>	82.07 <sup>0.66</sup>	88.15 <sup>0.30</sup>
ImageNet							
LT! UNI	<b>7</b>	34.30 <sup>0.19</sup>	19.02 <sup>0.26</sup>	33.57 <sup>0.33</sup>	<b>23.94</b> <sup>2.04</sup>	35.91 <sup>0.20</sup>	37.34 <sup>0.15</sup>
LT! UNI	<b>3</b>	34.30 <sup>0.19</sup>	17.28 <sup>0.48</sup>	32.34 <sup>0.41</sup>	<b>24.78</b> <sup>3.06</sup>	35.86 <sup>0.17</sup>	37.39 <sup>0.16</sup>

Table 4.2: Improve Estimates from Confusion Matrix. ' ' Accuracy ( std. dev.) after adaptation with new prior estimate based on confusion matrix (CM) inversion [16] and our proposed method from Section 3.1 (CM) and Section 3.4 (CM<sup>W</sup>). SCM denotes soft confusion matrix, NA denotes no adaptation, Oracle is an adaptation with ground truth priors. Results on CIFAR are averaged from 10 experiments, results on Places and ImageNet are averaged from 5 experiments. The best results are displayed in bold. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.

### 4.2.3 Comparing Regularization

Comparison of (S)CM<sup>L</sup> method with its regularized version (S)CM<sup>L-R</sup> is in Table 4.3. We show the results for (S)CM<sup>R</sup> with two different regularization constants  $\lambda = 0.001$  and  $\lambda = 1$ . For  $\lambda = 0.001$  the effect of regularization is very small and the results are very close to (S)CM<sup>L</sup>. With  $\lambda = 1$  the results start to be deteriorated in the case of UNI! LT shift. We evaluated the regularization also for bigger regularization constants and with increasing the results are getting even more deteriorated. This is due to the stronger effect of regularization, where the estimated test prior gets closer to the training prior. Note that the setting of our experiments disadvantage the methods with regularization because the purpose of regularization is to make the estimation of the test prior more stable when the label shift is small, while the shifts between the training and test distributions in our experiments are large. In several cases (S)CM<sup>R</sup> outperforms (S)CM<sup>L</sup> by a small margin. We conjecture this is because the likelihood close to the optimum doesn't have to correspond to the best accuracy due to an inaccurate estimate of the

#### 4. Experiments

D	C	NA	CM <sup>L</sup>	CM <sup>LR</sup> <sub>=0:001</sub>	CM <sup>LR</sup> <sub>=1</sub>	SCM <sup>L</sup>	SCM <sup>LR</sup> <sub>=0:001</sub>	SCM <sup>LR</sup> <sub>=1</sub>	Oracle
CIFAR100*									
LT! UNI	<b>7</b>	31.66 <sup>1.27</sup>	33.00 <sup>1.56</sup>	33.06 <sup>1.52</sup>	33.14 <sup>1.50</sup>	33.47 <sup>1.33</sup>	33.48 <sup>1.32</sup>	33.34 <sup>1.41</sup>	34.06 <sup>1.35</sup>
LT! UNI	<b>3</b>	31.67 <sup>1.27</sup>	32.41 <sup>1.50</sup>	32.42 <sup>1.50</sup>	33.34 <sup>1.47</sup>	33.42 <sup>1.51</sup>	33.38 <sup>1.48</sup>	33.48 <sup>1.50</sup>	34.40 <sup>1.38</sup>
UNI! LT	<b>7</b>	63.83 <sup>0.82</sup>	68.08 <sup>0.75</sup>	68.04 <sup>0.76</sup>	67.86 <sup>0.85</sup>	68.24 <sup>0.75</sup>	68.24 <sup>0.72</sup>	67.91 <sup>0.69</sup>	69.13 <sup>0.58</sup>
UNI! LT	<b>3</b>	63.83 <sup>0.82</sup>	69.10 <sup>0.98</sup>	69.07 <sup>1.00</sup>	69.02 <sup>0.85</sup>	69.40 <sup>0.77</sup>	69.39 <sup>0.79</sup>	69.19 <sup>0.80</sup>	70.65 <sup>0.75</sup>
Places365									
LT! UNI	<b>7</b>	25.14 <sup>0.14</sup>	27.77 <sup>0.45</sup>	27.77 <sup>0.45</sup>	27.69 <sup>0.46</sup>	28.47 <sup>0.14</sup>	28.54 <sup>0.10</sup>	28.55 <sup>0.15</sup>	32.99 <sup>0.46</sup>
LT! UNI	<b>3</b>	25.14 <sup>0.14</sup>	27.69 <sup>0.51</sup>	27.69 <sup>0.51</sup>	27.69 <sup>0.50</sup>	27.83 <sup>0.27</sup>	27.85 <sup>0.26</sup>	27.82 <sup>0.25</sup>	33.38 <sup>0.31</sup>
UNI! LT	<b>7</b>	58.17 <sup>1.01</sup>	81.64 <sup>0.63</sup>	81.67 <sup>0.66</sup>	80.29 <sup>0.67</sup>	82.04 <sup>0.63</sup>	82.09 <sup>0.64</sup>	80.55 <sup>1.15</sup>	88.14 <sup>0.27</sup>
UNI! LT	<b>3</b>	58.17 <sup>1.01</sup>	81.65 <sup>0.61</sup>	81.72 <sup>0.67</sup>	80.30 <sup>0.69</sup>	82.07 <sup>0.66</sup>	82.10 <sup>0.64</sup>	80.53 <sup>1.12</sup>	88.15 <sup>0.30</sup>
ImageNet									
LT! UNI	<b>7</b>	34.30 <sup>0.19</sup>	33.57 <sup>0.33</sup>	33.57 <sup>0.33</sup>	33.89 <sup>0.31</sup>	35.91 <sup>0.20</sup>	35.78 <sup>0.17</sup>	35.75 <sup>0.26</sup>	37.34 <sup>0.15</sup>
LT! UNI	<b>3</b>	34.30 <sup>0.19</sup>	32.34 <sup>0.41</sup>	32.34 <sup>0.41</sup>	32.74 <sup>0.42</sup>	35.86 <sup>0.17</sup>	35.59 <sup>0.14</sup>	35.63 <sup>0.15</sup>	37.39 <sup>0.16</sup>

Table 4.3: Compare Regularization.' Accuracy ( std. dev.) after adaptation by proposed method from Section 3.1 ((S)CM) and method proposed in Section 3.3 ((S)CM<sup>LR</sup>) making use of regularization. The (S)CM<sup>LR</sup> methods is evaluated with two di erent regularization constants = 0:001 and = 1. NA denotes no adaptation, Oracle is adaptation with ground truth priors. Results on CIFAR are averaged from 10 experiments, results on Places and ImageNet are averaged from 5 experiments. The best results are displayed in bold. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.

D	C	NA	CM <sup>LR</sup> <sub>=0:001</sub>	CM <sup>LR</sup> <sub>=1</sub>	SCM <sup>LR</sup> <sub>=0:001</sub>	SCM <sup>LR</sup> <sub>=1</sub>	RLLS	Oracle
CIFAR100*								
LT! UNI	<b>7</b>	31.66 <sup>1.27</sup>	33.06 <sup>1.52</sup>	33.14 <sup>1.50</sup>	33.48 <sup>1.32</sup>	33.34 <sup>1.41</sup>	32.75 <sup>1.40</sup>	34.06 <sup>1.35</sup>
LT! UNI	<b>3</b>	31.67 <sup>1.27</sup>	32.42 <sup>1.50</sup>	33.34 <sup>1.47</sup>	33.38 <sup>1.48</sup>	33.48 <sup>1.50</sup>	32.62 <sup>1.46</sup>	34.40 <sup>1.38</sup>
UNI! LT	<b>7</b>	63.83 <sup>0.82</sup>	68.04 <sup>0.76</sup>	67.86 <sup>0.85</sup>	68.24 <sup>0.72</sup>	67.91 <sup>0.69</sup>	68.02 <sup>0.77</sup>	69.13 <sup>0.58</sup>
UNI! LT	<b>3</b>	63.83 <sup>0.82</sup>	69.07 <sup>1.00</sup>	69.02 <sup>0.85</sup>	69.39 <sup>0.79</sup>	69.19 <sup>0.80</sup>	69.05 <sup>0.97</sup>	70.65 <sup>0.75</sup>
Places365								
LT! UNI	<b>7</b>	25.14 <sup>0.14</sup>	27.77 <sup>0.45</sup>	27.69 <sup>0.46</sup>	28.54 <sup>0.10</sup>	28.55 <sup>0.15</sup>	26.94 <sup>0.41</sup>	32.99 <sup>0.46</sup>
LT! UNI	<b>3</b>	25.14 <sup>0.14</sup>	27.69 <sup>0.51</sup>	27.69 <sup>0.50</sup>	27.85 <sup>0.26</sup>	27.82 <sup>0.25</sup>	27.03 <sup>0.40</sup>	33.38 <sup>0.31</sup>
UNI! LT	<b>7</b>	58.17 <sup>1.01</sup>	81.67 <sup>0.66</sup>	80.29 <sup>0.67</sup>	82.09 <sup>0.64</sup>	80.55 <sup>1.15</sup>	82.04 <sup>0.69</sup>	88.14 <sup>0.27</sup>
UNI! LT	<b>3</b>	58.17 <sup>1.01</sup>	81.72 <sup>0.67</sup>	80.30 <sup>0.69</sup>	82.10 <sup>0.64</sup>	80.53 <sup>1.12</sup>	82.04 <sup>0.69</sup>	88.15 <sup>0.30</sup>
ImageNet								
LT! UNI	<b>7</b>	34.30 <sup>0.19</sup>	33.57 <sup>0.33</sup>	33.89 <sup>0.31</sup>	35.78 <sup>0.17</sup>	35.75 <sup>0.26</sup>	34.69 <sup>0.14</sup>	37.34 <sup>0.15</sup>
LT! UNI	<b>3</b>	34.30 <sup>0.19</sup>	32.34 <sup>0.41</sup>	32.74 <sup>0.42</sup>	35.59 <sup>0.14</sup>	35.63 <sup>0.15</sup>	34.31 <sup>0.08</sup>	37.39 <sup>0.16</sup>

Table 4.4: Compare regularization.' Accuracy ( std. dev.) after adaptation with methods based on regularization, i.e. RLLS [19] and method proposed in Section 3.3 (S)CM<sup>LR</sup>. NA denotes no adaptation, Oracle is adaptation with ground truth priors. The (S)CM<sup>LR</sup> methods is evaluated with two di erent regularization constants = 0:001 and = 1. Results on CIFAR are averaged from 10 experiments, results on Places and ImageNet are averaged from 5 experiments. The best results are displayed in bold. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.



confusion matrix.

The methods leveraging regularization to prevent large shifts from training prior  $p_T(Y)$  are compared in Table 4.4. Concretely, we compare RLLS [19], which is estimating priors ratio  $w(Y)$  by minimizing  $l_2$ -norm with regularization term penalizing large shifts, against our method proposed in Section 3.3 with two different regularization constants  $\lambda = 0:001$  and  $\lambda = 1$ . The best performing method is  $SCM^R$  with  $\lambda = 0:001$ .

#### 4.2.4 Comparing Ratio Estimation

Having an estimate of the trained priors, Table 4.5 compares prior ratio estimation with BBSE, BBSE-S [20] and RLLS [19] against the best performing prior estimation method,  $SCM^L$ , and our methods ( $CM^W$  and  $SCM^W$ ) proposed to handle the negative weights in the estimates from BBSE and BBSE-S. The improvement of our method (S) $CM^W$  compared to BBSE and BBSE-S is apparent on datasets CIFAR100 and ImageNet, while on Places365 the result is deteriorated. The results indicate that it is better to estimate the new priors with  $SCM^L$  than to directly estimate the prior ratio.

#### 4.2.5 Methods for MLE and MAP Prior Estimation

Existing methods for maximum likelihood and maximum a-posteriori prior estimation are compared against the methods proposed in Sections 3.1 and 3.2 respectively in Table 4.6. Note that the methods maximize a different likelihood function: The EM algorithm of Saerens et al. [16] maximizes the likelihood of observed classifier outputs  $f(x_i)$ , while the proposed methods based on the confusion matrix ( $CM^L$ ) and the soft confusion matrix ( $SCM^L$ ) maximize the likelihood of classifiers decisions. The same difference in likelihood functions holds for the MAP approach of Sulc and Matas [17] and MAP estimate proposed in Section 3.2, but we use the same hyper-prior on  $p_E(Y)$  for all methods:  $Dir(\alpha = 3)$ . Although the (S) $CM^W$  is also based on maximum likelihood estimation, we do not compare this method here since it is outperformed by (S) $CM^L$  as experimentally shown in Section 4.2.4.

From the maximum likelihood estimators, the proposed  $SCM^L$  achieves the best results in most cases, except for Places365 "UNILT", where the EM algorithm performed slightly better. Similarly, the Maximum A-Posteriori version of the proposed method,  $SCM^M$  performs better than the existing MAP estimate [17] in most cases. As expected, the MAP estimation improves

#### 4. Experiments

D	C	NA	SCM <sup>L</sup>	CM <sup>W</sup>	SCM <sup>W</sup>	RLLS	BBSE	BBSE-S	Oracle
CIFAR100*									
LT! UNI	<b>7</b>	31.66 <sup>1.27</sup>	33.47 <sup>1.33</sup>	<u>32.77</u> <sup>1.40</sup>	32.75 <sup>1.28</sup>	32.75 <sup>1.40</sup>	31.28 <sup>1.58</sup>	31.92 <sup>1.62</sup>	34.06 <sup>1.35</sup>
LT! UNI	<b>3</b>	31.67 <sup>1.27</sup>	33.42 <sup>1.51</sup>	<u>33.11</u> <sup>1.29</sup>	32.98 <sup>1.35</sup>	32.62 <sup>1.46</sup>	26.47 <sup>1.87</sup>	29.06 <sup>2.65</sup>	34.40 <sup>1.38</sup>
UNI! LT	<b>7</b>	63.83 <sup>0.82</sup>	68.24 <sup>0.75</sup>	68.04 <sup>0.76</sup>	<u>68.24</u> <sup>0.72</sup>	68.02 <sup>0.77</sup>	67.95 <sup>0.96</sup>	68.12 <sup>0.90</sup>	69.13 <sup>0.58</sup>
UNI! LT	<b>3</b>	63.83 <sup>0.82</sup>	69.40 <sup>0.77</sup>	69.08 <sup>0.97</sup>	69.41 <sup>0.80</sup>	69.05 <sup>0.97</sup>	69.30 <sup>0.99</sup>	<u>69.51</u> <sup>0.98</sup>	70.65 <sup>0.75</sup>
Places365									
LT! UNI	<b>7</b>	25.14 <sup>0.14</sup>	28.47 <sup>0.14</sup>	23.14 <sup>2.72</sup>	19.75 <sup>2.99</sup>	<u>26.94</u> <sup>0.41</sup>	24.79 <sup>0.74</sup>	25.55 <sup>0.69</sup>	32.99 <sup>0.46</sup>
LT! UNI	<b>3</b>	25.14 <sup>0.14</sup>	27.83 <sup>0.27</sup>	23.16 <sup>3.59</sup>	18.80 <sup>1.32</sup>	<u>27.03</u> <sup>0.40</sup>	23.12 <sup>0.79</sup>	23.68 <sup>0.75</sup>	33.38 <sup>0.31</sup>
UNI! LT	<b>7</b>	58.17 <sup>1.01</sup>	82.04 <sup>0.63</sup>	81.49 <sup>0.51</sup>	81.57 <sup>0.32</sup>	<u>82.04</u> <sup>0.69</sup>	80.66 <sup>0.57</sup>	81.69 <sup>0.20</sup>	88.14 <sup>0.27</sup>
UNI! LT	<b>3</b>	58.17 <sup>1.01</sup>	82.07 <sup>0.66</sup>	81.52 <sup>0.57</sup>	81.64 <sup>0.34</sup>	<u>82.04</u> <sup>0.69</sup>	80.71 <sup>0.56</sup>	81.69 <sup>0.20</sup>	88.15 <sup>0.30</sup>
ImageNet									
LT! UNI	<b>7</b>	34.30 <sup>0.19</sup>	35.91 <sup>0.20</sup>	35.21 <sup>0.13</sup>	<u>35.28</u> <sup>0.15</sup>	34.69 <sup>0.14</sup>	30.77 <sup>0.31</sup>	31.31 <sup>0.90</sup>	37.34 <sup>0.15</sup>
LT! UNI	<b>3</b>	34.30 <sup>0.19</sup>	35.86 <sup>0.17</sup>	34.75 <sup>0.06</sup>	<u>35.00</u> <sup>0.17</sup>	34.31 <sup>0.08</sup>	26.89 <sup>0.49</sup>	28.05 <sup>1.69</sup>	37.39 <sup>0.16</sup>

Table 4.5: Estimate test priors or directly the prior ratio? Accuracy ( std. dev.) after adaptation with the priors estimated by SCM<sup>L</sup> or with the prior ratio estimated by BBSE [20] (without re-training), RLLS [ 19] (without re-training), CM<sup>W</sup> and SCM<sup>W</sup>. Results on CIFAR are averaged from 10 experiments, results on Places and ImageNet are averaged from 5 experiments. The best results among all methods are displayed in bold. The best results among methods estimating priors ratio are underlined. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.

D	C	NA	CM <sup>L</sup>	SCM <sup>L</sup>	EM	CM <sup>M</sup>	SCM <sup>M</sup>	MAP	Oracle
CIFAR100*									
LT! UNI	<b>7</b>	31.66 <sup>1.27</sup>	33.00 <sup>1.56</sup>	33.47 <sup>1.33</sup>	32.81 <sup>1.41</sup>	33.49 <sup>1.45</sup>	33.50 <sup>1.40</sup>	32.73 <sup>1.42</sup>	34.06 <sup>1.35</sup>
LT! UNI	<b>3</b>	31.67 <sup>1.27</sup>	32.41 <sup>1.50</sup>	33.42 <sup>1.51</sup>	29.43 <sup>1.59</sup>	33.99 <sup>1.43</sup>	34.13 <sup>1.53</sup>	24.46 <sup>12.43</sup>	34.40 <sup>1.38</sup>
UNI! LT	<b>7</b>	63.83 <sup>0.82</sup>	68.08 <sup>0.75</sup>	68.24 <sup>0.75</sup>	67.23 <sup>0.88</sup>	67.01 <sup>0.91</sup>	67.00 <sup>0.87</sup>	66.72 <sup>0.91</sup>	69.13 <sup>0.58</sup>
UNI! LT	<b>3</b>	63.83 <sup>0.82</sup>	69.10 <sup>0.98</sup>	69.40 <sup>0.77</sup>	69.17 <sup>0.91</sup>	68.42 <sup>0.84</sup>	68.38 <sup>0.73</sup>	68.30 <sup>0.77</sup>	70.65 <sup>0.75</sup>
Places365									
LT! UNI	<b>7</b>	25.14 <sup>0.14</sup>	27.77 <sup>0.45</sup>	28.47 <sup>0.14</sup>	28.02 <sup>0.92</sup>	28.02 <sup>0.24</sup>	27.68 <sup>0.13</sup>	25.22 <sup>0.13</sup>	32.99 <sup>0.46</sup>
LT! UNI	<b>3</b>	25.14 <sup>0.14</sup>	27.69 <sup>0.51</sup>	27.83 <sup>0.27</sup>	28.09 <sup>1.32</sup>	27.92 <sup>0.24</sup>	27.41 <sup>0.15</sup>	28.57 <sup>0.24</sup>	33.38 <sup>0.31</sup>
UNI! LT	<b>7</b>	58.17 <sup>1.01</sup>	81.64 <sup>0.63</sup>	82.04 <sup>0.63</sup>	82.63 <sup>0.31</sup>	76.13 <sup>0.56</sup>	73.27 <sup>0.46</sup>	76.97 <sup>0.45</sup>	88.14 <sup>0.27</sup>
UNI! LT	<b>3</b>	58.17 <sup>1.01</sup>	81.65 <sup>0.61</sup>	82.07 <sup>0.66</sup>	82.63 <sup>0.26</sup>	76.16 <sup>0.54</sup>	73.30 <sup>0.46</sup>	77.00 <sup>0.41</sup>	88.15 <sup>0.30</sup>
ImageNet									
LT! UNI	<b>7</b>	34.30 <sup>0.19</sup>	33.57 <sup>0.33</sup>	35.91 <sup>0.20</sup>	34.63 <sup>0.29</sup>	36.41 <sup>0.17</sup>	36.57 <sup>0.16</sup>	34.64 <sup>0.20</sup>	37.34 <sup>0.15</sup>
LT! UNI	<b>3</b>	34.30 <sup>0.19</sup>	32.34 <sup>0.41</sup>	35.86 <sup>0.17</sup>	27.26 <sup>2.25</sup>	36.18 <sup>0.12</sup>	36.80 <sup>0.14</sup>	20.65 <sup>18.77</sup>	37.39 <sup>0.16</sup>

Table 4.6: How to estimate new priors? Accuracy ( std. dev.) after adaptation to new priors estimated with different Maximum Likelihood and Maximum A Posteriori estimates. NA denotes no adaptation, Oracle is adaptation with ground truth priors. Results on CIFAR are averaged from 10 experiments, results on Places and ImageNet are averaged from 5 experiments. The best results are displayed in bold. The column denoted as D determines the training and the test distribution for a particular row, while the column denoted as C denotes if BCTS [47] calibration was applied.

upon MLE on the dense test distributions, favoured by the Dirichlet hyper-prior.

#### 4.2.6 Dependence on the Number of Test Samples

Figure 4.2 displays the accuracy on the uniformly distributed sets after the adaptation of classifiers trained on the CIFAR100-LT and Places365-LT datasets with different prior estimation methods, as a function of the number of test examples used for prior estimation. While the proposed SCM method achieves slightly higher accuracy with more samples, the EM algorithm works slightly better with a low number of samples. With an extremely low number of samples, the prior estimation should better be omitted. For the sake of clarity, we exclude the other proposed methods ((S)CM<sup>R</sup>, (S)CM<sup>W</sup>, (S)CM<sup>M</sup>) from the comparison since they are just extensions to the (S)CM method or perform worse compared to the (S)CM.

### 4.3 Confusion Matrices Illustrated on an Artificial Dataset

Let us consider an illustrative classification problem with two classes  $\{0, 1\}$ , generated from known normal distributions:  $p(x|Y = 0) = \mathcal{N}(-2; 2)$  and  $p(x|Y = 1) = \mathcal{N}(2; 2)$  with equal priors,  $p(Y = 0) = p(Y = 1) = 0.5$ . We use 3 different classifiers modeled by logistic regression in the form

$$f(x) = \frac{1}{1 + e^{-(ax+b)}} \quad (4.1)$$

The first classifier,  $f_1(x)$ , was trained with scikit-learn on 4 randomly generated samples. The other two are Bayes classifiers:  $f_c(x)$  is perfectly calibrated with parameters  $a = 1$  and  $b = 0$ ;  $f_o(x)$  is overconfident with parameters  $a = 2$  and  $b = 0$ . All three classifiers are illustrated together with their decision thresholds and the known posterior probabilities in Figure 4.3.

Following the Section 2.4.1, we denote  $\hat{C}_{dij}$  the confusion matrix estimated from top-1 predictions and  $\hat{C}_{dij}^{\text{soft}}$  the confusion matrix estimated using the softmax outputs, following the Equation (2.19). In this artificial example with known distributions, we can compute the true confusion matrix  $C_{dij}$  as follows

## 4. Experiments

Figure 4.2: How many samples do I need? Accuracy after adapting CIFAR100-LT! UNI (left) and Places365-LT! UNI (right) using #samples for prior estimation.

$$C_{1jj} = \int_1^{Z_1} p(x|Y = j) dx; \quad C_{0jj} = \int_1^{Z_t} p(x|Y = j) dx: \quad (4.2)$$

Figure 4.4 compares the distance of the two estimated confusion matrices from the true confusion matrix, depending on the size of the validation set. Note that the soft confusion matrix may not converge to the true confusion matrix even with a perfectly calibrated classifier, but it provides a better estimate in low-sample scenarios.

4.3. Confusion Matrices Illustrated on an Artificial Dataset

Figure 4.3: An illustrative 2-class example with known class posteriors (solid curves), outputs of 3 classifiers for given  $x$  (dashed lines) and their decision thresholds (vertical lines). Note that  $f_c(x)$  and  $f_o(x)$  are optimal Bayes classifiers minimizing the 0/1 loss.

#### 4. Experiments

(a) : Classifier  $f_t(x)$  trained on 4 samples:      (b) : Overconfident classifier  $f_o(x)$ :

"	# "	# "	# "	# "	# "	# "	# "
0:94 0:30	0:785 0:396	0:938 0:321	0:82 0:16	0:798 0:174	0:841 0:159	0:202 0:826	0:159 0:841
0:06 0:70	0:215 0:604	0:062 0:679	0:18 0:84	0:202 0:826	0:159 0:841	0:202 0:826	0:159 0:841
$\hat{C}_{d_{jy}}$	$\hat{C}_{d_{jy}}^{\text{soft}}$	$C_{d_{jy}}$	$\hat{C}_{d_{jy}}$	$\hat{C}_{d_{jy}}^{\text{soft}}$	$C_{d_{jy}}$	$\hat{C}_{d_{jy}}^{\text{soft}}$	$C_{d_{jy}}$

(c) : Calibrated classifier  $f_c(x)$ :

"	# "	# "	# "
0:82 0:16	0:764 0:215	0:841 0:159	0:159 0:841
0:18 0:84	0:236 0:785	0:159 0:841	0:159 0:841
$\hat{C}_{d_{jy}}$	$\hat{C}_{d_{jy}}^{\text{soft}}$	$C_{d_{jy}}$	$C_{d_{jy}}$

Figure 4.4: Top: The distance (sum of absolute differences) of estimated confusion matrix  $\hat{C}_{d_{jy}}$  and soft confusion matrix  $\hat{C}_{d_{jy}}^{\text{soft}}$  from the true confusion matrix  $C_{d_{jy}}$ , depending on the number of validation samples. The distance values are averaged over 50 trials. Bottom: Confusion matrices  $\hat{C}_{d_{jy}}$ ,  $\hat{C}_{d_{jy}}^{\text{soft}}$  estimated from 50 samples and the true confusion matrix  $C_{d_{jy}}$  computed from Equation (4.2).

## 4.4 Convergence Speed

The convergence of all algorithms on the test set of Places365-LT is displayed in the Figure 4.5. For the sake of clarity, we exclude  $(S)CM^R$ ,  $(S)CM^W$  methods from comparison in the Figure 4.5 since they are just extensions to the  $(S)CM^L$  method or perform worse compared to the  $(S)CM$ . Note that the optimization code for our methods is experimental and not optimized for run time. The results for RLLS are based on the authors' code<sup>3</sup>, using the cvxpy optimizer. Table 4.7 compares the runtimes of all algorithms on Cifar100-LT, Places365-LT and ImageNet-LT, using the termination conditions used in all our experiments and described in Section 4.1.2. While the proposed optimization in  $(S)CM^L$  and  $(S)CM^M$  takes longer than the baseline  $(S)CM$ , all methods converge within 0.5 seconds on Places365-LT. Even though the predictions are adapted on CPU and the classifier evaluation is computed on GPU, the time to adapt classifier to label shift is negligible compared to the time it takes to evaluate the classifier on the test set, which takes several minutes on ImageNet-LT.

Figure 4.5: Convergence of prior estimation on Places365-LT. The y-axis shows the classification accuracy after prior shift adaptation. Bold lines depict a particular method before the termination criterion is met. Dashed lines show the constant accuracy after the computation is finished. Note that  $(S)CM$ , BBSE and BBSE-S are single-step methods. Since RLLS uses a third party optimizer, we only report the total optimization time.

<sup>3</sup><https://github.com/Angie-Liu/labelshift/tree/5bbe517938f4e3f5bd14c2c105de973dcc2e0917>

#### 4. Experiments

	Cifar100-LT	Places365-LT	ImageNet-LT
CM	0.0038	0.0697	0.3486
SCM	0.0034	0.0555	0.2946
BBSE	0.0089	0.0678	0.3357
BBSE-S	0.0033	0.0506	0.2927
EM	0.0068	0.8293	2.2329
CM-L	0.2001	0.3104	1.3035
SCM-L	0.1828	0.2955	1.3595
CM-LR	0.2021	0.3349	1.3127
SCM-LR	0.1975	0.3197	1.3151
CM-W	0.1735	0.2985	1.2144
SCM-W	0.1665	0.2847	1.2456
MAP	0.0041	0.0916	0.3388
CM-M	0.2061	0.3355	1.3943
SCM-M	0.1998	0.3114	1.3510
RLLS	0.0268	0.1946	2.1280

Table 4.7: Run time (in seconds) of methods estimating test time prior  $p_E(Y)$  or prior ratio  $p_E(Y)=p_T(Y)$ . The run time was measured on a laptop with Intel® Core™ i7-6700HQ CPU @ 2.60GHz 8 and averaged over 10 runs.





## Chapter 5

### Conclusion

If we knew what it was we were doing, it would not be called research, would it?

Albert Einstein

This thesis reviews and compares the existing methods for adapting classifiers to label shift and proposes new methods based on the confusion matrix, likelihood maximization and maximum a-posteriori estimation. The proposed methods handle a known problem [1, 16, 18, 20] of inconsistent estimates of decision probabilities and of the confusion matrix, which can result in negative values in the estimated prior probabilities or in the estimated prior ratio. The (S)CM<sup>L</sup> method for test set prior estimation deals with the problem by maximization of the likelihood of classifier's decisions on the new unlabeled test set. The next method proposed in this thesis, (S)CM<sup>M</sup>, extends (S)CM<sup>L</sup> by adding Dirichlet hyper-prior on the test set label distribution, leading to a maximum a-posteriori formulation. The third method proposed in the thesis is (S)CM<sup>R</sup>, which extends (S)CM<sup>L</sup> by adding a regularization term to penalize large shifts. The last proposed method, (S)CM<sup>W</sup>, is estimating the prior ratio instead of the test set prior. As part of the (S)CM<sup>W</sup> method, this thesis proposed an algorithm for projection onto the prior ratio simplex by generalizing an algorithm for projection onto the probability simplex described by Wang et al. [48].

From the experimental analysis of the existing and the proposed methods for prior shift adaptation we can observe and conclude the following:

- Adapting classifier predictions typically performs better than re-training the classifier with sampling matching the shift, and it is significantly

## 5. Conclusion

computationally cheaper.

- The proposed (S)CM<sup>L</sup> method handles inconsistent estimates of decision probabilities  $p(D)$  and confusion matrix  $\hat{C}_{d|y}$ , and it consistently improves the results compared to the (S)CM baseline.
- From the compared maximum likelihood estimators, the proposed SCM achieves the best results in most cases.
- Extending (S)CM<sup>L</sup> with regularization, denoted (S)CM<sup>LR</sup>, does not lead to any advantage on datasets with large shifts. Investigation of the effects on datasets with smaller shifts is left for future work.
- The proposed Maximum A-Posteriori approach, SCM<sup>M</sup>, performs better than the existing MAP estimator [17].
- It is better to estimate the training and test priors separately than to directly estimate their ratio with BBSE, RLLS or the proposed (S)CM<sup>W</sup> method.
- Among the methods estimating priors ratio, the (S)CM<sup>W</sup> outperforms all other methods on CIFAR and ImageNet, while on Places365 the best result is achieved by RLLS.
- In a small sample scenario, the EM algorithm [6, 47] and BBSE-S perform the best among the compared methods. With an extremely low number of samples, prior estimation should better be omitted at all.
- Soft confusion matrix is a biased estimate of the classifier's confusion matrix, but for a small number of samples in the validation set it gives a better estimate than the confusion matrix computed by counting decisions  $\arg \max_k f(x)$ .
- All compared methods perform nearly real-time: The time to adapt classifier's predictions to label shift is negligible compared to the time it takes to make the predictions on the test set.
- The (S)CM baselines are faster than the proposed methods, since they compute matrix inversion in a single step, while the proposed methods are iterative. EM algorithm has very fast convergence on a dataset with a small number of classes (CIFAR100-LT), where it is also faster than the proposed methods. On the other hand, on datasets with a higher number of classes (Places365-LT and ImageNet-LT) the convergence of EM is about twice slower than the convergence of the proposed methods.

- Prior shift adaptation relies on a well-calibrated classifier, assumed in Eq. (2.14). In [47], BCTS improves prior shift adaptation of classifiers trained on uniform distribution, similarly to our UNI ! LT experiments. With class-specific parameters, BCTS may overfit to errors on the validation set. For classifiers trained on LT datasets, we show BCTS is not a reliable calibration method, as it often decreases the final recognition accuracy.

**Future Work.** This thesis reviewed previous work on classifier calibration, but the effects of calibration on adaptation to label shift were examined only briefly. Deeper understanding of the mutual influence between calibration and label shift adaptation will be needed for the future work to improve the accuracy of calibrated classifiers after label shift.

Investigation into how regularization affects the proposed (S)CM<sup>L</sup> method in small shift scenarios is left for future work, which will evaluate the methods on a wider range of dataset shifts. The assumption is that regularization will help when employed on smaller dataset shifts, as it prevents large deviations from the original distribution. Inspired by Azizzadenesheli et al. [19], the future work can attempt to find theoretical guarantees for our methods, which could give us better insight into how to set up the regularization constant in (S)CM<sup>LR</sup> or the hyperparameter of Dirichlet hyperprior in (S)CM<sup>M</sup> properly.

All methods in the thesis were solely focused on the classification task. Adaptation to dataset shifts in other tasks, such as regression, object detection or semantic segmentation, remains an interesting topic for future work.





## Bibliography

- [1] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach . Prentice Hall, 3 ed., 2010.
- [2] M. Mohri, A. Rostamizadeh, and A. Talwalkar, Foundations of Machine Learning. Adaptive Computation and Machine Learning, Cambridge, MA: MIT Press, 2 ed., 2018.
- [3] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics) Berlin, Heidelberg: Springer-Verlag, 2006.
- [4] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning MIT Press, 2016. <http://www.deeplearningbook.org> .
- [5] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE, vol. 86, no. 11, pp. 2278-2324, 1998.
- [6] H. Salehinejad, S. Sankar, J. Barfett, E. Colak, and S. Valaee, Recent advances in recurrent neural networks, 2018.
- [7] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, Graph neural networks: A review of methods and applications, 2021.
- [8] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, A comprehensive survey on transfer learning, 2020.
- [9] W. M. Kouw and M. Loog, An introduction to domain adaptation and transfer learning, 2019.

## 5. Conclusion

- [10] B. Zadrozny, Learning and evaluating classifiers under sample selection bias, in ICML, 2004.
- [11] M. Sugiyama et al., Direct importance estimation with model selection and its application to covariate shift adaptation, in NIPS, 2008.
- [12] S. Bickel, M. Brückner, and T. Scheffer, Discriminative learning under covariate shift, *Journal of Machine Learning Research*, vol. 10, no. 75, pp. 2137–2155, 2009.
- [13] J. Huang, A. Gretton, K. Borgwardt, B. Schölkopf, and A. Smola, Correcting sample selection bias by unlabeled data, in *Advances in Neural Information Processing Systems* (B. Schölkopf, J. Platt, and T. Hofmann, eds.), vol. 19, MIT Press, 2007.
- [14] K. Zhang et al., Domain adaptation under target and conditional shift, in ICML, 2013.
- [15] M. C. du Plessis and M. Sugiyama, Semi-supervised learning of class balance under class-prior change by distribution matching, *CoRR*, vol. abs/1206.4677, 2012.
- [16] M. Saerens, P. Latinne, and C. Decaestecker, Adjusting the outputs of a classifier to new a priori probabilities: A simple procedure, *Neural Comput.*, vol. 14, p. 2141, Jan. 2002.
- [17] M. Sulc and J. Matas, Improving cnn classifiers by estimating test-time priors, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [18] S. Vucetic and Z. Obradovic, Classification on data with biased class distribution, in *European Conference on Machine Learning* pp. 527–538, Springer, 2001.
- [19] K. Azizzadenesheli, A. Liu, F. Yang, and A. Anandkumar, Regularized learning for domain adaptation under label shifts., in ICLR, 2019.
- [20] Z. C. Lipton, Y.-X. Wang, and A. Smola, Detecting and correcting for label shift with black box predictors, 2018.
- [21] G. Forman, Quantifying counts and costs via classification, *Data Mining and Knowledge Discovery* vol. 17, pp. 164–206, Oct 2008.
- [22] K. Simonyan and A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2015.

- [23] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, Going deeper with convolutions, 2014.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, 2015.
- [25] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, Aggregated residual transformations for deep neural networks, 2017.
- [26] H. Zhang, C. Wu, Z. Zhang, Y. Zhu, H. Lin, Z. Zhang, Y. Sun, T. He, J. Mueller, R. Manmatha, M. Li, and A. Smola, Resnet: Split-attention networks, 2020.
- [27] M. Tan and Q. V. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, 2020.
- [28] V. Tra, J. Kim, S. A. Khan, and J.-M. Kim, Bearing fault diagnosis under variable speed using convolutional neural networks and the stochastic diagonal levenberg-marquardt algorithm, *Sensors* vol. 17, no. 12, 2017.
- [29] T. Kanamori, S. Hido, and M. Sugiyama, A least-squares approach to direct importance estimation, *Journal of Machine Learning Research* vol. 10, no. 48, pp. 1391 1445, 2009.
- [30] B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars, Unsupervised visual domain adaptation using subspace alignment, in *2013 IEEE International Conference on Computer Vision*, pp. 2960 2967, 2013.
- [31] M. Baktashmotlagh, M. T. Harandi, B. C. Lovell, and M. Salzmann, Domain adaptation on the statistical manifold, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [32] J. Ho man, T. Darrell, and K. Saenko, Continuous manifold based adaptation for evolving visual domains, 2014.
- [33] S. J. Pan, J. T. Kwok, and Q. Yang, Transfer learning via dimensionality reduction, in *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2, AAAI'08*, p. 677 682, AAAI Press, 2008.
- [34] S. J. Pan, X. Ni, J.-T. Sun, Q. Yang, and Z. Chen, Cross-domain sentiment classification via spectral feature alignment, in *Proceedings of*

## 5. Conclusion

the 19th International Conference on World Wide Web WWW '10, (New York, NY, USA), p. 751-760, Association for Computing Machinery, 2010.

- [35] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, Domain-adversarial training of neural networks, 2016.
- [36] R. Tachet, H. Zhao, Y.-X. Wang, and G. Gordon, Domain adaptation with conditional distribution matching and generalized label shift, 2020.
- [37] Y. Li, M. Murias, S. Major, G. Dawson, and D. E. Carlson, On target shift in adversarial domain adaptation, 2019.
- [38] H. Liu, M. Long, J. Wang, and M. Jordan, Transferable adversarial training: A general approach to adapting deep classifiers, in Proceedings of the 36th International Conference on Machine Learning (K. Chaudhuri and R. Salakhutdinov, eds.), vol. 97 of Proceedings of Machine Learning Research, pp. 4013-4022, PMLR, 09-15 Jun 2019.
- [39] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, Adversarial discriminative domain adaptation, 2017.
- [40] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial networks, 2014.
- [41] G. J. McLachlan, Discriminant Analysis and Statistical Pattern Recognition. Hoboken, NJ, USA: John Wiley & Sons, Inc, 1992-03-27.
- [42] M. Kull, M. Perelló-Nieto, M. Kängsepp, T. de Menezes e Silva Filho, H. Song, and P. A. Flach, Beyond temperature scaling: Obtaining well-calibrated multiclass probabilities with dirichlet calibration, CoRR, vol. abs/1910.12656, 2019.
- [43] J. Vaicenavicius, D. Widmann, C. R. Andersson, F. Lindsten, J. Roll, and T. B. Schön, Evaluating model calibration in classification, CoRR, vol. abs/1902.06977, 2019.
- [44] B. Zadrozny and C. Elkan, Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers, in Proceedings of the Eighteenth International Conference on Machine Learning ICML '01, (San Francisco, CA, USA), p. 609-616, Morgan Kaufmann Publishers Inc., 2001.



- [45] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, On calibration of modern neural networks, *CoRR*, vol. abs/1706.04599, 2017.
- [46] A. H. Murphy and R. L. Winkler, Reliability of Subjective Probability Forecasts of Precipitation and Temperature, *Journal of the Royal Statistical Society Series C* vol. 26, pp. 41–47, March 1977.
- [47] A. Alexandari, A. Kundaje, and A. Shrikumar, Maximum likelihood with bias-corrected calibration is hard-to-beat at label shift adaptation, in *ICML*, pp. 222–232, 2020.
- [48] W. Wang and M. Á. Carreira-Perpiñán, Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application, *ArXiv*, vol. abs/1309.1541, 2013.
- [49] K. Cao et al., Learning imbalanced datasets with label-distribution-aware margin loss, in *NeurIPS*, 2019.
- [50] Z. Liu et al., Large-scale long-tailed recognition in an open world, in *CVPR*, 2019.
- [51] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [52] Y. Idelbayev, Proper ResNet implementation for CIFAR10/CIFAR100 in PyTorch. [https://github.com/akamaster/pytorch\\_resnet\\_cifar10](https://github.com/akamaster/pytorch_resnet_cifar10). Accessed: 2021-03-23.



## Appendix A

### Code

#### A.1 Sample Code

We attach jupyter notebook with sample code to the thesis with the implementation of the proposed methods and algorithms used for comparison. Classifier's predictions precomputed on CIFAR100-LT are provided as well for easy demonstration. Packages required to run the code can be installed via pip or Conda using requirements.txt or env.yml files. The algorithms for adaptation are implemented in separated script algorithms.py. The sample code contains the following files:

```
sample_code
├── env.yml
├── requirements.txt
├── example.ipynb
├── algorithms.py
├── CIFAR100_LT_outputs
│   └── outputs.pth.tar
```

To run the code start jupyter notebook and run all cells.

#### A.2 Full Reproducibility Code

In addition to the sample code we also provide experimental code, used to generate all results in the main text. Use `.../experimental_code/requirements.txt` file to set up the python environment.

To train 10 classifiers on CIFAR100 on UNI and LT distributions run

## A. Code .....

following command (assuming current directory is the root of experimental code, i.e. `.../experimental_code/`):

```
python3 multi_train.py -c configs/train_configs/train_
config_Cifar100.json -s
```

For Places365:

```
python3 multi_train.py -c configs/train_configs/train_
config_Places365.json -s
```

And for ImageNet:

```
python3 multi_train.py -c configs/train_configs/train_
config_ImageNet.json -s
```

Note that you will need to set the path to the images of original CIFAR100, ImageNet and Places365 datasets. This is done in corresponding config file `.../experimental_code/configs/dataset_configs/{dataset}.json` by setting "root" tag.

To run evaluation with label shift adaptation you can run following command:

```
python3 multi_eval.py -c adaptation/eval_configs/eval.json -s
```

Before running this command you need to set paths to the saved classifier's outputs into the `eval.json` config file (or another one you used in the argument). The tag determined for this information is "outputs" , taking the paths as a list.



## Appendix B

### List of Attachments

- sample\_code.zip
- experimental\_code.zip