

---

# META LEARNING WITH LANGUAGE MODELS: CHALLENGES AND OPPORTUNITIES IN THE CLASSIFICATION OF IMBALANCED TEXT

---

**Apostol Vassilev**

National Institute of Standards and Technology  
Gaithersburg, USA  
apostol.vassilev@nist.gov

**Honglan Jin**

National Institute of Standards and Technology  
Gaithersburg, USA  
honglan.jin@nist.gov

**Munawar Hasan**

National Institute of Standards and Technology  
Gaithersburg, USA  
munawar.hasan@nist.gov

## ABSTRACT

Detecting out of policy speech (OOPS) content is important but difficult. While machine learning is a powerful tool to tackle this challenging task, it is hard to break the performance ceiling due to factors like quantity and quality limitations on training data and inconsistencies in OOPS definition and data labeling [1]. To realize the full potential of available limited resources, we propose a meta learning technique (MLT) that combines individual models built with different text representations. We analytically show that the resulting technique is numerically stable and produces reasonable combining weights. We combine the MLT with a threshold-moving (TM) technique to further improve the performance of the combined predictor on highly-imbalanced in-distribution and out-of-distribution datasets. We also provide computational results to show the statistically significant advantages of the proposed MLT approach.<sup>1</sup>

All authors contributed equally to this work.

**Keywords** Natural language processing · Out of policy speech detection · Meta learning · Deep learning · Large Language Models

## 1 Introduction

Out of policy speech (OOPS) has permeated social media with serious consequences for both individuals and society. Although it comprises a small fraction of the content generated daily on social media, sifting through the data to quickly identify and eliminate the toxic content is difficult. The scale of this problem has long passed a threshold that requires automated detection. Yet it remains to be a challenging problem for machine learning (ML) due to the way OOPS manifests itself in datasets: context-dependent, nuanced, non-colloquial language that may even be syntactically incorrect. Because the OOPS content of the dataset is usually only a small fraction of the overall size, there is a high imbalance between OOPS and in-policy text. Related to this, there are not many high-quality labeled datasets with consistent definitions of OOPS and in-policy content. The difficulties are exacerbated further by significant differences in the distributions of the datasets that the model has been trained on and the data it sees during deployment. When faced with all of these challenges, ML models applied to natural language processing (NLP) tasks quickly reach a performance ceiling that limits their usefulness for sensitive tasks, such as OOPS detection.

---

<sup>1</sup>DISCLAIMER: This paper is not subject to copyright in the United States. Commercial products are identified in order to adequately specify certain procedures. In no case does such identification imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the identified products are necessarily the best available for the purpose.

A recent comprehensive evaluation of the performance of NLP models on many tasks [2] has revealed a relationship between the accuracy and the access to the model, showing a consistent gap between the higher accuracy of non-open models and that of open models. This finding suggests troubling power dynamics related to language models. Attempts to close that gap by improving the weaker language models by finetuning them on outputs from a stronger proprietary model have not been successful [3].

We approach this problem differently with a meta-learning technique for combining multiple open NLP models with the goals of breaking through the ceiling of each individual model and attaining higher combined performance on a large number of OOPS detection tasks, at a par with the performance of the best non-open model on these tasks. We make three main contributions: **i) a meta learning technique (MLT)** with language models with statistically significant performance improvements on high-quality recent benchmark datasets specifically curated for many OOPS detection tasks, **ii) theoretical analysis of the MLT** that establishes it as computationally stable and results in reasonably valued combiner weights for the participating language models, backed by numerical results in full agreement with the theory, and **iii) combining MLT with threshold-moving (TM)**, further improves the performance on highly-imbalanced in-distribution and out-of-distribution datasets, thus closing the performance gap with the best non-open model. Our results also suggest that smaller models can outperform much larger models on tasks of interest.

This paper is organized as follows. We review related work to establish a baseline for our research. Then we introduce our approach based on MLT with language models and TM. Next, we provide extensive numerical results that demonstrate the advantages of our approach. We end with conclusions.

## 2 Related Work

The inception of attention mechanisms [4, 5, 6] stimulated the development of many modeling techniques for NLP problems. The transformer network, first introduced by [7], has since become the dominant technology used by the top performing modern NLP models, including large language models [8, 9, 10, 11]. Despite the enormous progress [12], NLP remains a hard target for ML, and the technology is still unable to fully capture text semantics in many specific contexts.

The task of OOPS detection in particular adds complexities like ethnic slurs, incorrect or shortened words, ambiguous grammar, sarcasm, context-dependent belligerence, and a lack of consensus about the definition of OOPS [13, 14, 15, 1]. Moreover, it is difficult to differentiate OOPS from offensive language [16]. [17] points out major issues like strong subjective interpretations that results in different subjective perceptions in a variety of systems, each trained on a different resource. [18] discusses the low generalizability of OOPS models. These difficulties are further exacerbated by the low-quality datasets available to researchers for OOPS detection. For example, [19] found that racial bias is present in most of the available datasets [20, 16, 21, 22]. Lack of agreement among the annotators and human errors add more complexity to already low-quality and highly imbalanced data. The absence of context awareness is yet another notable pitfall in such datasets. All of these factors have hindered the research into OOPS detection.

In an attempt to rectify some of these problems, [23] have created a high-quality dataset specifically curated for multiple OOPS detection tasks. While still highly-imbalanced, this dataset uses distinct primary categories, context-aware annotation, and high-quality regressive annotation. To enable more targeted assess of various model functionalities with a large suite of functional tests, [24] have developed a dataset and used it to asses the performance of open-source language models and some commercial models, revealing critical weaknesses in all of them. There are many other commercial models in use by the industry (e.g., Meta’s combination of linformer [25] and reinforcement integrity optimizer, Google Jigsaw’s Perspective Application Programming Interface (API) [26], Two Hat’s SiftNinja [27]), but their performance has not been independently verified. The work of [24] suggests that some skepticism about such unverified performance claims is warranted.

The idea of combining models to gain better performance has a long history in machine learning, starting with the linear stacking regression approach of [28] for tabular data. Later [29] extended the idea and presented a systematic approach to combining using Bayesian rules. The recent advancements in deep learning and available modern computational power have led to ensemble techniques [30], which deliver good performance in computer vision tasks [31, 32, 33, 34]. In contrast, the idea of ensemble techniques in NLP tasks has hardly been explored; see [6], [8] and [35]. One of the reasons for this is that the individual NLP models have distinct encoding, which leads to the creation of fundamentally distinct hyperspaces that in turn may result in disappointing performance. Recently, a promising technique for combining language models has been proposed and applied to sentiment analysis - a binary classification task, on a well-balanced dataset [36].

High data imbalance and transfer learning, when the training dataset and the data with which the model is used in deployment have different distributions, pose serious challenges to machine learning. If left unchecked, these challenges

can lead to a severe performance degradation with dire consequences in critical applications, such as OOPS detection. Many techniques for dealing with these challenges have emerged in the literature: oversampling and undersampling [37], the weighted class approach [23, 24, 38], and threshold-moving (TM) [39, 40, 41, 42]. Oversampling and undersampling require modifications to the original datasets [37], which is either undesirable or impossible in many applications. The analysis in [38] shows that the effect of class weighting tends to be prominent initially but vanishes as the model training progresses, effectively making the model more skewed towards the majority class. Moreover, the weighted class method presents many challenges computationally. It assigns the class weights before model training using a class ratio distribution. This heuristic assignment is not guaranteed to produce an optimal result. To search for optimal class weights, the models must be retrained every time a class weight is changed. Repeated model training requires not only significant computing resources but also a long time. TM, on the other hand, is a method that maximizes a selected performance measure without retraining. Additionally, TM can be trained on a different performance measure of interest easily, as it is a plug-in method that works during the post-training stage. Due to all of these factors, we utilize TM to address the class imbalance problem.

### 3 Methodology

Imbalanced data distribution is prevalent in many application domains, including OOPS detection. Metrics like the F1 score are a better measure than accuracy for uneven data distribution [43]. We choose macro-averaged F1 scores to report our performance because we treat all classes equally. Additionally, We report accuracy when applicable, as an absolute improvement should ideally be made on both F1 and accuracy, regardless of data distribution. For brevity, in this paper, we will refer to macro-averaged F1 as "*macroF1*".

In multi-label classification, the F1 score for class  $c$  is calculated as follows:

$$F_c = \frac{2 * (precision_c * recall_c)}{(precision_c + recall_c)} = \frac{2TP_c}{(2TP_c + FP_c + FN_c)}$$

where  $TP_c, FP_c, FN_c$  represents the counts of true positives, false positives, and false negatives for class  $c$ , respectively.

Then, the macro-averaged F1 score is obtained by calculating the average of the F1 scores of each class.

Throughout this paper, we will denote a multi-label classifier or classification as M-C, and a binary classifier or classification as B-C.

First, we introduce an MLT architecture to combine predictions from multiple models. We assume that the number of texts in a corpus  $\mathbb{D}$  is large. Let there be  $n_c$  classes of texts in  $\mathbb{D}$ , labeled  $1, \dots, n_c$ , for some integer  $n_c > 2$ . Here, we define a predictor  $F$  in terms of a neural network with dense layers and a sigmoid. The neural network computes the weights  $\{w_i\}_{i=1}^K$  for combining the participating models  $M_1, \dots, M_K$ . Let  $\mathbf{y}_i(\tau^j)$  be the probability estimate computed by the  $i$ -th model on the  $j$ -th text in  $\mathbb{D}$ .

We define the MLT predictor as a combination of  $K > 1$  individual predictors:

$$\mathbf{y}(\tau) = \sum_{i=1}^K w_i \mathbf{y}_i(\tau), \quad \forall \tau \in \mathbb{D}, \quad (1)$$

$$\text{where either } w_i \geq 0, \forall i; \quad \text{or } w_i \leq 0, \forall i. \quad (2)$$

The real-valued functions  $\mathbf{y}_i$  with a range of interval  $[0, 1]$  have values that correspond to the probability of being assigned to a class  $I^{(C)}$  and are assigned to it with respect to the class threshold test with some value  $t \in (0, 1)$ :

$$\tau \in \begin{cases} \mathbf{I}^{(C)} & , \text{ if } \sigma(\mathbf{y}(\tau)) \geq t, \\ \mathbf{I}^{(\bar{C})} & , \text{ otherwise.} \end{cases} \quad (3)$$

Here  $\sigma(x)$  is the sigmoid function. Because  $\mathbf{y}_i(\tau)$  is, with ranges, shifted with respect to the domain of the sigmoid, to get accurate classification one needs to shift the range of  $\mathbf{y}$ ,

$$\sigma_b(\mathbf{y}(\tau)) = \sigma(\mathbf{y}(\tau) + b), \quad (4)$$

for some  $b > 0$  if  $w_i \leq 0, \forall i$  or  $b < 0$  if  $w_i \geq 0, \forall i$ .

Secondly, we employ the threshold-moving (TM) technique to post-process predictions from MLT, aiming to alleviate data imbalance issues [37, 40]. We refer to this combination as MLT-plus-TM. TM involves utilizing a trained threshold

value to map probabilities to class labels, thereby optimizing selected performance metrics such as the *macroF1* score of the model.

Subsequently, we construct classification systems, denoted as MLT-plus-TM, for both multi-label classifiers (M-C) and binary classifiers (B-C). Both systems utilize Gold labeled data from Contextual Abuse Dataset (CAD) [23].

For both classification tasks, we begin by constructing five individual models using five different embeddings. Subsequently, we combine the results using the approach outlined in (1)-(4), followed by applying the threshold-moving (TM) technique in both M-C and B-C.

To establish a notion about the reasonableness of the combiner weight values in (1), we show that the following theorem holds.

**Theorem 1.** *Let  $W = \sum_{i=1}^K w_i$  with  $w_i$  computed by the procedure (1)-(4). Then for any class  $C \in \mathbb{D}$ ,  $W$  is either close to 1 or -1.*

*Proof.*<sup>2</sup>

For a real-valued function  $\mathbf{y}_i(\tau)$  let us define

$$\|\mathbf{y}(\tau)\|_t^2 = \sum_{\tau \in \mathbb{D}} \mathbf{I}_t^{(C)}(\mathbf{y}(\tau))^2, \quad (5)$$

where  $\mathbf{I}_t^{(C)}$  is the assigned class for  $\mathbf{y}(\tau)$  with respect to the threshold  $t$  according to (3).

Let  $\mathbf{u}$  be defined on the texts in the corpus  $\mathbb{D}$ , so that  $\mathbf{I}_t^{(C)}(\mathbf{u}(\tau))$  is equal to the label assigned to each text  $\tau$  in  $\mathbb{D}$  that also belong to class  $C$ . In other words,  $\mathbf{u}$  is the true label on each class. Similarly, we define

$$\|\mathbf{y}(\tau) - \mathbf{z}(\tau)\|_t^2 = \sum_{\tau \in \mathbb{D}} (\mathbf{I}_t^{(C)}(\mathbf{y}(\tau)) - \mathbf{I}_t^{(C)}(\mathbf{z}(\tau)))^2. \quad (6)$$

We assume that each predictor  $\mathbf{y}_i$  is decent, i.e.,  $\|\mathbf{u} - \mathbf{y}_i\|_t$  is small, compared to  $\|\mathbf{u}\|_t$ . Let

$$\hat{\mathbf{y}}(\tau) = \frac{1}{|W|} \mathbf{y}(\tau) = \sum_{i=1}^K \frac{w_i}{|W|} \mathbf{y}_i(\tau)$$

be the interpolation predictor constructed as a linear combination of  $\mathbf{y}_i$  with coefficients that sum up to  $\pm 1$ . If the individual predictors are good then the interpolation predictor  $\hat{\mathbf{y}}$  is also good, i.e.,  $\|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$  is small relative to  $\|\mathbf{u}\|_t$ .

Let  $\mathbb{L}_t(x)$  be a linear approximation of  $\sigma(x)$  for some constant  $t > 0$ , such that  $\mathbb{L}_t(x)$  minimizes  $\|\mathbb{L}_t(x) - \sigma(x)\|_t$ . Note that any straight line passing through the point  $(\ln(\frac{t}{1-t}), t)$  and having the same slope as  $\sigma'(\ln(\frac{t}{1-t}))$  satisfies  $\|\mathbb{L}_t(x) - \sigma(x)\|_t = 0$ . Take

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t = \|\mathbf{u} - W\mathbf{u} + W\mathbf{u} - \sigma_b(\mathbf{y})\|_t$$

and consider the case  $b < 0$  and  $w_i \geq 0$  for  $\forall i$  in (4). Considering the two cases for  $W$  ( $W \leq 1$  and  $W > 1$ ), applying the triangle inequality and using the interpolation predictor  $\hat{\mathbf{y}}(\tau)$  and the linear approximation  $\mathbb{L}_t(x)$ , one gets

$$W \geq \frac{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t}. \quad (7)$$

and

$$W \leq \frac{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t}. \quad (8)$$

Next, consider the case  $b > 0$  and  $w_i \leq 0$  for  $\forall i$  in (4). As before, we consider the two cases for  $w$  ( $1 + W > 0$  and  $W < -1$ ). Again, with the help of the triangle inequality, the interpolation predictor  $\hat{\mathbf{y}}(\tau)$ , and the linear approximation  $\mathbb{L}_t(x)$ , one gets

$$W \leq -\frac{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t} \quad (9)$$

and

$$W \geq -\frac{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t} \quad (10)$$

Combining (7), (8), (9), and (10) completes the proof.  $\square$

<sup>2</sup>A detailed proof is provided in Appendix 6.

Figure 1: **Dataset for M-C:** In the figure below, we see the CAD train, dev and test splits with number of samples for each label. It is clear from the figure that the majority of the samples are neutral.

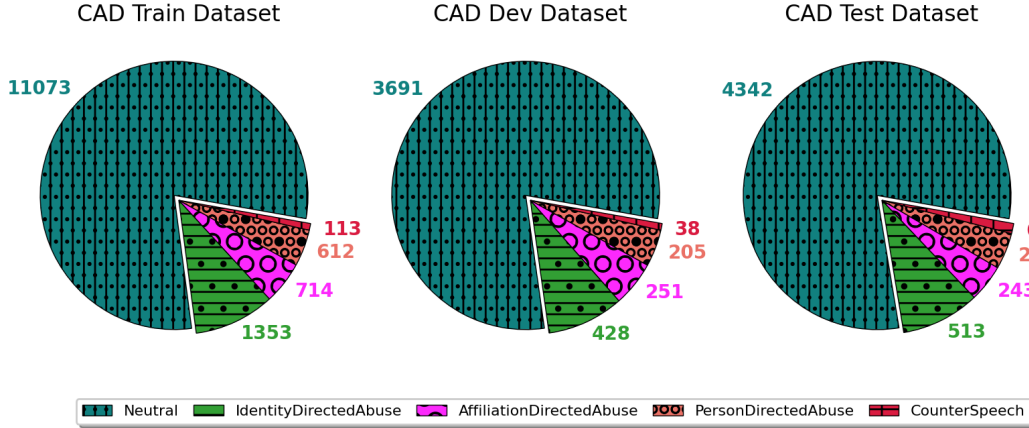
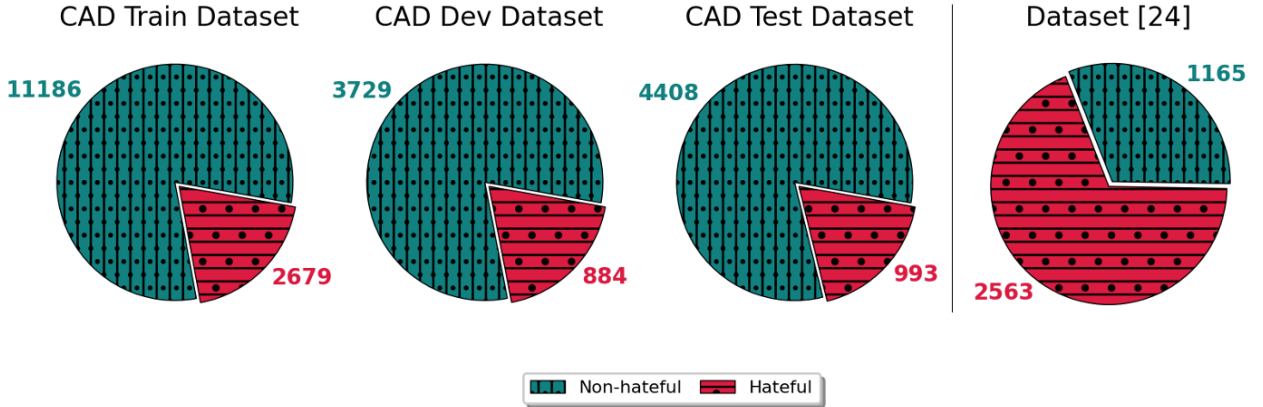


Figure 2: **Dataset for B-C:** In the figure below, we see the CAD train, dev and test dataset with number of samples for "Non-hateful" and "Hateful" category. The ratio between "Non-hateful" to "Hateful" category is approximately 80 % to 20 %. For the dataset [24], the ratio is approximately 32 % to 68 %.



The result of Theorem 1 shows that MLT tends to balance the contributions of the individual models to the joint prediction, which is important for obtaining better combined performance than that of the individual models. We illustrate this with computational examples in Section 4.

Finally, we assess the performance of MLT-plus-TM, trained with CAD data, on a distinct high-quality human-labeled dataset [24].

The dataset [24] encompasses an extensive array of functional tests within both the "Hateful" and "Non-hateful" categories. We have confirmed with the authors of CAD [23] and [24] that these two datasets were independently constructed without any overlap between them. However, it's noteworthy that both datasets adhere to the same well-defined OOPS taxonomy, rendering them suitable for rigorous assessment of our binary classifier ( B-C) through transfer learning.

Significantly, the two datasets exhibit distinctly different class distributions. In CAD, the ratio between "Hateful" and "Non-hateful" instances stands at approximately 20 % to 80 %. This distribution aligns more closely with real-world scenarios where "Non-hateful" content predominates. In contrast, the distribution in [24] is skewed in the opposite direction, with a ratio of 68.8 % "Non-hateful" to 31.2 % "Hateful". It's important to note that the dataset [24] is designed to uncover model weaknesses, intentionally featuring a considerably larger proportion of challenging cases than they occur naturally. Consequently, this skewed distribution in [24] is not reflective of a real-world class distribution. In situations where such distribution discrepancies exist between training and test data, the threshold-moving (TM)

technique might shift its "favor" in a direction contrary to its intended effect. We offer illustrative examples and delve into this topic further in Section 4.2.

### 3.1 MLT architecture

The MLT combiner architecture consists of the following components: (i) individual base models characterized by the embeddings they use, and ii) the method for combining the models.

#### 1. Individual base models

- **BERT** [8]: It is one of the state-of-the-art model for solving various tasks in the area of natural language processing. We use an instance from [44] pre-trained on real-world text from Wikipedia (en.wikipedia.org) and BooksCorpus [45].
- **BERTweet** [46]: Based on BERT, it is the first public large-scale pre-trained language model for English Tweets. It was pre-trained on 850M English Tweets (cased).
- **Bigbird** [47]: It is a sparse-attention-based transformer model that extends BERT to much longer sequences.
- **Bloom** [48]: It emerged as world's largest open multi-lingual language model. Bloom has 176B parameters and is able to generate text in 46 languages.
- **XLNet** [9]: It emerged as an extension to Transformer-XL [49], since it borrows the recurrence mechanism from Transformer-XL, to establish long-term dependencies and, hence, XLNet is a popular choice in the area of long-range NLP tasks.

All five individual models were created using the `AutoModelForSequenceClassification` from the Hugging Face Transformers library [50]. This is a generic model class designed for loading pre-trained transformer models. It additionally incorporates a classification head on top of the model outputs, specifically utilized for tasks involving classification. The same `CADtrain` data were used as training data; however, they are utilized at distinct level of granularity for M-C and B-C.

#### 2. Combiner Method

We subsequently construct M-C and B-C MLT combiners by amalgamating five individual models. In both classifications, the predicted probabilities on `CADdev` data from the individual models are used as input within a neural network architecture. The hidden layer has only one node responsible for combining the input values. The resulting output entails a consolidated, singular predicted probability value, post application of sigmoid activation, in accordance with equation (1)-(4).

In B-C MLT, only a single combiner on "Hateful" class is necessary. The outcome corresponds to the merged probability of the "Hateful" class.

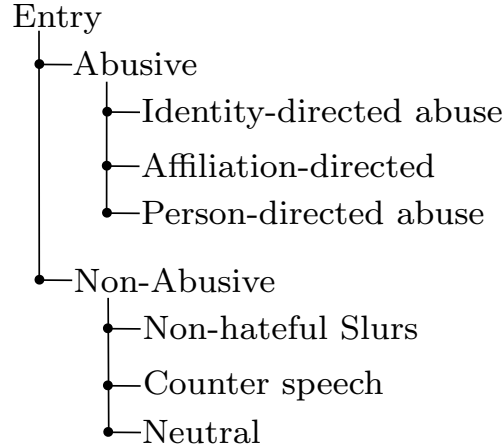
Multi-label classification algorithms can be categorized into two different groups [51]: (i) problem transformation methods, and (ii) algorithm adaptation methods. Binary relevance (BR) [52, 53] is a popular problem transformation approach that trains binary classifiers, one for each class in the datasets. One vs. the rest (OVR) is a typical type of the BR approach. BR has faced criticism due to its assumption of label independence within a multi-label dataset, disregarding potential correlations among labels. Nevertheless, despite this constraint, the BR method remains a straightforward and efficient strategy for tackling multi-label classification challenges, alongside the methods in the second group that directly handle multi-label data.

Within the CAD dataset [23], only 1.94% of entries in the training set contain more than one primary category. The training data have a label cardinality of 1.02. Unlike many other multi-label classifications where entries with multiple categories constitute a significant portion of the data, the label correlations within CAD are not dominantly strong.

In M-C MLT, we extend the aforementioned classification task approaches to construct our MLT combiners. We explored both strategies in M-C MLT: a BR-like method where we decompose the  $N$  columns and create separate combiners for each column, and a neural network approach to build a single combiner that learns a direct mapping from probabilities across all  $N$  columns to the output of combined probabilities for all  $N$  columns in simultaneously.

Throughout our experiments, both methods yielded comparable performance. Despite the fact that the second method captures label correlations, the BR method, which assumes label independence, performs comparatively well due to the label structure of the CAD dataset. We choose the BR method

Figure 3: **Class hierarchy** (Adapted from [13])



as our final choice in MLT due to its simplicity, efficiency, and its seamless integration with training threshold cut-off values later on – a method also based on BR principles.

In the BR method, we employ an *N-separate combiners approach*: we decompose the  $N$  columns and construct an individual combiner for each of them. Every combiner merges the predicted probabilities from the five base models for a particular column into a unified result. This process is iterated for all  $N = 5$  columns within the CAD dataset, yielding the following combiners: "Neutral" combiner, "Identity-directed abusive" combiner, "Affiliation-directed abuse" combiner, "Person-directed abuse" combiner, and "Counter Speech" combiner. The composition of verdicts from all binary combiners constitutes the multi-label output.

### 3.2 Data Composition

CAD is a high-quality labeled dataset with a granular level and with a great coverage of both "Abusive" and "Non-abusive" sub-categories, cf., Figure 3. We use five sub-categories in our M-C, i.e., we exclude "Non-hateful Slurs". "Neutral" entries dominate, accounting for 79.8 % of the data. "Identity-directed Abuse" accounts for 9.9 %, "Affiliation-directed Abuse" (5.0 %), "Person-directed Abuse" (4.0 %) and "Counter Speech" (0.8 %). "Non-hateful Slurs" accounts for only 0.5 % and is removed from consideration [23]. The distribution of the five sub-categories that we use from CAD dataset in M-C, for each of the `train`, `dev` and `test` splits is shown in Figure 1.

In the case of B-C, for the CAD dataset, the top-level entries of "Abusive" and "Non-abusive" are mapped to "Hateful" and "Non-hateful" classes, i.e., for each of the five sub-categories, we use "Neutral", and "Counter Speech" as the "Non-hateful" category and "Identity-directed Abuse", "Affiliation-directed Abuse" and "Person-directed Abuse" as the "Hateful" category, respectively (cf., Figure 2). After the mapping, the ratio between the two classes is about 20 % to 80 % for all the three `train`, `dev` and `test` splits.

CAD is used in all aspects of training for both M-C and B-C including individual models, MLT combining weights, and threshold cut-off values for TM.

The independent dataset [24] is used to test above trained systems. The ratio between "Hateful" to "Non-hateful" categories in the dataset [24] is around 68 % to 32 % (cf., Figure 2). Unlike in CAD where most of samples are "Non-hateful" as in real-world situation, the dataset [24] has a much larger portion of "Hateful" class, but it is deliberately crafted that way for testing system performance purpose.

Both CAD and [24] are high quality and share consistent OOPS interpretation.

### 3.3 Threshold-moving (TM)

Models trained to minimize the overall loss function with imbalanced data tend to exhibit bias in favor of the majority class, using the default 0.5 threshold to map probabilities to category labels would more likely assign minority class data to the majority class. Threshold-moving (TM) is utilized to counter this bias by enhancing the likelihood of correctly classifying minority data. An optimal threshold value is trained by optimizing F1 on the positive class (minority class) "Hateful" in B-C after MLT.

---

**Algorithm 1** Optimal threshold training in TM

---

**Input:** classes, step size, base and max threshold

```
1:  $\mathbf{I}^{(C)} \leftarrow$  number of classes  $\triangleright$  Initialization
2:  $\delta \leftarrow$  step size,  $\alpha \leftarrow$  base threshold,  $\beta \leftarrow$  max threshold
3:  $\mathbb{D} \leftarrow$  predicted probabilities from individual model
4: execute Meta Learning Technique on  $\mathbb{D}$   $\triangleright$  MLT
5:  $\mathcal{OptTh} \leftarrow \phi$   $\triangleright$  Stores optimal threshold for each class
6: for  $i = 1$  to  $\mathbf{I}^{(C)}$  do
7:    $\mathcal{DICT}^{F_1} \leftarrow \phi$ 
8:    $\mathbb{D}'[i] \leftarrow$  build data from  $\mathbb{D}$  corresponding to  $i^{th}$  class
9:   for  $\Delta = \alpha$  to  $\beta$  do  $\triangleright$  TM
10:     $\gamma \leftarrow F_1$  score for  $\mathbb{D}'[i]$  using  $\Delta$  as threshold
11:     $\mathcal{DICT}^{F_1}[\Delta] \leftarrow \gamma$ 
12:     $\Delta \leftarrow \Delta + \delta$ 
13:   end for
14:    $\mathcal{OptTh}[i] \leftarrow \text{key}(\max(\mathcal{DICT}^{F_1}))$ 
15: end for
16: Return ( $\mathcal{OptTh}$ )  $\triangleright$  Return optimal threshold
```

---

In the case of M-C, employing the default threshold of 0.5 would assign the predicted class as the one with the highest probability equal to or greater than 0.5. Various TM strategies are available. The simplest approach is training a single global threshold value for all classes to achieve the best performance, known as the score-cut strategy (S-Cut). The default threshold value of 0.5 is a special instance of S-Cut. S-Cut returns a label set where scores exceed or equal the trained threshold value. The rank-cut strategy (R-cut) aims to return a set of precisely  $n$  labels with the highest scores. However, in our application, most samples possess only one class, with only a few having more than one. Thus, R-cut is not a suitable candidate. A more complex strategy, called CS-Cut, exists. The class-specific score-cut (CS-Cut) [54] optimizes a vector of thresholds, each corresponding to a label. It then employs the same thresholding operation as S-Cut. CS-cut proves useful when certain labels are harder to classify than others, as is the case in this application. Less dominant classes are more challenging to classify accurately. Hence, we adopt this method, as depicted in Algorithm 1. CS-cut can provide higher accuracy but is more susceptible to overfitting compared to S-Cut.

Within CS-Cut, multiple variants exist, ranging from simple to more intricate approaches. In CS-Cut for M-C thresholding, we adopt a binary relevance (BR)-like method. BR is among the most common thresholding methods in multi-label classification [55]. It’s both simple and highly efficient, serving as a strong baseline for comparison with more advanced methods. Given our use of BR in MLT, training threshold values and applying them through TM on MLT output are streamlined and straightforward.

As elaborated in Section 3.1 for M-C, five MLT combiners generate predicted probabilities for five classes. Post-MLT, we train a threshold value for TM using the MLT output for each class, similar to B-C. We repeat this process for all classes. The resultant final threshold values for all classes constitute a tuple of five values ranging between 0 and 1. This TM approach assumes no class dependency in M-C regarding the effect of a threshold value in one class on those from other classes. Therefore, extending TM training from B-C to M-C is straightforward. This assumption of class independence greatly simplifies the task of identifying suitable threshold values for M-C. The macro-averaged F1 score for M-C is the average of F1 scores calculated in B-C for each class.

TM alters the decision-making dynamics between the minority and majority classes and impacts evaluation results. Depending on the most important for the application, finding the optimal threshold values becomes necessary to maximize performance with respect to the chosen metric, or to satisfy lower bound constraints on other metrics while optimizing the primary one.

In our experiments, TM improves the *macroF1* score in both M-C (0.439 to 0.467) and B-C (0.714 to 0.722) on CAD test data, at the expense of a slight accuracy decline. This aligns with expectations, as further discussed in Section 4.3. We also applied TM trained on CAD to a different dataset [24] in B-C. Through this transfer learning, we achieve performance improvements in both *macroF1* (0.688 to 0.69) and accuracy (0.717 to 0.741). More extensive details and analysis are provided in Section 4.3.

## 4 Computational Results

In this section, we provide computational results with analysis. We use macro-averaged F1 (*macroF1*) score as the main metric in performance evaluation because datasets in the OOPS domain are highly imbalanced, and we want to



Table 1: CAD M-C scores on the test set

Model(#)	Accuracy	<i>macroF1</i>
<b>DistilBERT Baseline 1 (n=5)</b>	0.769 (0.005)	0.44 (0.007)
<b>BERT Baseline 2 (n=5)</b>	0.762 (0.005)	0.455 (0.006)
<b>Bigbird (n=2)</b>	0.806(0.0076)	0.418(0.0086)
<b>BERT (n=2)</b>	0.806(0.0041)	0.424(0.0025)
<b>BERTweet (n=2)</b>	0.806(0.0043)	0.441(0.0024)
<b>Bloom (n=2)</b>	0.785(0.0075)	0.304(0.0035)
<b>XLNet (n=2)</b>	0.802(0.0039)	0.439(0.0014)
<b>MLT combiner (n=32)</b>	0.821(0.0031)	0.439(0.0049)
<b>MLT-plus-TM (n=32)</b>	0.780(0.0096)	0.467(0.0095)

treat all classes equally, regardless of their presence volume in the datasets. We also provide accuracy in reporting as a reference, but it should be considered a secondary metric.

All results are reported as means with standard deviations from multiple runs ( $n$  represents the number of runs in all tables). In case of individual models, two runs from an individual model are used in performance evaluation and in joining combiners. In case of combiners (MLT and MLT-plus-TM runs), if each individual model has two runs, combining five individual models with all possible permutations create 32 runs.

Three sets of results are presented: M-C (Table 1) and B-C (Table 3) on CAD test, and B-C (Table 4) on dataset [24] using transfer learning. For both M-C and B-C on CAD, five individual models are trained using the CAD training dataset. Subsequently, predictions are obtained on the CAD dev dataset and used to train combining weights through the MLT structure explained in Section 3.1. Following the MLT step, threshold cut-off values are trained (detailed in Section 3.3) using the output from MLT. The cut-off values are then applied through TM on the predictions from the CAD test dataset to obtain the final MLT-plus-TM results.

For M-C, we use the reported performance on the same tasks in [24] as baselines to compare with our results.

In B-C transfer learning, we utilize the trained MLT combining weights and TM cut-off values from CAD to generate MLT and MLT-plus-TM results on dataset [24]. Finally, we compare our results with the reported ones in [24] and provide observations and analysis.

All experiments reported here were developed in Python 3.8 with the TensorFlow 2.9.1 [56] library and PyTorch Lightning (version 1.7.7). The source code for reproducing the results reported here along with appropriate meta data is available at <https://data.nist.gov/od/id/mds2-3074>.

The individual participating open models, **BigBird** (M1), **BERT** (M2), **BERTweet** (M3), **Bloom** (M4), and **XLNet** (M5), were trained on a professional Graphics Processing Unit (GPU) cluster with 8 NVIDIA Tesla V100 (32 GB each). The experiments for **MLT** and **TM** were carried on a 2017 MacBook Pro with 3.1 GHz Quad-Core Intel Core i7 and 16 GB RAM, *without* Graphics Processing Unit (GPU) acceleration.

#### 4.1 Performance of M-C

Table 1 presents the performance at all category levels. Among the five participating individual models, BERTweet demonstrates the best overall performance, achieving the highest accuracy and *macroF1* scores. MLT surpasses all individual models in accuracy and only slightly lags behind BERTweet in *macroF1* score. MLT-plus-TM further improves the *macroF1* score to the highest level, albeit at the expense of a slight decrease in accuracy. However, in imbalanced classification, relying solely on accuracy can be misleading as it may not adequately account for minority classes. On the other hand, *macroF1* score is a more meaningful metric in an imbalanced classification settings. Our MLT-plus-TM outperforms the baseline in both *macroF1* score (0.467 vs. 0.455) and accuracy (0.780 vs. 0.769) and achieves the highest *macroF1* score in all runs.

Table 2 presents the performance at a per-category level. MLT-plus-TM achieves the highest F1 score in all categories, except for ‘Counter Speech’. However, it is worth noting that the ‘Counter Speech’ category has a very low prevalence (0.8%), which makes the performance in this category less reliable and should be interpreted with caution.

#### 4.2 Performance of B-C

Table 3 displays the performance on CAD test. Once again, BERTweet emerges as the top-performing individual model. But MLT outperforms all individual models in terms of both *macroF1* score and accuracy. MLT-plus-TM further

Table 2: CAD M-C Scores per category on the test set. Scores: P=Precision, R=Recall, F1=F1

Model(#)	Neutral N=4342			Identity-directed N=513			Affiliation-directed N=243			Person-directed N=237			Counter Speech N=66		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
DistilBERT Baseline 1 (n=5)	0.88	0.917	0.898	0.414	0.473	0.441	0.368	0.45	0.405	0.359	0.404	0.38	0.083	0.073	0.076
BERT Baseline 2 (n=5)	0.883	0.922	0.902	0.411	0.51	0.455	0.368	0.481	0.416	0.356	0.488	0.411	0.107	0.088	0.091
Bigbird (n=2)	0.940	0.871	0.904	0.298	0.621	0.402	0.327	0.404	0.361	0.317	0.481	0.379	0.030	0.084	0.044
BERT (n=2)	0.933	0.876	0.903	0.322	0.581	0.414	0.333	0.497	0.396	0.285	0.467	0.354	0.038	0.179	0.053
BERTweet (n=2)	0.926	0.882	0.904	0.356	0.547	0.429	0.377	0.466	0.416	0.382	0.441	0.409	0.030	0.129	0.046
Bloom (n=2)	0.964	0.845	0.901	0.157	0.527	0.233	0.191	0.367	0.245	0.122	0.239	0.142	0.0000	1.0000	0.0000
XLNet (n=2)	0.927	0.876	0.901	0.386	0.517	0.442	0.329	0.470	0.387	0.325	0.469	0.383	0.068	0.109	0.083
MLT combiner (n=32)	0.871	0.960	0.913	0.678	0.338	0.451	0.573	0.343	0.429	0.609	0.302	0.403	1.0000	0.0000	0.0000
MLT-plus-TM (n=32)	0.866	0.968	0.914	0.504	0.478	0.488	0.418	0.476	0.443	0.471	0.427	0.441	0.269	0.046	0.050

enhances the *macroF1* score to its highest level, albeit with a slight decrease in accuracy. Both M-C and B-C on CAD test exhibit very similar patterns in TM.

Table 4 presents the performance on dataset [24], which represents a transfer learning scenario involving an out-of-distribution target dataset. Our objective is to assess how well MLT and TM trained from CAD generalize to dataset [24]. We believe that models trained with CAD data tend to capture the real-world imbalanced distribution between the 'Hateful' and 'Non-hateful' classes, which may introduce some bias towards predicting the majority "Non-hateful" class. Therefore, TM is applied to counterbalance this bias. It is important to note that dataset [24] deliberately has a much larger portion of the 'Hateful' class, resulting in an opposite class distribution compared to CAD. This deliberate design choice serves the purpose of testing and evaluation.

Table 3: B-C performance on CAD test data. Scores: Accuracy (Recall) and F1 with Std by test case label

Model(#)	Hateful N=899		Non-hateful N=4408		Total N=5307	
	Accuracy(Recall)	F1	Accuracy(Recall)	F1	Accuracy(Recall)	<i>macroF1</i>
Bigbird (n=2)	0.572(0.0047)	0.460(0.0001)	0.882(0.0002)	0.911(0.0006)	0.847(0.0008)	0.685(0.0003)
BERT (n=2)	0.546(0.0145)	0.492(0.0056)	0.891(0.0004)	0.907(0.0024)	0.843(0.0037)	0.700(0.0040)
BERTweet (n=2)	0.580(0.0421)	0.510(0.0260)	0.895(0.0099)	0.911(0.0056)	0.850(0.0068)	0.711(0.0102)
Bloom (n=2)	0.499(0.0430)	0.372(0.0653)	0.870(0.0131)	0.899(0.0082)	0.827(0.0096)	0.635(0.0286)
XLNet (n=2)	0.531(0.0260)	0.489(0.0028)	0.892(0.0029)	0.904(0.0048)	0.839(0.0067)	0.697(0.0010)
MLT (n=32)	0.430(0.0172)	0.510(0.0081)	0.948(0.0055)	0.918(0.0016)	0.860(0.0021)	0.714(0.0038)
MLT-plus-TM(n=32)	0.500(0.0304)	0.533(0.0080)	0.923(0.0122)	0.912(0.0039)	0.851(0.0052)	0.722(0.0031)

Table 4: B-C performance on dataset [24]. Scores: Accuracy (Recall) and F1 with Std by test case label

Model(#)	Hateful N=2563		Non-hateful N=1165		Total N=3728	
	Accuracy(Recall)	F1	Accuracy(Recall)	F1	Accuracy(Recall)	<i>macroF1</i>
B-D	0.755	0.738	0.36	0.379	0.632	0.559
B-F	0.655	0.694	0.485	0.432	0.602	0.563
Google-P	0.895	0.841	0.482	0.563	0.766	0.702
Bigbird (n=2)	0.847(0.0021)	0.559(0.0120)	0.394(0.0033)	0.535(0.0014)	0.547(0.0070)	0.547(0.0067)
BERT (n=2)	0.795(0.0028)	0.766(0.0183)	0.506(0.0233)	0.539(0.0012)	0.690(0.0164)	0.653(0.0097)
BERTweet (n=2)	0.805(0.0099)	0.810(0.0374)	0.613(0.0924)	0.573(0.0018)	0.739(0.0363)	0.692(0.0196)
Bloom (n=2)	0.696(0.0147)	0.156(0.0475)	0.313(0.0011)	0.466 (0.0058)	0.346(0.0099)	0.311(0.0208)
XLNet (n=2)	0.776(0.0128)	0.731(0.0249)	0.455(0.0147)	0.498(0.0185)	0.651(0.0166)	0.614(0.0032)
MLT (n=32)	0.743(0.0321)	0.783(0.0149)	0.660(0.0325)	0.593(0.0065)	0.717(0.0128)	0.688(0.0081)
MLT-plus-TM(n=32)	0.833(0.0346)	0.815(0.0106)	0.540(0.0597)	0.565(0.0255)	0.741(0.0087)	0.690(0.0109)

For ease of comparison, we include the following model performances reported in [24] in Table 4: two pre-trained BERT models denoted as B-D and B-F, as well as one commercial non-open model - Google Jigsaw's Perspective - denoted as Google-P. It is important to note that these models have been trained on different data; however, they are all tested on the same suite [24].

Table 4 illustrates that MLT slightly underperforms when compared to the best individual model, BERTweet, in transfer learning. We provide further analysis on this in Section 4.3.1. However, MLT-plus-TM achieves the best performance in terms of accuracy and performs nearly as well as the best model in terms of *macroF1* score.

At the category level, in the "Hateful" category (Table 4), MLT-plus-TM outperforms all individual participating models in terms of F1 score but falls slightly behind the Google-P model (0.815 vs. 0.841). In the "Non-hateful" category

(Table 4), MLT outperforms all other models, including the Google-P model (0.593 vs. 0.563), in terms of F1 score. MLT-plus-TM also marginally outperforms the Google-P model (0.565 vs. 0.563).

Overall, based on Table 4, we can observe that both MLT (0.688) and MLT-plus-TM (0.690) exhibit performance almost comparable to that of the Google-P model (0.702) in *macroF1*.

Table 5: MLT M-C performance in *macroF1* scores order on CAD test

Model Combination	<i>macroF1</i>	Accuracy
M1, M2, M3, M4, M5	0.439	0.821
M1, M2, M3, M5	0.438	0.821
M1, M3, M4, M5	0.432	0.820
M1, M2, M3, M4	0.431	0.821
M1, M2, M3	0.431	0.824
M2, M3, M4, M5	0.429	0.819
M1, M3, M5	0.429	0.821
M1, M2, M4, M5	0.428	0.820
M2, M3, M5	0.427	0.819
M1, M2, M5	0.425	0.822
M2, M3, M4	0.410	0.819
M3, M4, M5	0.407	0.819
M1, M3, M4	0.406	0.816
M2, M3	0.406	0.817
M1, M3	0.404	0.823
M3, M4	0.403	0.803
M3, M5	0.402	0.820
M1, M2, M4	0.401	0.820
M2, M4, M5	0.401	0.820
M1, M4, M5	0.399	0.814
M1, M5	0.392	0.814
M2, M5	0.391	0.805
M1, M2	0.391	0.826
M2, M4	0.384	0.805
M4, M5	0.380	0.803
M1, M4	0.368	0.806
<b>Average score</b>	<b>0.410</b>	<b>0.817</b>

### 4.3 Inference from Combiner Performance

Section 4.1 and 4.2 demonstrate that MLT-plus-TM significantly improves performance, surpassing all individual models and achieving best overall performance. Even when we transfer the learned MLT combining weights and threshold cut-off values from CAD to the dataset [24], where the data distribution differs substantially from the training data, the final MLT-plus-TM remains at least on par with the best participating model.

We also provide MLT and MLT-plus-TM results for all possible model combinations (5, 6, 7, 8, 9, and 10). The best combiner for M-C on the CAD test is (M2, M3, M5) with a *macroF1* score of 0.469, compared to 0.467 in all combinations of (M1, M2, M3, M4, M5). In B-C on the CAD test, the best combination is (M2, M3) with a *macroF1* score of 0.726, while all combinations yield 0.722. For B-C on dataset [24] transfer learning, the best combination is (M1, M3) with a *macroF1* score of 0.709, whereas all combinations yield 0.690. These results clearly demonstrate that all combinations fall short of achieving the highest possible performance. Nevertheless, we still report the performance by combining all five models (M1, M2, M3, M4, M5).

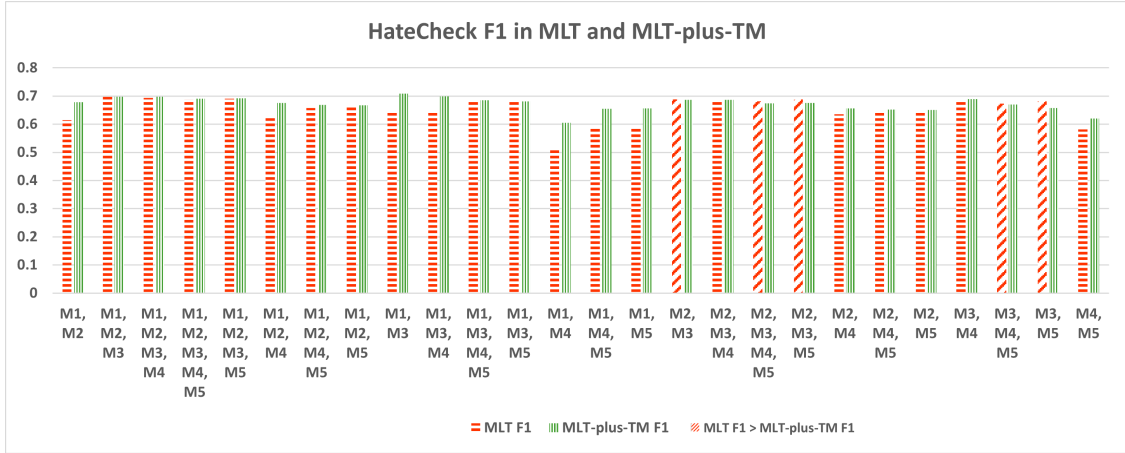
#### 4.3.1 MLT performance analysis

The results(1, 3) indicate that MLT achieves the best accuracy, surpassing the best individual model, BERTweet (0.821 vs. 0.806 in M-C and 0.86 vs. 0.85 in B-C). A plausible explanation for this is that CAD is highly imbalanced, and MLT may further enhance the advantage of the majority class, resulting in higher accuracy. However, when considering *macroF1*, which treats all classes equally and places greater emphasis on the minority class, MLT may not necessarily lead to a higher *macroF1* score. Although MLT may not directly boost the *macroF1* score, the higher accuracy achieved

Table 6: **MLT-plus-TM M-C performance in *macroF1* scores order on CAD test**

Model Combination	<i>macroF1</i>	Accuracy
M2, M3, M5	0.469	0.778
M1, M2, M3, M4, M5	0.467	0.780
M1, M3, M4, M5	0.465	0.780
M1, M2, M3, M5	0.464	0.784
M2, M3, M4, M5	0.464	0.784
M3, M4, M5	0.462	0.783
M1, M3, M5	0.461	0.787
M1, M2, M4, M5	0.459	0.786
M1, M2, M3, M4	0.459	0.786
M3, M5	0.459	0.789
M1, M2, M3	0.457	0.787
M1, M3	0.457	0.782
M2, M3	0.456	0.783
M1, M3, M4	0.456	0.783
M1, M2, M5	0.454	0.790
M2, M3, M4	0.452	0.787
M1, M4, M5	0.451	0.771
M2, M4, M5	0.449	0.781
M2, M5	0.446	0.782
M3, M4	0.446	0.794
M1, M5	0.445	0.764
M1, M2, M4	0.443	0.780
M1, M2	0.440	0.776
M4, M5	0.437	0.770
M2, M4	0.428	0.782
M1, M4	0.420	0.776
<b>Average score</b>	<b>0.453</b>	<b>0.782</b>

Figure 4: **Combiner Performance F1 scores on dataset [24] using transfer learning: MLT B-C performance and MLT-plus-TM B-C performance for all the possible combinations.** MLT B-C performance is taken from table 9, and corresponding MLT-plus-TM B-C performance is taken from table 10.



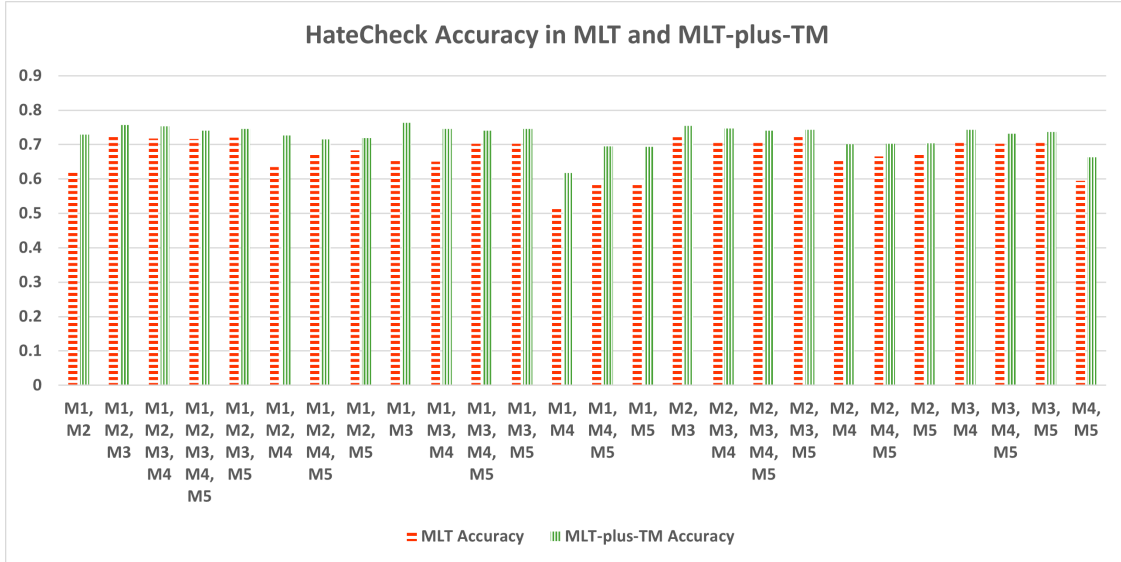
through MLT leaves more room for improvement in the *macroF1* score through TM. Consequently, MLT-plus-TM achieves the best performance, surpassing that of Bertweet (0.467 vs. 0.441 in M-C and 0.722 vs. 0.711 in B-C) in our targeted metric, macro-averaged F1 (*macroF1*). With MLT-plus-TM, both M-C and B-C on CAD achieve the best performance in *macroF1*.

The goal of MLT is to train optimized combining weights to achieve better performance than that of the individual models on a given dataset. However, when such weights are applied to a different dataset, such as in a transfer learning

Table 7: **MLT B-C performance in *macroF1* scores order on CAD test**

Model Combination	<i>macroF1</i>	Accuracy
M2, M3, M5	0.715	0.859
M1, M2, M3, M4, M5	0.714	0.860
M1, M2, M3, M5	0.714	0.860
M2, M3, M4, M5	0.713	0.860
M1, M2, M3, M4	0.709	0.859
M1, M2, M3	0.709	0.859
M2, M4, M5	0.709	0.857
M1, M3, M4, M5	0.709	0.859
M2, M5	0.708	0.858
M2, M3	0.708	0.860
M1, M2, M4, M5	0.708	0.857
M1, M3, M5	0.708	0.859
M3, M4	0.706	0.857
M3, M4, M5	0.706	0.857
M2, M3, M4	0.706	0.859
M1, M3, M4	0.706	0.860
M1, M2, M5	0.705	0.856
M1, M3	0.705	0.860
M3, M5	0.704	0.857
M1, M4, M5	0.704	0.858
M1, M2, M4	0.701	0.854
M1, M5	0.699	0.856
M2, M4	0.697	0.849
M1, M2	0.697	0.855
M4, M5	0.694	0.850
M1, M4	0.688	0.853
<b>Average score</b>	<b>0.706</b>	<b>0.857</b>

Figure 5: **Combiner Performance Accuracy scores on dataset [24] using transfer learning: MLT B-C performance and MLT-plus-TM B-C performance for all the possible combinations.** MLT B-C performance is taken from Table 9, and corresponding MLT-plus-TM B-C performance is taken from Table 10.



scenario, the dynamics of the combination effect may shift. In Table 4 on dataset [24], MLT slightly underperforms

Table 8: MLT-plus-TM B-C performance in *macroF1* scores order on CAD test

Model Combination	<i>macroF1</i>	Accuracy
M2, M3	0.726	0.848
M1, M3, M5	0.724	0.845
M1, M2, M3, M5	0.723	0.852
M1, M3, M4, M5	0.723	0.847
M1, M2, M3, M4	0.723	0.847
M1, M2, M3	0.722	0.845
M1, M2, M3, M4, M5	0.722	0.851
M2, M3, M4	0.722	0.848
M2, M3, M5	0.721	0.847
M2, M3, M4, M5	0.720	0.847
M1, M3	0.720	0.843
M3, M4, M5	0.719	0.847
M1, M3, M4	0.719	0.846
M3, M5	0.717	0.839
M1, M2, M4, M5	0.717	0.844
M1, M2, M5	0.716	0.842
M2, M4, M5	0.714	0.842
M3, M4	0.713	0.845
M1, M4, M5	0.712	0.838
M1, M5	0.712	0.840
M2, M5	0.712	0.844
M1, M2, M4	0.711	0.833
M1, M2	0.708	0.833
M2, M4	0.705	0.837
M1, M4	0.701	0.833
M4, M5	0.699	0.831
<b>Average score</b>	<b>0.716</b>	<b>0.843</b>

when compared with the best individual model, Bertweet (0.717 vs. 0.739 in accuracy and 0.688 vs. 0.692 in *macroF1*). This indicates that the combining weights trained on CAD are not optimal for dataset [24].

When we examining individual model performance in Table 3 and 4, we notice that the inference on CAD test and dataset [24] does not always follow the same pattern. Although BERTweet (M3) has the best performance and Bloom (M5) is the weakest in both datasets, in Table 3 on CAD test, BERTweet (M3), Bigbird (M1), BERT (M2) and XLNet (M5) are very close in performance. However, in Table 4 on dataset [24], Bertweet stands out from the rest. Following BERTweet are BERT and XLNet. Bigbird performs significantly worse than the three better models, and Bloom is at the bottom. The performance gap among individual models on dataset [24] is much wider than on CAD. Bigbird and Bloom show the sharpest performance drop on dataset [24]. The difference in the pattern of relative model performance between CAD and dataset [24] shows that the optimal combining weights for CAD may not be optimal on dataset [24], which is not surprising in transfer learning.

In MLT-plus-TM, we use a separate mechanism, TM, to boost final performance in *macroF1* after MLT.

#### 4.3.2 TM performance analysis

As shown in Figure 2, the ratio between "Non-hateFul" to "Hateful" is 80 % to 20 % on CAD. We train threshold cut-off values to detect more instances of minority class, "Hateful", in order to maximize the overall *macroF1* score. It's important to note that the trained cut-off value is lower than the default 0.5 threshold. This adjustment increases the number of true positives (more instances of "Hateful" instances are correctly classified as such), but it can also decrease the number of true negatives (instances of "Non-hateful" are correctly classified as such). Consequently, the effect of TM effect on accuracy depends on the combined count of true positives and true negatives. If this count increases, accuracy increases; otherwise, accuracy decreases. Considering that 80 % of data are labeled as "Non-hateful" in the CAD dataset, the decrease in true negatives may be much larger than the increase in true negatives. As a result, the overall accuracy after applying TM may decrease in CAD.

Table 9: **MLT B-C performance in *macroF1* scores order on dataset [24]**

<b>Model Combination</b>	<b><i>macroF1</i></b>	<b>Accuracy</b>
M1, M2, M3	0.697	0.723
M1, M2, M3, M4	0.693	0.718
M1, M2, M3, M5	0.690	0.720
M1, M2, M3, M4, M5	0.688	0.717
M2, M3	0.688	0.722
M2, M3, M5	0.688	0.726
M3, M4	0.683	0.710
M2, M3, M4	0.681	0.712
M1, M3, M5	0.681	0.702
M1, M3, M4, M5	0.681	0.703
M2, M3, M4, M5	0.681	0.713
M3, M5	0.681	0.712
M3, M4, M5	0.673	0.703
M1, M2, M5	0.661	0.684
M1, M2, M4, M5	0.657	0.678
M2, M5	0.646	0.674
M1, M3	0.644	0.654
M1, M3, M4	0.642	0.651
M2, M4, M5	0.641	0.666
M2, M4	0.635	0.659
M1, M2, M4	0.626	0.635
M1, M2	0.615	0.622
M1, M4, M5	0.586	0.590
M1, M5	0.585	0.589
M4, M5	0.582	0.594
M1, M4	0.515	0.516
<b>Average score</b>	<b>0.652</b>	<b>0.673</b>

This effect is observed in both M-C and B-C on CAD. In M-C (Table 5, 6), the *macroF1* score for all combiners increases from 0.410 to 0.458, but the averaged accuracy decreases from 0.817 to 0.782. In B-C (Table 7, 8), the *macroF1* score increases from 0.706 to 0.716, while the averaged accuracy decreases from 0.875 to 0.843. This trade-off is acceptable because *macroF1* is a more important metric as in this highly imbalanced data situation.

The effect of applying TM in a dataset with an opposite class distribution compared to the training data is more complex.

In dataset [24], when we apply threshold cut-off value trained from CAD on the 'Hateful' class, which is lower than the 0.5 cut-off, it increases the chance of detecting "Hateful" instances more easily. Consequently, more "Hateful" samples are detected, resulting in an increase in the number of true positives for the "Hateful" class. While the number of false negatives for the "Non-hateful" class may also increase, and the number of true negatives decreases, since the majority of the data is "Hateful" (68 %), the increase in true positives in "Hateful" is likely to be larger than the decrease in true negatives in "Non-hateful". As a result, the overall accuracy is more likely to improve. This is opposite to what happens in CAD and is not unexpected. Figure 5 illustrates this effect: in all 26 combiners, MLT-plus-TM achieves higher accuracy scores than MLT.

The case of *macroF1* is more delicate. The precision for the "Hateful" class may increase, while the precision for the "Non-hateful" class may decrease. The recall for both classes may also be affected. The *macroF1* score is the average of the F1 scores for both classes, so it will depend on how these changes in precision and recall balance out. Figure 4 shows that in all 26 combiners, 21 combiners achieve higher *macroF1* scores after TM. Only 5 (represented by the upward diagonal strips pattern bars) produce slightly lower *macroF1* scores.

Overall, the threshold cut-off values trained in CAD work very well on dataset [24], despite the opposite class composition of the two datasets. From Table 9 and 10, both *macroF1* and accuracy scores are improved after applying TM using the CAD trained threshold values: average gain in *macroF1* for the 26 combiners is from 0.652 to 0.672, and the average gain in accuracy is from 0.673 to 0.725. The reason why TM works well even when applied to a dataset with an opposite class composition may be that the models trained on CAD have an implicit bias toward the majority class, and TM helps to offset this bias. When the offset brings the predictions closer to the ground truth, it results in

Table 10: **MLT-plus-TM B-C performance in *macroF1* scores order on dataset [24]**

Model Combination	<i>macroF1</i>	Accuracy
M1, M3	0.709	0.763
M1, M3, M4	0.699	0.745
M1, M2, M3, M4	0.698	0.753
M1, M2, M3	0.697	0.757
M1, M2, M3, M5	0.692	0.745
M1, M2, M3, M4, M5	0.690	0.741
M3, M4	0.689	0.743
M2, M3	0.687	0.754
M2, M3, M4	0.686	0.747
M1, M3, M4, M5	0.685	0.741
M1, M3, M5	0.681	0.745
M1, M2	0.678	0.729
M1, M2, M4	0.676	0.727
M2, M3, M5	0.675	0.743
M2, M3, M4, M5	0.674	0.740
M3, M4, M5	0.670	0.731
M1, M2, M4, M5	0.668	0.715
M1, M2, M5	0.667	0.719
M3, M5	0.658	0.737
M1, M5	0.656	0.694
M2, M4	0.656	0.701
M1, M4, M5	0.655	0.695
M2, M4, M5	0.652	0.703
M2, M5	0.651	0.704
M4, M5	0.620	0.663
M1, M4	0.605	0.617
<b>Average score</b>	<b>0.672</b>	<b>0.725</b>

performance gains. Compared to MLT, MLT-plus-TM improves both accuracy (0.741 vs 0.717) and *macroF1* (0.69 vs 0.688) in out-of-distribution transfer learning on dataset [24] (4).

### 4.3.3 Reinforcing the Weak Model

The combiner is not necessarily composed of the best participating models. We try to develop an intuition about how superior and inferior models can combine together and obtain a better metric overall. First, we look at the overall metric results, then, we consider how the weights of the combiner adjust themselves among the superior and the inferior models.

Bloom (M4) has the weakest performance among the five participating models. When M4 is compared with M1, M2, M3 or M5, its *macroF1* score and accuracy results are inferior for both the CAD *test* and dataset [24]. Though the results of M4 are inferior, it may still be possible that the hyperspace of M4 might carry some useful information that could support other models in establishing the correct inference. This is where MLT plays a big role, the combinations gives M4 some opportunity to shine in the company of other superior models. Similarly, the case with M1 and M5 when compared to M2 and M3. Further, when M4 is combined with any other model, there is a considerable increment in both the metrics (*macroF1* score and accuracy), compared to M4 alone; while when compared with the individual metrics of the participating models, there is always a small increment. One of the major reasons for this observation is the higher weight of the other participating models compared to M4.

If we analyze the weights of the participating models in the combiner, we find that the models that are inferior to their participating peers, tend to have a lower weight. For example, in the case of a combiner with models M1 and M4, the weight of model M1 is 1.246, while the weight of model M4 is 0.885. For a three model combiner with models M1, M2 and M4, the weight of model M1 is 0.692, model M2 is 0.755, and model M4 is 0.421, respectively. In the case of a combiner with four models (M1, M2, M3 and M4), we have the weights as 0.566, 0.615, 0.697 and 0.341 for models M1, M2, M3 and M4, respectively. Similar weight distribution is seen when all the models participate for the combiner, i.e., 0.327, 0.367, 0.496, 0.225 and 0.419 for M1, M2, M3, M4 and M5, respectively. Further, from these examples one can see that the weights tend to reflect the individual strength of the models. The inferences and the experimental results



are in agreement with *Vassilev et.al.* [36], where the authors show that in a combination of superior and inferior model, the inferior model becomes a supporting agent to the combiner, resulting into overall improvements of metrics when compared against the metrics of individual models. Overall, combining more models results in better performance.

Table 11: B-C Accuracy scores on dataset [24] 29 functionalities

	Functionality	B-D	B-F	Google-P	BIGBIRD	BERT	BERTWEET	BLOOM	XLNET	MLT_1	MLT-plus-TM_1	MLT_2	MLT-plus-TM_2	MLT_3	MLT-plus-TM_3
F1	derog_neg_emote_h	0.886	0.907	0.986	0.257	0.657	0.729	0.086	0.579	0.621	0.779	0.679	0.814	0.600	0.893
F2	derog_neg_attrib_h	0.886	0.843	0.957	0.436	0.864	0.921	0.107	0.764	0.871	0.964	0.907	0.964	0.850	0.993
F3	derog_dehum_h	0.914	0.807	0.986	0.414	0.914	0.921	0.029	0.814	0.886	0.921	0.900	0.936	0.886	0.971
F4	derog_impl_h	0.714	0.614	0.850	0.207	0.443	0.507	0.100	0.586	0.471	0.614	0.493	0.643	0.436	0.729
F5	threat_dir_h	0.872	0.759	1.000	0.474	0.767	0.669	0.023	0.812	0.759	0.850	0.789	0.910	0.737	0.947
F6	threat_norm_h	0.914	0.836	1.000	0.357	0.807	0.786	0.079	0.664	0.786	0.879	0.821	0.914	0.757	0.943
F7	slur_h	0.604	0.410	0.660	0.722	0.861	0.938	0.000	0.771	0.924	0.958	0.931	0.972	0.910	1.000
F8	slur_homonym_h	0.667	0.700	0.633	0.767	0.600	0.533	0.800	0.633	0.633	0.533	0.600	0.467	0.667	0.467
F9	slur_reclaimed_h	0.395	0.333	0.284	0.383	0.358	0.506	0.667	0.099	0.383	0.259	0.370	0.272	0.370	0.160
F10	profanity_h	0.829	0.729	1.000	0.593	0.914	0.886	0.229	0.821	0.871	0.921	0.879	0.929	0.857	0.964
F11	profanity_nh	0.990	1.000	0.980	0.950	0.920	0.950	0.990	0.850	0.930	0.900	0.930	0.900	0.930	0.880
F12	ref_subs_clause_h	0.871	0.807	0.993	0.379	0.821	0.857	0.471	0.843	0.807	0.914	0.836	0.921	0.779	0.957
F13	ref_subs_sent_h	0.857	0.707	1.000	0.474	0.857	0.782	0.534	0.722	0.797	0.880	0.812	0.872	0.789	0.925
F14	negate_pos_h	0.850	0.607	0.964	0.286	0.686	0.621	0.057	0.593	0.571	0.771	0.621	0.757	0.564	0.836
F15	negate_neg_h	0.128	0.120	0.038	0.789	0.421	0.511	0.940	0.549	0.556	0.338	0.534	0.346	0.586	0.211
F16	phrase_question_h	0.807	0.750	0.993	0.500	0.614	0.714	0.000	0.500	0.664	0.793	0.693	0.764	0.636	0.836
F17	phrase_opinion_h	0.857	0.759	0.985	0.466	0.827	0.805	0.045	0.669	0.759	0.887	0.797	0.910	0.737	0.940
F18	ident_neutral_h	0.206	0.587	0.841	0.929	0.802	0.897	1.000	0.849	0.929	0.778	0.897	0.770	0.952	0.651
F19	ident_pos_h	0.217	0.529	0.540	0.937	0.651	0.608	0.963	0.455	0.704	0.534	0.683	0.492	0.725	0.360
F20	counter_quote_h	0.266	0.329	0.156	0.803	0.491	0.376	0.694	0.647	0.601	0.405	0.566	0.376	0.624	0.254
F21	counter_ref_h	0.291	0.298	0.184	0.730	0.475	0.560	0.745	0.617	0.589	0.404	0.589	0.390	0.610	0.284
F22	target_obj_h	0.877	0.846	0.954	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
F23	target_indiv_h	0.277	0.554	0.846	0.846	0.600	0.646	1.000	0.646	0.692	0.662	0.662	0.615	0.692	0.569
F24	target_group_h	0.355	0.597	0.629	0.839	0.532	0.613	0.919	0.694	0.710	0.597	0.645	0.548	0.710	0.452
F25	spell_char_swap_h	0.699	0.586	0.887	0.571	0.609	0.556	0.105	0.541	0.669	0.744	0.677	0.782	0.624	0.842
F26	spell_char_del_h	0.593	0.479	0.743	0.557	0.671	0.707	0.079	0.514	0.700	0.821	0.693	0.836	0.679	0.900
F27	spell_space_del_h	0.681	0.511	0.801	0.383	0.674	0.830	0.071	0.652	0.738	0.872	0.759	0.872	0.723	0.929
F28	spell_space_add_h	0.439	0.376	0.740	0.110	0.266	0.376	0.121	0.249	0.243	0.335	0.272	0.376	0.231	0.468
F29	spell_leet_h	0.480	0.439	0.682	0.607	0.561	0.705	0.046	0.526	0.647	0.786	0.665	0.786	0.607	0.867

#### 4.4 Performance in 29 functionalities

The test suite [24] provides 29 functional tests, as shown in Table 11. We modified original code to include our performance numbers for comparison. In MLT and MLT-plus-TM, we include three runs to provide insights into more general performance. Each test evaluates a specific functionality and is associated with a gold standard label ("Hateful" or "Non-hateful"). The table also provides performance across functional tests for the models mentioned above. We present our results using the same tests in Table 11. The 'Non-hateful' functionality lines are highlighted in gray. As expected, MLT-plus-TM improves performance in all "Hateful" functionalities compared to MLT alone, while it has the opposite effect on all "Non-hateful" functionalities. Bloom is strongly biased towards classifying all cases as "Non-hateful", resulting in high accuracy on "Non-hateful" cases but misclassifying most "Hateful" cases.

#### 4.5 Performance on target groups

We also present our models' performance on target groups, alongside the results of models reported in [24], as displayed in Table 12 using the modified original code. We adhere to the Centers for Disease Control and Prevention (CDC) guidelines for Preferred Terms for Select Population Groups & Communities, available at [https://www.cdc.gov/healthcommunication/Preferred\\_Terms.html](https://www.cdc.gov/healthcommunication/Preferred_Terms.html), when referring to some target group names. Similar to section 4.4, we also include three runs for MLT and MLT-plus-TM. MLT-plus-TM demonstrates a more balanced accuracy spread across all seven target groups, thereby reducing performance biases among them and closing the gap with Google-P's performance.

Table 12: Model accuracy on test cases by target group

Target group	B-D	B-F	Google-P	BIGBIRD	BERT	BERTWEET	BLOOM	XLNET	MLT_1	MLT-plus-TM_1	MLT_2	MLT-plus-TM_2	MLT_3	MLT-plus-TM_3
Group 1	0.349	0.523	0.805	0.634	0.715	0.710	0.287	0.713	0.767	0.781	0.755	0.770	0.748	0.796
Group 2	0.691	0.694	0.808	0.644	0.677	0.808	0.323	0.610	0.748	0.789	0.760	0.760	0.732	0.767
Group 3	0.739	0.743	0.808	0.646	0.772	0.786	0.373	0.724	0.798	0.777	0.789	0.767	0.800	0.767
Group 4	0.698	0.722	0.805	0.458	0.720	0.689	0.337	0.613	0.694	0.755	0.734	0.774	0.672	0.791
Group 5	0.710	0.371	0.798	0.418	0.656	0.477	0.233	0.496	0.523	0.634	0.568	0.696	0.506	0.736
Group 6	0.722	0.736	0.796	0.354	0.755	0.784	0.323	0.762	0.803	0.784	0.791	0.786	0.796	0.767
Group 7	0.705	0.589	0.805	0.447	0.515	0.698	0.235	0.553	0.591	0.686	0.637	0.713	0.589	0.739

Group 1: Females  
Group 2: Gender Transitioned persons  
Group 3: LGBTQ persons  
Group 4: African American persons  
Group 5: People with disabilities  
Group 6: The Muslim community  
Group 7: Immigrants/Migrants

## 5 Conclusion

We have introduced a meta learning technique (MLT) for OOPS detection that combines individual language models to improve performance. Further, we have combined MLT with a threshold-moving (TM) technique to further boost the performance of the combined predictor. The experiments with open models confirm that our proposed methodology is numerically stable and is able to produce superior results on HS detection as compared to traditional methods. Thus, our technique helps to close the performance gap between non-open and open models [2] on the tasks considered in this paper. Additionally, we also establish theoretical bounds on the combiner weight coefficients to show that MLT behaves well computationally, which is supported by the experimental results. Overall, we think that our proposal is suited for a wide range of problems spanning various application domains (e.g., other NLP tasks and multimodal classification).

## References

- [1] Paula Fortuna and Sérgio Nunes. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30, 2018.
- [2] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- [3] Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. The false promise of imitating proprietary llms, 2023.
- [4] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [6] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- [7] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [9] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Nee-lakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [11] Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.
- [12] Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. Deep learning-based text classification: A comprehensive review. *ACM Comput. Surv.*, 54(3), apr 2021.
- [13] Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. Challenges and frontiers in abusive content detection. In *Proceedings of the Third Workshop on Abusive Language Online*. Association for Computational Linguistics, 2019.
- [14] Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899*, 2017.
- [15] Patricia Rossini. Beyond incivility: Understanding patterns of uncivil and intolerant discourse in online political talk. *Communication Research*, 49(3):399–425, 2020.
- [16] Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 512–515, 2017.
- [17] Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. Resources and benchmark corpora for hate speech detection: a systematic review. *Language Resources and Evaluation*, 55(2):477–523, 2021.

- [18] Wenjie Yin and Arkaitz Zubiaga. Towards generalisable hate speech detection: a review on obstacles and solutions. *PeerJ Computer Science*, 7:e598, 2021.
- [19] Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. Racial bias in hate speech and abusive language detection datasets. *arXiv preprint arXiv:1905.12516*, 2019.
- [20] Zeerak Waseem. Are You a Racist or Am I Seeing Things? Annotator Influence on Hate Speech Detection on Twitter. In *Proceedings of the first workshop on NLP and computational social science*, pages 138–142, 2016.
- [21] Jennifer Golbeck, Zahra Ashktorab, Rashad O Banjo, Alexandra Berlinger, Siddharth Bhagwan, Cody Buntain, Paul Cheakalos, Alicia A Geller, Rajesh Kumar Gnanasekaran, Raja Rajan Gunasekaran, et al. A large labeled corpus for online harassment research. In *Proceedings of the 2017 ACM on web science conference*, pages 229–233, 2017.
- [22] Antigoni Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Twelfth International AAAI Conference on Web and Social Media*, 2018.
- [23] Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini, Rebekah Tromble, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, et al. Introducing cad: the contextual abuse dataset. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2289–2303. Association for Computational Linguistics, 2021.
- [24] Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet B Pierrehumbert. Hatecheck: Functional tests for hate speech detection models. *arXiv preprint arXiv:2012.15606*, 2020.
- [25] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [26] Google Jigsaw Perspective. Perspective api. <https://perspectiveapi.com/>, 2022.
- [27] Two Hat. Shift ninja. <https://www.twohat.com/blog/introducing-sift-ninja/>, 2022.
- [28] Leo Breiman. Stacked regressions. *Machine learning*, 24(1):49–64, 1996.
- [29] Josef Kittler, Mohamad Hatef, Robert PW Duin, and Jiri Matas. On combining classifiers. *IEEE transactions on pattern analysis and machine intelligence*, 20(3):226–239, 1998.
- [30] David Opitz and Richard Maclin. Popular ensemble methods: An empirical study. *Journal of artificial intelligence research*, 11:169–198, 1999.
- [31] Ashnil Kumar, Jinman Kim, David Lyndon, Michael Fulham, and Dagan Feng. An ensemble of fine-tuned convolutional neural networks for medical image classification. *IEEE journal of biomedical and health informatics*, 21(1):31–40, 2016.
- [32] Rahul Paul, Lawrence Hall, Dmitry Goldgof, Matthew Schabath, and Robert Gillies. Predicting nodule malignancy using a cnn ensemble approach. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.
- [33] Fábio Perez, Sandra Avila, and Eduardo Valle. Solo or ensemble? choosing a cnn architecture for melanoma classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [34] Benedetta Savelli, Alessandro Bria, Mario Molinara, Claudio Marrocco, and Francesco Tortorella. A multi-context cnn ensemble for small lesion detection. *Artificial Intelligence in Medicine*, 103:101749, 2020.
- [35] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*, 2018.
- [36] Apostol Vassilev, Munawar Hasan, and Honglan Jin. Can you tell? ssnet - a biologically-inspired neural network framework for sentiment classifiers. In Giuseppe Nicosia, Varun Ojha, Emanuele La Malfa, Gabriele La Malfa, Giorgio Jansen, Panos M. Pardalos, Giovanni Giuffrida, and Renato Umeton, editors, *Machine Learning, Optimization, and Data Science*, pages 357–382, Cham, 2022. Springer International Publishing.
- [37] Sotiris Kotsiantis, Dimitris Kanellopoulos, Panayiotis Pintelas, et al. Handling imbalanced datasets: A review. *GESTS international transactions on computer science and engineering*, 30(1):25–36, 2006.
- [38] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, pages 872–881. PMLR, 2019.

- [39] Jason Brownlee. A gentle introduction to threshold-moving for imbalanced classification. *Machine Learning Mastery*, 2020.
- [40] Guillem Collell, Drazen Prelec, and Kaustubh R Patil. A simple plug-in bagging ensemble based on threshold-moving for classifying binary and multiclass imbalanced data. *Neurocomputing*, 275:330–340, 2018.
- [41] Ouadie Gharroudi, Haytham Elghazel, and Alex Aussem. Ensemble multi-label classification: a comparative study on threshold selection and voting methods. In *2015 IEEE 27th international conference on tools with artificial intelligence (ICTAI)*, pages 377–384. IEEE, 2015.
- [42] Rong-En Fan and Chih-Jen Lin. A study on threshold selection for multi-label classification. *Department of Computer Science, National Taiwan University*, pages 1–23, 2007.
- [43] Yanmin Sun, Andrew KC Wong, and Mohamed S Kamel. Classification of imbalanced data: A review. *International journal of pattern recognition and artificial intelligence*, 23(04):687–719, 2009.
- [44] TensorFlowHub. Bert-cased. [https://tfhub.dev/tensorflow/bert\\_multi\\_cased\\_L-12\\_H-768\\_A-12/2](https://tfhub.dev/tensorflow/bert_multi_cased_L-12_H-768_A-12/2), 2022.
- [45] Wolfram Neural Net Repository. Bookcorpus dataset. <https://resources.wolframcloud.com/NeuralNetRepository/resources/BERT-Trained-on-BookCorpus-and-English-Wikipedia-Data>, 2019.
- [46] Mohiuddin Md Abdul Qudar and Vijay Mago. Tweetbert: a pretrained language representation model for twitter text analysis. *arXiv preprint arXiv:2010.11091*, 2020.
- [47] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in neural information processing systems*, 33:17283–17297, 2020.
- [48] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [49] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- [50] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- [51] Konstantinos Trohidis, Grigorios Tsoumakas, George Kalliris, and Ioannis Vlahavas. Multi-label classification of music by emotion. *EURASIP Journal on Audio, Speech, and Music Processing*, 2011(1):4, September 2011.
- [52] Grigorios Tsoumakas, Ioannis Manousos Katakis, and Ioannis P. Vlahavas. Mining multi-label data. In *Data Mining and Knowledge Discovery Handbook*, 2010.
- [53] Min-Ling Zhang, Yu-Kun Li, Xu-Ying Liu, and Xin Geng. Binary relevance for multi-label learning: an overview. *Frontiers of Computer Science*, 12(2):191–202, April 2018.
- [54] Karol Draszawka and Julian Szymański. Thresholding strategies for large scale multi-label text classifier. In *2013 6th International Conference on Human System Interactions (HSI)*, pages 350–355, 2013.
- [55] Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 26(8):1819–1837, 2014.
- [56] Google Brain Team. Open source library for ml models. <https://www.tensorflow.org/>, 2022.

## 6 Appendix: a detailed proof of Theorem 1

*Proof.* Here we provide a detailed proof of Theorem 1 in the paper. We note that all equation numbers are local for this appendix, unless explicitly mentioned to be those in the paper. With the notation and assumptions from Section 3.1 let

$$\hat{\mathbf{y}}(\tau) = \frac{1}{|W|} \mathbf{y}(\tau) = \sum_{i=1}^K \frac{w_i}{|W|} \mathbf{y}_i(\tau) \quad (11)$$

be the interpolation predictor constructed as a linear combination of  $\mathbf{y}_i$  with coefficients that sum up to  $\pm 1$ . If the individual predictors are good, then the interpolation predictor  $\hat{\mathbf{y}}$  is also good, i.e.,  $\|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$  is small relative to  $\|\mathbf{u}\|_t$ .

Let  $\mathbb{L}_t(x)$  be a linear approximation of  $\sigma(x)$  for some constant  $t > 0$ , such that  $\mathbb{L}_t(x)$  minimizes  $\|\mathbb{L}_t(x) - \sigma(x)\|_t$ . Note that any straight line passing through the points  $(\ln(\frac{t}{1-t}), t)$  and having the same slope as  $\sigma'(\ln(\frac{t}{1-t}))$  satisfies  $\|\mathbb{L}_t(x) - \sigma(x)\|_t = 0$ .

First, consider the case  $b < 0$  and  $w_i \geq 0$  for  $\forall i$  in (4).

Then,

$$\begin{aligned} \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t &= \|\mathbf{u} - W\mathbf{u} + W\mathbf{u} - \sigma_b(\mathbf{y})\|_t = \\ &= \|(1 - W)\mathbf{u} + W(\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y}))\|_t. \end{aligned}$$

Applying the triangle inequality, we get

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq |1 - W| \|\mathbf{u}\|_t - W \|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t. \quad (12)$$

Next, consider the case  $W \leq 1$ . Then, from inequality (12)

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq (1 - W) \|\mathbf{u}\|_t - W \|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t.$$

Let  $\mathbb{L}_t$  be as defined above and

$$\mathbb{L}_{t,b}(x) = \mathbb{L}_t(x - b). \quad (13)$$

Then,

$$\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t = \|\mathbf{u} - \frac{1}{W}\mathbb{L}_{t,b}(\mathbf{y}) + \frac{1}{W}\mathbb{L}_{t,b}(\mathbf{y}) - \frac{1}{W}\sigma_b(\mathbf{y})\|_t$$

Thus,

$$\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t \leq \|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t + \frac{1}{W} \|\mathbb{L}_{t,b}(\mathbf{y}) - \sigma_b(\mathbf{y})\|_t.$$

Note that

$$\frac{1}{W} \|\mathbb{L}_{t,b}(\mathbf{y}) - \sigma_b(\mathbf{y})\|_t = 0.$$

From here we get

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq (1 - W) \|\mathbf{u}\|_t - W (\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t).$$

This implies that

$$W \geq \frac{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t + \|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t}.$$

Note that  $\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t = \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$ , because by construction  $\mathbf{I}_t(\mathbb{L}_{t,b}(\hat{\mathbf{y}})) = \mathbf{I}_t(\sigma_b(\hat{\mathbf{y}}))$ . Hence,

$$W \geq \frac{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t}. \quad (14)$$

Note also that  $\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t$  is small, especially with respect to the size of  $\|\mathbf{u}\|_t$ . Similarly, by the definition of  $\hat{\mathbf{y}}$  in (11),  $\|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$  is small.

Next, consider the case  $W > 1$ . Then, inequality (12) implies

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq (W - 1)\|\mathbf{u}\|_t - W\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t. \quad (15)$$

Introducing  $\mathbb{L}_{t,b}(x)$  as in the previous case, we get

$$\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t = \|\mathbf{u} - \frac{1}{W}\mathbb{L}_{t,b}(\mathbf{y}) + \frac{1}{W}\mathbb{L}_{t,b}(\mathbf{y}) - \frac{1}{W}\sigma_b(\mathbf{y})\|_t$$

Thus,

$$\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t \leq \|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t + \frac{1}{W}\|\mathbb{L}_{t,b}(\mathbf{y}) - \sigma_b(\mathbf{y})\|_t.$$

As we observed above,  $W^{-1}\|\mathbb{L}_{t,b}(\mathbf{y}) - \sigma_b(\mathbf{y})\|_t = 0$  and  $\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t = \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$ . Hence,

$$\|\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y})\|_t \leq \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t.$$

Substituting this into inequality (15) gives

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq (W - 1)\|\mathbf{u}\|_t - W\|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t.$$

From here we get

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t + \|\mathbf{u}\|_t \geq W(\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t).$$

Note that  $\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t$  and  $\|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$  are small relative to  $\|\mathbf{u}\|_t$ , which means that  $\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t > 0$  and not too far from  $\|\mathbf{u}\|_t$ . This implies that

$$W \leq \frac{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t}. \quad (16)$$

Next, consider the case  $b > 0$  and  $w_i \leq 0$  for  $\forall i$  in (4). As before,

$$\begin{aligned} \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t &= \|\mathbf{u} - W\mathbf{u} + W\mathbf{u} - \sigma_b(\mathbf{y})\|_t = \\ &= \|(1 - W)\mathbf{u} + W(\mathbf{u} - \frac{1}{W}\sigma_b(\mathbf{y}))\|_t, \end{aligned}$$

and from this we get

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq |1 + W| \|\mathbf{u}\|_t - |W| \|\mathbf{u} + \frac{1}{W}\sigma_b(\mathbf{y})\|_t \quad (17)$$

$$= |1 + W| \|\mathbf{u}\|_t - |W| \|\mathbf{u} + \frac{1}{W}\mathbb{L}_{t,b}(\mathbf{y})\|_t \quad (18)$$

$$= |1 + W| \|\mathbf{u}\|_t + W \|\mathbf{u} - \frac{1}{|W|}\mathbb{L}_{t,b}(\mathbf{y})\|_t \quad (19)$$

$$= |1 + W| \|\mathbf{u}\|_t + W \|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t. \quad (20)$$

Furthermore, consider the case  $1 + W > 0$ . Then from (20) we get

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq (1 + W)\|\mathbf{u}\|_t + W\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t.$$

Thus, using  $\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t = \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$ ,

$$W \leq -\frac{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t} \quad (21)$$

Now, let  $W < -1$ . Then,

$$\|\mathbf{u} - \sigma_b(\mathbf{y})\|_t \geq -(1 + W)\|\mathbf{u}\|_t + W\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t.$$

Form here using again  $\|\mathbf{u} - \mathbb{L}_{t,b}(\hat{\mathbf{y}})\|_t = \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t$ , we get

$$W \geq -\frac{\|\mathbf{u}\|_t + \|\mathbf{u} - \sigma_b(\mathbf{y})\|_t}{\|\mathbf{u}\|_t - \|\mathbf{u} - \sigma_b(\hat{\mathbf{y}})\|_t} \quad (22)$$

Combining (14), (16), (21), and (22) completes the proof.  $\square$