

# Consistent Estimators for Learning to Defer to an Expert

by

Hussein Mozannar

B.E., American University of Beirut (2019)

Submitted to the Institute for Data, Systems, and Society  
in partial fulfillment of the requirements for the degree of  
Master of Science Degree in Social and Engineering Systems

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

© 2023 Hussein Mozannar. This work is licensed under a CC BY-SA 4.0.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license..

Authored by: Hussein Mozannar  
Institute for Data, Systems, and Society  
May 1, 2023

Certified by: David Sontag  
Professor of Electrical Engineering and Computer Science

Accepted by: Fotini Christia  
Chair, Social and Engineering Systems Program



# Consistent Estimators for Learning to Defer to an Expert

by

Hussein Mozannar

Submitted to the Institute for Data, Systems, and Society  
on May 1, 2023, in partial fulfillment of the  
requirements for the degree of  
Master of Science Degree in Social and Engineering Systems

## Abstract

Learning algorithms are often used in conjunction with expert decision makers in practical scenarios, however, this fact is largely ignored when designing these algorithms. In this thesis, we explore how to learn predictors that can either predict or choose to defer the decision to a downstream expert. Given only samples of the expert's decisions, we give a procedure based on learning a classifier and a rejector and analyze it theoretically. Our approach is based on a novel reduction to cost sensitive learning where we give a consistent surrogate loss for cost sensitive learning that generalizes the cross entropy loss. We show the effectiveness of our approach on a variety of experimental tasks.

Thesis Supervisor: David Sontag

Title: Professor of Electrical Engineering and Computer Science



# Contents

<b>Abstract</b>	<b>3</b>
<b>1 Introduction</b>	<b>11</b>
1.1 Overview . . . . .	11
1.2 Related Work . . . . .	13
<b>2 Problem Formulation</b>	<b>15</b>
<b>3 Proposed Surrogate Loss</b>	<b>19</b>
<b>4 Theoretical analysis</b>	<b>25</b>
4.1 Failure of Confidence Scores Method . . . . .	25
4.2 Inconsistency of mixtures of experts loss and Realizable-consistency . . . . .	27
4.3 Generalization Bound For Joint Learning . . . . .	29
<b>5 Experiments</b>	<b>31</b>
5.1 Synthetic Data . . . . .	31
5.2 CIFAR-10 . . . . .	33
5.3 CIFAR-100 . . . . .	36
5.4 CIFAR10H and limited expert data . . . . .	36
5.5 Hate Speech and Offensive Language Detection . . . . .	38
5.6 Synthetic Experts on CheXpert . . . . .	39
<b>6 Conclusion</b>	<b>49</b>
<b>A Appendix</b>	<b>51</b>
A.1 Practitioner’s guide to our approach . . . . .	51
A.1.1 General implementation . . . . .	51
A.1.2 Choice of $\alpha$ . . . . .	53
A.2 Experimental Details and Results . . . . .	54
A.2.1 CIFAR-10 . . . . .	54
A.2.2 CIFAR-10H . . . . .	56
A.2.3 CIFAR-100 . . . . .	57
A.2.4 Hate Speech experiments . . . . .	58

A.2.5	Baseline Implementation . . . . .	59
A.2.6	CheXpert Experiments . . . . .	62
A.3	Deferred Proofs and Derivations . . . . .	63
A.3.1	Section 3 . . . . .	63
A.3.2	Section 4 . . . . .	71

# List of Figures

2-1	The expert deferral pipeline, the rejector first $r(x)$ decides who between the classifier $h(x)$ and expert $M(z)$ should predict and then whoever makes the final prediction incurs a specific cost. . . . .	16
4-1	Setting of two groups, red and blue, the task is binary classification with labels $\{o, +\}$ , the expert fits the red majority group, hence the classifier should attempt to fit the blue group with the rejector (black line) separating the groups.	26
5-1	Left figure shows overall system accuracy of our method and baselines ( $k$ is the number of classes the expert can predict) and right figure compares the accuracy on the non-deferred examples versus the coverage for every $k$ . . .	34
5-2	Varying training set size when training with expert $k = 5$ for Confidence baseline and our method $L_{CE}^1$ . . . . .	35
5-3	Plot of AU-ROC of the ROC curve (a) for each level of coverage (0 coverage means only the expert predicting and 1 coverage is only the classifier predicting) and of the area under the precision-recall curve (AU-PR) (b) for each of the 5 tasks comparing our method with the baselines on the training derived test set for the toy expert with $q = 0.7, p = 1$ . We report the maximum AU-ROC and AU-PR achieved on each task, error bars are standard deviations derived from 10 runs (averaging over the expert's randomness). . . . .	42
5-4	Left figure shows the average of the maximum achievable AU-ROC for the 5 tasks (average over the tasks) when the changing the size of the training data (as a % of the original set) and right figure shows the same for AU-PR . . .	44
5-5	Left figure (a) shows the chest X-ray of a patient with Cardiomegaly, the right figure (b) shows that same X-ray but now with the left part hidden which is used as input to the models. . . . .	45
5-6	Left figure shows the average of the maximum achievable AU-ROC for the 5 tasks (average over the tasks) when the changing the scale of the noise (size of rectangular region) and right figure shows the same for AU-PR. . . . .	46
5-7	Plots of AU-ROC and AU-PR as we vary coverage when training and testing with chest X-rays that have their left section hidden. The expert model is $q = 0.7$ and $p = 1$ . . . . .	47

A-1	Comparison of the developed method $L_{CE}^1$ on CIFAR-100 versus the confidence baseline. $k$ is the number of classes the expert can predict . . . . .	57
A-2	Plot of classifier AU-ROC on non-deferred examples versus coverage for (a) for systems learned with 100% of training data (b) and learned with 10% of training data. Noise at low coverage is due to reduced data size. . . . .	62



# List of Tables

5.1	Comparison of our methods with the confidence score baseline, oracle baseline and our implementation of [43] method. We compute a 95% confidence interval for the average difference between the baselines and our method. . . . .	32
5.2	Comparing our proposed methods on CIFAR10H and a baseline based on confidence scores recording system accuracy, coverage and classifier accuracy on non-deferred examples. . . . .	37
5.3	Average difference in AU-ROC across all coverage and difference between maximum achievable AU-ROC between our method and the Confidence baseline for each of the 5 tasks and different toy expert probabilities $p$ and $q$ ; each entry is (average difference $\pm$ standard deviation; difference of maximums). The difference between our method and the ModelConfidence is roughly twice the values noted in table 5.3, only at Expert (0.7, 0.7) does Confidence and ModelConfidence achieve the same performance since the expert has uniform error over the domain. . . . .	43
A.1	Accuracy of the expert on deferred examples shown for the methods and baselines proposed with varying expert competence (k) on CIFAR-10. . . . .	55
A.2	Accuracy of the classifier on non-deferred examples shown for our method $L_{CE}^1$ and baseline Confidence with varying training set size for expert $k = 5$ on CIFAR-10. . . . .	56
A.3	Accuracy of the CIFAR10H [53] expert on each of the 10 classes . . . . .	56
A.4	Complete results of table 5.2 comparing our proposed approaches and baseline.	56
A.5	Accuracy of the expert on deferred examples shown for the methods and baselines proposed with varying expert competence (k) on CIFAR-100. . . . .	58
A.6	Detailed results for our method and baselines on the hate speech detection task [15]. sys: system accuracy, class: classifier accuracy, disc: system discrimination, AAE-biased: Expert 2 that has higher error rate for AAE group, non-AAE biased: Expert 3 that has higher error for non AAE tweets . . . . .	59
A.7	System accuracy of our implementation of [43] and our method and baselines with varying expert competence (k) on CIFAR-10. . . . .	61



# Chapter 1

## Introduction

### 1.1 Overview

Machine learning systems are now being deployed in settings to complement human decision makers such as in healthcare [30, 54], risk assessment [27] and content moderation [39]. These models are either used as a tool to help the downstream human decision maker: judges relying on algorithmic risk assessment tools [28] and risk scores being used in the ICU [23], or instead these learning models are solely used to make the final prediction on a selected subset of examples [43, 54]. A current application of the latter setting is Facebook’s and other online platform’s content moderation approach [58, 33]: an algorithm is used to filter easily detectible inappropriate content and the rest of the examples are screened by a team of human moderators. Another motivating application arises in health care settings, for example, deep neural networks can outperform radiologists in detecting pneumonia from chest X-rays [31], however, many obstacles are limiting complete automation, an intermediate step to automating this task will be the use of models as triage tools to complement radiologist expertise. Our focus in this work is to give theoretically sound approaches for machine learning models that can either predict or defer the decision to a downstream expert to complement and augment their capabilities.

The learned model should adapt to the underlying human expert in order to achieve better performance than deploying the model or expert individually. In situations where we have limited data or model capacity, the gains from allowing the model to focus on regions where the expert is less accurate are expected to be more significant. However, even when data or model capacity are not a concern, the expert may have access to side information unavailable to the learner due to privacy concerns for example, the hard task is then to identify when we should defer without having access to this side-information. We will only assume in this work that we are allowed access to samples of the expert’s decisions or to costs of deferring, we believe that this is a reasonable assumption that can be achieved in practical settings. Inspired by the literature on rejection learning [13], our approach will be to learn two functions: a classifier that can predict the target and a rejector that decides whether the classifier or the expert should predict.

We start by formulating a natural loss function for the combined machine-expert system in section 2 and show a reduction from the expert deferral setting to cost sensitive learning. With this reduction in hand, we are able to give a novel convex surrogate loss that upper bounds our system loss and that is furthermore consistent in section 3. This surrogate loss settles the open problem posed by [50] for finding a consistent loss for multiclass rejection learning. Our proposed surrogate loss and approach requires only adding an additional output layer to existing model architectures and changing the loss function, hence it necessitates minimal to no added computational costs. In section 4, we show the limitations of approaches in the literature from a consistency point-of-view and then provide generalization bounds for minimizing the empirical loss. To show the efficacy of our approach, we give experimental evidence on image classification datasets CIFAR-10 and CIFAR-100 using synthetic and human experts based on CIFAR10H [53], on a hate speech and offensive language detection task [15], and on classification of chest X-rays with synthetic experts in section 5. To summarize, the contributions of this thesis are the following:

- We formalize the expert deferral setup and analyze it theoretically giving a generalization bound for solving the empirical problem.

- We propose a novel convex consistent surrogate loss  $L_{CE}$  (3.4) for expert deferral easily integrated into current learning pipelines.
- We provide a detailed experimental evaluation of our method and baselines from the literature on image and text classification tasks.

Note that this thesis is based on the ICML 2020 paper I published in the first year of my PhD [49]. It is part of a series of work I have built on during my PhD which I mention briefly in the related work, my PhD thesis will bridge these work together in the next year.

## 1.2 Related Work

Learning with a reject option, *rejection learning*, has long been studied starting with [11] who investigated the trade-off between accuracy and the rejection rate. The framework of rejection learning assumes a constant cost  $c$  of deferring and hence the problem becomes to predict only if one is  $1 - c$  confident. Numerous works have proposed surrogate losses and uncertainty estimation methods to solve the problem [4, 56, 50, 34]. [13, 12] proposed a different approach by learning two functions: a classifier and a rejection function, and analyzed the approach by giving a kernel-based algorithm in the binary setting. [50] tried to extend their approach to the multiclass setting but failed to give a consistent surrogate loss and hence resorted to confidence based methods.

Recent work has started to explore models that defer to downstream experts, [43] considers an identical framework to the one considered here however their approach does not allow the model to adapt to the underlying expert and the loss used is not consistent and requires an uncertainty estimate of the expert decisions. On the other hand, [16] gives an approximate procedure to learn a linear model that picks a subset of the training data on which to defer and uses a nearest neighbor algorithm to defer on new examples, the approach used is only feasible for small dataset sizes and does not generalize beyond ridge regression. [54] considers

binary classification with expert deferral, their approach is to learn a classifier ignoring the expert and obtain uncertainty estimates for both the expert and classifier and then defer based on which is higher, we detail the limitations of this approach in section 4. Concurrent work [59] learns a model with the mixtures of expert loss first introduced in [43] and defers based on estimated model and expert confidence as in [54]. Work on AI-assisted decision making has focused on the reverse setting considered here: the expert chooses to accept or reject the decision of the classifier instead of a learned rejector [2, 1]. Additionally, the fairness in machine learning community has started to consider the fairness impact of having downstream decision makers [43, 8, 27, 20] but in slightly different frameworks than the ones considered here and work has started to consider deferring in reinforcement learning [44].

A related framework to our setting is selective classification [21] where instead of setting a cost for rejecting to predict one sets a constraint on the probability of rejection; here is no assumed downstream expert. Approaches range from deferring based on confidence scores [25], learning a deep network with two heads, one for predicting and the other for deferring [26], and learning with portfolio theory inspired loss functions [62]. Finally, our work bears a resemblance to active learning with weak (the expert) and strong labelers (the ground truth) [61].

In follow-up work to this thesis with collaborator Mohammad-Amin Charusaie, we showed how one can derive a family of surrogate losses for cost-sensitive learning using a construction from surrogate losses for multi-class learning [9]. We also show a theoretical worst case bound between joint learning and separate learning, this makes the motivation in section 4.1 concrete. However, in [47], we show that the surrogate loss presented in this thesis fails to learn complementary predictions in simple linear data settings. This is partially due to the fact that learning a classifier-rejector pair jointly is computationally hard, and thus we build non-convex approaches that can exactly minimize the error of the human-AI system. The notion of deferring has been very fruitful in other parts of my work where a human can choose to defer to the classifier instead of an AI rejector deciding who should predict [19, 48, 46].

# Chapter 2

## Problem Formulation

We are interested in predicting a target  $Y \in \mathcal{Y} = \{1, \dots, K\}$  based on covariates  $X \in \mathcal{X}$  where  $X, Y \sim \mathbf{P}$ . We assume that we have query access to an expert  $M$  that has access to a domain  $\mathcal{Z}$  that may contain additional information than  $\mathcal{X}$  to classify instances according to the target  $\mathcal{Y}$ . Querying the expert implies deferring the decision which incurs a cost  $l_{exp}(x, y, m)$  that depends on the target  $y$ , covariate  $x$ , and the expert's prediction  $m$ . On the other hand, predicting without querying the expert implies that a classifier makes the final decision and incurs a cost  $l(x, y, \hat{y})$  where  $\hat{y}$  is the prediction of the classifier. Our goal is to build a predictor  $\hat{Y} : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$  that can either predict or defer the decision to the expert denoted by  $\perp$ . We can now formulate a natural system loss function  $L$  for the system consisting of the classifier in conjunction with the expert:

$$L(\hat{Y}) = \mathbb{E}_{(x,y) \sim \mathbf{P}, m \sim M|(x,y)} \left[ \underbrace{l(x, y, \hat{Y}(x))}_{\text{classifier cost}} \overbrace{\mathbb{I}_{\hat{Y}(x) \neq \perp}}^{\text{predict}} + \underbrace{l_{exp}(x, y, m)}_{\text{expert cost}} \overbrace{\mathbb{I}_{\hat{Y}(x) = \perp}}^{\text{defer}} \right] \quad (2.1)$$

Our strategy for learning the predictor  $\hat{Y}$  will be to learn two separate functions  $h : \mathcal{X} \rightarrow \mathcal{Y}$

(classifier) and  $r : \mathcal{X} \rightarrow \{0, 1\}$  (rejector) and hence we write our loss as:

$$L(h, r) = \mathbb{E}_{(x,y) \sim \mathbf{P}, m \sim M|(x,y)} [ l(x, y, h(x))\mathbb{I}_{r(x)=0} + l_{\text{exp}}(x, y, m)\mathbb{I}_{r(x)=1} ] \quad (2.2)$$

Figure 2-1 illustrates our expert deferral setting with its different components. The above

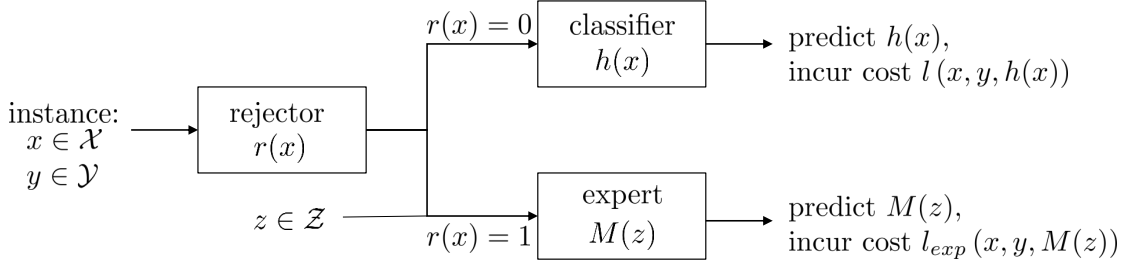


Figure 2-1: The expert deferral pipeline, the rejector first  $r(x)$  decides who between the classifier  $h(x)$  and expert  $M(z)$  should predict and then whoever makes the final prediction incurs a specific cost.

formulation is a generalization of the learning with rejection framework studied by [13] as by setting  $l_{\text{exp}}(x, y, m) = c$  for a constant  $c > 0$  the two objectives coincide. In [43], the loss proposed assumes that the classifier and expert costs are the logistic loss between the target and their predictions in the binary target setting.

While our treatment extends to general forms of expert and classifier costs, we will pay particular attention in our theoretical analysis when the costs are the misclassification error with the target. Formally, we define a 0–1 loss version of our system loss:

$$L_{0-1}(h, r) = \mathbb{E}_{(x,y) \sim \mathbf{P}, m \sim M|(x,y)} [ \mathbb{I}_{h(x) \neq y} \mathbb{I}_{r(x)=0} + \mathbb{I}_{m \neq y} \mathbb{I}_{r(x)=1} ] \quad (2.3)$$

One may also assume a constant additive cost function  $c(x)$  for querying the expert depending on the instance  $x$  making  $l_{\text{exp}}(x, y, m) = \mathbb{I}_{m \neq y} + c(x)$ ; such additive costs can be easily integrated into our analysis.

Our approach will be to cast this problem as a *cost sensitive learning* problem over an augmented label space that includes the action of deferral. Let the random costs  $\mathbf{c} \in \mathbb{R}_+^{K+1}$



where for  $i \in [K]$ ,  $c(i)$  is the  $i$ 'th component of  $\mathbf{c}$  represents the cost of predicting  $i \in \mathcal{Y}$  while  $c[K + 1]$  represents the cost of deferring to the expert. The goal of this setup is to learn a predictor  $h : \mathcal{X} \rightarrow [K + 1]$  minimizing the cost sensitive loss  $\tilde{L}(h) := \mathbb{E}[c(h(x))]$ . For example, giving an instance  $(x, y)$ , our loss (2.2) is obtained by setting  $c(i) = l(x, y, i)$  for  $i \in [K]$  and  $c(K + 1) = l_{\text{exp}}(x, y, m)$ .

For the majority of this thesis we assume access to samples  $S = \{(x_i, y_i, m_i)\}_{i=1}^n$  where  $\{(x_i, y_i)\}_{i=1}^n$  are drawn i.i.d. from the unknown distribution  $\mathbf{P}$  and  $m_i$  is drawn from the distribution of the random variable  $M|(X = x_i, Y = y_i)$  and access to the realizations of  $l_{\text{exp}}$  and  $l$  when required.



# Chapter 3

## Proposed Surrogate Loss

It is clear that the system loss function (2.2) is not only non-convex but also computationally hard to optimize. The usual approach in machine learning is to formulate upper bounding convex surrogate loss functions and optimize them in hopes of approximating the minimizers of the original loss [3]. Work from rejection learning [13, 50] suggested learning two separate functions  $h$  and  $r$  and provided consistent convex surrogate loss functions only for the binary setting. We extend their proposed surrogates for our expert deferral setting for binary labels with slight modifications in appendix A.3. Consistency is used to prove that a proposed surrogate loss is a good candidate and is often treated as a necessary condition. The issue with the proposed surrogates in [13] for rejection learning is that when extended to the multiclass setting, it is impossible for them to be consistent as was shown by [50]. Aside the consistency issue, [50] found that simple baselines can outperform the proposed losses in practice.

The construction of our proposed surrogate loss for the multiclass expert deferral setting will be motivated in two ways, the first is through a novel reduction to cost sensitive learning, and the second is inspired by the Bayes minimizer for the 0–1 system loss (2.3). Let  $g_i : \mathcal{X} \rightarrow \mathbb{R}$  for  $i \in [K + 1]$  and define  $h(x) = \arg \max_{i \in [K+1]} g_i$ , motivated by the success of the cross

entropy loss, our proposed surrogate for cost-sensitive learning  $\tilde{L}_{CE}$  takes the following form:

$$\tilde{L}_{CE}(g_1, \dots, g_{K+1}, x, c(1), \dots, c(K+1)) = - \sum_{i=1}^{K+1} \left( \max_{j \in [K+1]} c(j) - c(i) \right) \log \left( \frac{\exp(g_i(x))}{\sum_k \exp(g_k(x))} \right) \quad (3.1)$$

The loss  $\tilde{L}_{CE}$  is a novel surrogate loss for cost sensitive learning that generalizes the cross entropy loss when the costs correspond to multiclass misclassification. The following proposition shows that the loss is consistent, meaning its minimizer over all measurable functions agrees with the Bayes solution.

**Proposition 1.**  *$\tilde{L}_{CE}$  is convex in  $\mathbf{g}$  and is a consistent loss function for  $\tilde{L}$ :*

let  $\tilde{\mathbf{g}} = \arg \inf_{\mathbf{g}} \mathbb{E} \left[ \tilde{L}_{CE}(\mathbf{g}, \mathbf{c}) | X = x \right]$ , then:  $\arg \max_{i \in [K+1]} \tilde{\mathbf{g}}_i = \arg \min_{i \in [K+1]} \mathbb{E}[c(i) | X = x]$

Proof of Proposition 1 can be found in Appendix A.3;  $\tilde{L}_{CE}$  is a simpler consistent alternative to the surrogates derived in [10] for cost sensitive learning.

Now we consider when the system loss function is  $L_{0-1}$  (2.3), our approach is to treat deferral as a new class and construct a new label space  $\mathcal{Y}^\perp = \mathcal{Y} \cup \perp$  and a corresponding distribution  $\mathbb{P}(Y^\perp | X = x)$  such that minimizing the misclassification loss on this new space will be equivalent to minimizing our system loss  $L_{0-1}$ . The Bayes optimal classifier on  $\mathcal{Y}^\perp$  is clearly  $h^\perp = \arg \max_{y^\perp \in \mathcal{Y}^\perp} \mathbb{P}(\mathcal{Y}^\perp = y^\perp | X = x)$ , and we need it to match the decision of the Bayes solution  $h^B, r^B$  of  $L_{0-1}$  (2.3):

$$h^B, r^B = \arg \inf_{h, r} L_{0-1}(h, r) \quad (3.2)$$

where the infimum is over all measurable functions. Denote by  $\eta_y(x) = \mathbb{P}(Y = y | X = x)$ , it is clear that for  $x \in \mathcal{X}$  the best classifier is the same as the Bayes solution for standard classification since if we don't defer we have to do our best. Now we only reject the classifier if its expected error is higher than the expected error of the expert which we formalize in the below proposition:

**Proposition 2.** *The minimizers of the loss  $L_{0-1}$  (2.3) are defined point-wise for all  $x \in \mathcal{X}$  as:*

$$\begin{aligned} h^B(x) &= \arg \max_{y \in \mathcal{Y}} \eta_y(x) \\ r^B(x) &= \mathbb{I}_{\max_{y \in \mathcal{Y}} \eta_y(x) \leq \mathbb{P}(Y=M|X=x)} \end{aligned} \quad (3.3)$$

Proof of the above proposition can be found in Appendix A.3 and equation (3.3) give us sufficient conditions for consistency to check our proposed loss. Let  $g_y : \mathcal{X} \rightarrow \mathbb{R}$  for  $y \in \mathcal{Y}$  and define  $h(x) = \arg \max_{y \in \mathcal{Y}} g_y$ , similarly let  $g_{\perp} : \mathcal{X} \rightarrow \mathbb{R}$  and define  $r(x) = \mathbb{I}_{\max_{y \in \mathcal{Y}} g_y(x) \leq g_{\perp}}$  the proposed surrogate loss for  $L_{0-1}$  (2.2) in the multiclass setting is then:

$$L_{CE}(h, r, x, y, m) = -\log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) - \mathbb{I}_{m=y} \log \left( \frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \quad (3.4)$$

The proposed surrogate  $L_{CE}$  is in fact consistent and upper bounds  $L_{0-1}$  as the following theorem demonstrates.

**Theorem 1.** *The loss  $L_{CE}$  is convex in  $\mathbf{g}$ , upper bounds  $L_{0-1}$  and is consistent:*

$$\inf_{h, r} \mathbb{E}_{x, y, m} [L_{CE}(h, r, x, y, m)]$$

*is attained at  $(h_{CE}^*, r_{CE}^*)$  such that  $h^B(x) = h_{CE}^*(x)$  and  $r^B(x) = r_{CE}^*(x)$  for all  $x \in \mathcal{X}$ .*

*Proof Sketch.* Please refer to appendix A.3 for the detailed proof. First the infimum over functions  $h, r$  can be replaced by a point-wise infimum as:

$$\inf_{h, r} \mathbb{E}_{x, y, m} [L_{CE}(h, r, x, y, m)] = \mathbb{E}_x \inf_{h(x), r(x)} \mathbb{E}_{y|x} \mathbb{E}_{m|x, y} [L_{CE}(h(x), r(x), x, y, m)]$$

Now let us expand the inner expectation:

$$\begin{aligned} \mathbb{E}_{y|x} \mathbb{E}_{m|x,y} [L_{SH}(h(x), r(x), x, y, m)] &= - \sum_{y \in \mathcal{Y}} \eta_y(x) \log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \\ &\quad - \mathbb{P}(Y = M | X = x) \log \left( \frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \end{aligned} \quad (3.5)$$

For ease of notation denote the RHS of equation (3.5) as  $L_{CE}(g_1, \dots, g_{|\mathcal{Y}|}, g_{\perp})$ , note that it is a convex function, hence we will take the partial derivatives with respect to each argument and set them to 0. For any  $g_{\perp}$  and  $i \in \mathcal{Y}$  we have :

$$\frac{\exp(g_i^*(x))}{\sum_{y' \in \tilde{\mathcal{Y}}} \exp(g_{y'}(x))} = \frac{\eta_i(x)}{1 + \mathbb{P}(Y = M | X = x)} \quad (3.6)$$

The optimal  $h^*$  for any  $g_{\perp}$  should satisfy equation (3.6) for every  $i \in \mathcal{Y}$ . Plugging  $h^*$  and taking the derivative with respect to  $g_{\perp}$  we get:

$$\frac{\exp(g_{\perp}^*(x))}{\sum_{y' \in \mathcal{Y}} \exp(g_{y'}^*(x))} = \frac{\mathbb{P}(Y = M | X = x)}{1 + \mathbb{P}(Y = M | X = x)}$$

since exponential is an increasing function we get that the optimal  $h^*$  and  $r^*$  in fact agrees with the Bayes solution.  $\square$

When the costs  $c(1), \dots, c(K+1)$  are in accordance with our expert deferral setting the loss  $\tilde{L}_{CE}$  reduces to  $L_{CE}$ . Now stepping back and looking more closely at our loss  $L_{CE}$ , we can see that the loss on examples where the expert makes a mistake becomes the cross entropy loss with the target. On the other hand, when the expert agrees with the target, the learner faces two opposing decisions whether to defer or predict the target. We can encourage or hinder the action of deferral by modifying the loss with an additional parameter  $\alpha \in \mathbb{R}^+$  as

$L_{CE}^\alpha(h, r, x, y, m)$ :

$$\begin{aligned}
L_{CE}^\alpha(h, r, x, y, m) = & - (\alpha \cdot \mathbb{I}_{m=y} + \mathbb{I}_{m \neq y}) \log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \\
& - \mathbb{I}_{m=y} \log \left( \frac{\exp(g_\perp(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right)
\end{aligned} \tag{3.7}$$

Note that  $L_{CE}^1 = L_{CE}$ . The effect of  $\alpha$  is to re-weight examples where the expert is correct to discourage the learner from fitting them and instead focus on examples where the expert makes a mistake. In practice, one would treat  $\alpha$  as an additional hyperparameter to optimize for.





# Chapter 4

## Theoretical analysis

In this chapter we focus on the zero-one system loss function  $L_{0-1}$  and try to understand previous proposed solutions in the literature in comparison with our method from a theoretical perspective.

### 4.1 Failure of Confidence Scores Method

Let us first remind ourselves of the Bayes solution for the system loss:

$$h^B(x) = \arg \max_{y \in \mathcal{Y}} \eta_y(x), \quad r^B(x) = \mathbb{I}_{\max_{y \in \mathcal{Y}} \eta_y(x) \leq \mathbb{P}(Y=M|X=x)}$$

The form of the Bayes solution above suggests a very natural approach: 1) learn a classifier minimizing the misclassification loss with the target and obtain confidence scores for predictions, 2) obtain confidence scores for expert agreement with the target, this can be done by learning a model where the target is whether the expert agrees with the task label and extracting confidence scores from this model [55], and finally 3) compare who between the classifier and the expert is more confident and accordingly defer. We refer to this as the

confidence score method (Confidence), this approach leads to a consistent estimator for both the rejector and classifier and was proposed by [54].

In fact this is the standard approach in rejection learning [4, 56, 50], a host of different methods exist for estimating a classifier’s confidence on new examples including trust scores [34], Monte-Carlo dropout for neural networks [24] among many others. However, the key pitfall of this method in the expert deferral setup it that it does not allow  $h$  to adapt to the expert’s strengths and weaknesses. When we restrict our search space to a limited class of functions  $\mathcal{H}$  and  $\mathcal{R}$  this approach can easily fail. We now give a toy example where learning the classifier independently fails which motivates the need to jointly learn both the classifier and rejector.

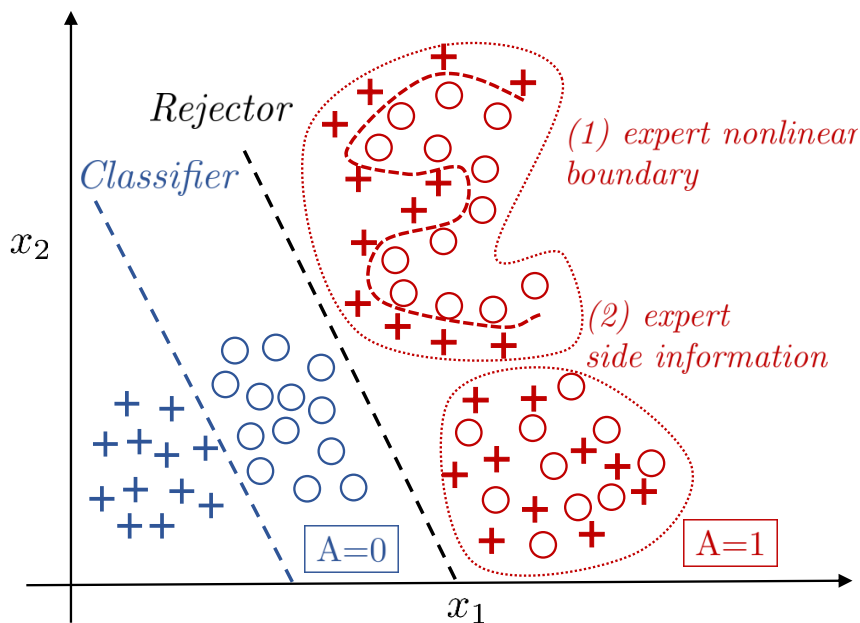


Figure 4-1: Setting of two groups, red and blue, the task is binary classification with labels  $\{o, +\}$ , the expert fits the red majority group, hence the classifier should attempt to fit the blue group with the rejector (black line) separating the groups.

Assume that there exists two sub-populations in the data denoted  $A = 1$  and  $A = 0$  where  $\mathbb{P}(A = 1) \geq \mathbb{P}(A = 0)$  from which  $X \in \mathbb{R}^d$  is generated from and conditional on the target and population,  $X|(Y = y, A = 0)$  is normally distributed according to  $\mathcal{N}(\mu_{y,0}, \Sigma)$  and

$X|Y = y, A = 1$ ) consists of two clusters: cluster (1) is normally distributed but the means are not well separated and cluster (2) is only separable by a complex non-linear boundary; the data is illustrated in Figure 4-1. Finally we assume the expert to be able to perfectly classify group  $A = 1$ , on cluster (1) the expert is able to compute the complex nonlinear boundary and on cluster (2) the expert has side-information  $Z$  that allows him to separate the classes which is not possible from only  $X$ . Our hypothesis spaces  $\mathcal{H}$  and  $\mathcal{G}$  will be the set of all  $d$ -dimensional hyperplanes. If we start by learning  $h$ , then the resulting hyperplane will try to minimize the average error across both groups, this will likely result into a hyperplane that separates neither group as the data is not linearly separable, especially on group  $A = 1$ . If we assume that the boundary between the groups is linear as shown, then we can achieve the error of the Bayes solution within our hypothesis space: the optimal behavior in this setting is clearly to have  $h$  fit group  $A = 0$ , note here the Bayes solution corresponds to a hyperplane via linear discriminant analysis for 2 classes on  $A = 0$ , and the rejector  $r$  separating the groups as illustrated in Figure 4-1. This example illustrates the complexities of this setting, due to model capacity there are significant gains to be achieved from adapting to the expert by focusing only group  $A = 0$ . Setting aside model capacity, the nonlinear boundary of cluster (1) is sample intensive to learn as we only have access to finite data. Finally, cluster (2) cannot be separated even with infinite data, the side information of the expert is needed, and so the hard task is to identify the region of cluster (2). This serves to illustrate the complexities of the setup and the importance of learning the classifier and rejector jointly.

## 4.2 Inconsistency of mixtures of experts loss and Realizable-consistency

So far we have focused on classification consistency to verify the soundness of proposed approaches, however, we usually have specific hypothesis classes  $\mathcal{H}, \mathcal{R}$  in mind, and if the Bayes predictor is not in our class then consistency might not guarantee much [5]. For

example, for binary classification with half-spaces, any predictor learned with a convex surrogate loss can have arbitrarily high error if the best half-space has non-zero error [5]. The previous example illustrated in Figure 4-1 shows an the mode of failure that exists in the expert deferral setup even in the realizable setting. Therefore, a more relevant requirement to our example is that the minimizers of a proposed surrogate loss and the original loss agree for given hypothesis classes in the *realizable* setting; this is formally defined with the below notion.

**Definition 1** (realizable  $(\mathcal{H}, \mathcal{R})$ -consistency). A surrogate loss  $L_{surr}$  is realizable  $(\mathcal{H}, \mathcal{R})$ -consistent if for all distributions  $\mathbf{P}$  and experts  $M$  for which there exists  $h^*, r^* \in \mathcal{H} \times \mathcal{R}$  that have zero error  $L(h^*, r^*) = 0$ , we have  $\forall \epsilon > 0, \exists \delta > 0$  such that if  $(\hat{h}, \hat{r})$  satisfies

$$\left| L_{surr}(\hat{h}, \hat{r}) - \inf_{h \in \mathcal{H}, r \in \mathcal{R}} L_{surr}(h, r) \right| \leq \delta, \text{ then: } L(\hat{h}, \hat{r}) \leq \epsilon$$

A similar notion was introduced for classification by [40] and by [13] for rejection learning, however here we have the the added dimension of the expert.

Note that the expert deferral setting considered here can be thought of as a hard mixture of two experts problem where one of the experts is fixed [35, 57, 43]. This observation motivates a natural mixture of experts type loss, let  $g_y : \mathcal{X} \rightarrow \mathbb{R}$  for  $y \in \mathcal{Y}$ ,  $h(x) = \arg \max_{y \in \mathcal{Y}} g_y$ ,  $r_i : \mathcal{X} \rightarrow \mathbb{R}$  for  $i \in \{0, 1\}$  and  $r(x) = \arg \max_{i \in \{0, 1\}} r_i(x)$ , the mixture of experts loss is defined as:

$$L_{mix}(\mathbf{g}, \mathbf{r}, x, y, m) = -\log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y}} \exp(g_{y'}(x))} \right) \frac{\exp(r_0(x))}{\sum_{i \in \{0, 1\}} \exp(r_i(x))} + \mathbb{I}_{m \neq y} \frac{\exp(r_1(x))}{\sum_{i \in \{0, 1\}} \exp(r_i(x))} \quad (4.1)$$

The above loss extends [43] approach to the multiclass setting. As the next proposition demonstrates,  $L_{mix}$  is in general *not* classification consistent, however, it is realizable  $(\mathcal{H}, \mathcal{R})$ -consistent for classes closed under scaling which include linear models and neural networks.

**Proposition 3.**  $L_{mix}$  is realizable  $(\mathcal{H}, \mathcal{R})$ -consistent for classes closed under scaling but is not classification consistent.

Proof of proposition 3 can be found in Appendix A.3. Note that integrating more information about  $M$  in  $L_{mix}$  would not make the loss consistent, the inconsistency arises from the parameterization in  $\mathbf{g}$ , setting the classifier loss to simply be  $\mathbb{I}_{h(x) \neq y}$  would make  $L_{mix}$  consistent at the cost of losing the convexity and differentiability in  $\mathbf{g}$ . While  $L_{mix}$  is indeed realizable consistent however it is not convex in both  $\mathbf{g}$  and  $\mathbf{r}$ , hence it is not clear how to efficiently optimize it. Setting aside computational feasibilities, it is also not immediately clear which between consistency and realizable  $(\mathcal{H}, \mathcal{R})$ -consistency will be more practically relevant. In our experimental chapter we show how the mismatch between the model and expert loss and their actual errors causes this method to learn the incorrect behavior which hints that classification consistency is crucial.

### 4.3 Generalization Bound For Joint Learning

In this subsection we analyze the sample complexity to jointly learn a rejector and classifier. The goal is to find the minimizer of the empirical version of our system loss when our hypothesis space for  $h$  and  $r$  are  $\mathcal{H}, \mathcal{R}$  respectively:

$$\hat{h}^*, \hat{r}^* = \arg \min_{h \in \mathcal{H}, r \in \mathcal{G}} L_{0-1}^S(h, r) := \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{h(x_i) \neq y_i} \mathbb{I}_{r(x_i)=0} + \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i)=1} \quad (4.2)$$

By going after the system loss directly, we can approximate the population minimizers  $h^*, r^*$  over  $\mathcal{H} \times \mathcal{R}$  of  $L_{0-1}$  (2.3). The optimum  $h^*$  may not necessarily coincide with the optimal minimizer of the misclassification loss with the target which is why learning jointly is critical. We now give a generalization bound for our empirical minimization procedure for a binary target.

**Theorem 2.** For any expert  $M$  and data distribution  $\mathbf{P}$  over  $\mathcal{X} \times \mathcal{Y}$ , let  $0 < \delta < \frac{1}{2}$ , then

with probability at least  $1 - \delta$ , the following holds for the empirical minimizers  $(\hat{h}^*, \hat{r}^*)$ :

$$\begin{aligned}
L_{0-1}(\hat{h}^*, \hat{r}^*) &\leq L_{0-1}(h^*, r^*) + \mathfrak{R}_n(\mathcal{H}) + \mathfrak{R}_n(\mathcal{R}) + \mathfrak{R}_{n\mathbb{P}(M \neq Y)/2}(\mathcal{R}) \\
&\quad + 2\sqrt{\frac{\log \frac{2}{\delta}}{2n}} + \frac{\mathbb{P}(M \neq Y)}{2} \exp\left(-\frac{n\mathbb{P}(M \neq Y)}{8}\right)
\end{aligned} \tag{4.3}$$

Proof of the above theorem can be found in Appendix A.3. We can see that the performance of our empirical minimizer is controlled by the Rademacher complexity  $\mathfrak{R}_n(\mathcal{R})$  and  $\mathfrak{R}_n(\mathcal{H})$  of both the classifier and rejector model classes and the error of the expert. Note that when  $\mathbb{P}(M \neq Y) = 0$  we recover the bound proved in Theorem 1 [13] for rejection learning when  $c = 0$ ; this gives evidence that deferring to an expert is a more sample intensive problem than rejection learning. Both our loss  $L_{CE}$  and the confidence scores approach lead to consistent estimators, however, as we will later show in our experiments, one differentiating factor will be that of sample complexity. We can already see in the bound (4.3), that we pay the complexity of the rejector and classifier model classes, however, our approach combines the rejector and classifier in one model to avoid these added costs.

# Chapter 5

## Experiments

We provide code to reproduce our experiments <sup>1</sup>, we have built an extended code base after this thesis in <sup>2</sup>. In Appendix A.1 we give a detailed guide on implementing our method. Additional experimental details and results are left to Appendix A.2.

### 5.1 Synthetic Data

As a first toy example to showcase that our proposed loss  $L_{CE}^\alpha$  is able to adapt to the underlying expert behavior, we perform experiments in a Gaussian mixture setup akin to the example in chapter 4. The covariate space is  $\mathcal{X} = \mathbb{R}^d$  and target  $\mathcal{Y} = \{0, 1\}$ , we assume that there exists two sub-populations in the data denoted  $A = 1$  and  $A = 0$ . Furthermore,  $X|(Y = y, A = a)$  is normally distributed according to  $\mathcal{N}(\mu_{y,a}, \Sigma_{y,a})$ . The expert follows the Bayes solution for group  $A = 1$  which here corresponds to a hyperplane. Our hypothesis spaces  $\mathcal{H}$  and  $\mathcal{R}$  will be the set of all  $d$ -dimensional hyperplanes.

**Setup:** We perform 200 trials where on each trial we generate: random group proportions  $\mathbb{P}(A = 1) \sim U(0, 1)$  fixing  $\mathbb{P}(Y = 1|A = a) = 0.5$ , random means and variances for each

---

<sup>1</sup><https://github.com/clinicalml/learn-to-defer>

<sup>2</sup>[https://github.com/clinicalml/human\\_ai\\_deferral](https://github.com/clinicalml/human_ai_deferral)

Gaussian component  $X|Y = y, A = a \sim \mathcal{N}(\mu_{y,a}, \Sigma_{y,a})$  where  $\mu_{y,a} \sim U(0, 10)^d$  and similarly for the diagonal components of  $\Sigma_{y,a}(i, i) \sim U(0, 10)$  keeping non-diagonal components 0 with dimension  $d = 10$ ; we generate in total 1000 samples each for training and testing. We compare against oracle behavior and two baselines: 1) An oracle baseline (Oracle) that trains only on  $A = 0$  data and trains the rejector to separate the groups with knowledge of group labels and 2) the confidence score baseline (Confidence) that trains a linear model on all the data and then trains a different linear model on all the data where labels are the expert’s agreement with the target and finally compares which of the two is more confident according to the probabilities assigned by the corresponding models and 3) our implementation of the approach in [43] (MixOfExp).

**Results:** We train a multiclass logistic regression model with our loss  $L_{CE}^\alpha$  with  $\alpha \in \{0, 0.5, 1\}$  and record in table 5.1 the difference in accuracy between our method and baselines for the best performing  $\alpha$ . We can see that our method with  $\alpha = 0$  outperforms the confidence baseline by 6.39 on average in classification accuracy and matches the oracle method with 0.22 positive difference which shows the success of our method. When trained with loss  $L_{CE}^1$  or  $L_{CE}^5$  the model matches the confidence baseline, the reason being is that with  $\alpha \neq 0$  the model will still try to fit the target  $Y$  but the model class here is not rich enough to allow the model to reasonably fit the target and adapt to the expert.

Table 5.1: Comparison of our methods with the confidence score baseline, oracle baseline and our implementation of [43] method. We compute a 95% confidence interval for the average difference between the baselines and our method.

DIFFERENCE IN SYSTEM ACCURACY	AVERAGE	95% INTERVAL
$L_{CE}^0$ -CONFIDENCE [54]	6.39	[3.71,9.06]
$L_{CE}^0$ -ORACLE	0.22	[-1.71,2.15]
$L_{CE}^0$ - MIXOFEXP [43]	2.01	[0.14,4.06]



## 5.2 CIFAR-10

As our first real data experimental evaluation we conduct experiments on the celebrated CIFAR-10 image classification dataset [38] consisting of  $32 \times 32$  color images drawn from 10 classes split into 50,000 train and 10,000 test images.

**Synthetic Expert.** We simulate multiple synthetic experts of varying competence in the following way: let  $k \in [10]$ , then if the image belongs to the first  $k$  classes the expert predicts perfectly, otherwise the expert predicts uniformly over all classes. The classifier and expert costs are assumed to be the misclassification costs.

**Base Network.** Our base network for classification will be the Wide Residual Networks (WideResNets) [60] which with data augmentation and hyperparameter tuning can achieve a 96.2% test accuracy. Since our goal is not to achieve better accuracies but to show the merit of our approach for a given fixed model, we disadvantage the model by not using data augmentation and a smaller network size. The WideResNet with 28 layers minimizing the cross-entropy loss achieves 90.47% test accuracy with training until fitting the data in 200 epochs; this will be our benchmark model. We use SGD with momentum and a cosine annealing learning rate schedule.

**Proposed Approach:** Following chapter 4, we parameterize  $h$  and  $r$  (specifically  $g_{\perp}$ ) by a WideResNet with 11 output units where the first 10 units represent  $h$  and the 11<sup>th</sup> unit is  $g_{\perp}$  and minimize the proposed surrogate  $L_{CE}^{\alpha}$  (3.4). We also experimented with having  $h$  be a WideResNet with 10 output units and  $g_{\perp}$  a WideResNet with a single output unit and observed identical results. We show results for  $\alpha \in \{0.5, 1\}$ .

**Baselines:** We compare against three baselines. The first baseline trains the rejector to recognize if the image is in the first  $k$  classes and accordingly defers, we call this baseline "LearnedOracle"; this rejector is a learned implementation of what the optimal rejector should do. The second baseline is the confidence score method [54] and the third is the mixture-of-experts loss of [43], details of the implementation of this final baseline are left to

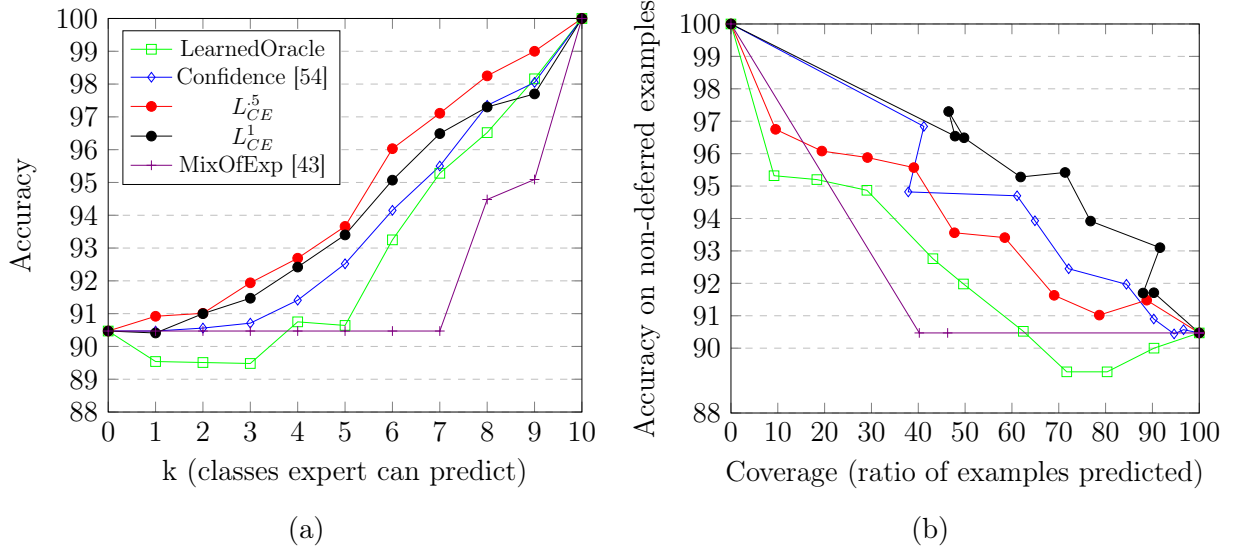


Figure 5-1: Left figure shows overall system accuracy of our method and baselines ( $k$  is the number of classes the expert can predict) and right figure compares the accuracy on the non-deferred examples versus the coverage for every  $k$

#### Appendix A.2.5.

**Results.** In figure 5-5a we plot the accuracy of the combined algorithm and expert system versus  $k$ , the number of classes the expert can predict perfectly. We can see that the model trained with  $L_{CE}^{0.5}$  and  $L_{CE}^1$  outperforms the baselines by 1.01% on average for the confidence score baseline and by 1.94 on average for LearnedOracle. To look more closely at the behavior of our method, we plot in figure 5-1b the accuracy on the non-deferred examples versus the coverage, the fraction of the examples non-deferred, for each  $k$ . We can see that that the model trained with  $L_{CE}^1$  dominates all other baselines giving better coverage and accuracy for the classifier’s predictions. This gives evidence that our loss allows the model to only predict when it is highly confident.

#### Why do we outperform the baselines?

1) *Sample complexity:* The Confidence baseline [54] requires training two networks while ours only requires one, when data is limited our approach gives significant improvements in comparison. We experiment with increasing training set sizes while keeping the test set

fixed and training our model with  $L_{CE}^1$  and the Confidence baseline. Figure 5-2 plots system accuracy versus training set size when training with expert  $k = 5$ . We can see when data is limited our approach massively improves on the baseline, for example with 2000 training points, Confidence achieves 62.33% accuracy while our method achieves 70.12%, a 7.89 point increase.

2) *Taking into consideration both expert and model confidence*: the LearnedOracle baseline ignores model confidence entirely and only focuses on the region where the expert is correct. While this is the behavior of the Bayes classifier in this setup, when dealing with a limited model class and limited data, this no longer is the correct behavior. For this reason, our model outperforms the LearnedOracle baseline.

3) *Consistency*: the mixtures of experts loss of [43] fails in this setup and learns never to defer. The reason is that when training, the loss of the classifier will converge to zero and validation classifier accuracy will still improve in the mean-time, however the loss of the expert remains constant, thus we never defer.




Figure 5-2: Varying training set size when training with expert  $k = 5$  for Confidence baseline and our method  $L_{CE}^1$ .

### 5.3 CIFAR-100

We repeat the experiments described above on the CIFAR-100 dataset [38]. A 28 layer WideResNet achieves a 79.28 % test accuracy when training with data augmentation (random crops and flips). The simulated experts also operate in a similar fashion, for  $k \in \{10, 20, \dots, 100\}$ , if the image is in the first  $k$  classes, the expert predicts the correct label with probability 0.94 to simulate SOTA performance on CIFAR-100 with 93.8% test accuracy [37], otherwise the expert predicts uniformly at random.

Compared against the confidence score baseline, the model trained with  $L_{CE}^1$  outperforms it by a 1.60 difference in test accuracy for  $30 \leq k \leq 90$  on average and otherwise performs on par. This again gives evidence for the efficacy of our method, full experimental results are available in appendix A.2.3.

### 5.4 CIFAR10H and limited expert data

Obtaining expert labels for entire datasets may in fact be prohibitively expensive as standard dataset sizes have grown into million of points [17]. Therefore it is more realistic to expect that the expert has labeled only a fraction of the data. In the following experiments we assume access to fully labeled data  $S_l = \{(x_i, y_i, m_i)\}_{i=1}^m$  and data without expert labels  $S_u = \{(x_i, y_i)\}_{i=m+1}^n$ . The goal again is to learn a classifier  $h$  and rejector  $r$  from the two datasets  $S_l$  and  $S_u$ .

**Data:** To experiment in settings where we have limited expert data, we use the dataset CIFAR10H [53] initially developed to improve model robustness. CIFAR10H contains for each data point in the CIFAR-10 test set fifty crowdworker annotations recorded as counts for each of the 10 classes. The training set of CIFAR-10 will constitute  $S_u$ , and we randomly split the test set in half where one half constitutes  $S_l$  and the other is for testing; we randomize the splitting over 10 trials.

**Expert:** We simulate the behavior of an average human annotator by sampling from the class counts for each data point. The performance of our simulated expert has an average classification accuracy of 95.22 with a standard deviation of 0.18 over 100 runs. The performance of the expert is non uniform over the classes, for example on the class *cat* the expert has 91.0% accuracy while on *horse* a 97.8% accuracy.

**Proposed Approach:** Our method will be to learn  $f_m : \mathcal{X} \rightarrow \{0, 1\}$  to predict whether the expert errs from data  $\tilde{S}_l = \{(x_i, \mathbb{I}_{y_i \neq m_i})\}_{i=1}^m$ , using  $f_m$  we label  $S_u$  with the expert disagreement labels to use in our loss function and obtain  $\hat{S}_u$ . Note since our loss function does not care which label the expert predicts but whether he errs or not, our task simplifies to binary classification instead of classification over the target  $\mathcal{Y}$ . Finally we train using our loss  $L_{CE}$  on  $\hat{S}_u \cup S_l$ ; we refer to our method as " $L_{CE}$  impute"

Table 5.2: Comparing our proposed methods on CIFAR10H and a baseline based on confidence scores recording system accuracy, coverage and classifier accuracy on non-deferred examples.

METHOD	SYSTEM	COVERAGE	CLASSIFIER
$L_{CE}$ IMPUTE	<b>96.29</b> $\pm$ 0.25	51.67 $\pm$ 1.46	<b>99.2</b> $\pm$ 0.08
$L_{CE}$ 2-STEP	96.03 $\pm$ 0.21	60.81 $\pm$ 0.87	98.11 $\pm$ 0.22
CONFIDENCE [54]	95.09 $\pm$ 0.40	<b>79.48</b> $\pm$ 5.93	96.09 $\pm$ 0.42

**Results.** We compare against a confidence score baseline where we train a classifier on  $S_u$  and then model the expert on  $S_l$ . Results are shown in table 5.2 and we can see that our method outperforms the confidence method by 1.2 points on system accuracy and an impressive 3.1 on data points where the classifier has to predict. To show the effect of imputing expert labels on  $S_u$ , we train first our model using  $L_{CE}$  on  $S_u$  and then fine tune to learn deferral on  $S_l$ , we refer to this as " $L_{CE}$  2-step". It is possible that further approaches inspired by SOTA methods in semi supervised learning methods give further improvements [51, 6].

## 5.5 Hate Speech and Offensive Language Detection

We conduct experiments on the dataset created by [15] consisting of 24,783 tweets annotated as hate speech, offensive language or neither. We create a synthetic expert that has differing error rates according to the demographic of the tweet’s author as described in what follows.

**Expert.** [7] developed a probabilistic language model that can identify if a tweet is in African-American English (AAE), this model was used by [14] to audit for racial bias in classifiers. We use the same model and predict that a tweet is in AAE if the probability predicted is higher than 0.5. Our expert model is as follows: if the tweet is in AAE then with probability  $p$  we predict the correct label and otherwise predict uniformly at random. On the other hand if the tweet is not in AAE, we predict with probability  $q$  the correct label. We experiment with 3 different expert probabilities for  $p$  and  $q$ : 1) a fair expert with  $\{p = 0.9, q = 0.9\}$ , 2) a biased expert towards AAE tweets  $\{p = 0.75, q = 0.9\}$  and 3) a biased expert towards non AAE tweets  $\{p = 0.9, q = 0.75\}$ .

**Our Approach.** For our model we use the CNN developed in [36] for text classification with 100 dimensional Glove embeddings [52] and 300 filters of sizes  $\{3, 4, 5\}$  using dropout. This CNN achieves a 89.5% average accuracy on the classification task, comparable to the 91% achieved by [15] with a feature heavy linear model. We randomly split the dataset with a 60, 10, 30% split into a training, validation and test set respectively; we repeat the experiments for 5 random splits. We used a grid search over the validation set to find  $\alpha$ .

**Results.** We compare against two baselines: the first is Confidence, the second is an oracle baseline that trains first a model on the classification task and then implements the Bayes rejector  $r^B(x)$  equipped with the knowledge of  $p, q$  and the tweet’s demographic group. Both our model trained with  $L_{CE}^1$  and the confidence score baseline achieve similar accuracy and coverage with the oracle baseline performing only slightly better across the three experts. For the AAE biased expert, our model trained with  $L_{CE}^1$  achieves  $92.91 \pm 0.17$  system accuracy, Confidence  $92.42 \pm 0.40$  and Oracle  $93.22 \pm 0.11$ . This suggests that both approaches are

performing optimally in this setting.

**Racial Bias.** A major concern in this setting is whether the end to end system consisting of the classifier and expert is discriminatory. We define the discrimination of a predictor as the difference in the false positive rates of AAE tweets versus non AAE tweets where false positives indicate tweets that were flagged as hate speech or offensive when they were not. Surprisingly, the confidence score baseline with the fair expert doubles the discrimination of the overall system compared to the classifier acting on it’s own: the classifier has a discrimination of 0.226 on all the test data, the fair expert a discrimination of 0.03 while the confidence score baseline has a discrimination of 0.449. This again reiterates the established fact that fairness does not compose [20]. In fact, the end-to-end system can be less discriminatory even if the individual components are more discriminatory, for the second expert that has higher error rates on non AAE tweets with discrimination of 0.084, the discrimination of the confidence score method reduces to 0.151. While our method does not achieve significantly lower discrimination than the baseline, however integrating fairness constraints for the end-to-end system becomes easier as we can adapt the classifier. Complete experimental results can be found in Appendix A.2.4.

## 5.6 Synthetic Experts on CheXpert

**Task.** CheXpert is a large chest radiograph dataset that contains over 224 thousand images of 65,240 patients automatically labeled for the presence of 14 observations using radiology reports [31]. In addition to the automatically labeled training set, [31] make publicly accessible a validation set of 200 patients labeled by a consensus of 3 radiologists and hide a further testing set of 500 patients labeled by 8 radiologists. We focus here on the detection of only the 5 observations that make up the "competition tasks" [31]: Atelectasis, Cardiomegaly, Consolidation, Edema, and Pleural Effusion. This is a multi-task problem, we have 5 separate binary tasks, we will learn to defer on an individual task basis.

**Expert.** We create a simulated expert as follows: if the chest X-ray contains support devices (the presence of support devices is part of the label) then the expert is correct with probability  $p$  on all tasks independently and if the X-ray does not contain support devices, then the expert is correct with probability  $q$ . We vary  $q \in \{0.5, 0.7\}$  and  $p \in \{0.7, 0.8, 0.9, 1\}$  to obtain different experts, we let  $p \geq q$  as one can think that a patient that has support devices might have a previous medical history that the expert is aware of and can use as side-information.

**Data.** We use the downsampled resolution version of CheXpert [31] and split the training data set with an 80-10-10 split on a patient basis for training, validation and testing respectively, no patients are shared among the splits. Images are normalized and resized to be compatible with pre-trained ImageNet models, we use data augmentation in the form of random resized crops, horizontal flips and random rotations of up to 15 degrees while training. Note that a small subset of the training data has an uncertainty label "U" instead of a binary label that implies that the automatic annotator is uncertain, we ignore these points on a task basis while training and testing.

**Baselines.** We implement two baselines: a threshold confidence baseline that learns a threshold to maximize system AU-ROC on just the confidence of the classifier model to defer (ModelConfidence), this is the post-hoc thresholding method in [43], and the Confidence baseline [54]. We use temperature scaling [29] to ensure calibration of all baselines on the validation set.

**Model.** Following [31], we use the DenseNet121 architecture for our model with pre-trained weights on ImageNet, the loss for the baseline models is the average of the binary cross entropy for each of the tasks. We train the baseline models using Adam for 4 epochs. For our approach we train for 3 epochs using the cross entropy loss and then train for one epoch using  $L_{CE}^\alpha$  with  $\alpha$  chosen to maximize the area under the receiver operating characteristic curve (AU-ROC) of the combined system on the validation set for each of the 5 tasks (each task is treated separately). We also observe similar results if we train for the first three epochs with  $L_{CE}^1$  and then train for one epoch with a validated choice of  $\alpha$ .



**Experimental setup.** In a clinical setting there might be a cost associated to querying a radiologist, this then imposes a constraint on how often we can query the radiologist i.e. our model’s coverage (fraction of examples where algorithm predicts). We constrain our method and the baselines to achieve  $c\%$  coverage for  $c \in [100]$  to simulate the spectrum between complete automation and none.


We achieve this for our method by first sorting the test set based on  $g_{\perp}(x) - \max(g_0(x), g_1(x)) := q(x)$  across all patients  $x$  in the test set, then to achieve coverage  $c$ , we define  $\tau = q(x_c)$  where  $q(x_c)$  is the  $c$ ’th percentile of the outputs  $q(x)$ , then we let  $r(x) = 1 \iff q(x) \geq \tau$ . The definition of  $\tau$  ensures that we obtain exactly  $c\%$  coverage.

For ModelConfidence we achieve this by letting  $q(x) = 1 - \max(g_0(x), g_1(x))$  ( $g$  is the result of a separate trained model than the one for our method), this is the natural classifier’s probability of error from the softmax output, and for the Confidence we let  $q(x)$  be the difference between the radiologists confidence and the classifier’s confidence.

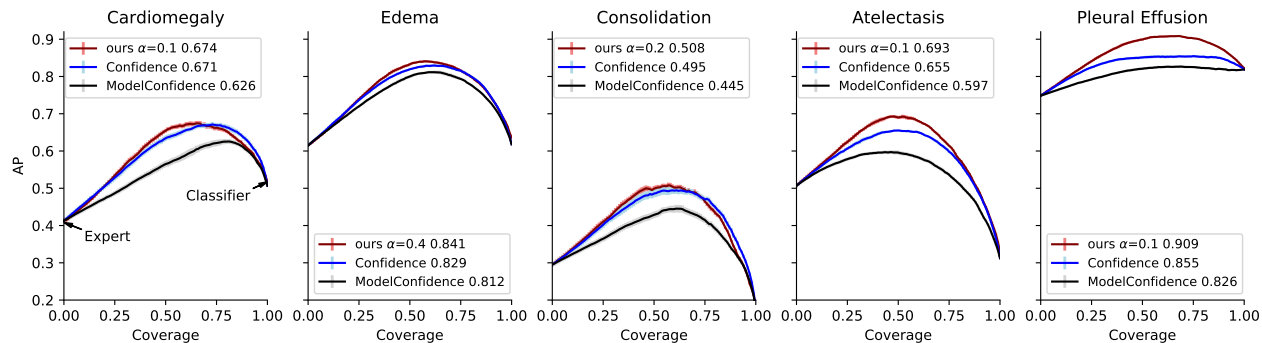
**Results.** In Figure 5-3a we plot the overall system (expert and algorithm combined) AU-ROC for each desired coverage for the methods and in Figure 5-3b we plot the overall system area under the precision-recall curve (AU-PR) versus the coverage; this is for the expert with  $q = 0.7$  and  $p = 1$ . We can see that the curve for our method dominates the baselines over the entire coverage range for both AU-ROC and AU-PR, moreover the curves are concave and we can achieve higher performance by combining expert and algorithm than using both separately. Our method is able to achieve a higher maximum AU-ROC and AU-PR than both baselines: the difference between the maximum attainable AU-ROC of our method and Confidence is 0.043, 0.029, 0.016, 0.022 and 0.025 respectively for each of the five tasks. There is a clear hierarchy between the 3 compared methods: our method dominates Confidence and Confidence in turn dominates ModelConfidence, in fact ModelConfidence is a special case of the Confidence baseline, since the expert does not have uniform performance over the domain there are clear gains in modeling the expert.

This hierarchy continues to hold as we change the expert behavior as we vary the probabilities

$p$  and  $q$ , in Table 5.3 we show for each of the 5 tasks the difference between the average AU-ROC across all coverages (average value of the curves shown in Figure 5-3a) for our method and the Confidence baseline for different expert probabilities and the difference between the maximum achievable AU-ROC. A positive average difference serves to show the degree of dominance of our method over the Confidence baseline, note that the difference alone cannot imply dominance of the curves however dominance is still observed. Our method improves on the baselines as the difference between  $q$  and  $p$  increases, this difference encodes the non-uniformity of the expert behavior over the domain.



(a) AU-ROC vs coverage for expert  $q = 0.7, p = 1$ , maximum AU-ROC is noted.



(b) AU-PR vs coverage for expert  $q = 0.7, p = 1$ , maximum AU-PR is noted.

Figure 5-3: Plot of AU-ROC of the ROC curve (a) for each level of coverage (0 coverage means only the expert predicting and 1 coverage is only the classifier predicting) and of the area under the precision-recall curve (AU-PR) (b) for each of the 5 tasks comparing our method with the baselines on the training derived test set for the toy expert with  $q = 0.7, p = 1$ . We report the maximum AU-ROC and AU-PR achieved on each task, error bars are standard deviations derived from 10 runs (averaging over the expert's randomness).

Table 5.3: Average difference in AU-ROC across all coverage and difference between maximum achievable AU-ROC between our method and the Confidence baseline for each of the 5 tasks and different toy expert probabilities  $p$  and  $q$ ; each entry is (average difference  $\pm$  standard deviation; difference of maximums). The difference between our method and the ModelConfidence is roughly twice the values noted in table 5.3, only at Expert (0.7, 0.7) does Confidence and ModelConfidence achieve the same performance since the expert has uniform error over the domain.

EXPERT ( $p, q$ )	CARDIOMEGALY	EDEMA	CONSOLIDATION	ATELECTASIS	PLEURAL EFFUSION
(0.5,0.7)	0.032 $\pm$ 0.024; 0.002	0.015 $\pm$ 0.012; 0.007	0.015 $\pm$ 0.008; 0.007	0.017 $\pm$ 0.009; 0.007	0.007 $\pm$ 0.003 ;0.007
(0.5,0.9)	0.032 $\pm$ 0.017; 0.014	0.026 $\pm$ 0.016; 0.024	0.010 $\pm$ 0.005; 0.015	0.016 $\pm$ 0.008; 0.026	0.012 $\pm$ 0.010; 0.004
(0.5,1)	0.022 $\pm$ 0.012; 0.029	0.013 $\pm$ 0.009; 0.019	0.007 $\pm$ 0.008; 0.012	0.013 $\pm$ 0.006; 0.020	0.010 $\pm$ 0.008; 0.012
(0.7,0.7)	0.024 $\pm$ 0.018; 0.005	0.011 $\pm$ 0.009; 0.010	0.011 $\pm$ 0.010; 0.009	0.006 $\pm$ 0.006; 0.008	0.001 $\pm$ 0.001; 0.003
(0.7,0.9)	0.032 $\pm$ 0.020; 0.024	0.010 $\pm$ 0.007; 0.010	0.007 $\pm$ 0.007; 0.017	0.014 $\pm$ 0.008; 0.017	0.010 $\pm$ 0.006; 0.006
(0.7,1)	0.027 $\pm$ 0.014; 0.042	0.016 $\pm$ 0.010; 0.027	0.007 $\pm$ 0.007; 0.019	0.013 $\pm$ 0.007; 0.022	0.014 $\pm$ 0.010; 0.027
(0.8,1)	0.017 $\pm$ 0.009; 0.023	0.011 $\pm$ 0.008; 0.012	0.001 $\pm$ 0.004; 0.007	0.012 $\pm$ 0.006; 0.009	0.010 $\pm$ 0.006; 0.018

**Sample Complexity** Training data for chest X-rays is a valuable resource that may not be abundantly available when trying to deploy a machine learning model in a new clinical setting where for example the imaging mechanism may differ. It is important to see the effectiveness of the proposed approaches when training data size is limited, this furthermore helps us understand the comparative sample complexity of our method versus the baselines.

**Experimental details.** We restrict the training data size for our model and baselines while keeping the same validation and testing data as previously; the validation data is used only for calibration of models and optimizing over choice of  $\alpha$ . We train using the same procedure as before and report the maximum achievable AU-PR and AU-ROC. The expert we defer to is the synthetic expert described above with  $q = 0.7$  and  $p = 1$ .

**Results.** In Figure 5-4 we plot the average of the maximum achievable AU-ROC 5-4a and AU-PR 5-4b across the 5 tasks for the different methods as we vary the the training set size. We observe that our method consistently outperforms the baselines and continues to take advantage of further data as the baselines performance starts to saturate. If we look at the AU-ROC and AU-PR of the expert on deferred examples, we observe negligible differences as the training set size increases for each method, however if we look at classifier performance on the non-deferred examples, we start to observe a significant difference in AU-ROC and

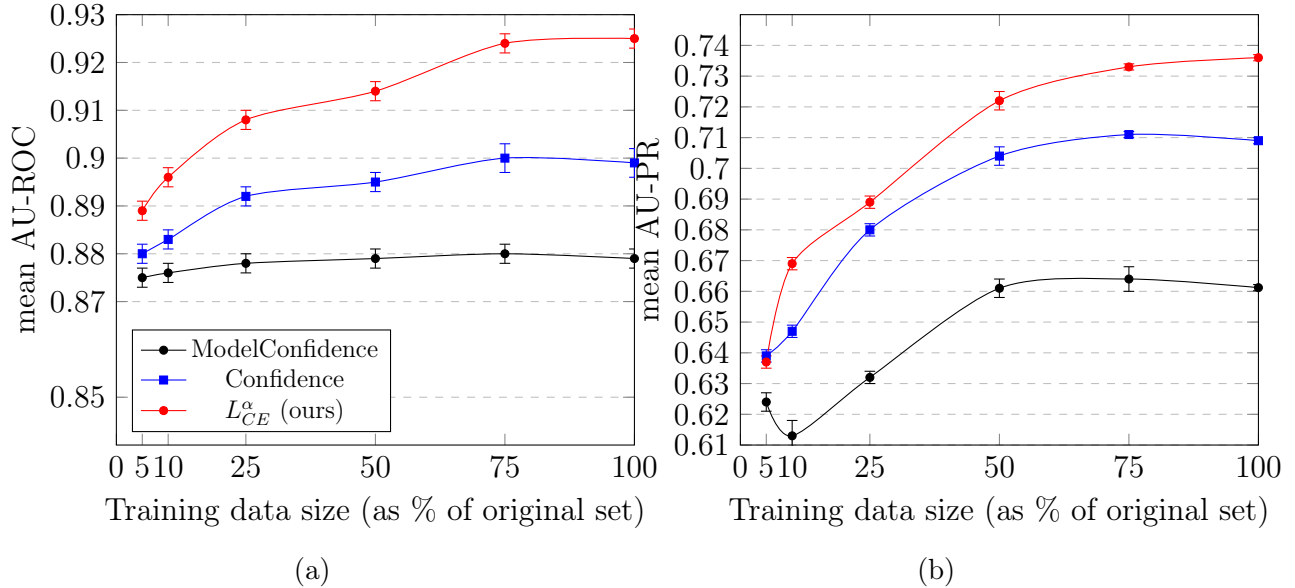


Figure 5-4: Left figure shows the average of the maximum achievable AU-ROC for the 5 tasks (average over the tasks) when the changing the size of the training data (as a % of the original set) and right figure shows the same for AU-PR

AU-PR for our method while the baselines lag behind. In Figure A-2 (found in Appendix A.2.6) we plot the classifier AU-ROC on non-deferred examples versus the coverage for each of the 5 tasks, we can see for example on Cardiomegaly, our method at full training data obtains an AU-ROC that is at least 0.2 points greater than that of ModelConfidence at coverage levels less than 50%. One expects ModelConfidence to achieve the best performance when looking at non-deferred examples, and this in fact is true when we look at accuracy, however for AU-ROC, what happens is that the ModelConfidence baseline never defers on negative predicted examples due to the class imbalance which makes the model very confident in it's negative predictions. Thus, any positive labeled example that the model mistakenly labels as negative with high confidence will cause the AU-ROC to be reduced at low coverage levels. This also allows us to see that our method, and to an extent the Confidence baseline, make very different deferral decisions that factor in the expert.

**Impact of input noise** In our previous experimental setup, the input domain of the classifier  $\mathcal{X}$ , the chest X-ray, is assumed to be sufficient to perfectly predict the label  $Y$  as



(a) Original X-ray



(b) X-ray with hidden left part

Figure 5-5: Left figure (a) shows the chest X-ray of a patient with Cardiomegaly, the right figure (b) shows that same X-ray but now with the left part hidden which is used as input to the models.

our golden standard is the prediction of expert radiologists from just looking at the X-ray. Therefore, given enough training data and a sufficiently rich model class, a learned classifier from  $\mathcal{X}$  will be able to perfectly predict the target and won't need to defer to any expert to achieve better performance. In this set of experiments, we perform two studies: the first we hide the left part of the chest X-ray on both training and testing examples to obtain a new input domain  $\tilde{\mathcal{X}}$ . This now limits the power of any learned classifier even in the infinite data regime as the left part of the X-ray may hide crucial parts of the input. Figure 5-5 shows this noise applied to a patient's X-ray, the size of the rectangular region was chosen to cover one side of the chest area, we later experiment with varying the scale of the noise. In the second experiment, we train with the original chest X-rays but evaluate with noisy X-rays with noise in the form of erasing a randomly placed rectangular region of the X-ray. This second experiment is meant to illustrate the robustness of the different methods to input noise.

**Experimental details.** The noise in the second set of experiments consists of a 2:1 (height:width) randomly located rectangular region of scale (area) that we vary from 0.1 to 0.66. The expert is the synthetic expert model with  $q = 0.7$  and  $p = 1$ .

**Results.** In Figure 5-7 we plot the AU-ROC and AU-PR of the different methods as we vary coverage when training and testing while hiding the left section of the X-rays. We can first observe that the maximum achievable performance for the different methods is significantly reduced, however the gap between the different methods is still observed. In Figure 5-6 we plot the average maximum AU-ROC and AU-PR across the 5 tasks as we vary the area of the rectangular region. While the performance of all the methods degrade with the scale of the noise, the gap between the methods remains constant in terms of AU-PR but diminishes in terms of AU-ROC as the performance of the baselines remains steady.






Figure 5-6: Left figure shows the average of the maximum achievable AU-ROC for the 5 tasks (average over the tasks) when the changing the scale of the noise (size of rectangular region) and right figure shows the same for AU-PR.



(a) AU-ROC vs coverage when hiding left part of X-ray.



(b) AU-PR vs coverage when hiding left part of X-ray.

Figure 5-7: Plots of AU-ROC and AU-PR as we vary coverage when training and testing with chest X-rays that have their left section hidden. The expert model is  $q = 0.7$  and  $p = 1$ .





# Chapter 6

## Conclusion

In this thesis, we explored a framework where the learning model can choose to defer to an expert or predict. We analyzed the framework theoretically and proposed a novel surrogate loss via a reduction to multiclass cost sensitive learning. Through experiments on image and text classification tasks, we showcased that our approach not only achieves better accuracy than confidence score baselines but does so with better sample complexity and computational cost. We hope that our method will inspire machine learning practitioners to integrate downstream decision makers into their learning algorithms. Future work will explore how to defer in settings where we have limited expert data, learning from biased expert data and learning to defer to multiple experts simultaneously.

## Acknowledgements

This work was supported by NSF CAREER award #1350965.



# Appendix A

## Appendix

### A.1 Practitioner’s guide to our approach

#### A.1.1 General implementation

Given a dataset of tuples  $S = \{(x_i, y_i, m_i)\}_{i=1}^n$  where  $x_i$  represents the covariates,  $y_i$  is the target and  $m_i$  are the expert labels, we want to construct a classifier  $h : \mathcal{X} \rightarrow \mathcal{Y}$  and rejector function  $r : \mathcal{X} \rightarrow \{-1, 1\}$ . Our method for predicting on a new example  $x \in \mathcal{X}$  given expert context  $z \in \mathcal{Z}$  that only the expert can observe, a function class  $\mathcal{H}$  where  $h \in \mathcal{H} : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|+1}$  (an example would be the set of deep networks with  $|\mathcal{Y}| + 1$  output units) , and an expert  $M : \mathcal{Z} \rightarrow \mathcal{Y}$  is summarized below in Algorithm 1.

---

**Algorithm 1:** Our proposed method for prediction on a new example  $x \in \mathcal{X}$  with expert input  $z \in \mathcal{Z}$

---

**Input:** training data  $S = \{(x_i, y_i, m_i)\}_{i=1}^n$ , function class  $\mathcal{H}$ , example  $x$ , Expert  $M$

and expert input  $z$

$g_1, \dots, g_{|\mathcal{Y}|}, g_{\perp} \leftarrow \arg \min_{\mathbf{g} \in \mathcal{H}} \sum_{i \in S} L_{CE}^{\alpha}(\mathbf{g}, x_i, y_i, m_i)$

prediction = 0

$r(x) \leftarrow \text{sign}(-\max_{y \in \mathcal{Y}} g_y(x) + g_{\perp}(x))$

**if**  $r(x) = 0$  **then**

$h(x) \leftarrow \arg \max_{y \in \mathcal{Y}} g_y(x)$   
    prediction  $\leftarrow h(x)$

**else**

$m \leftarrow M(z)$  (expert query)  
    prediction  $\leftarrow m$

**end**

**Return:** prediction

---

Where the loss  $L_{CE}^{\alpha}$  used in algorithm is the following:

$$L_{CE}^{\alpha}(h, r, x, y, m) = -(\alpha \cdot \mathbb{I}_{m=y} + \mathbb{I}_{m \neq y}) \log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \\ - \mathbb{I}_{m=y} \log \left( \frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right)$$

Practically, integrating an expert decision maker into a machine learning model amounts to two modifications in training: increasing the output size of the function class in consideration by an additional output unit representing deferral and training with the loss  $L_{CE}^{\alpha}$  instead of the cross entropy loss. We show how to implement  $L_{CE}^{\alpha}$  in PyTorch below:

```
def deferral_loss_L_CE(outputs, target, expert, k_classes, alpha):
    , , ,

    outputs: model outputs
    target: target labels
```

```

expert: expert agreement labels for batch
k_classes: cardinality of target Y
'''
batch_size = outputs.size()[0]
defer_position =
outputs = torch.nn.functional.softmax(outputs, dim=1)
loss = -expert*torch.log2(outputs[range(batch_size), k_classes])
      - (alpha*expert + (1-expert)) *
      torch.log2(outputs[range(batch_size), labels])
return torch.sum(loss)/batch_size

```

### A.1.2 Choice of $\alpha$

The choice of the hyperparameter  $\alpha$  has sizable influence on system performance. Naive validation over  $\alpha$  requires re-training on the entire training set from scratch over the search space. We find that a simple validation strategy often works as well as re-training from scratch especially in scenarios where there is little gain in adapting to the expert but there are major gains in being able to defer correctly.

The strategy first requires splitting the training set into two sets  $S_{T_1}$  and  $S_{T_2}$  where  $S_{T_1}$  is larger than  $S_{T_2}$  (e.g. an 80-20 split), access to a validation set  $S_V$  and a set of possible values  $\mathcal{A}$  for  $\alpha$  (an evenly spaced grid over  $[0, 10]$  is more than sufficient). The strategy then proceeds in two steps:

- **Step 1:** Train on  $S_{T_1}$  with  $L_{CE}^1$  (i.e. setting  $\alpha = 1$ ) to maximize system performance on  $S_V$ . One may find more success instead training on  $S_{T_1}$  with the *cross-entropy loss* (however with the model having an extra output) to maximize *classifier* performance on  $S_V$  rather than system performance. Call the resulting model of this first step  $M_1$

- **Step 2:** For each  $\alpha \in \mathcal{A}$ , fine-tune on  $S_{T_2}$  starting from model  $M_1$  to maximize system performance on  $S_V$  measuring it with the rejector  $r(x) = \mathbb{I}\{-\max_{y \in \mathcal{Y}} g_y(x) + g_\perp(x) \geq \tau\}$  where the threshold  $\tau$  is chosen to maximize performance on  $S_V$  post-hoc. The resulting model  $M'_1$  and  $\tau^*$  that obtains best system performance across all choices of  $\alpha$  and choices of  $\tau$  is the final model.
- **Inference time:** Use the rejector defined by  $r(x) = \mathbb{I}\{-\max_{y \in \mathcal{Y}} g_y(x) + g_\perp(x) \geq \tau^*\}$  and proceed as in Algorithm 1.

*Note* that system performance here refers to metrics measured with respect to the machine+expert system with deferral while classifier performance refers to metrics measured as if the system never deferred.

## A.2 Experimental Details and Results

All experiments were run on a Linux system with an NVIDIA T4 GPU on PyTorch 1.4.0.

### A.2.1 CIFAR-10

**Implementation Details.** We employ the implementation in <https://github.com/xternalz/WideResNet-pytorch> for the Wide Residual Networks. To train, we run SGD with an initial learning rate of 0.1, Nesterov momentum at 0.9 and weight decay of 5e-4 with a cosine annealing learning rate schedule [41]. We train for a total of 200 epochs for all experiments, at this point the network has perfectly fit the training set, we found that early stopping based on a validation set did not make any difference and similarly training for more than 200 epochs also did not hurt test accuracy.

**Expert Accuracy.** In Table A.1 we show the accuracy of the expert on the deferred examples versus the classes the expert can predict  $k$ . We can see that our method  $L_{CE}^5$  has higher expert accuracy than all other baselines except at  $k = 1, 2$  where coverage is very high. This contrasts with Figure 5-1b that shows the classifier accuracy on non-deferred accuracy where  $L_{CE}^5$  had lower accuracy for each expert level compared to Confidence and  $L_{CE}^1$ . Hence there is a clear trade-off between choosing the hyper-parameter  $\alpha < 1$  and  $\alpha = 1$ . For  $\alpha < 1$ , the model will prefer to always defer to the expert if it is correct, this is advantageous in this setup as the expert is perfect on a subset of the data and uniformly random on the other. However, for  $\alpha = 1$ , the model will compare the confidence of the expert and the model essentially performing the computation of the Bayes rejector  $r^B$  as shown by the consistency of the loss  $L_{CE}^1$ ; note that for  $\alpha \neq 1$  the loss  $L_{CE}$  is no longer consistent.

Table A.1: Accuracy of the expert on deferred examples shown for the methods and baselines proposed with varying expert competence ( $k$ ) on CIFAR-10.

METHOD / EXPERT (K)	1	2	3	4	5	6	7	8	9	10
$L_{CE}^1$	73.65	86.01	73.66	87.41	88.81	94.7	96.67	98.72	98.65	100
$L_{CE}^5$	86.44	90.96	<b>92.65</b>	<b>91.67</b>	<b>93.71</b>	<b>96.32</b>	<b>97.61</b>	98.77	<b>99.24</b>	<b>100</b>
CONFIDENCE	<b>87.5</b>	<b>92.74</b>	88.88	88.3	92.8	94.56	96.76	<b>98.89</b>	98.89	100
ORACLEREJECT	85.3	90.49	88.23	91.13	89.33	93.61	95.45	96.82	98.45	100

**Increasing data size.** In table A.2 we show the accuracy of the classifier and the coverage of the system for our method compared to the baseline Confidence for expert  $k = 5$ . We can see that when data is limited, our method retains high classification accuracy for the classifier versus the baseline. This is due in fact to the low coverage of our method compared to Confidence, as data size grows the coverage our method increases as now the classifier’s performance improves and the system can now safely defer to it more often. On the other hand, the baseline remains at almost constant coverage, not adapting to growing data sizes.

Table A.2: Accuracy of the classifier on non-deferred examples shown for our method  $L_{CE}^1$  and baseline Confidence with varying training set size for expert  $k = 5$  on CIFAR-10.

METHOD / DATA SIZE (THOUSANDS)	1	2	3	5	8	10	20	50
$L_{CE}^1$ (CLASSIFIER)	<b>62.84</b>	<b>71.51</b>	<b>72.63</b>	<b>75.03</b>	80.1	<b>82.11</b>	86.44	<b>95.42</b>
CONFIDENCE (CLASSIFIER)	50.31	59	66.3	70.12	<b>80.33</b>	78.67	<b>87.01</b>	92.45
$L_{CE}^1$ (COVERAGE)	25.7	35.87	40.42	49.62	46.38	46.51	50	71.35
CONFIDENCE (COVERAGE)	<b>69.32</b>	<b>72.93</b>	<b>71.99</b>	<b>75.05</b>	<b>73.09</b>	<b>65.9</b>	<b>74.16</b>	<b>72.12</b>

## A.2.2 CIFAR-10H

**Class-wise Accuracy of Expert.** Table A.3 shows the average accuracy of the synthetic CIFAR10H [53] expert on each of the 10 classes. We can see that the expert has very different accuracies for the classes which gives an opportunity for an improvement.

**Results.** Table A.4 shows full experimental results for the CIFAR-10H results.

Table A.3: Accuracy of the CIFAR10H [53] expert on each of the 10 classes

CLASS	1	2	3	4	5	6	7	8	9	10
ACCURACY	95.15	97.23	94.75	91.58	90.51	94.90	96.22	97.91	97.33	96.74

Table A.4: Complete results of table 5.2 comparing our proposed approaches and baseline.

METHOD	SYSTEM ACCURACY	COVERAGE	CLASSIFIER ACCURACY	EXPERT ACCURACY
$L_{CE}$ IMPUTE	<b>96.29</b> ±0.25	51.67±1.46	<b>99.2</b> ± 0.08	<b>93.18</b> ± 0.48
$L_{CE}$ 2-STEP	96.03±0.21	60.81±0.87	98.11 ± 0.22	92.77 ± 0.58
CONFIDENCE	95.09±0.40	<b>79.48</b> ±5.93	96.09 ± 0.42	90.94 ± 1.34



### A.2.3 CIFAR-100

**Results.** In figure A-1 we plot the accuracy of the combined algorithm and expert system versus  $k$ , the number of classes the expert can predict. We can see that our method dominates the baseline over all  $k$ . In table A.5 we show expert, classifier and system accuracy along with coverage of both methods. Our approach  $L_{CE}^1$  obtains both better expert and classifier accuracy however gets lower coverage than Confidence.




Figure A-1: Comparison of the developed method  $L_{CE}^1$  on CIFAR-100 versus the confidence baseline.  $k$  is the number of classes the expert can predict

Table A.5: Accuracy of the expert on deferred examples shown for the methods and baselines proposed with varying expert competence (k) on CIFAR-100.

METHOD / EXPERT (K)	10	20	30	40	50	60	70	80	90	100
$L_{CE}^1$ (SYSTEM)	<b>78.67</b>	<b>79.43</b>	<b>81.02</b>	<b>82.09</b>	<b>83.8</b>	<b>85.15</b>	<b>87.58</b>	<b>90.23</b>	<b>91.81</b>	94.59
CONFIDENCE (SYSTEM)	78.48	79.37	79.67	80.75	81.92	83.67	85.15	88.63	90.31	<b>94.74</b>
$L_{CE}^1$ (COVERAGE)	89.19	82.44	84.79	71.66	74.52	65.72	62.23	59.37	52.15	49.07
CONFIDENCE (COVERAGE)	<b>99.17</b>	<b>95.47</b>	<b>93.96</b>	<b>86.64</b>	<b>86.71</b>	<b>80.67</b>	<b>79.56</b>	<b>75.36</b>	<b>72.39</b>	<b>63.32</b>
$L_{CE}^1$ (CLASSIFIER)	<b>82.35</b>	<b>84.03</b>	<b>84.07</b>	<b>85.29</b>	<b>86.44</b>	<b>87.78</b>	<b>90.13</b>	<b>91.89</b>	<b>92.4</b>	94.59
CONFIDENCE (CLASSIFIER)	78.99	80.66	81.79	84.75	84.62	87.30	88.75	90.97	92.07	<b>94.97</b>
$L_{CE}^1$ (EXPERT)	<b>47.36</b>	<b>57.8</b>	<b>68.87</b>	<b>73.99</b>	<b>76.06</b>	<b>79.65</b>	<b>83.37</b>	<b>87.79</b>	<b>91.16</b>	<b>94.57</b>
CONFIDENCE (EXPERT)	18.07	52.09	51.49	54.79	64.4	68.55	71.13	82.11	85.70	94.30

## A.2.4 Hate Speech experiments

**Implementation details.** We train all models with Adam for 15 epochs and select the best performing model on the validation set.

**Results.** Table A.6 shows complete results of our method, baselines, expert and classifier. The performance of our method and the baselines all achieve comparable results.

Table A.6: Detailed results for our method and baselines on the hate speech detection task [15]. sys: system accuracy, class: classifier accuracy, disc: system discrimination, AAE-biased: Expert 2 that has higher error rate for AAE group, non-AAE biased: Expert 3 that has higher error for non AAE tweets

METHOD/EXPERT	FAIR			AAE-BIASED		
	SYS	CLASS	DISC	SYS	CLASS	DISC
$L_{CE}^1$ (OURS)	93.36 $\pm$ 0.16	<b>95.60</b> $\pm$ 0.44	<b>0.294</b> $\pm$ 0.03	92.91 $\pm$ 0.17	<b>94.67</b> $\pm$ 0.61	<b>0.37</b> $\pm$ 0.06
CONFIDENCE	93.22 $\pm$ 0.11	94.49 $\pm$ 0.12	0.45 $\pm$ 0.02	92.42 $\pm$ 0.40	94.56 $\pm$ 0.40	0.41 $\pm$ 0.02
ORACLE	<b>93.57</b> $\pm$ 0.11	94.87 $\pm$ 0.22	0.32 $\pm$ 0.02	<b>93.22</b> $\pm$ 0.11	94.49 $\pm$ 0.12	0.449 $\pm$ 0.024
EXPERT	89.76	–	0.031	84.28	–	0.071
CLASSIFIER	88.26	88.26	0.226	88.26	88.26	0.226

METHOD/EXPERT	NON-AAE BIASED		
	SYS	CLASS	DISC
$L_{CE}^1$ (OURS)	90.42 $\pm$ 0.38	<b>94.04</b> $\pm$ 0.81	0.231 $\pm$ 0.04
CONFIDENCE	90.60 $\pm$ 0.13	93.68 $\pm$ 0.24	0.15 $\pm$ 0.03
ORACLE	<b>91.09</b> $\pm$ 0.12	92.57 $\pm$ 0.15	<b>0.15</b> $\pm$ 0.02
EXPERT	80.4	–	0.084
CLASSIFIER	88.26	88.26	0.226

## A.2.5 Baseline Implementation

**Description of [43] approach.** A different approach to our method, is to try directly to approximate the system loss (2.1), this was the road taken by [43] in their differentiable model method. Let us introduce the loss used in [43]:

$$L(h, r, M) = \mathbb{E}_{(x,y) \sim \mathbf{P}, m \sim M|(x,y)} [(1 - r(x, h(x)))l(y, h(x)) + r(x, h(x))l(y, m)] \quad (\text{A.1})$$

where  $h : \mathcal{X} \rightarrow \Delta^{|\mathcal{Y}|-1}$  (classifier),  $r : \mathcal{X} \times \Delta^{|\mathcal{Y}|-1} \rightarrow \{0, 1\}$  (rejector) and the expert  $M : \mathcal{Z} \rightarrow \Delta^{|\mathcal{Y}|-1}$ . [43] considers only binary labels and uses the logistic loss for  $l(., .)$  and

thus requires the expert to produce uncertainty estimates for its predictions instead of only a label; we can extend this to the multiclass setting by using the cross entropy loss for  $l$ . It is clear that the loss (A.1) is non-convex in  $r$ , hence to optimize it [43] estimates the gradient through the Concrete relaxation [42, 32]. However, in the code of [43] found at <https://github.com/dmadras/predict-responsibly>, the authors replace  $r(x)$  by its estimated probability from its model. [43] considers an additional parameter  $\gamma_{defer}$  found in the code, however it is not clear what effect this parameter has as we found its description in the paper did not match the code. In detail, let  $r_0, r_1 : \mathcal{X} \rightarrow \mathbb{R}$  and  $r(x) = \arg \max_{i \in \{0,1\}} r_i$ , the loss [43] considers is:

$$\tilde{L}(h, r, M) = \mathbb{E}_{(x,y) \sim \mathbf{P}, m \sim M|(x,y)} \left[ \frac{\exp(r_0(x))}{\exp(r_0(x)) + \exp(r_1(x))} l(y, h(x)) + \frac{\exp(r_1(x))}{\exp(r_0(x)) + \exp(r_1(x))} l(y, m) \right] \quad (\text{A.2})$$

All terms in loss (A.2) are on the same scale which is crucial for the model to train well. We explicitly have two functions  $r_0$  and  $r_1$  defining  $r$  even though  $r$  is binary; this is for ease of implementation.

Another key detail of [43] approach, is that the classifier is independently trained of the rejector by stopping the gradient from  $r$  to backpropagate through  $h$ . This no longer allows  $h$  to adapt to the expert,  $h$  is trained with the cross entropy loss on its own concurrently with  $r$ .

**CIFAR-10 details.** In our CIFAR-10 setup, the dataset  $S$  contains only the final prediction  $m$  of the expert  $M$ , thus to compute  $l(y, m)$  we set  $l(y, m) = -\log(1 - \epsilon)$  if  $y = m$  and  $l(y, m) = -\log(\frac{1}{|y|})$  if  $y \neq m$  (simulating a uniform prediction in accordance with our expert behavior) with  $\epsilon = 10^{-12}$ . One could instead train a network to model the expert’s prediction, we found this approach to fail as there is a big amount of noise in the labels caused by the expert’s random behavior.

**Results on CIFAR-10.** For expert  $k < 8$ , we found that the [43] baseline to almost never defer to the expert and when  $k = 8, 9$  at the end of training (200 epochs) the rejector never defers but the optimal system is found in the middle of training ( $\sim 100$  epochs). The optimal


systems achieve 46.27 and 40.22 coverage, 98.81 and 98.89 expert accuracy on deferred examples and 89.38 and 89.40 classifier accuracy on non-deferred examples respectively for  $k = 8, 9$ . The classifier alone for the optimal systems achieve  $\sim 86$  classification accuracy on all of the validation set for both experts, notice that there is not much difference between the classification accuracy on all the data and non-deferred examples, while for our method and other baselines there is a considerable increase. This indicates that the rejector is only looking at the expert loss and ignoring the classifier

What is causing this behavior is that as the classifier  $h$  trains, it’s loss  $l(y, h(x))$  eventually goes to 0, however the loss of the expert  $l(y, m)$  is either 0 or equal to  $-\log(0.1)$ , hence the rejector will make the easier decision to never defer. At initial epochs, we have a non-trivial rejector as the classifier  $h$  is still learning, and the coverage progressively grows till 100% over training. Essentially, what [43] approach is trying to do is choosing between the lower cost between expert and classifier: a cost-sensitive learning problem at it’s heart. Therefore, one can use the losses developed here to tackle the problem better; we leave this to future investigations. Another potential fix is to learn the classifier and rejector on two different data sets.


Table A.7: System accuracy of our implementation of [43] and our method and baselines with varying expert competence ( $k$ ) on CIFAR-10.

METHOD / SYSTEM ACCURACY ( $k$ )	1	2	3	4	5	6	7	8	9	10
$L_{CE}^5$	<b>90.92</b>	<b>91.01</b>	<b>91.94</b>	<b>92.69</b>	<b>93.66</b>	<b>96.03</b>	<b>97.11</b>	<b>98.25</b>	<b>99</b>	<b>100</b>
$L_{CE}^1$	90.41	91.00	91.47	92.42	93.4	95.06	96.49	97.30	97.70	100
CONFIDENCE	90.47	90.56	90.71	91.41	92.52	94.15	95.5	97.35	98.05	100
ORACLE REJECT	89.54	89.51	89.48	90.75	90.64	93.25	95.28	96.52	98.16	100
[43]	90.40	90.40	90.40	90.40	90.40	90.40	90.40	94.48	95.09	100

## A.2.6 CheXpert Experiments



(a) classifier AU-ROC on non-deferred examples vs coverage for expert  $q = 0.7, p = 1$  with 100% of training data.



(b) classifier AU-ROC on non-deferred examples vs coverage for expert  $q = 0.7, p = 1$  with 10% of training data.

Figure A-2: Plot of classifier AU-ROC on non-deferred examples versus coverage for (a) for systems learned with 100% of training data (b) and learned with 10% of training data. Noise at low coverage is due to reduced data size.

## A.3 Deferred Proofs and Derivations

### A.3.1 Section 3

#### Binary Setting

As we eluded to in the body of the paper, we can extend the losses introduced by [13] to our setting for binary labels. Let  $\mathcal{Y} = \{-1, +1\}$  and  $r, h : \mathcal{X} \rightarrow \mathbb{R}$  where we defer if  $r(x) \leq 0$ , for generality we assume  $l_{exp}(x, y, m) = \max(c, \mathbb{I}_{m \neq y})$  as this allows to treat rejection learning as an immediate special case. Following the derivation in [13], let  $u \rightarrow \phi(-u)$  and  $u \rightarrow \psi(-u)$  be two convex function upper bounding  $\mathbb{I}_{u \leq 0}$  and let  $\alpha, \beta > 0$ , then:

$$\begin{aligned}
 L_c(h, r, x, y, m) &= \mathbb{I}_{h(x)y \leq 0} \mathbb{I}_{r(x) > 0} + \max(c, \mathbb{I}_{m \neq y}) \mathbb{I}_{r(x) \leq 0} \\
 &\leq \max \left\{ \mathbb{I}_{\max\{h(x)y, -r(x)\} \leq 0}, \max(c, \mathbb{I}_{m \neq y}) \mathbb{I}_{r(x) \leq 0} \right\} \\
 &\stackrel{(a)}{\leq} \max \left\{ \mathbb{I}_{\frac{\alpha}{2}(h(x)y - r(x)) \leq 0}, \max(c, \mathbb{I}_{m \neq y}) \mathbb{I}_{\beta r(x) \leq 0} \right\} \\
 &\stackrel{(b)}{\leq} \max \left\{ \phi \left( \frac{-\alpha}{2}(h(x)y - r(x)) \right), \max(c, \mathbb{I}_{m \neq y}) \psi(-\beta r(x)) \right\} \tag{A.3}
 \end{aligned}$$

$$\leq \phi \left( \frac{-\alpha}{2}(h(x)y - r(x)) \right) + \max(c, \mathbb{I}_{m \neq y}) \psi(-\beta r(x)) \tag{A.4}$$

step (a) is by noting that  $\max(a, b) \geq \frac{a+b}{2}$ , step (b) since  $\phi(u)$  and  $\psi(u)$  upper bound  $\mathbb{I}_{u \leq 0}$ . Both the right hand sides of equations (A.3) and (A.4) are convex functions of both  $h$  and  $r$ . When  $\phi$  and  $\psi$  are both the exponential loss we obtain the following loss with  $\beta(x, y, m) : \mathcal{X} \times \mathcal{Y}^2 \rightarrow \mathbb{R}^+$ :

$$L_{SH}(h, r, x, y, m) := \exp \left( \frac{\alpha}{2}(r(x) - h(x)y) \right) + (c + \mathbb{I}_{m \neq y}) \exp(-\beta(x, y, m)r(x))$$

we will see that it will be necessary that  $\beta$  is no longer constant for the loss to be consistent while in the standard case it sufficed to have  $\beta$  constant [13]. The following proposition shows that for an appropriate choice of  $\beta$  and  $\alpha$  we can make  $L_{SH}$  consistent.

**Proposition 4.** Let  $c(x) = c - c\mathbb{P}(Y \neq M|X = x) + \mathbb{P}(Y \neq M|X = x)$ , for  $\alpha = 1$  and  $\beta = \sqrt{\frac{1-c(x)}{c(x)}}$ ,  $\inf_{h,r} \mathbb{E}_{x,y,m}[L_{SH}(h, r, x, y, m)]$  is attained at  $(h_{SH}^*, r_{SH}^*)$  such that  $\text{sign}(h^B) = \text{sign}(h_{SH}^*)$  and  $\text{sign}(r^B) = \text{sign}(r_{SH}^*)$ .

*Proof.* Denote  $\eta(x) = \mathbb{P}(Y = 1|X = x)$  and  $q(x, y) = \mathbb{P}(M = 1|X = x, Y = y)$ , we have:

$$\begin{aligned} \inf_{h,r} \mathbb{E}_{x,y,m}[L_{SH}(h, r, x, y, m)] &= \inf_{h,r} \mathbb{E}_x \mathbb{E}_{y|x} \mathbb{E}_{m|x,y}[L_{SH}(h, r, x, y, m)] \\ &= \mathbb{E}_x \inf_{h(x), r(x)} \mathbb{E}_{y|x} \mathbb{E}_{m|x,y}[L_{SH}(h(x), r(x), x, y, m)] \end{aligned}$$

Now we will expand the inner expectation:

$$\begin{aligned} &\mathbb{E}_{y|x} \mathbb{E}_{m|x,y}[L_{SH}(h(x), r(x), x, y, m)] \tag{A.5} \\ &= \eta(x)q(x, 1)(\exp\left(\frac{\alpha}{2}(r(x) - h(x))\right) + c \exp(-\beta r(x))) \\ &+ (1 - \eta(x))q(x, -1)(\exp\left(\frac{\alpha}{2}(r(x) + h(x))\right) + (1) \exp(-\beta r(x))) \\ &+ \eta(x)(1 - q(x, 1))(\exp\left(\frac{\alpha}{2}(r(x) - h(x))\right) + (1) \exp(-\beta r(x))) \\ &+ (1 - \eta(x))(1 - q(x, -1))(\exp\left(\frac{\alpha}{2}(r(x) + h(x))\right) + c \exp(-\beta r(x))) \end{aligned}$$

The Bayes optimal solution for our original loss in the binary setting is:

$$\begin{aligned} h^B(x) &= \eta(x) - \frac{1}{2} \\ r^B(x) &= |\eta(x) - \frac{1}{2}| - \left(\frac{1}{2} - c - \mathbb{P}(M \neq Y|X = x)\right) \end{aligned}$$

**Case 1:** if  $\eta(x) = 0$ , writing  $v = r(x)$ ,  $u = h(x)$  then term (A.5) becomes:

$$q(x, -1)(\exp\left(\frac{\alpha}{2}(v + u)\right) + 1 \exp(-\beta v)) + (1 - q(x, -1))(\exp\left(\frac{\alpha}{2}(v + u)\right) + c \exp(-\beta v))$$

then to minimize the above it is necessary that the optimal solutions are such that  $u^* < 0$ ,  $v^* > 0$  which agree with the sign of the original Bayes solution.



**Case 2:** if  $\eta(x) = 1$ , then term (A.5) becomes:

$$q(x, 1)(\exp\left(\frac{\alpha}{2}(v - u)\right) + c \exp(-\beta v)) + (1 - q(x, 1))(\exp\left(\frac{\alpha}{2}(v - u)\right) + (1) \exp(-\beta v))$$

then to minimize the above it is necessary that the optimal solutions are such that  $u^* > 0, v^* > 0$  which agree with the sign of the original Bayes solution.

**Case 3:**  $\eta(x) \in (0, 1)$ , for ease of notation denote the RHS of equation (A.5) as  $L_\psi(u, v)$ , note that  $L_\psi(u, v)$  is a convex function of both  $u$  and  $v$  and therefore to find the optimal solution it suffices to take the partial derivatives with respect to each and set them to 0.

For  $u$ :

$$\begin{aligned} \frac{\partial_\psi(u, v)}{\partial u} &= 0 \\ \iff -\eta(x) \frac{\alpha}{2} \exp\left(\frac{\alpha}{2}(v - u^*)\right) + (1 - \eta(x)) \exp\left(\frac{\alpha}{2}(v + u^*)\right) &= 0 \\ \iff -\eta(x) \frac{\alpha}{2} \exp\left(\frac{-\alpha}{2}u^*\right) + (1 - \eta(x)) \frac{\alpha}{2} \exp\left(\frac{\alpha}{2}u^*\right) &= 0 \\ \iff u^* &= \frac{1}{\alpha} \log\left(\frac{\eta(x)}{1 - \eta(x)}\right) \end{aligned}$$

we note that  $u^*$  has the same sign as the minimizer of the exponential loss and hence has the same sign as  $h^B(x)$ .

Plugging  $u^*$  and taking the derivative with respect to  $v$ :

$$\begin{aligned}
\frac{\partial \psi(u^*, v)}{\partial v} &= 0 \\
\iff \eta(x) \frac{\alpha}{2} \exp\left(\frac{\alpha}{2}(v^* - u^*)\right) + (1 - \eta(x)) \exp\left(\frac{\alpha}{2}(v^* + u^*)\right) \\
&- \beta c(\eta(x)q(x, 1) + (1 - \eta(x))(1 - q(x, -1))) \exp(-\beta v^*) \\
&- (1 - \eta(x))q(x, -1)\beta \exp(-\beta v^*) - \eta(x)(1 - q(x, 1))\beta \exp(-\beta v^*) = 0 \\
\iff \eta(x) \frac{\alpha}{2} \exp\left(\frac{\alpha}{2}(v^* - u^*)\right) + (1 - \eta(x)) \exp\left(\frac{\alpha}{2}(v^* + u^*)\right) \\
&- \beta(c - c\mathbb{P}(M \neq Y|X = x) + \mathbb{P}(M \neq Y|X = x)) \exp(-\beta v^*) = 0
\end{aligned}$$

Appealing to the proof of Theorem 1 in [12] we obtain that:

$$v^* = \frac{1}{\alpha/2 + \beta} \log \left( \frac{c(x)\beta}{\alpha} \sqrt{\frac{1}{\eta(x)(1 - \eta(x))}} \right)$$

Furthermore by the proof of Theorem 1 in [12], the sign of  $v^*$  matches that of  $r^B(x)$  if and only if:

$$\frac{\beta}{\alpha} = \sqrt{\frac{1 - c(x)}{c(x)}}$$

□

## Multiclass setting

**Proposition 1.**  $\tilde{L}_{CE}$  is convex and is a consistent loss function for  $\tilde{L}$ :

let  $\tilde{\mathbf{g}} = \arg \inf_{\mathbf{g}} \mathbb{E} [\tilde{L}_{CE}(\mathbf{g}, \mathbf{c})|X = x]$ , then:  $\arg \max_{i \in [K+1]} \tilde{\mathbf{g}}_i = \arg \min_{i \in [K+1]} \mathbb{E}[c(i)|X = x]$

*Proof.* Writing the expected loss:

$$\inf_{\mathbf{g}} \mathbb{E}_{x, \mathbf{c}}[\tilde{L}_{CE}(\mathbf{g}, x, \mathbf{c})] = \inf_{\mathbf{g}} \mathbb{E}_x \mathbb{E}_{\mathbf{c}|x}[\tilde{L}_{CE}(\mathbf{g}, x, \mathbf{c})] = \mathbb{E}_x \inf_{\mathbf{g}(x)} \mathbb{E}_{\mathbf{c}|x}[\tilde{L}_{CE}(\mathbf{g}(x), x, \mathbf{c})]$$

Now we will expand the inner expectation:

$$\mathbb{E}_{\mathbf{c}|x}[\tilde{L}_{CE}(\mathbf{g}(x), x, \mathbf{c})] = - \sum_{y \in [K+1]} \mathbb{E}[\max_j c(j) - c(y) | X = x] \log \left( \frac{\exp(g_y(x))}{\sum_k \exp(g_k(x))} \right)$$

The loss  $\tilde{L}_{CE}$  is convex in the predictor, so it suffices to differentiate with respect to each  $g_y$  for  $y \in \mathcal{Y}^\perp$  and set to 0.

$$\begin{aligned} \frac{\partial L_{CE}}{\partial g_y^*} &= 0 \\ \iff \mathbb{E}[\max_j c(j) - c(y) | X = x] - \frac{\exp(g_y^*(x))}{\sum_k \exp(g_k(x))} \sum_{i \in [K+1]} \mathbb{E}[\max_j c(j) - c(i) | X = x] &= 0 \\ \iff \frac{\exp(g_y^*(x))}{\sum_k \exp(g_k(x))} &= \frac{\mathbb{E}[\max_j c(j) - c(y) | X = x]}{\sum_{i \in [K+1]} \mathbb{E}[\max_j c(j) - c(i) | X = x]} \end{aligned}$$

From this we can deduce:

$$\begin{aligned} h(x) &= \arg \max_{y \in [K+1]} g_y^*(x) = \arg \max_{y \in [K+1]} \frac{\exp(g_y^*(x))}{\sum_{y \in [K+1]} \exp(g_y^*(x))} \\ &= \arg \max_{y \in [K+1]} \frac{\mathbb{E}[\max_j c(j) | X = x] - \mathbb{E}[c(y) | X = x]}{\sum_{i \in [K+1]} \mathbb{E}[\max_j c(j) - c(i) | X = x]} \\ &= \arg \min_{y \in [K+1]} \mathbb{E}[c(y) | X = x] = \tilde{h}^B(x) \end{aligned}$$

□

**Proposition 2.** *The minimizers of the loss  $L_{0-1}$  (2.3) are defined point-wise for all  $x \in \mathcal{X}$  as:*

$$\begin{aligned} h^B(x) &= \arg \max_{y \in \mathcal{Y}} \eta_y(x) \\ r^B(x) &= \mathbb{I}_{\max_{y \in \mathcal{Y}} \eta_y(x) \leq \mathbb{P}(Y=M|X=x)} \end{aligned} \tag{A.6}$$

*Proof.* When we don't defer, the loss incurred by the model is the misclassification loss in the

standard multiclass setting and hence by standard arguments [22] we can define  $h^B$  point-wise regardless of  $r$ :

$$h^B(x) = \arg \inf_h \mathbb{E}_y [\mathbb{I}_{h \neq y}] = \arg \max_{y \in \mathcal{Y}} \eta_y(x)$$

Now for the rejector, we should only defer if the expected loss of having the expert predict is less than the error of the classifier  $h^B$  defined above, define  $r^B : \mathcal{X} \rightarrow \{0, +1\}$  as:

$$\begin{aligned} r^B(x) &= \mathbb{I}_{\mathbb{E}[\mathbb{I}_{M \neq Y} | X=x] \leq \mathbb{E}[\mathbb{I}_{h^B(x) \neq Y} | X=x]} \\ &= \mathbb{I}_{\mathbb{P}(Y \neq M) \leq (1 - \max_{y \in \mathcal{Y}} \eta_y(x))} \\ &= \mathbb{I}_{\mathbb{P}(Y=M) \geq \max_{y \in \mathcal{Y}} \eta_y(x)} \end{aligned}$$

□

**Theorem 2.** *The loss  $L_{CE}$  is a convex upper bound of  $L_{0-1}$  and is consistent:*

*$\inf_{h,r} \mathbb{E}_{x,y,m} [L_{CE}(h, r, x, y, m)]$  is attained at  $(h_{CE}^*, r_{CE}^*)$  such that  $h^B(x) = h_{CE}^*(x)$  and  $r^B(x) = r_{CE}^*(x)$  for all  $x \in \mathcal{X}$ .*

*Proof.* The fact that  $L_{CE}$  is convex is immediate as  $\mathbb{I}_{m=y} \geq 0$  and the cross entropy loss is convex.

Now we show that  $L_{CE}$  is an upper bound of  $L_{0-1}$ :

$$\begin{aligned} L_{0-1}(h, r, x, y, m) &= \mathbb{I}_{h(x) \neq y} \mathbb{I}_{r(x)=0} + \mathbb{I}_{m \neq y} \mathbb{I}_{r(x)=1} \\ &\stackrel{(a)}{\leq} -\log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) - \mathbb{I}_{m=y} \log \left( \frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \end{aligned} \quad (\text{A.7})$$

To justify inequality (a), consider first if  $r(x) = 0$ , then if  $\mathbb{I}_{h(x) \neq y} = 1$  we know that  $\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \leq \frac{1}{2}$  giving  $-\log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \right) \geq 1$ , moreover all the terms in the RHS of (a) are always positive.

On the other hand if  $r(x) = 1$ , then again  $\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))} \leq \frac{1}{2}$  as we decided to reject and since also giving  $-\log\left(\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \geq 1$ . Finally note that  $L_{0-1}(h, r, x, y, m) \leq 1$ .

We will now show that the optimal rejector minimizing the upper bound (A.7) is in fact consistent.

Denote  $q_m(x, y) = \mathbb{P}(M = m | X = x, Y = y)$  and  $\eta_y(x) = \mathbb{P}(Y = y | X = x)$ , we have:

$$\begin{aligned} \inf_{h,r} \mathbb{E}_{x,y,m} [L_{CE}(h, r, x, y, m)] &= \inf_{h,r} \mathbb{E}_x \mathbb{E}_{y|x} \mathbb{E}_{m|x,y} [L_{CE}(h, r, x, y, m)] \\ &= \mathbb{E}_x \inf_{h(x), r(x)} \mathbb{E}_{y|x} \mathbb{E}_{m|x,y} [L_{CE}(h(x), r(x), x, y, m)] \end{aligned}$$

Let us expand the inner expectation:

$$\begin{aligned} &\mathbb{E}_{y|x} \mathbb{E}_{m|x,y} [L_{CE}(h(x), r(x), x, y, m)] \\ &= \mathbb{E}_{y|x} \left[ -\log\left(\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) - \sum_{m \in \mathcal{Y}} \mathbb{I}_{m=y} \log\left(\frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \right] \\ &= -\sum_{y \in \mathcal{Y}} \eta_y(x) \log\left(\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \\ &\quad - \sum_{y \in \mathcal{Y}} \eta_y(x) \sum_{m \in \mathcal{Y}} q_m(x, y) \mathbb{I}_{m=y} \log\left(\frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \\ &\stackrel{(a)}{=} -\sum_{y \in \mathcal{Y}} \eta_y(x) \log\left(\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) - \sum_{y \in \mathcal{Y}} \eta_y(x) q_y(m, y) \log\left(\frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \\ &\stackrel{(b)}{=} -\sum_{y \in \mathcal{Y}} \eta_y(x) \log\left(\frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \\ &\quad - \mathbb{P}(Y = M | X = x) \log\left(\frac{\exp(g_{\perp}(x))}{\sum_{y' \in \mathcal{Y} \cup \perp} \exp(g_{y'}(x))}\right) \tag{A.8} \end{aligned}$$

In step (a) all terms that differed on  $y$  and  $m$  disappear, in step (b) we have:

$$\sum_{y \in \mathcal{Y}} \eta_y(x) q_y(m, y) = \sum_{y \in \mathcal{Y}} \mathbb{P}(M = y, Y = y | X = x) = \mathbb{P}(Y = M | X = x)$$

For ease of notation denote the RHS of equation (A.8) as  $L_{CE}(g_1, \dots, g_{|\mathcal{Y}|}, g_\perp)$ , note that it is a convex function, hence we will take the partial derivatives with respect to each argument and set them to 0.

For any  $g_\perp$ , and for  $i \in \mathcal{Y}$  we have :

$$\begin{aligned} \frac{\partial L_{CE}(g_1^*, \dots, g_{|\mathcal{Y}|}^*, g_\perp)}{\partial g_i^*} &= 0 \\ \iff \frac{\exp(g_i^*(x))}{\sum_{y' \in \tilde{\mathcal{Y}}} \exp(g_{y'}^*(x))} &= \frac{\eta_i(x)}{1 + \mathbb{P}(Y = M|X = x)} \end{aligned} \quad (\text{A.9})$$

The optimal  $h^*$  for any  $g_\perp$  should satisfy equation (A.9) for every  $i \in \mathcal{Y}$ , however since exponential is an increasing function we get that the optimal  $h^*$  in fact agrees with the Bayes solution as:

$$\begin{aligned} \arg \max_{y \in \mathcal{Y}} g_y^*(X) &= \arg \max_{y \in \mathcal{Y}} \frac{\exp(g_y^*(x))}{\sum_{y \in \mathcal{Y}} \exp(g_y^*(x)) + \exp(g_\perp(x))} \\ &= \arg \max_{y \in \mathcal{Y}} \frac{\eta_y(x)}{1 + \mathbb{P}(Y = M|X = x)} = h^B(x) \end{aligned}$$

Plugging  $h^*$  and taking the derivative with respect to the optimal  $g_\perp^*$ :

$$\begin{aligned} \frac{\partial L_{CE}(g_1^*, \dots, g_{|\mathcal{Y}|}^*, g_\perp^*)}{\partial g_\perp^*} &= 0 \\ \iff \frac{\exp(g_\perp^*(x))}{\sum_{y' \in \mathcal{Y}} \exp(g_{y'}^*(x))} &= \frac{\mathbb{P}(Y = M|X = x)}{1 + \mathbb{P}(Y = M|X = x)} \end{aligned}$$

Note note that  $r^*(x) = 1$  only if  $\mathbb{P}(Y = M|X = x) \geq \max_{y \in \mathcal{Y}} \eta_y(x)$  which agrees with  $r^B(x)$  □

### A.3.2 Section 4

**Proposition 3.**  $L_{mix}$  is realizable  $(\mathcal{H}, \mathcal{R})$ -consistent for classes closed under scaling but is not classification consistent.

*Proof.* We first prove that  $L_{mix}$  is realizable  $(\mathcal{H}, \mathcal{R})$ -consistent. Let  $\mathbf{P}$  and  $M$  be such that there exists  $h^*, r^* \in \mathcal{H} \times \mathcal{R}$  that have zero error  $L(h^*, r^*) = 0$ . Assume that  $(\hat{h}, \hat{r})$  satisfy

$$\left| \mathbb{E}[L_{mix}(\hat{h}, \hat{r}, x, y, m)] - \inf_{h \in \mathcal{H}, r \in \mathcal{R}} \mathbb{E}[L_{mix}(h, r, x, y, m)] \right| \leq \delta$$

Let  $u > 0$ , we have:

$$\begin{aligned} & \mathbb{E}[L(\hat{h}, \hat{r}, x, y, m)] \\ & \leq 2\mathbb{E}[L_{mix}(\hat{h}, \hat{r}, x, y, m)] \quad (\text{factor of 2 is upper bound}) \\ & \leq 2\mathbb{E}[L_{mix}(uh^*, ur^*, x, y, m)] + 2\delta \quad (\text{by assumption and closed under scaling}) \\ & = 2\mathbb{E}[L_{mix}(uh^*, ur^*, x, y, m)|r^* = 1]\mathbb{P}(r^* = 1) + 2\mathbb{E}[L_{mix}(uh^*, ur^*, x, y, m)|r^* = 0]\mathbb{P}(r^* = 0) \\ & + 2\delta \\ & = 2\mathbb{E}\left[-\log\left(\frac{\exp(ug_y(x))}{\sum_{y' \in \mathcal{Y}} \exp(ug_{y'}(x))}\right) \frac{\exp(ur_0(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))}\right. \\ & \left. + \mathbb{I}_{m \neq y} \frac{\exp(ur_1(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))} \middle| r^* = 0\right] \mathbb{P}(r^* = 0) \end{aligned} \tag{A.10}$$

$$\begin{aligned} & + 2\mathbb{E}\left[-\log\left(\frac{\exp(ug_y(x))}{\sum_{y' \in \mathcal{Y}} \exp(ug_{y'}(x))}\right) \frac{\exp(ur_0(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))}\right. \\ & \left. + \mathbb{I}_{m \neq y} \frac{\exp(ur_1(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))} \middle| r^* = 1\right] \mathbb{P}(r^* = 1) + 2\delta \end{aligned} \tag{A.11}$$

Let us examine each term in the RHS of (A.11), when  $r^* = 1$  we have  $r_1(x) > r_0(x)$  hence:

$$\lim_{u \rightarrow \infty} \frac{\exp(ur_0(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))} = 0$$

Furthermore it must be that  $\mathbb{I}_{m \neq y} = 0$  as we decided to defer.

When  $r^* = 0$ , we have  $r_0(x) \geq r_1(x)$  hence:

$$\lim_{u \rightarrow \infty} \frac{\exp(ur_1(x))}{\sum_{i \in \{0,1\}} \exp(ur_i(x))} = 0$$

moreover we have  $h^*(x) = y$  by optimality of  $(h^*, r^*)$  (as we did not defer) and realizability thus:

$$\lim_{u \rightarrow \infty} \log \left( \frac{\exp(ug_y(x))}{\sum_{y' \in \mathcal{Y}} \exp(ug_{y'}(x))} \right) = 0$$

We can conclude that taking the limit as  $u \rightarrow \infty$  on the RHS of (A.11) and applying the monotone convergence theorem (swap of expectation and limit) we get:

$$\mathbb{E}[L(\hat{h}, \hat{r}, x, y, m)] \leq 2\delta$$

taking  $\delta = \epsilon/2$  completes the proof.

We now move to looking at the Bayes solution of  $L_{mix}$ , denote  $q_m(x, y) = \mathbb{P}(M = m | X = x, Y = y)$ , we have:

$$\inf_{h, r} \mathbb{E}_{x, y, m} [L_{mix}(h, r, x, y, m)] = \mathbb{E}_x \inf_{h(x), r(x)} \mathbb{E}_{y|x} \mathbb{E}_{m|x, y} [L_{mix}(h(x), r(x), x, y, m)]$$

Let us expand the inner expectation:

$$\begin{aligned} \mathbb{E}_{y|x} \mathbb{E}_{m|x, y} [L_{mix}(h(x), r(x), x, y, m)] = & \tag{A.12} \\ & - \sum_{y \in \mathcal{Y}} \eta_y(x) \log \left( \frac{\exp(g_y(x))}{\sum_{y' \in \mathcal{Y}} \exp(g_{y'}(x))} \right) \frac{\exp(r_0(x))}{\sum_{i \in \{0,1\}} \exp(r_i(x))} \\ & + \mathbb{P}(Y \neq M | X = x) \frac{\exp(r_1(x))}{\sum_{i \in \{0,1\}} \exp(r_i(x))} \end{aligned}$$

Denote the RHS of (A.12) by  $L_{mix}(g_1, \dots, g_{|\mathcal{Y}|}, r_0, r_1)$ , it is a convex function in  $g_i$  for all



$i \in \mathcal{Y}$ , consider any  $r_0, r_1$ , we have :

$$\frac{\partial L_{mix}(g_1^*, \dots, g_{|\mathcal{Y}|^*}, r_0, r_1)}{\partial g_i^*} = 0 \iff \frac{\exp(g_i^*(x))}{\sum_{y' \in \mathcal{Y}} \exp(g_{y'}^*(x))} = \eta_i(x) \quad (\text{A.13})$$

Since the optimal  $h^*$  for any  $r_0, r_1$  *does not depend* on the form of  $r_0$  and  $r_1$  we conclude that (A.13) gives the optimal choice of  $h$ . We now need to find the optimal choice of  $r_0(x)$  and  $r_1(x)$  to minimize  $L_{mix}(g_1^*, \dots, g_{|\mathcal{Y}|^*}, r_0, r_1)$  which takes the following form:

$$L_{mix}(g_1^*, \dots, g_{|\mathcal{Y}|^*}, r_0, r_1) = H(h^B(x)) \frac{\exp(r_0(x))}{\sum_{i \in \{0,1\}} \exp(r_i(x))} + \mathbb{P}(Y \neq M | X = x) \frac{\exp(r_1(x))}{\sum_{i \in \{0,1\}} \exp(r_i(x))}$$

where  $H(X)$  is the Shannon entropy of the random variable  $X$ , here by  $H(h^B(x))$  we refer to the entropy of the probabilistic form of  $h^B(x)$  according to (A.13). Clearly the optimal  $r_0^*$  and  $r_1^*$  have the following behavior for a given  $x \in \mathcal{X}$ :

$$\begin{cases} r_0(x) = \infty, r_1(x) = -\infty & \text{if } H(h^B(x)) < \mathbb{P}(Y \neq M | X = x) \\ r_0(x) = -\infty, r_1(x) = \infty & \text{if } H(h^B(x)) \geq \mathbb{P}(Y \neq M | X = x) \end{cases}$$

This does not have the form of  $r^B(x)$ , as this rejector compares the entropy of  $h^B(x)$  instead of it's confidence to the probability of error of the expert which will not always be in accordance.  $\square$

**Theorem 2.** *For any expert  $M$  and data distribution  $\mathbf{P}$  over  $\mathcal{X} \times \mathcal{Y}$ , let  $0 < \delta < \frac{1}{2}$ , then with probability at least  $1 - \delta$ , the following holds for the empirical minimizers  $(\hat{h}^*, \hat{r}^*)$ :*

$$\begin{aligned} L_{0-1}(\hat{h}^*, \hat{r}^*) &\leq L_{0-1}(h^*, r^*) + \mathfrak{R}_n(\mathcal{H}) + \mathfrak{R}_n(\mathcal{R}) + \mathfrak{R}_{n\mathbb{P}(M \neq Y)/2}(\mathcal{R}) \\ &\quad + 2\sqrt{\frac{\log(\frac{2}{\delta})}{2n}} + \frac{\mathbb{P}(M \neq Y)}{2} \exp\left(-\frac{n\mathbb{P}(M \neq Y)}{8}\right) \end{aligned}$$

*Proof.* Let  $\mathcal{L}_{\mathcal{H}, \mathcal{R}}$  be the family of functions defined as  $\mathcal{L}_{\mathcal{H}, \mathcal{R}} = \{(x, y, m) \rightarrow L(h, r, x, y, m); h \in$

$\mathcal{H}, r \in \mathcal{R}$  with  $L(h, r, x, y, m) := \mathbb{I}_{h(x) \neq y} \mathbb{I}_{r(x) = -1} + \mathbb{I}_{m \neq y} \mathbb{I}_{r(x) = 1}$ . Let  $\mathfrak{R}_n(\mathcal{L}_{\mathcal{H}, \mathcal{R}})$  be the Rademacher complexity of  $\mathcal{L}_{\mathcal{H}, \mathcal{R}}$ , then since  $L(h, r, x, y, m) \in [0, 1]$ , by the standard Rademacher complexity bound (Theorem 3.3 in [45]), with probability at least  $1 - \delta/2$  we have:

$$L_{0-1}(\hat{h}^*, \hat{r}^*) \leq L_{0-1}^S(\hat{h}^*, \hat{r}^*) + 2\mathfrak{R}_n(\mathcal{L}_{\mathcal{H}, \mathcal{R}}) + \sqrt{\frac{\log(\frac{2}{\delta})}{2n}}$$

We will now relate the complexity of  $\mathcal{L}_{\mathcal{H}, \mathcal{R}}$  to the individual classes:

$$\begin{aligned} \mathfrak{R}_n(\mathcal{L}_{\mathcal{H}, \mathcal{R}}) &= \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{h(x_i) \neq y_i} \mathbb{I}_{r(x_i) = -1} + \epsilon_i \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i) = 1} \right] \\ &\stackrel{(a)}{\leq} \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{h(x_i) \neq y_i} \mathbb{I}_{r(x_i) = -1} \right] \\ &\quad + \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i) = 1} \right] \\ &\stackrel{(b)}{\leq} \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{h(x_i) \neq y_i} \right] + \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{r(x_i) = -1} \right] \\ &\quad + \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i) = 1} \right] \\ &\leq \frac{1}{2} \mathfrak{R}_n(\mathcal{H}) + \frac{1}{2} \mathfrak{R}_n(\mathcal{R}) + \mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i) = 1} \right] \end{aligned} \tag{A.14}$$

step (a) follows as the supremum is a subadditive function, step (b) is the application of Lemma 2 in [18] to  $\mathbb{E}_\epsilon \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{h(x_i) \neq y_i} \mathbb{I}_{r(x_i) = -1} \right]$  which says that the Rademacher complexity of a product of two indicators functions is upper bounded by the sum of the complexities of each class, now we will take a closer look at the last term in the RHS of inequality (A.14). Denote  $n_m^S = \sum_{j \in S} \mathbb{I}_{y_j \neq m_j}$  and define the random variable  $S_m = \{i : y_i \neq$

$m_i\}$ , we have that  $n_m^S \sim \text{Binomial}(n, \mathbb{P}(M \neq Y))$  and  $\mathbb{E}[n_m^S | S_m] = n\mathbb{P}(M \neq Y)$ , hence:

$$\begin{aligned}
& \mathbb{E} \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{m_i \neq y_i} \mathbb{I}_{r(x_i)=1} \right] \\
&= \mathbb{E} \left[ \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{I}_{r(x_i)=1} \right] \\
&= \mathbb{E} \left[ \frac{n_m^S}{m} \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{n_m^S} \sum_{i=1}^{n_m^S} \epsilon_i \mathbb{I}_{r(x_i)=1} \right] \text{ (by relabeling)} \\
&\stackrel{(a)}{=} \mathbb{E} \left[ \mathbb{E}_\epsilon \left[ \frac{n_m^S}{m} \sup_{(h,r) \in \mathcal{H} \times \mathcal{R}} \frac{1}{n_m^S} \sum_{i=1}^{n_m^S} \epsilon_i \mathbb{I}_{r(x_i)=1} \middle| S_m \right] \right] \\
&\stackrel{(b)}{=} \mathbb{E} \left[ \frac{n_m^S}{m} \hat{\mathfrak{R}}_{S_m}(\mathcal{R}) \right] \\
&\stackrel{(c)}{=} \mathbb{P}(n_m^S < \frac{n\mathbb{P}(A)}{2}) \mathbb{E} \left[ \frac{n_m^S}{m} \hat{\mathfrak{R}}_{S_m}(\mathcal{R}) \middle| n_m^S < \frac{n\mathbb{P}(A)}{2} \right] \\
&\quad + \mathbb{P}(n_m^S \geq \frac{n\mathbb{P}(A)}{2}) \mathbb{E} \left[ \frac{n_m^S}{m} \hat{\mathfrak{R}}_{S_m}(\mathcal{R}) \middle| n_m^S \geq \frac{n\mathbb{P}(A)}{2} \right] \\
&\stackrel{(d)}{\leq} \frac{\mathbb{P}(M \neq Y)}{2} \exp \left( -\frac{n\mathbb{P}(M \neq Y)}{8} \right) + \mathfrak{R}_{n\mathbb{P}(M \neq Y)/2}(\mathcal{R})
\end{aligned}$$

In step (a) we conditioned on the dataset  $S_m$ , in step (b) we used the definition of the empirical Rademacher complexity  $\hat{\mathfrak{R}}_{S_m}(\mathcal{R})$  on  $S_m$ , step (c) we introduce the event  $A = \{M \neq Y\}$ , step (d) follows from a Chernoff bound on  $n_m^S$  and since the Rademacher complexity is bounded by 1 and is non-increasing with respect to sample size.

We can now proceed with inequality (A.14):

$$\mathfrak{R}_n(\mathcal{L}_{\mathcal{H}, \mathcal{R}}) \stackrel{(a)}{\leq} \frac{1}{2} \mathfrak{R}_n(\mathcal{H}) + \frac{1}{2} \mathfrak{R}_n(\mathcal{R}) + \frac{\mathbb{P}(M \neq Y)}{2} \exp \left( -\frac{n\mathbb{P}(M \neq Y)}{8} \right) + \mathfrak{R}_{n\mathbb{P}(M \neq Y)/2}(\mathcal{R})$$

step (a) follows as the Rademacher complexity of indicator functions based on a certain class is equal to half the Rademacher complexity of the class [45].

The final step is to note by Hoeffding's inequality we have with probability at least  $1 - \delta/2$ :

$$L^S(h^*, r^*) \leq L(h^*, r^*) + \sqrt{\frac{\log(\frac{2}{\delta})}{2n}}$$

Now since  $(\hat{h}^*, \hat{r}^*)$  are the empirical minimizers we have that  $L^S(\hat{h}^*, \hat{r}^*) \leq L^S(h^*, r^*)$ , collecting all the inequalities we obtain the following generalization bound with probability at least  $1 - \delta$ :

$$\begin{aligned} L(\hat{h}^*, \hat{r}^*) &\leq L(h^*, r^*) + \mathfrak{R}_n(\mathcal{H}) + \mathfrak{R}_n(\mathcal{R}) + 2\sqrt{\frac{\log(\frac{2}{\delta})}{2n}} \\ &\quad + \frac{\mathbb{P}(M \neq Y)}{2} \exp\left(-\frac{n\mathbb{P}(M \neq Y)}{8}\right) + \mathfrak{R}_{n\mathbb{P}(M \neq Y)/2}(\mathcal{R}) \end{aligned}$$

□

# Bibliography

- [1] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. Optimizing ai for teamwork. *arXiv preprint arXiv:2004.13102*, 2020.
- [2] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S Weld, Walter S Lasecki, and Eric Horvitz. Updates in human-ai teams: Understanding and addressing the performance/-compatibility tradeoff. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2429–2437, 2019.
- [3] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [4] Peter L Bartlett and Marten H Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(Aug):1823–1840, 2008.
- [5] Shai Ben-David, David Loker, Nathan Srebro, and Karthik Sridharan. Minimizing the misclassification error rate using a surrogate convex loss. *arXiv preprint arXiv:1206.6442*, 2012.
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*, 2019.
- [7] Su Lin Blodgett, Lisa Green, and Brendan O’Connor. Demographic dialectal variation in social media: A case study of African-American English. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130, Austin, Texas, November 2016. Association for Computational Linguistics.
- [8] Ran Canetti, Aloni Cohen, Nishanth Dikkala, Govind Ramnarayan, Sarah Scheffler, and Adam Smith. From soft classifiers to hard decisions: How fair can we be? In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 309–318. ACM, 2019.
- [9] Mohammad-Amin Charusaie, Hussein Mozannar, David Sontag, and Samira Samadi. Sample efficient learning of predictors that complement humans. In *International Conference on Machine Learning*, pages 2972–3005. PMLR, 2022.

- [10] Minmin Chen, Ramki Gummadi, Chris Harris, and Dale Schuurmans. Surrogate objectives for batch policy optimization in one-step decision making. In *Advances in Neural Information Processing Systems*, pages 8825–8835, 2019.
- [11] C Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- [12] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Boosting with abstention. In *Advances in Neural Information Processing Systems*, pages 1660–1668, 2016.
- [13] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with rejection. In *International Conference on Algorithmic Learning Theory*, pages 67–82. Springer, 2016.
- [14] Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. Racial bias in hate speech and abusive language detection datasets. *arXiv preprint arXiv:1905.12516*, 2019.
- [15] Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Eleventh international aaai conference on web and social media*, 2017.
- [16] Abir De, Paramita Koley, Niloy Ganguly, and Manuel Gomez-Rodriguez. Regression under human assistance. *arXiv preprint arXiv:1909.02963*, 2019.
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [18] Giulia DeSalvo, Mehryar Mohri, and Umar Syed. Learning with deep cascades. In *International Conference on Algorithmic Learning Theory*, pages 254–269. Springer, 2015.
- [19] Michael Doron, Hussein Mozannar, and David Sontag. Machine teaching with generative models for human learning.
- [20] Cynthia Dwork and Christina Ilvento. Fairness under composition. *arXiv preprint arXiv:1806.06122*, 2018.
- [21] Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11(May):1605–1641, 2010.
- [22] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- [23] Joseph Futoma, Sanjay Hariharan, Katherine Heller, Mark Sendak, Nathan Brajer, Meredith Clement, Armando Bedoya, and Cara O’Brien. An improved multi-output gaussian process rnn with real-time validation for early sepsis detection. In *Machine Learning for Healthcare Conference*, pages 243–254, 2017.

- [24] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [25] Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. In *Advances in neural information processing systems*, pages 4878–4887, 2017.
- [26] Yonatan Geifman and Ran El-Yaniv. Selectivenet: A deep neural network with an integrated reject option. *arXiv preprint arXiv:1901.09192*, 2019.
- [27] Ben Green and Yiling Chen. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 90–99. ACM, 2019.
- [28] Ben Green and Yiling Chen. The principles and limits of algorithm-in-the-loop decision making. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019.
- [29] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR. org, 2017.
- [30] Kanza Hamid, Amina Asif, Wajid Abbasi, Durre Sabih, et al. Machine learning with abstention for automated liver disease diagnosis. In *2017 International Conference on Frontiers of Information Technology (FIT)*, pages 356–361. IEEE, 2017.
- [31] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 590–597, 2019.
- [32] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [33] Shagun Jhaver, Iris Birman, Eric Gilbert, and Amy Bruckman. Human-machine collaboration for content regulation: The case of reddit automoderator. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(5):1–35, 2019.
- [34] Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. To trust or not to trust a classifier. In *Advances in neural information processing systems*, pages 5541–5552, 2018.
- [35] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.
- [36] Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.

- [37] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Large scale learning of general visual representations for transfer. *arXiv preprint arXiv:1912.11370*, 2019.
- [38] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009.
- [39] Daniel Link, Bernd Hellengrath, and Jie Ling. A human-is-the-loop approach for semi-automated content moderation. In *ISCRAM*, 2016.
- [40] Phil Long and Rocco Servedio. Consistency versus realizable h-consistency for multiclass classification. In *International Conference on Machine Learning*, pages 801–809, 2013.
- [41] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [42] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- [43] David Madras, Toni Pitassi, and Richard Zemel. Predict responsibly: Improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*, pages 6150–6160, 2018.
- [44] Vahid Balazadeh Meresht, Abir De, Adish Singla, and Manuel Gomez-Rodriguez. Learning to switch between machines and humans. *arXiv preprint arXiv:2002.04258*, 2020.
- [45] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [46] Hussein Mozannar, Gagan Bansal, Adam Fournay, and Eric Horvitz. Reading between the lines: Modeling user behavior and costs in ai-assisted programming. *arXiv preprint arXiv:2210.14306*, 2022.
- [47] Hussein Mozannar, Hunter Lang, Dennis Wei, Prasanna Sattigeri, Subhro Das, and David Sontag. Who should predict? exact algorithms for learning to defer to humans. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.
- [48] Hussein Mozannar, Arvind Satyanarayan, and David Sontag. Teaching humans when to defer to a classifier via exemplars. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5323–5331, 2022.
- [49] Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, pages 7076–7087. PMLR, 2020.



- [50] Chenri Ni, Nontawat Charoenphakdee, Junya Honda, and Masashi Sugiyama. On possibility and impossibility of multiclass classification with rejection. *arXiv preprint arXiv:1901.10655*, 2019.
- [51] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, pages 3235–3246, 2018.
- [52] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [53] Joshua C Peterson, Ruairidh M Battleday, Thomas L Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9617–9626, 2019.
- [54] Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*, 2019.
- [55] Maithra Raghu, Katy Blumer, Rory Sayres, Ziad Obermeyer, Bobby Kleinberg, Sendhil Mullainathan, and Jon Kleinberg. Direct uncertainty prediction for medical second opinions. In *International Conference on Machine Learning*, pages 5281–5290, 2019.
- [56] Harish G Ramaswamy, Ambuj Tewari, Shivani Agarwal, et al. Consistent algorithms for multiclass classification with an abstain option. *Electronic Journal of Statistics*, 12(1):530–554, 2018.
- [57] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- [58] James Vincent. Ai won’t relieve the misery of facebook’s human moderators, February 2019.
- [59] Bryan Wilder, Eric Horvitz, and Ece Kamar. Learning to complement humans. *arXiv preprint arXiv:2005.00582*, 2020.
- [60] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [61] Chicheng Zhang and Kamalika Chaudhuri. Active learning from weak and strong labelers. In *Advances in Neural Information Processing Systems*, pages 703–711, 2015.
- [62] Liu Ziyin, Zhikang Wang, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency, and Masahito Ueda. Deep gamblers: Learning to abstain with portfolio theory. *arXiv preprint arXiv:1907.00208*, 2019.