
Efficient Online ML API Selection for Multi-Label Classification Tasks

Lingjiao Chen¹ Matei Zaharia¹ James Zou^{1,2}

Abstract

Multi-label classification tasks such as OCR and multi-object recognition are a major focus of the growing machine learning as a service industry. While many multi-label APIs are available, it is challenging for users to decide which API to use for their own data and budget, due to the heterogeneity in their prices and performance. Recent work has shown how to efficiently select and combine single-label APIs to optimize performance and cost. However, its computation cost is exponential in the number of labels, and is not suitable for settings like OCR. In this work, we propose FrugalMCT, a principled framework that adaptively selects the APIs to use for different data in an online fashion while respecting the user’s budget. It allows combining ML APIs’ predictions for any single data point, and selects the best combination based on an accuracy estimator. We run systematic experiments using ML APIs from Google, Microsoft, Amazon, IBM, Tencent, and other providers for tasks including multi-label image classification, scene text recognition and named entity recognition. Across these tasks, FrugalMCT can achieve over 90% cost reduction while matching the accuracy of the best single API, or up to 8% better accuracy while matching the best API’s cost.

1. Introduction

Many machine learning users are starting to adopt machine learning as a service (MLaaS) APIs to obtain high-quality predictions. One of the most common tasks these APIs target is multi-label classification. For example, one can use Google’s computer vision API (Goo) to tag an image with a wide range of possible labels for \$0.0015, or Mi-

crosoft’s API (Mic) for \$0.0010. Another example is to extract all text strings from an image for \$0.005 via iFLYTEK’s API (Ifi) or \$0.021 via Tencent’s API (Ten). In practice, these APIs also provide different performance on different types of input data (e.g., English vs Chinese text). The heterogeneity in APIs’ performance and prices makes it hard for users to decide which API, or combination of APIs, to use for their own datasets and budgets.

Recent work (Chen et al., 2020) proposed FrugalML, an algorithmic framework that adaptively decides which APIs to call for a data point to optimize accuracy and cost. Their approach learns a fast decision rule for each possible output label that can significantly improve cost-performance over the individual APIs. However, FrugalML requires a large amount of training data and involves solving a non-convex optimization problem with complexity exponential in the number of distinct labels. This prevents it from being used for tasks with large number of labels, such as multi-label classification. Furthermore, FrugalML ignores correlation between different APIs’ predictions, potentially limiting its accuracy. For example, APIs A and B may output $\{person, car\}$ and $\{car, bike\}$ separately for an image whose true keywords are $\{person, car, bike\}$. FrugalML would select one of the two label sets, but combining them results in the true label set and thus higher accuracy. Thus, this paper aims to solve these significant limitations and address the question: *how do we design efficient ML API selection strategies for multi-label classification tasks to maximize accuracy within a budget?*

We propose FrugalMCT, a principled framework that learns the relative strengths of different combinations of multi-label classification APIs and efficiently selects the optimal combinations of APIs to call for different data and budget constraints. As shown in Fig. 1 (a), FrugalMCT directly estimates the accuracy of each API combination on a particular input based on the features and predicted labels of that input. Then it uses a fast service selector based on the estimated accuracy to balance accuracy and budget. For example, we might first call API A on an input. If A returns *person* and *teddy bear* and the accuracy predictor gives relatively high estimated accuracy (Fig. 1 (c)), then we stop and report $\{person, teddy bear\}$ as the label set. If A returns *person* and *tennis racket*, and we predict that combining it with API B’s output gives a much higher ac-

¹Department of Computer Sciences, Stanford University, Stanford, USA ²Department of Biomedical Data Science, Stanford University, Stanford, USA. Correspondence to: Lingjiao Chen <lingjiao@stanford.edu>.

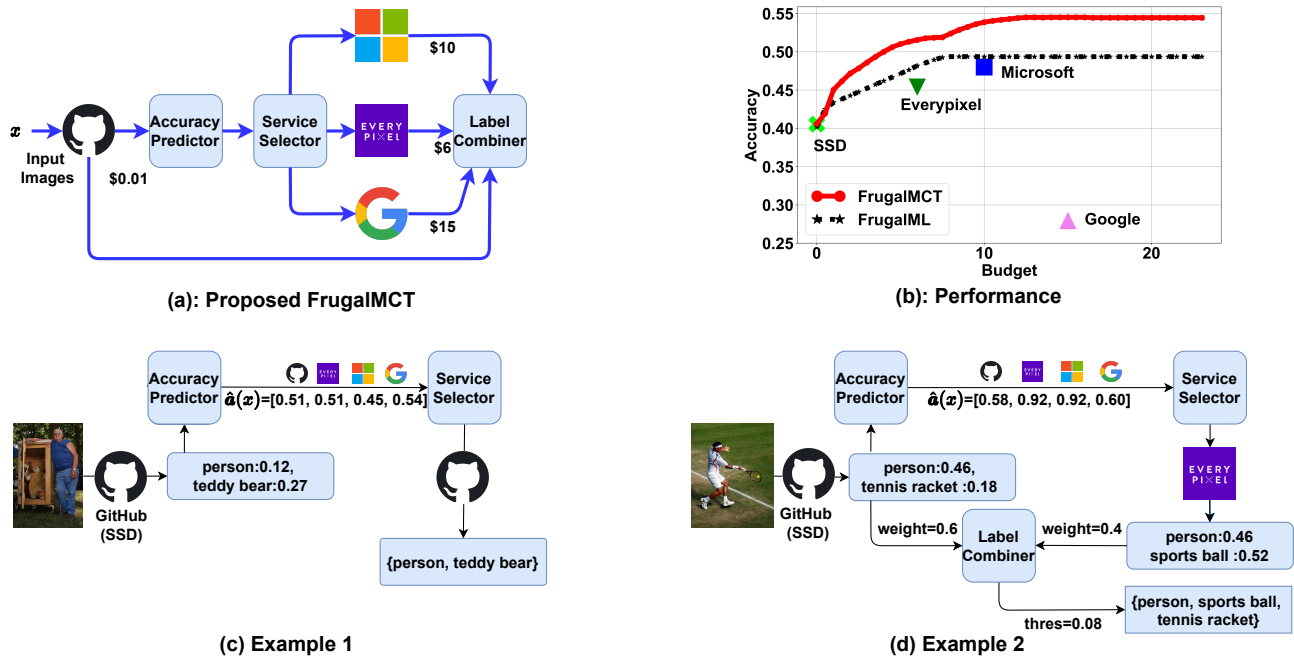


Figure 1. Demonstration of FrugalMCT. (a): FrugalMCT workflow. (b): Performance of FrugalMCT on COCO, a multi-label image dataset, using real commercial ML APIs. (c), (d): Examples of FrugalMCT’s behavior on different inputs. In (c), FrugalMCT estimates that the accuracy of a cheap open source SSD model from GitHub is high, and thus directly returns its predictions. In (d), FrugalMCT estimates that combining SSD’s results with the Everypixel API has a much higher estimated accuracy, and thus it invokes EveryPixel and combines its results with SSD’s results.

curacy, then we invoke API B and combine their prediction to obtain $\{person, sports ball, tennis racket\}$ (Fig. 1 (d)).

Contributions. FrugalMCT is an end-to-end approach that integrates the selection of APIs and the combination of their outputs for individual user queries. It leverages our key new finding that *current commercial APIs have complementary strengths and weaknesses, and that we can reliably predict which APIs are likely to work well for a new query based on easy-to-generate metadata about its input*. Based on this API accuracy predictor, FrugalMCT then leverages an efficient online algorithm to determine which combination of APIs to call for different user queries. We show that the online algorithm enjoys an accuracy provably close to the offline method as well as a small computational cost. All components in FrugalMCT are trainable, making it easy to customize for different applications. To our knowledge, FrugalMCT is the first work on how to effectively select and combine multi-label ML APIs.

Empirically, FrugalMCT produces substantially better prediction performance than individual APIs and than FrugalML adapted for multi-label tasks (Fig. 1 (b)). Extensive experiments with real commercial APIs on several tasks, including multi-label image classifications, scene text recognition, and named entity recognition, show that FrugalMCT typically provides over 60% (as high as 98%)

cost reduction when aiming to match the best commercial API’s performance. Also, when targeting the same cost as the best commercial API, FrugalMCT can improve performance up to 8%. As a dataset contribution, we have also released¹ our dataset of 295,212 samples annotated by commercial multi-label APIs as the largest dataset and resource for studying multi-label ML prediction APIs.

2. Related Work

MLaaS: With the growing importance and adoption of MLaaS APIs (Ama; Ten; Goo; IBM; Mic), existing research has largely focused on evaluating individual APIs for their performance (Yao et al., 2017), robustness (Hosseini et al., 2017), biases (Koenecke et al., 2020), performance estimation (Chen et al., 2021), pricing (Chen et al., 2019), and applications (Buolamwini & Gebru, 2018; Hosseini et al., 2019; Reis et al., 2018). Recent work on FrugalML (Chen et al., 2020) studies API calling strategies for single label classification. While their approach’s computational complexity is exponential in the number of labels, FrugalMCT’s complexity does not depend on the number of labels, making it suitable for multi-label prediction APIs. In addition, FrugalML selects only one API per user query, while FrugalMCT considers the combination of multiple

¹<https://github.com/lchen001/FrugalMCT>

APIs’ output for each input data. This improves the overall accuracy (as shown in Sec 5), but also creates unique optimization challenges that we solve.

Ensembles for multi-label classification: Ensemble learning is a natural approach to combine different predictors’ output. Several ensemble methods have been developed, such as using pruned sets (Read et al., 2008), classifier chains (Read et al., 2011), and random subsets (Tsoumakas & Vlahavas, 2007), with applications in image annotations (Xu et al., 2011), document classification (Chen et al., 2017), and speech categorization (Liu et al., 2019). Moyano et al. (2018) provide a detailed survey of this area. Almost all of these ensemble methods require joint training of the base classifiers, but MLaaS APIs are black box to the users. Also, while ensemble methods focus only on improving accuracy, FrugalMCT explicitly considers the cost of each API and enforces a budget constraint.

Model cascades: A series of works (Viola & Jones, 2001a;b; Sun et al., 2013; Cai et al., 2015; Wang et al., 2011; Xu et al., 2014; Chen et al., 2018; Kumar et al., 2018; Chen et al., 2018) explores cascades (a sequence of models) to balance the quality and runtime of inference. Model cascades use a *single* predicted quality score to avoid calling computationally expensive models, but FrugalMCT strategies utilize both *quality scores and predicted label sets* to select an expensive add-on service. While cascades do not explicitly specify inference speed, FrugalMCT allows users to explicitly incorporate different budget requirements. Designing such strategies requires solving a significantly harder optimization problem, e.g., choosing how to divide the available budget between classes (§4), but also improves performance substantially over using the quality score alone (§5).

AutoML for multi-label classification: AutoML (Thorn-ton et al., 2013) automates the customization of ML pipelines, including the selection, combination, and parametrization of the learning algorithms. There is a rich literature of AutoML techniques for standard single label tasks, and fewer methods on multi-label predictions (Wever et al., 2021) (e.g. genetic algorithms (de Sá et al., 2017) and a neural network-based search scheme (Pakrashi & Namee, 2019)). We refer interested readers to a recent survey (Wever et al., 2021) for more details. Applying AutoML to use multiple ML APIs is underexplored, and FrugalMCT can be viewed as the first AutoML approach designed for automating the selection of multiple multi-label ML APIs. While most AutoML systems exclusively focus on prediction performance, FrugalMCT optimizes accuracy and cost jointly, which is desirable for cost-sensitive API users.

Multiple choice knapsack and integer programming: Many resource allocation problems can be modeled as mul-

iple choice knapsack problem (MCKP) (Pamela H. Vance & Toth, 1993), such as keyword bidding (Zhou & Naroditskiy, 2008) and quality of service control (Lee et al., 1999). While NP-hard (Sinha & Zoltners, 1979), various approximations have been proposed for MCKP, such as branch and bound (Pamela H. Vance & Toth, 1993), convex hull relaxation (Akbar et al., 2006) and bi-objective transformation (Bednarczuk et al., 2018). Inherently an integer linear programming (ILP) problem, MCKP can also be tackled by ILP solvers, motivated by online adwords searching (Devanur & Hayes, 2009), resource allocation (Devanur & Hayes, 2019) and general linear programming (Li et al., 2020). The service selector of FrugalMCT can be viewed as a MCKP with the same item cost vector per item group, which we leverage to obtain a customized fast and online solver. Our goal is to not develop novel MCKP solver, but to efficiently adapt ILP methods as a subroutine of our end-to-end FrugalMCT to tackle a practical new application.

3. Preliminaries

Notation. We denote matrices and vectors in bold, and scalars, sets, and functions in standard script. Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, we let $\mathbf{A}_{i,j}$ denote its entry at location (i, j) . $\mathbb{1}(\cdot)$ represents the indicator function.

Multi-label classification Tasks. Throughout this paper, we focus on multi-label classification tasks: assigning a label set $Y \subseteq \mathcal{Y}$ to any data point $x \in \mathcal{X}$. In contrast to basic supervised learning, in multi-label learning each data point is associated with a set of labels instead of a single label. Many MLaaS APIs target such tasks. Consider, for example, image tagging, where \mathcal{X} is a set of images and \mathcal{Y} is the set of all tags. Example label sets could be $\{person, car\}$ or $\{bag, train, sky\}$.

MLaaS Market. Consider a MLaaS market consisting of K different ML services for some multi-label tasks. For a data point x , the k th service returns to the user a set of labels with their quality scores, denoted by $Y_k(x) \subseteq \mathcal{Y} \times [0, 1]$. For example, one API for multi-label image classification might produce $Y_k(x) = \{(person, 0.8), (car, 0.7)\}$, indicating the label *person* with confidence 0.8 and *car* with confidence 0.7. Let the vector $\mathbf{c} \in \mathbb{R}^K$ denote the unit cost of all services. E.g., $c_k = 0.01$ means that users need to pay \$0.01 every time they call the k th service.

4. FrugalMCT Framework

In this section, we present FrugalMCT, a framework to adaptively select ML APIs for multi-label classification tasks within a budget. All proofs are left to the appendix. We generalize the scheme in Figure 1 (a) to K ML services. As shown in Figure 2, FrugalMCT contains three

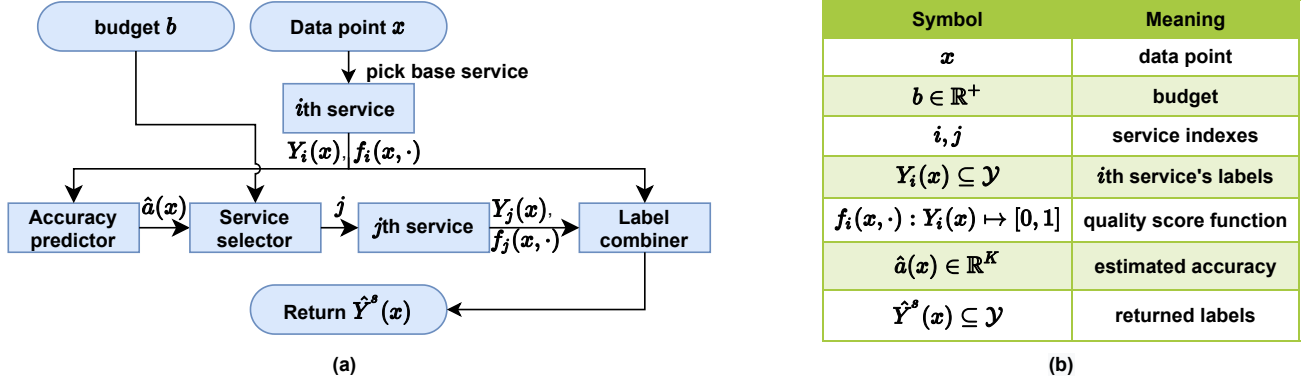


Figure 2. Overview of FrugalMCT. (a) shows how it works: Given a data point, FrugalMCT first invokes a base service. An accuracy predictor then estimates the performance of different APIs. Next, an add-on service is selected based on the predicted accuracy and budget. Finally, the add-on and base services’ predictions are combined to return FrugalMCT’s prediction. (b) lists notation.

main components: an accuracy estimator, a service selector, and a label combiner.

Given a data point x , it first calls some base service, denoted by $base$, which is one of the K APIs, and obtains $Y_{base}(x)$. Often, $base$ is a cheap or free service, such as an inexpensive open source model; we discuss how to choose $base$ out of multiple options in Section 4.4. Next, an accuracy predictor produces a vector $\hat{a}(x) \in [0, 1]^K$, whose k th value estimates the accuracy of the label set produced by the label combiner using base’s and k th API’s outputs. The service selector $s(\cdot) : \mathcal{X} \mapsto [K]$ then decides if and which *add-on service* needs to be invoked. Finally, a label combiner generates a label set by combining the predictions from the base and add-on APIs. Take Figure 1 (d) as an example. The image is first passed to the GitHub model, which produces $\{(person, 0.46), (tennis racket, 0.18)\}$, by which the accuracy predictor predicts the accuracy of the label set generated by combining each API’s output with GitHub model’s. The service selector then decides to further invoke Everypixel, which gives $\{(person, 0.46), (sports ball, 0.52)\}$. Finally, the label combiner uses both APIs’ output for the final prediction.

FrugalMCT allows users to customize the accuracy predictor and the label combiner, depending on the applications. For example, for the image tagging problem, one might use image features (e.g., brightness and contrast) to build the accuracy predictor, while word embeddings can be more useful for named entity recognition. In the following sections, we explain the key of accuracy predictor, API selector and the label combiner in more detail.

4.1. Accuracy prediction

The accuracy predictor $\hat{a}(\cdot)$ can be obtained by two steps. The first step is to generate a feature vector for every data point in the training dataset $\mathbb{X}^{Tr} \triangleq \{x_1^{Tr}, x_2^{Tr}, \dots, x_N^{Tr}\}$.

Generally the feature vector can be any embedding of the data point x and base service prediction $Y_{base}(x)$. In this paper we adopt a simple approach: if the label set \mathcal{Y} is bounded, a $|\mathcal{Y}|$ dimensional vector is generated using one hot encoding on $Y_{base}(x)$. For example, given $\mathcal{Y} = \{person, car, bike\}$ and $Y_{base}(x) = \{(person, 0.8), (car, 0.7)\}$, the generated feature vector is $[0.8, 0.7, 0]$. For unbounded \mathcal{Y} , word embedding is used to generate a vector for every predicted label, and the sum of them (weighted by their quality values) becomes the corresponding feature vector.

The next step is to train the accuracy predictor. For each $x_n^{Tr} \in \mathbb{X}^{Tr}$, as its true label sets and prediction from each API are available, we can construct its true accuracy vector $\mathbf{a}(x_n^{Tr}) \in [0, 1]^K$, whose k th element is the accuracy of the label produced by the label combiner using base and k th service predictions. Then we can train some regressor (e.g., random forest) to map the feature vector to the accuracy vector. We use standard multi-label accuracy² (Zhang & Zhou, 2014) as a concrete metric. FrugalMCT can as easily use another metric such as F1-score, precision or subset accuracy.

4.2. The API selection problem

A core subroutine of FrugalMCT is the API selector s : given a budget b and the estimated accuracy $\hat{a}(x)$, which service should be invoked? Let $\mathbb{X} \triangleq \{x_1, x_2, \dots, x_N\}$ be the entire unlabeled dataset to be classified, and $S \triangleq \{1, 2, \dots, K\}^{\mathbb{X}}$ be the set of all functions mapping each data point in \mathbb{X} to an API. Let $base$ be the index of the base service. For any $s \in S$, $s(x) = base$ implies no add-on API is needed, and $s(x) = k \neq base$ implies k th API is invoked. Our goal is to find some $s \in S$ to maximize the estimated accuracy while satisfying the budget constraint,

² $\frac{\|Y \cap Y'\|}{\|Y \cup Y'\|}$ where Y/Y' is the true/predicted label set.

formally stated as below.

Definition 4.1. Let $\mathbf{Z}_{n,k}^*$ be the optimal solution to the budget aware API selection problem

$$\begin{aligned} & \max_{\mathbf{Z} \in \mathbb{R}^{N \times K}}: \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k} \hat{\mathbf{a}}_k(x_n) \\ & \text{s.t.} \quad \frac{1}{N} \sum_{n=1}^N \sum_{k=1, k \neq \text{base}}^K \mathbf{Z}_{n,k} \mathbf{c}_k + \mathbf{c}_{\text{base}} \leq b; \quad (4.1) \\ & \quad \sum_{k=1}^K \mathbf{Z}_{n,k} = 1, \forall n; \mathbf{Z}_{n,k} \in \{0, 1\}, \forall n, k \end{aligned}$$

Then the optimal FrugalMCT strategy is given by $s^*(x_n) \triangleq \arg \max_k \mathbf{Z}_{n,k}^*$.

Here, the objective quantifies the average accuracy, the first constraint models the budget requirement, and the last two constraints enforces only one add-on API is picked for each data point. Base service is needed for every data point and thus its cost \mathbf{c}_{base} appears for every n in the budget constraint. Note that Problem 4.1 is a MCKP (and thus integer linear program) and NP-hard in general.

4.3. An online algorithm for FrugalMCT

In many time-sensitive applications, the input data x_n (as well as the accuracy vector $\hat{\mathbf{a}}(x_n)$) comes sequentially, and the API needs to be selected before observing the future data. The selection process also needs to be fast.

To tackle this challenge, we present an efficient online algorithm, which requires $O(K)$ computations per round and gives a provably near-optimal solution. The key idea is to explicitly balance between accuracy and cost at every iteration. Specifically, for a given data point x_n and $p \in \mathbb{R}$, let us define a strategy $s^p(x_n) \triangleq \arg \max_k \hat{\mathbf{a}}_k(x_n) - p \mathbf{c}_k \mathbb{1}_{k \neq \text{base}}$ and break ties by picking k with smallest cost. Here, p is a parameter to balance between accuracy $\hat{\mathbf{a}}(x_n)$ and cost \mathbf{c} . When $p = 0$, $s^p(x_n)$ selects the API with highest estimated accuracy. When p is large enough $s^p(x_n)$ enforces to pick the base API. In fact, larger value of p implies more weights on cost and smaller p favors more the accuracy. Let $r(s) \triangleq \frac{1}{N} \sum_{n=1}^N \hat{\mathbf{a}}_{s(x_n)}(x_n)$ denote the average accuracy achieved by a strategy s . We can show, interestingly, an appropriate choice of p leads to small average accuracy loss.

Theorem 4.2. Assume the probability density of $\hat{\mathbf{a}}(x)$ is a continuous function on $[0, 1]^K$. Then with probability 1, there exists p^* such that s^{p^*} satisfies budget constraint, and $r(s^{p^*}) \geq r(s^*) - \frac{1}{N}$.

In words, $s^{p^*}(x_n)$ gives a solution to the API selection problem with accuracy loss at most $\frac{1}{N}$. In practice, $\hat{\mathbf{a}}(x)$

is continuous for standard ML models of accuracy predictors (e.g., logistic regressors) and thus the assumption holds. In addition, it is computationally efficient: at iteration n , it only requires computing $\hat{\mathbf{a}}_k(x_n) - p \mathbf{c}_k \mathbb{1}_{k \neq \text{base}}$ for $k = 1, 2, \dots, K$, which takes only $O(K)$ computations.

The remaining question is how to obtain p^* . As we cannot see the future data to compute p^* , a natural idea is to estimate it using the training dataset. More precisely, given the training dataset $\{x_1^{Tr}, x_2^{Tr}, \dots, x_{N^{Tr}}^{Tr}\}$, let $\hat{p}, \hat{\mathbf{q}}$ be the optimal solution to the following problem

$$\begin{aligned} & \min_{p, \mathbf{q}} (1 - \delta)(b - \mathbf{c}_{\text{base}})p + \sum_{n=1}^{N^{Tr}} \mathbf{q}_n, \\ & \text{s.t.} \quad \frac{\mathbf{c}_k \cdot \mathbb{1}_{k \neq \text{base}} \cdot p}{N^{Tr}} + \mathbf{q}_n \geq \frac{\hat{\mathbf{a}}_k(x_n^{Tr})}{N^{Tr}}, \forall n, k \\ & \quad p \geq 0, \mathbf{q} \in \mathbb{R}^{N^{Tr}}, \mathbf{q} \geq 0 \end{aligned} \quad (4.2)$$

where $\delta \in (0, 1)$ is a small buffer to ensure that we don't exceed the budget (in practice we set $\delta \leq 0.01$). Technically, Problem 4.2 is the dual problem to the linear programming by relaxing the integer constraint in Problem 4.1 on the training dataset with budget $(1 - \delta)b$, and \hat{p} corresponds to the near-optimal strategy for the training dataset. If the training and testing datasets are from the same distribution, then a small δ can ensure with high probability, \hat{p} is slightly less than p^* and thus $s^{\hat{p}}$ satisfies the budget constraint. Given \hat{p} , one can use $s^{\hat{p}}$ to select the APIs in an online fashion. The details are given in Algorithm 1.

Algorithm 1 FrugalMCT Online API Selection Algorithm.

Input : $\mathbf{c}, b, \{x_1^{Tr}, x_2^{Tr}, \dots, x_{N^{Tr}}^{Tr}\}, \{x_1, x_2, \dots, x_N\}$

Output : FrugalMCT online API selector $s^o(\cdot)$

- 1: Compute \hat{p} by solving Problem 4.2 and set $b^r = N(b - \mathbf{c}_{\text{base}})$.
 - 2: At iteration $n = 1, 2, \dots, N$:
 - 3: $s^o(x_n) = \begin{cases} s^{\hat{p}}(x_n) & \text{if } b^r - \mathbf{c}_{s^{\hat{p}}(x_n)} \geq 0 \\ \text{base} & \text{o/w} \end{cases}$
 - 4: $b^r = b^r - \mathbf{c}_{s^{\hat{p}}(x_n)} \mathbb{1}_{s^{\hat{p}}(x_n) \neq \text{base}}$
-

Here, b^r is used to ensure the generated solution is always feasible. The following theorem gives the performance guarantee of the online solution.

Theorem 4.3. If $\delta = \Theta\left(\sqrt{\frac{\log N/\epsilon}{N}} + \sqrt{\frac{\log N^{Tr}/\epsilon}{N^{Tr}}}\right)$ and the probability density of $\hat{\mathbf{a}}(x)$ is a continuous function on $[0, 1]^K$, then s^o satisfies the budget constraint, and with probability at least $1 - \epsilon$, $r(s^o) \geq r(s^*) - O\left(\sqrt{\frac{\log N/\epsilon}{N}} + \sqrt{\frac{\log N^{Tr}/\epsilon}{N^{Tr}}}\right)$.

Roughly speaking, s^o leads to an accuracy loss at most $O\left(\sqrt{\frac{\log N}{N}} + \sqrt{\frac{\log N^{Tr}}{N^{Tr}}}\right)$ compared to the optimal offline

strategy. For large training and testing datasets, such an accuracy loss is often negligible, which is also verified by our experiments on real world datasets.

4.4. Base service selection and label combination

Now we describe how the base service is selected and how the label combiner works. The base service can be picked by an offline searching process. More precisely, for each possible base service, we train a FrugalMCT strategy and evaluate its performance on a validation dataset, and pick the base service corresponding to the highest performance.

The label combiner contains two phases. First, a new label set associated with its quality function is produced. The label set is simply the union of that from the base service and add-on service. The quality score is a weighted sum of the score from both APIs, controlled by a hyperparameter w . For example, suppose the base predicts $\{(person, 0.8), (car, 0.7)\}$ and the add-on predicts $\{(car, 0.5), (bike, 0.4)\}$. Given $w = 0.3$, new confidence for $person$ is $0.3 \times 0.8 = 0.24$, for car is $0.3 \times 0.7 + 0.7 \times 0.5 = 0.46$, and for $bike$ is $0.7 \times 0.4 = 0.28$. Thus the combined set is $\{(person, 0.24), (car, 0.46), (bike, 0.28)\}$. Next, a threshold θ is applied to remove labels with low confidence. For example, given $\theta = 0.25$, the label $person$ would be removed, and the final predicted label set becomes $\{car, bike\}$. The parameters w and θ are global hyperparameters for each dataset, and can be obtained by an efficient searching algorithm to maximize the overall performance. The details are left to Appendix A.

5. Experiments

We compare the accuracy and incurred costs of FrugalMCT to that of real world ML services for various tasks. Our goal is to (i) understand when and why FrugalMCT can reduce cost without hurting accuracy, (ii) investigate the trade-offs between accuracy and cost achieved by FrugalMCT, and (iii) assess the effect of training data size and accuracy predictors on FrugalMCT’s performance.

Tasks, ML Services, and Datasets. We focus on three common ML tasks in different domains: multi-label image classification (*MIC*), scene text recognition (*STR*), and named entity recognition (*NER*). *MIC* aims at obtaining all keywords associated with an image, *STR* seeks to recognize all texts in an image, and *NER* desires to extract all entities in a text paragraph. The ML services used for each task and their prices are summarized in Table 1. For each task we use three datasets, summarized in Table 2. More details can be found in Appendix C.

Accuracy Predictors. Except when explicitly noted, we use a random forest regressor as the accuracy predictor for

Table 1. ML services used for each task. Price unit: USD/10,000 queries. A publicly available (and thus free) GitHub model is also used per task: a single shot detector (SSD) (*SSD*) pretrained on Open Images V4 (Kuznetsova et al., 2020) for *MIC*, a convolutional recurrent neural network (PP-OCR) (*Pad*) pretrained on an industrial dataset (Du et al., 2020) for *STR*, and a convolutional neural network (spaCy (*Spa*)) pretrained on OntoNotes (Weischedel et al., 2017) for *NER*.

| Task | ML Service | Price | ML Service | Price |
|------|--------------------------|-------|----------------------------|-------|
| MIC | SSD (<i>SSD</i>) | <0.01 | Everyapixel (<i>Eve</i>) | 6 |
| | Microsoft (<i>Mic</i>) | 10 | Google (<i>Goo</i>) | 15 |
| STR | PP-OCR (<i>Pad</i>) | <0.01 | Google (<i>Goo</i>) | 15 |
| | iFLYTEK (<i>Ifi</i>) | 50 | Tencent (<i>Ten</i>) | 210 |
| NER | spaCy (<i>Spa</i>) | <0.01 | Amazon (<i>Ama</i>) | 3 |
| | Google (<i>GoN</i>) | 10 | IBM (<i>IBM</i>) | 30 |

Table 2. Dataset Statistics.

| Task | Dataset | Size | # Labels | Dist Labels |
|------|---------|--------|----------|-------------|
| MIC | PASCAL | 11540 | 16682 | 20 |
| | MIR | 25000 | 92909 | 24 |
| | COCO | 123287 | 357662 | 80 |
| STR | MTWI | 9742 | 867727 | 4404 |
| | ReCTS | 20000 | 555286 | 4134 |
| | LSVT | 30000 | 1878682 | 4852 |
| NER | CONLL | 10898 | 43968 | 9910 |
| | ZHNER | 16915 | 147164 | 4375 |
| | GMB | 47830 | 116225 | 14376 |

all the datasets. For *MIC* and *STR* datasets, we map each possible label to an index, and create a feature vector whose k th element is base service’s quality score for the label corresponding to k . If a label is not predicted, the corresponding value is 0. For *NER* datasets, we map each predicted label to a 96-dimensional vector using a word embedding from spaCy (*Spa*), and then use the sum weighted by their corresponding quality scores as the feature vector. The accuracy predictor is then trained on half of the datasets using the feature vectors generated as above. Interestingly, we found we are able to accurately predict which commercial API is best for each instance using relatively simple features. This makes the approach more broadly applicable. We will study the effects of accuracy predictors later in this section.

Multi-label Image Classification: A Case Study. Let us start with multi-label image classification on the COCO dataset (Lin et al., 2014). We set budget $b = 6$, the price of Everyapixel, the cheapest commercial API (except the open source model from GitHub). For comparison, we also use

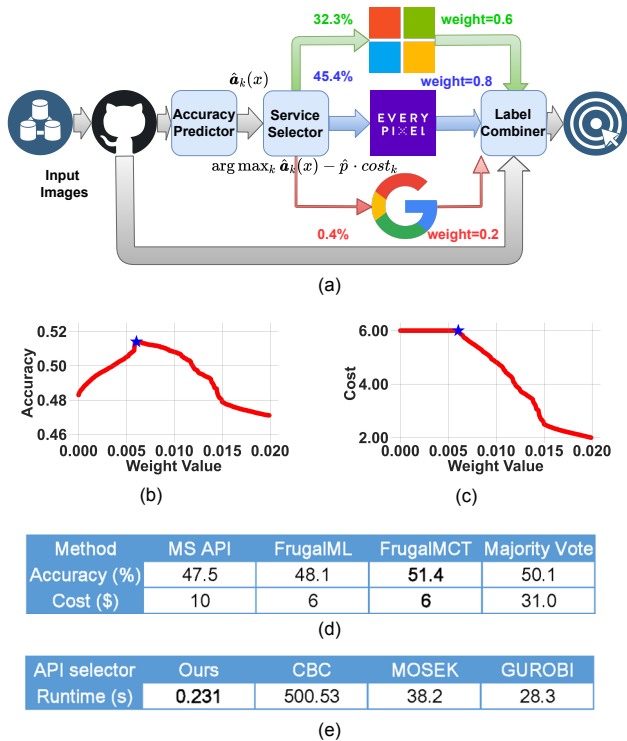


Figure 3. A FrugalMCT strategy learned on the dataset COCO. (a) shows that FrugalMCT reduces cost by mostly calling the Everyapixel API (45.4%) or the GitHub API (22.1%) only. (b) and (c) show how the accuracy and cost vary with weight p . The blue point corresponds to 0.006, the learned \hat{p} . (d) shows the accuracy and cost of FrugalMCT, FrugalML, Microsoft API, and majority vote. (e) gives the runtime performance of our (online) API selector and three commercial ILP solvers.

the average quality score over all predicted labels as the confidence score and adapt FrugalML (Chen et al., 2020) with the same budget (= 6) as another baseline .

Figure 3 demonstrates the learned FrugalMCT strategy. As shown in Figure 3 (a), the learned FrugalMCT reduces the cost by mostly using the Everyapixel API (45%, 6\$) and occasionally calling Microsoft API (32%, 10\$), and rarely invoking the Google API (0.4%, 15\$). Note that its performance depends on the threshold value \hat{p} . As shown in Figure 3 (b) and (c), for small thresholds, FrugalMCT tends to call the more accurate and expensive APIs. However, it runs out of budget quickly, and for many data points only base service can be used, leading to low accuracy. For large thresholds, FrugalMCT tends to call cheaper but less accurate APIs, failing to fully use the budget and thus causing low accuracy too. The \hat{p} value learned by FrugalMCT (blue point in Figure 3 (b) and (c)) produces the optimal accuracy given the budget. Figure 3 (d) shows that FrugalMCT’s accuracy (0.514) is higher than that of the best ML service (MS, 0.475) and majority vote (Maj 0.501), while its cost is

much lower. This is primarily because FrugalMCT learns when the cheaper APIs perform better and call them aptly. FrugalMCT also outperforms FrugalML by exploiting the label combination. This is due to (i) that FrugalML cannot utilize the label information due to explosion of complexity, and (ii) that the label combiner in FrugalMCT gives higher accuracy than both the base and add-on APIs.

To understand the efficiency of FrugalMCT’s API selector, we compare it with three ILP solvers, namely, CBC, MOSEK, and GUROBI. CBC (Forrest & Lougee-Heimer, 2005) is an integer linear programming package developed based on cutting and branch. MOSEK (Andersen & Andersen, 2000) was originally developed for sparse programming and then extended for general mixed integer programming. On the other hand, the focus of GUROBI (Bixby, 2007) is parallelism optimization in integer programming. As shown in Figure 3 (e), the API selector of FrugalMCT (Alg. 1) is several orders of magnitude faster than those commercial ILP solvers. This is because it leverages the specific structure of Problem 4.1.

Table 3. End-to-end runtime comparison on COCO.

| Runtime | FrugalMCT | FrugalML | Majority Vote |
|-----------|-----------|----------|---------------|
| Training | 60s | 6627s | N/A |
| Inference | 1.25s | 1.24s | 1.92s |

The end-to-end runtime comparison of FrugalMCT with FrugalML and an ensemble approach (majority vote) is given in Table 3 . Majority vote does not need training, but its inference time is high due to calling all ML APIs. FrugalMCT enjoys a similar inference time with FrugalML but a 100x smaller training time.

Table 4. Cost savings achieved by FrugalMCT that reaches same accuracy as the best commercial API. On average the cost saving across the evaluated datasets is 73%.

| Task | Dataset | Acc (%) | Best API \$ | Our \$ | Save |
|------|---------|---------|-------------|--------|------|
| MIC | PASCAL | 74.8 | 10 | 1.4 | 86% |
| | MIR | 41.2 | 10 | 4.2 | 58% |
| | COCO | 47.5 | 10 | 3 | 70% |
| STR | MTWI | 67.9 | 210 | 30 | 86% |
| | ReCTS | 61.3 | 210 | 78 | 63% |
| | LSVT | 53.8 | 210 | 67 | 68% |
| NER | CONLL | 52.6 | 3 | 1.5 | 50% |
| | ZHNER | 61.3 | 30 | 0.7 | 98% |
| | GMB | 50.1 | 30 | 4.1 | 80% |

Analysis of Cost Savings. Next, we evaluate how much cost can be saved by FrugalMCT to reach the highest accuracy produced by a single API on different tasks. As shown

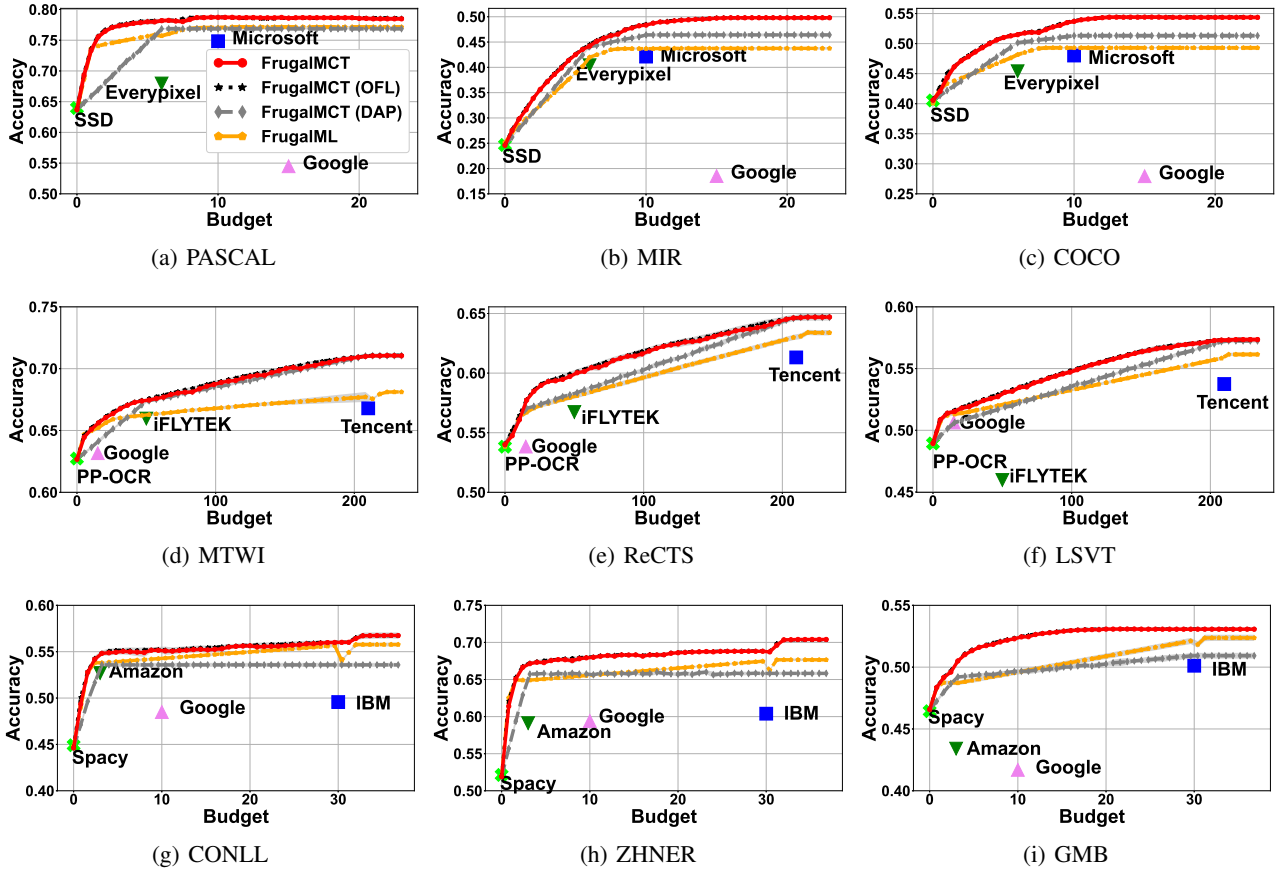


Figure 4. Accuracy cost trade-offs. The offline FrugalMCT (black) observes the full data and then make decisions. The online FrugalMCT (red) matches the offline performance in all the experiments. DAP (grey) is an oblation of FrugalMCT where a dummy accuracy predictor is used. FrugalML (orange) is the previous state-of-the-art method. The task of row 1, 2, 3 is *MIC*, *STR*, and *NER*.

in Table 4, FrugalMCT can typically save more than 60% of the cost. Interestingly, the cost saving can be up to 98% on the dataset ZHNER. This is probably because (i) the accuracy estimator enables the API selector to identify when the base service’s prediction is reliable and to avoid unnecessarily calling add-on services, and (ii) when add-on API is invoked, the apt combination of the base and add-on services leads to a high accuracy improvement.

Accuracy and Cost Trade-offs. Now we dive deeply into the accuracy and cost trade-offs achieved by FrugalMCT, shown in Figure 4. We compare with two ablations: “Offline”, where the full data is observed before making decision, “DAP”, where a dummy accuracy predictor is used, which, for each API, always returns its mean accuracy on the training dataset. We also compared with an adapted version of the previous state-of-the-art for single label task, FrugalML. To adapt it to multi-label tasks, we use the average quality score over all predicted labels as a single score, and cluster all labels into a “superclass”.

Compared to any single API, FrugalMCT allows users to pick any point in its trade-off curve and offers substantial more flexibility. In addition, FrugalMCT often achieves higher accuracy than any ML services it calls. For example, on COCO and ZHNER, more than 5% accuracy improvement can be reached with the same cost of the best API. Note that FrugalMCT also outperforms FrugalML with the same budget. This is primarily because FrugalMCT (i) utilizes a more principled way to use the features (learning an accuracy estimator) than FrugalML (directly using the label info), and (ii) adopts a label combiner designed for multi-label tasks. Ensemble methods such as majority votes (in the appendix C) produce accuracy similar to FrugalMCT, but their cost is much higher. Note that there is little performance difference between the online FrugalMCT strategy and the offline approach, due to the carefully designed on-line algorithm. This directly supports our theory.

Effects of Accuracy Predictors. The accuracy predictors play an important role in FrugalMCT’s performance. As Table 5 shows, FrugalMCT provides nontrivial accu-

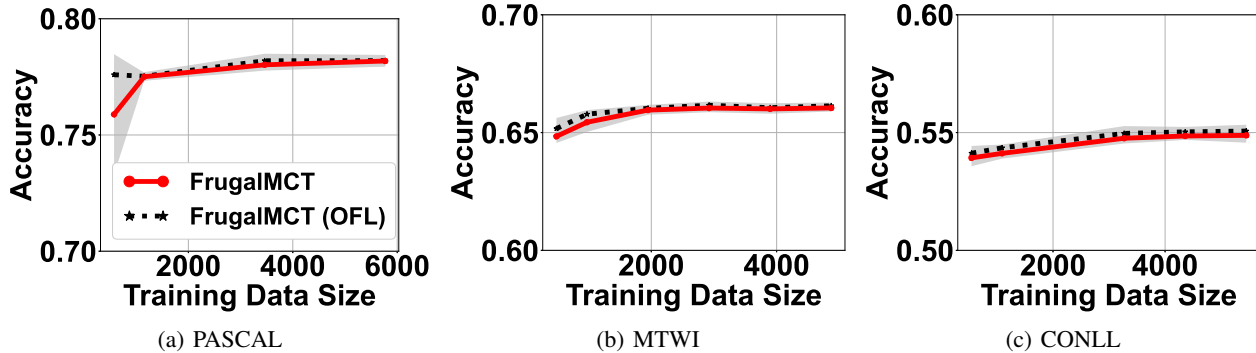


Figure 5. Testing accuracy v.s. training data size. The fixed budget is 6, 15, 3, respectively.

Table 5. Performance of FrugalMCT’s accuracy predictor. Root mean square error (RMSE) quantifies the standard deviation of the differences between the predicted and the true accuracy.

| Data | RMSE | Data | RMSE | Data | RMSE |
|--------|------|-------|------|------|------|
| PASCAL | 0.28 | MIR | 0.22 | COCO | 0.24 |
| MTWI | 0.17 | ReCTS | 0.22 | LSVT | 0.19 |
| CONLL | 0.29 | ZHNER | 0.31 | GMB | 0.28 |

accuracy estimates which enables its success. It’s interesting to note that the accuracy predictor doesn’t need to be perfect for FrugalMCT to do well; for example, the root mean square error (RMSE) of the accuracy predictor is 0.28 on PASCAL (and 0.29 on CONLL), but FrugalMCT still produces consistently better accuracy than FrugalML. We also evaluated FrugalMCT’s performance when the accuracy predictors are obtained via two AutoML toolkits, autsklearn (Feurer et al., 2015) and Auto-PyTorch (Mendoza et al., 2019) instead of random forest, and observe a similar performance.

Effects of Training Sample Size. Finally we study how the training dataset size affects FrugalMCT’s performance. As shown in Figure 5, across different tasks, a few thousand training samples are typically sufficient to learn the optimal FrugalMCT strategy. This is usually more efficient than training a customized ML model from scratch. It also only takes a few minutes to train those FrugalMCT strategies, which is much faster than training a model from scratch. This is useful in latency-critical applications.

Training cost of FrugalMCT. Both dollar cost and computation time of training are often much smaller than ML APIs inference cost. This is because (i) training is a one-time cost and (ii) FrugalMCT requires a small number of label annotations (a few thousands see Figure 5). Consider the image tagging task as an example: the dollar cost of calling all APIs is $\$0.0006 + \$0.001 + \$0.0015 = \0.0031

per image. Labeling for (say) five thousands images takes $\$0.0031 \times 5000 = \15.5 . Training a FrugalMCT strategy on half of the COCO dataset takes 59.5s on the experiment machine. This is much cheaper than calling the selected APIs after at large scale (e.g., millions of images).

6. Conclusion

In this paper, we presented FrugalMCT, an algorithmic framework to adaptively select and combine ML APIs for multi-label classification tasks within a budget constraint. FrugalMCT integrates forecasts of API’s accuracy with online constrained optimization to create an end-to-end algorithm with strong empirical performance and theoretical guarantees. How to efficiently use multi-label APIs is an important problem in practice for the large number of ML users who have chosen to rely on commercial prediction APIs, and has not been studied heavily in the ML literature. This work helps MLaaS users improve the overall accuracy and cost of their applications. Extensive empirical evaluation using real commercial APIs shows that FrugalMCT significantly improves both cost and accuracy. To encourage more research on MLaaS, we also release the dataset used to develop FrugalMCT, consisting of 295,212 samples annotated by commercial multi-label prediction APIs. The dataset and our code can be accessed from <https://github.com/lchen001/Frugalmct>.

Acknowledgement

This work was supported in part by a Google PhD Fellowship, a Sloan Fellowship, NSF CCF 1763191, NSF CAREER AWARD 1651570 and 1942926, NIH P30AG059307, NIH U01MH098953, grants from the Chan-Zuckerberg Initiative, Sutherland, and affiliate members and other supporters of the Stanford DAWN project, including Ant Financial, Meta, Google, Infosys, NEC, and VMware as well as Cisco and SAP. We also thank anonymous reviewers for helpful discussion and feedback.

References

- Amazon Comprehend API. <https://aws.amazon.com/comprehend>. [Accessed Oct-2020].
- Everypixel Image Tagging API. <https://labs.everypixel.com/api>. [Accessed Oct-2020].
- Google NLP API. <https://cloud.google.com/natural-language>. [Accessed Oct-2020].
- Google Vision API. <https://cloud.google.com/vision>. [Accessed Oct-2020].
- IBM NLP API. <https://www.ibm.com/cloud/watson-natural-language-understanding>. [Accessed Oct-2020].
- iFLYTEK Text Recognition API. <https://global.xfyun.cn/products/wordRecg>. [Accessed Oct-2020].
- Microsoft computer vision API. <https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision>. [Accessed Oct-2020].
- PaddleOCR, a text recognition tool from GitHub. <https://github.com/PaddlePaddle/PaddleOCR>. [Accessed Oct-2020].
- SSD, a multi-label image classification tool from GitHub/TensorflowHub. https://tfhub.dev/google/openimages_v4/ssd/mobilenet_v2/1. [Accessed Oct-2020].
- spaCy, a named entity recognition tool from GitHub. <https://github.com/explosion/spaCy>. [Accessed Oct-2020].
- Tencent Text Recognition API. <https://intl.cloud.tencent.com/product/ocr>. [Accessed Oct-2020].
- ZHNER dataset. <https://github.com/zjy-ucas/ChineseNER/tree/master/data>. [Accessed Oct-2020].
- Akbar, M. M., Rahman, M. S., Kaykobad, M., Manning, E. G., and Shoja, G. C. Solving the multidimensional multiple-choice knapsack problem by constructing convex hulls. *Comput. Oper. Res.*, 33:1259–1273, 2006.
- Andersen, E. D. and Andersen, K. D. The mosek interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High performance optimization*, pp. 197–232. Springer, 2000.
- Bednarczuk, E. M., Miroforidis, J., and Pyzel, P. A multi-criteria approach to approximate solution of multiple-choice knapsack problem. *Comput. Optim. Appl.*, 70(3): 889–910, 2018.
- Bixby, B. The gurobi optimizer. *Transp. Re-search Part B*, 41(2):159–178, 2007.
- Bos, J. The groningen meaning bank. In *JSSP*, pp. 2, 2013.
- Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FAT*, volume 81, pp. 77–91, 2018.
- Cai, Z., Saberian, M. J., and Vasconcelos, N. Learning complexity-aware cascades for deep pedestrian detection. In *ICCV*, pp. 3361–3369, 2015.
- Chen, G., Ye, D., Xing, Z., Chen, J., and Cambria, E. Ensemble application of convolutional and recurrent neural networks for multi-label text categorization. In *IJCNN*, 2017.
- Chen, L., Koutris, P., and Kumar, A. Towards model-based pricing for machine learning in a data marketplace. In *Proceedings of the 2019 International Conference on Management of Data*, pp. 1535–1552, 2019.
- Chen, L., Zaharia, M., and Zou, J. FrugalML: How to use ML prediction apis more accurately and cheaply. In *NeurIPS*, 2020.
- Chen, L., Zaharia, M., and Zou, J. How did the model change? efficiently assessing machine learning api shifts. In *International Conference on Learning Representations*, 2021.
- Chen, X., Liew, J., Xiong, W., Chui, C., and Ong, S. H. Focus, segment and erase: An efficient network for multi-label brain tumor segmentation. In *ECCV*, pp. 674–689, 2018.
- de Sá, A. G. C., Pappa, G. L., and Freitas, A. A. Towards a method for automatically selecting and configuring multi-label classification algorithms. In *GECCO*, pp. 1125–1132, 2017.
- Devanur, N. R. and Hayes, T. P. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *EC*, pp. 7178, 2009.
- Devanur, N. R. and Hayes, T. P. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. *J. ACM*, 66(1):1–41, 2019.
- Du, Y., Li, C., Guo, R., Yin, X., Liu, W., Zhou, J., Bai, Y., Yu, Z., Yang, Y., Dang, Q., and Wang, H. PP-OCR: A practical ultra lightweight OCR system. *CoRR*, 2020.

- Everingham, M., Eslami, S. M. A., Gool, L. V., Williams, C. K. I., Winn, J. M., and Zisserman, A. The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vis.*, 111(1):98–136, 2015.
- Feurer, M., Klein, A., Eggenberger, K., Springenberg, J., Blum, M., and Hutter, F. Efficient and robust automated machine learning. In *NIPS*. 2015.
- Forrest, J. and Lougee-Heimer, R. Cbc user guide. In *Emerging theory, methods, and applications*, pp. 257–277. INFORMS, 2005.
- He, M., Liu, Y., Yang, Z., Zhang, S., Luo, C., Gao, F., Zheng, Q., Wang, Y., Zhang, X., and Jin, L. Contest on robust reading for multi-type web images. In *ICPR*, pp. 7–12, 2018.
- Hosseini, H., Xiao, B., and Poovendran, R. Google’s cloud vision API is not robust to noise. In *ICMLA*, 2017.
- Hosseini, H., Xiao, B., and Poovendran, R. Studying the live cross-platform circulation of images with computer vision API: An experiment based on a sports media event. *IJC*, 13:1825–1845, 2019.
- Huiskes, M. J. and Lew, M. S. The MIR flickr retrieval evaluation. In *MIR*, pp. 39–43, 2008.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Touns, C., Rickford, J. R., Jurafsky, D., and Goel, S. Racial disparities in automated speech recognition. *PNAS*, 117(14):7684–7689, 2020.
- Kumar, P., Grewal, M., and Srivastava, M. M. Boosted cascaded convnets for multilabel classification of thoracic diseases in chest radiographs. pp. 546–552, 2018.
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J. R. R., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., Duerig, T., and Ferrari, V. The open images dataset V4. *IJCV*, 128(7):1956–1981, 2020.
- Lee, C., Lehoczky, J. P., Rajkumar, R., and Siewiorek, D. P. On quality of service optimization with discrete qos options. In *RTAS*, pp. 276, 1999.
- Li, X., Sun, C., and Ye, Y. Simple and fast algorithm for binary integer and online lp. In *NeurIPS*, 2020.
- Lin, T., Maire, M., Belongie, S. J., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: common objects in context. In *ECCV*, volume 8693, pp. 740–755, 2014.
- Liu, H., Burnap, P., Alorainy, W., and Williams, M. L. Fuzzy multi-task learning for hate speech type identification. In *WWW*, pp. 3006–3012, 2019.
- Mendoza, H., Klein, A., Feurer, M., Springenberg, J. T., Urban, M., Burkart, M., Dippel, M., Lindauer, M., and Hutter, F. Towards automatically-tuned deep neural networks. In *Automated Machine Learning - Methods, Systems, Challenges*, pp. 135–149. 2019.
- Moyano, J. M., Galindo, E. L. G., Cios, K. J., and Ventura, S. Review of ensembles of multi-label classifiers: Models, experimental study and prospects. *Inf. Fusion*, 44: 33–45, 2018.
- Pakrashi, A. and Namee, B. M. CascadeML: An automatic neural network architecture evolution and training algorithm for multi-label classification. In *AI*, volume 11927, pp. 3–17, 2019.
- Pamela H. Vance, S. M. and Toth), P. Knapsack problems: Algorithms and computer implementations. *SIAM Rev.*, 35(4):684–685, 1993.
- Read, J., Pfahringer, B., and Holmes, G. Multi-label classification using ensembles of pruned sets. In *ICDM*, pp. 995–1000, 2008.
- Read, J., Pfahringer, B., Holmes, G., and Frank, E. Classifier chains for multi-label classification. *Mach. Learn.*, 85(3):333–359, 2011.
- Reis, A., Paulino, D., Filipe, V., and Barroso, J. Using online artificial vision services to assist the blind - an assessment of microsoft cognitive services and google cloud vision. In *WorldCIST*, pp. 174–184, 2018.
- Sang, E. F. T. K. and Meulder, F. D. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *CoNLL*, pp. 142–147, 2003.
- Sinha, P. and Zoltners, A. A. The multiple-choice knapsack problem. *Oper. Res.*, 27(3):503–515, 1979.
- Sun, Y., Wang, X., and Tang, X. Deep convolutional network cascade for facial point detection. In *CVPR*, pp. 3476–3483, 2013.
- Sun, Y., Karatzas, D., Chan, C. S., Jin, L., Ni, Z., Chng, C. K., Liu, Y., Luo, C., Ng, C. C., Han, J., Ding, E., and Liu, J. ICDAR 2019 competition on large-scale street view text with partial labeling. In *ICDAR*, pp. 1557–1562, 2019.
- Thornton, C., Hutter, F., Hoos, H. H., and Leyton-Brown, K. Auto-WEKA: Combined selection and hyperparameter optimization of classification algorithms. In *KDD*, pp. 847–855, 2013.
- Tsoumakas, G. and Vlahavas, I. P. Random k -labelsets: An ensemble method for multilabel classification. In *ECML*, volume 4701, pp. 406–417, 2007.

- Viola, P. and Jones, M. Robust real-time object detection. In *International Journal of Computer Vision*, 2001a.
- Viola, P. A. and Jones, M. J. Fast and robust classification using asymmetric adaboost and a detector cascade. In *NIPS*, pp. 1311–1318, 2001b.
- Wang, L., Lin, J. J., and Metzler, D. A cascade ranking model for efficient ranked retrieval. In *SIGIR*, 2011.
- Weischedel, R., Hovy, E., Marcus, M., and Palmer, M. Ontonotes : A large training corpus for enhanced processing. 2017.
- Wever, M., Tornede, A., Mohr, F., and Hüllermeier, E. Automl for multi-label classification: Overview and empirical evaluation. *IEEE TPAMI*, 43(9):3037–3054, 2021.
- Xu, X., Jiang, Y., Peng, L., Xue, X., and Zhou, Z. Ensemble approach based on conditional random field for multi-label image and video annotation. In *MM*, pp. 1377–1380, 2011.
- Xu, Z. E., Kusner, M. J., Weinberger, K. Q., Chen, M., and Chapelle, O. Classifier cascades and trees for minimizing feature evaluation cost. *J. Mach. Learn. Res.*, 15(1): 2113–2144, 2014.
- Yao, Y., Xiao, Z., Wang, B., Viswanath, B., Zheng, H., and Zhao, B. Y. Complexity vs. performance: empirical analysis of machine learning as a service. In *IMC*, pp. 384–397, 2017.
- Zhang, M. and Zhou, Z. A review on multi-label learning algorithms. *IEEE TKDE*, 26(8):1819–1837, 2014.
- Zhang, R., Yang, M., Bai, X., Shi, B., Karatzas, D., Lu, S., Jawahar, C. V., Zhou, Y., Jiang, Q., Song, Q., Li, N., Zhou, K., Wang, L., Wang, D., and Liao, M. ICDAR 2019 robust reading challenge on reading chinese text on signboard. In *ICDAR*, pp. 1577–1581, 2019.
- Zhou, Y. and Naroditskiy, V. Algorithm for stochastic multiple-choice knapsack problem and application to keywords bidding. In *WWW*, pp. 1175–1176, 2008.

Outline. The appendix is organized as follows. We present missing technical details in Section A. The proofs are provided in Section B. Finally, Section C gives detailed experiment setups and additional empirical results.

A. Technical Details

Additional technical details are presented here.

Label combiner parameter search. Recall that the label combiner requires two parameters: the combining weight $w \in [0, 1]$ and the quality score threshold $\theta \in [0, 1]$. We adopt a simple grid search approach to select w and θ . More precisely, we first create a parameter candidate set $PCS \triangleq \{w_0, w_1, w_2, \dots, w_M\} \times \{\theta_0, \theta_1, \theta_2, \dots, \theta_M\}$, where $w_m = \frac{m}{M}$ and $\theta_i = \frac{m}{M}$. Next, for each $(w, \theta) \in PCS$, we evaluate the performance of combining the base service and the k th service using (w, θ) , and select the parameter that gives the highest accuracy. Note that this involves M^2 number of label combinations for each $k \in [K]$. In practice, we have found that $M = 10$ is sufficient to obtain a good combiner.

δ selection in Algorithm 1. A naive approach is to set a small constant value, say, $\delta = 0.01$. To obtain a more accurate strategy, we can adopt a search algorithm to select the best δ value based on the evaluation the performance on a validation dataset. More precisely, we first create a constant set CS . Then for each $\alpha \in CS$, let $\delta = \alpha \frac{\log N}{N}$, and then solve Problem 4.2 to obtain the parameter \hat{p} , evaluate the performance on a validation dataset. Finally, we select the $\alpha \in CS$ that achieves the highest accuracy on the validation dataset. In practice, we have found that $CS = \{-10, -9, -8, \dots, 0, 1, 2, \dots, 10\}$ is sufficient to obtain a highly accurate solution.

B. Proofs

For ease of notations, let us introduce $\hat{b}_k \triangleq b - \mathbf{c}_{base}$ and $\hat{\mathbf{c}}_k \triangleq \mathbf{c}_k \cdot \mathbb{1}_{k \neq base}$ first. Then we can rewrite the API selection problem (Problem 4.1) as

$$\begin{aligned} \max_{\mathbf{Z} \in \mathbb{R}^{N \times K}}: & \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k} \hat{\mathbf{a}}_k(x_n) \\ \text{s.t.} & \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k} \hat{\mathbf{c}}_k \leq \hat{b} \\ & \sum_{k=1}^K \mathbf{Z}_{n,k} = 1, \forall n; \mathbf{Z}_{n,k} \in \{0, 1\}, \forall n, k \end{aligned} \quad (\text{B.1})$$

Its corresponding linear programming simply becomes

$$\begin{aligned} \max_{\mathbf{Z} \in \mathbb{R}^{N \times K}}: & \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k} \hat{\mathbf{a}}_k(x_n) \\ \text{s.t.} & \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k} \hat{\mathbf{c}}_k \leq \hat{b} \\ & \sum_{k=1}^K \mathbf{Z}_{n,k} = 1, \forall n; \mathbf{Z}_{n,k} \in [0, 1], \forall n, k \end{aligned} \quad (\text{B.2})$$

We will analyze some useful properties for those two problems first, and then prove the desired results for the original API selection problem on top of those properties.

B.1. Helpful Lemmas

Before proving the desired results, let us also provide a few generic lemmas.

Lemma B.1. *Let $\mathbf{A} \in \mathbb{R}^{N_1 \times N_2}$ be a fixed matrix and $\boldsymbol{\beta} \in \mathbb{R}^{N_1}$ be a random vector. If $\boldsymbol{\beta}$ is supported on $[0, 1]^{N_1}$ with a continuous density function, then with probability 1,*

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \boldsymbol{\beta}\|_0 \geq N_1 - N_2.$$

Proof. If $N_1 \leq N_2$ then the above inequality obviously holds. Suppose $N_1 > N_2$. We prove this by contradiction. Assume the inequality does not hold. Then there exists some \mathbf{x}' , such that with probability larger than 0,

$$\|\mathbf{A}\mathbf{x}' - \boldsymbol{\beta}\|_0 < N_1 - N_2.$$

That is to say, at least $N_1 - (N_1 - N_2) + 1 = N_2 + 1$ many equations in $\mathbf{A}\mathbf{x}' = \boldsymbol{\beta}$ can be forced to 0. Let U be the set of those $N_2 + 1$ indexes. Then formally we have

$$\mathbf{A}_U \mathbf{x}' = \boldsymbol{\beta}_U.$$

That is to say, with probability larger than 0, $\boldsymbol{\beta}_U$ is in the subspace formed by the columns of \mathbf{A}_U .

On the other hand, we can show that for any set of indexes V with $|V| = N_2 + 1$, $\boldsymbol{\beta}_V$ lies in the subspace formed by the columns of \mathbf{A}_V with probability 0, which gives a contradiction. To see this, let us start by considering a fixed set of indexes V . Let Ω_V denote the subspace formed by \mathbf{A}_V and $p_{\boldsymbol{\beta}_V}(\cdot)$ be the density function of $\boldsymbol{\beta}_V$. The density function of $\boldsymbol{\beta}$ is continuous and thus $p_{\boldsymbol{\beta}_V}(\cdot)$ is also continuous. The support of $\boldsymbol{\beta}$ is in $[0, 1]^{N_1}$, and thus the support of $\boldsymbol{\beta}_V$ is in $[0, 1]^{N_2+1}$ (since $|V| = N_2 + 1$ by definition). That is to say, $p_{\boldsymbol{\beta}_V}(\cdot)$ is a continuous function on a compact set. Therefore, $p_{\boldsymbol{\beta}_V}(\cdot)$ must be bounded, i.e., there exists a constant p^{sup} such that $p_{\boldsymbol{\beta}_V}(\cdot) \leq p^{\text{sup}}$. Hence we have

$$\Pr[\boldsymbol{\beta}_V \in \Omega_V] = \int_{\mathbf{x}_V \in \Omega_V} p_{\boldsymbol{\beta}_V}(\mathbf{x}) d\mathbf{x} \leq \int_{\mathbf{x}_V \in \Omega_V} p^{\text{sup}} d\mathbf{x} = p^{\text{sup}} \int_{\mathbf{x}_V \in \Omega_V} 1 d\mathbf{x}$$

where the first equation is by definition of the random variable $\boldsymbol{\beta}_V$, the inequality is by increasing the density function to its upper bound p^{sup} , and the last equation simply moves the constant out of the integral. In addition, Ω_V is a $N_2 + 1$ dimensional space spanned by N_2 vectors, which implies that its measure in \mathbb{R}^{N_2+1} is 0, i.e., $\int_{\mathbf{x}_V \in \Omega_V} 1 d\mathbf{x} = 0$. Thus, we have just shown that

$$\Pr[\boldsymbol{\beta}_V \in \Omega_V] \leq p^{\text{sup}} \int_{\mathbf{x}_V \in \Omega_V} 1 d\mathbf{x} = 0$$

Probability is non-negative, and thus $\Pr[\boldsymbol{\beta}_V \in \Omega_V] = 0$ (for a fixed V). Note that the size of V is $N_2 + 1$ and there are in total N_1 possible indexes. Thus, there are $\binom{N_1}{N_2+1}$ many possible choices of V . Applying union bound, we have for any V , $\Pr[\boldsymbol{\beta}_V \in \Omega_V] = 0$. A contradiction. The assumption is incorrect, and thus we must have

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \boldsymbol{\beta}\|_0 \geq N_1 - N_2.$$

□

Lemma B.2. *Let f be a function defined on $\Omega_{\mathbf{z}}$. Assume there exists a set $\Omega_{\mathbf{z},1} \subseteq \Omega_{\mathbf{z}}$, such that for any $\mathbf{z} \in \Omega_{\mathbf{z}}$, there exists $\mathbf{z}' \in \Omega_{\mathbf{z},1}$, such that $\|f(\mathbf{z}) - f(\mathbf{z}')\| \leq \Delta$. Then we have*

$$\|f(\mathbf{z}^*) - f(\mathbf{z}_1^*)\| \leq \Delta,$$

where $\mathbf{z}^* = \arg \max_{\mathbf{z} \in \Omega_{\mathbf{z}}} f(\mathbf{z})$, $\mathbf{z}_1^* = \arg \max_{\mathbf{z} \in \Omega_{\mathbf{z},1}} f(\mathbf{z})$.

Proof. By assumption, there exists a $\mathbf{z}' \in \Omega_{\mathbf{z},1}$, such that

$$\|f(\mathbf{z}^*) - f(\mathbf{z}')\| \leq \Delta$$

which implies

$$f(\mathbf{z}') \geq f(\mathbf{z}^*) - \Delta$$

Noting that \mathbf{z}_1^* is the optimal solution on $\Omega_{\mathbf{z},1}$ and \mathbf{z}' is a feasible solution, we have

$$f(\mathbf{z}_1^*) \geq f(\mathbf{z}')$$

Combining the above two inequalities, we have

$$f(\mathbf{z}_1^*) \geq f(\mathbf{z}^*) - \Delta$$

On the other hand, since $\Omega_{\mathbf{z},1} \subseteq \Omega_{\mathbf{z}}$, \mathbf{z}_1^* is a feasible solution on $\Omega_{\mathbf{z}}$, and thus we have

$$f(\mathbf{z}_1^*) \leq f(\mathbf{z}^*) \leq f(\mathbf{z}^*) + \Delta$$

Combing those two inequalities we have

$$\|f(\mathbf{z}_1^*) - f(\mathbf{z}^*)\| \leq \Delta$$

which completes the proof. □

Lemma B.3. Let X_1, X_2, \dots, X_{N_1} and $X'_1, X'_2, \dots, X'_{N_2}$ be two i.i.d. samples from the same distribution which lies in $[x_{\inf}, x_{\sup}]$. Then we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N_2} \sum_{n=1}^{N_2} X'_n - \frac{1}{N_1} \sum_{n=1}^{N_1} X_n \right\| \leq (x_{\sup} - x_{\inf}) \left[\sqrt{\frac{\log 4 - \log \epsilon}{2N_2}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N_1}} \right].$$

Proof. We can apply the Hoeffding's inequality for both sequences separately, and we can obtain with probability $1 - \epsilon$,

$$\left\| \frac{1}{N_1} \sum_{n=1}^{N_1} X_n - \mathbb{E}[X_1] \right\| \leq (x_{\sup} - x_{\inf}) \sqrt{\frac{\log 2 - \log \epsilon}{2N_1}}$$

and with probability $1 - \epsilon$

$$\left\| \frac{1}{N_2} \sum_{n=1}^{N_2} X'_n - \mathbb{E}[X_1] \right\| \leq (x_{\sup} - x_{\inf}) \sqrt{\frac{\log 2 - \log \epsilon}{2N_2}}$$

Now applying union bound, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N_1} \sum_{n=1}^{N_1} X_n - \mathbb{E}[X_1] \right\| \leq (x_{\sup} - x_{\inf}) \sqrt{\frac{\log 4 - \log \epsilon}{2N_1}}$$

and $1 - \epsilon$

$$\left\| \frac{1}{N_2} \sum_{n=1}^{N_2} X'_n - \mathbb{E}[X_1] \right\| \leq (x_{\sup} - x_{\inf}) \sqrt{\frac{\log 4 - \log \epsilon}{2N_2}}$$

Now applying the triangle inequality, we have

$$\left\| \frac{1}{N_2} \sum_{n=1}^{N_2} X'_n - \frac{1}{N_1} \sum_{n=1}^{N_1} X_n \right\| \leq (x_{\sup} - x_{\inf}) \left[\sqrt{\frac{\log 4 - \log \epsilon}{2N_2}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N_1}} \right]$$

which completes the proof. \square

Lemma B.4. Let f_1, f_2, g_1, g_2 be functions defined on $\Omega_{\mathbf{z}}$, such that $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} |(f_1(\mathbf{z}) - f_2(\mathbf{z}))| \leq \Delta_1$ and $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} \|g_2(\mathbf{z}) - g_1(\mathbf{z})\| \leq \Delta_2$. Suppose

$$\begin{aligned} \mathbf{z}_1^* &= \arg \max_{\mathbf{z} \in \Omega_{\mathbf{z}}} f_1(\mathbf{z}) \\ &\text{s.t. } g_1(\mathbf{z}) \leq 0 \end{aligned}$$

and

$$\begin{aligned} \mathbf{z}_2^* &= \arg \max_{\mathbf{z} \in \Omega_{\mathbf{z}}} f_2(\mathbf{z}) \\ &\text{s.t. } g_2(\mathbf{z}) \leq \Delta_2, \end{aligned}$$

then we must have

$$\begin{aligned} f_1(\mathbf{z}_2^*) &\geq f_1(\mathbf{z}_1^*) - 2\Delta_1 \\ g_1(\mathbf{z}_2^*) &\leq 2\Delta_2. \end{aligned}$$

Proof. Note that $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} |(f_1(\mathbf{z}) - f_2(\mathbf{z}))| \leq \Delta_1$ implies $f_1(\mathbf{z}) \geq f_2(\mathbf{z}) - \Delta_1$ for any $\mathbf{z} \in \Omega_{\mathbf{z}}$. Specifically,

$$f_1(\mathbf{z}_2^*) \geq f_2(\mathbf{z}_2^*) - \Delta_1$$

Noting $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} \|g_2(\mathbf{z}) - g_1(\mathbf{z})\| \leq \Delta_2$, we have $g_2(\mathbf{z}_1^*) \leq g_1(\mathbf{z}_1^*) - \Delta_2 \leq -\Delta_2$, where the last inequality is due to $g_1(\mathbf{z}_1^*) \leq 0$ by definition. Since, \mathbf{z}_1^* is a feasible solution to the second optimization problem, and the optimal value must be no smaller than the value at \mathbf{z}_1^* . That is to say,

$$f_2(\mathbf{z}_2^*) \geq f_2(\mathbf{z}_1^*)$$

Hence we have

$$f_1(\mathbf{z}_2^*) \geq f_2(\mathbf{z}_2^*) - \Delta_1 \geq f_2(\mathbf{z}_1^*) - \Delta_1$$

In addition, $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} |(f_1(\mathbf{z}) - f_2(\mathbf{z}))| \leq \Delta_1$ implies $f_2(\mathbf{z}) \geq f_1(\mathbf{z}) - \Delta_1$ for any $\mathbf{z} \in \Omega_{\mathbf{z}}$. Thus, we have $f_2(\mathbf{z}_1^*) \geq f_1(\mathbf{z}_1^*) - \Delta_1$ and thus

$$f_1(\mathbf{z}_2^*) \geq f_2(\mathbf{z}_1^*) - \Delta_1 \geq f_1(\mathbf{z}_1^*) - 2\Delta_1$$

By $\max_{\mathbf{z} \in \Omega_{\mathbf{z}}} |g_1(\mathbf{z}) - g_2(\mathbf{z})| \leq \Delta_2$, we must have $g_1(\mathbf{z}_2^*) \leq g_2(\mathbf{z}_2^*) + \Delta_2 \leq 2\Delta_2$, where the last inequality is by definition of \mathbf{z}' , which completes the proof. \square

B.2. Proof of Theorem 4.2

Proof. We give a constructive proof via explicitly giving the value of p^* . In fact, let p^* and \mathbf{q}^* be the optimal solution to

$$\begin{aligned} \min_{p, \mathbf{q}} \quad & \hat{b}p + \sum_{n=1}^N \mathbf{q}_n \\ \text{s.t.} \quad & \frac{1}{N} \hat{\mathbf{c}}_k p + \mathbf{q}_n \geq \frac{1}{N} \hat{\mathbf{a}}_k(x_n) \\ & p, \mathbf{q} \geq 0 \end{aligned} \tag{B.3}$$

Then our goal is to show that for this constructed p^* , s^{p^*} is a feasible solution to Problem 4.1 and $r(s^{p^*}) \geq r(s^*) - \frac{1}{N}$ with probability 1 (Since probabilistic statement is only introduced in Lemma B.1 whose result holds with probability 1, and we only apply it finite times, we will omit the probabilistic statement for the rest of the proof for simplicity). To achieve this, let us construct a $N \times K$ matrix

$$\tilde{\mathbf{Z}}_{n,k}^{p^*} \triangleq \mathbb{1}_{s^{p^*}(x_n)=k}$$

It is not hard to see that $s^{p^*}(x_n) = \arg \max_k \tilde{\mathbf{Z}}_{n,k}^{p^*}$. By construction of s^{p^*} , feasibility of s^{p^*} to Problem 4.1 is equivalent to feasibility of $\tilde{\mathbf{Z}}^{p^*}$ to Problem B.1. By construction of s^* and s^{p^*} , $r(s^{p^*}) \geq r(s^*) - \frac{1}{N}$ is equivalent to

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{a}}_k(x_n) \geq \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^* \hat{\mathbf{a}}_k(x_n) - \frac{1}{N}.$$

Therefore, our goal becomes showing the feasibility of $\tilde{\mathbf{Z}}^{p^*}$ and the above inequality. By construction of $\tilde{\mathbf{Z}}^{p^*}$, the natural constraints ($\tilde{\mathbf{Z}}_{n,k}^{p^*} \in \{0, 1\}$ and $\sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} = 1, \forall n$) are obviously satisfied. Thus, we only need to show $\tilde{\mathbf{Z}}^{p^*}$ satisfies the budget constraint and the above inequality. To show those two results, let us introduce another variable $\mathbf{Z}^{*,LP}$, which represents a sparse optimal solution to the relaxed version of Problem B.1 (i.e., Problem B.2). The proof idea is then (roughly) to show (i) that $\tilde{\mathbf{Z}}^{p^*}$ is actually close to $\mathbf{Z}^{*,LP}$, (ii) that $\mathbf{Z}^{*,LP}$ satisfies the budget constraint and gives an estimated accuracy as high as that of the optimal solution \mathbf{Z}^* , and (iii) that the difference between $\tilde{\mathbf{Z}}^{p^*}$ and $\mathbf{Z}^{*,LP}$ does not break the budget constraints and only decreases the estimated accuracy by $1/N$. Combining the three points finishes the proof. Now we formalize this idea.

Step 1: We first show that $\tilde{\mathbf{Z}}^{p^*}$ and $\mathbf{Z}^{*,LP}$ are close to each other.

Lemma B.5. *Let $\mathbf{Z}^{*,LP}$ be an optimal solution to Problem B.2. Then there exists some constant n' , such that $\tilde{\mathbf{Z}}_{n,\cdot}^{p^*} = \mathbf{Z}_{n,\cdot}^{*,LP}, \forall n \neq n'$.*

Proof. Note that Problem B.3 is the dual problem to Problem B.2. We can write the complementary slackness constraints as follows

$$\mathbf{Z}_{n,k}^{*,LP} \left(\frac{1}{N} \hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N} \hat{\mathbf{a}}_k(x_n) \right) = 0, \forall n, k$$

Now let us construct the matrix

$$\mathbf{A} = \begin{bmatrix} \frac{1}{N} \hat{\mathbf{c}}, & \mathbf{1}, & \mathbf{0}, & \cdots, & \mathbf{0} \\ \frac{1}{N} \hat{\mathbf{c}}, & \mathbf{0}, & \mathbf{1}, & \cdots, & \mathbf{0} \\ \vdots, & \vdots, & \cdots, & \ddots, & \vdots \\ \frac{1}{N} \hat{\mathbf{c}}, & \mathbf{0}, & \mathbf{0}, & \cdots, & \mathbf{1} \end{bmatrix} \in \mathbb{R}^{NK \times (N+1)}$$

and the vector

$$\boldsymbol{\beta} = \frac{1}{N} \begin{bmatrix} \hat{\mathbf{a}}(x_1) \\ \hat{\mathbf{a}}(x_2) \\ \vdots \\ \hat{\mathbf{a}}(x_N) \end{bmatrix} \in \mathbb{R}^{NK}.$$

Then by Lemma B.1, $\min_{\mathbf{x}} \|\mathbf{Ax} - \boldsymbol{\beta}\|_0 \geq NK - N - 1$. Specifically, if $\mathbf{x} = [p^*, \mathbf{q}^{*T}]^T$, then we should have $\|\mathbf{Ax} - \boldsymbol{\beta}\|_0 \geq NK - N - 1$. Note that each row of $\mathbf{Ax} - \boldsymbol{\beta}$ corresponds to $\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n)$, and thus we effectively have

$$\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) \neq 0$$

for at least $NK - N - 1$ choices of n, k . In other words, among all possible choices of n, k , at most $N + 1$ many of them satisfies

$$\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) = 0$$

Furthermore, note that the constraint $\sum_{k=1}^K \mathbf{z}_k^{*LP}(x_n) = 1$ ensures that for any n , there must exist at least one k' such that $\mathbf{Z}_{n,k}^{*,LP} \neq 0$ and thus $\frac{1}{N}\hat{\mathbf{c}}_{k'} p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_{k'}(x_n) = 0$. By the pigeonhole principle, we can conclude that for all n except one (denoted by n'), exactly one equation in $\{\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) = 0\}_k$ can be satisfied.

Now let us fix any $n \neq n'$. Then there exists some k' , such that $\frac{1}{N}\hat{\mathbf{c}}_{k'} p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_{k'}(x_n) = 0$, and for any $k \neq k'$, $\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) > 0$ (due to the natural constraint in Problem B.3). That is to say, for any $k \neq k'$,

$$\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) > 0 = \frac{1}{N}\hat{\mathbf{c}}_{k'} p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_{k'}(x_n)$$

Multiplying N and rearranging the terms gives

$$\hat{\mathbf{a}}_{k'}(x_n) - \hat{\mathbf{c}}_{k'} p^* > \hat{\mathbf{a}}_k(x_n) - \hat{\mathbf{c}}_k p^*$$

That is to say, k' is the unique solution to $\max_k \hat{\mathbf{a}}_k(x_n) - \hat{\mathbf{c}}_k p^*$. By definition of $\tilde{\mathbf{Z}}^{p^*}$, we have $\tilde{\mathbf{Z}}_{n,k'}^{p^*} = 1$ and $\tilde{\mathbf{Z}}_{n,k}^{p^*} = 0, \forall k \neq k'$. Meanwhile, for any $k \neq k'$, by the slackness constraint, since $\frac{1}{N}\hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* - \frac{1}{N}\hat{\mathbf{a}}_k(x_n) > 0$, we must have $\mathbf{Z}_{n,k}^{*,LP} = 0$. The natural constraint in Problem B.2 requires $\sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} = 1$. Thus, we have $\mathbf{Z}_{n,k'}^{*,LP} = \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} - \sum_{k \neq k'} \mathbf{Z}_{n,k}^{*,LP} = 1$.

That is to say, for any $n \neq n'$, we always have $\tilde{\mathbf{Z}}_{n,\cdot}^{p^*} = \mathbf{Z}_{n,\cdot}^{*,LP}$, which completes the proof. \square

Step 2: Now we can show $\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{a}}_k(x_n) \geq \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{a}}_k(x_n) - \frac{1}{N}$. To see this, by Lemma B.5, $\tilde{\mathbf{Z}}_{n,\cdot}^{p^*} = \mathbf{Z}_{n,\cdot}^{*,LP}, \forall n \neq n'$, we must have

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{a}}_k(x_n) - \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{a}}_k(x_n) = \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{a}}_k(x'_n) - \frac{1}{N} \sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{a}}_k(x'_n)$$

As $\hat{\mathbf{a}}_k(x'_n)$ is bounded in $[0, 1]$, we have

$$\frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{a}}_k(x'_n) - \frac{1}{N} \sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{a}}_k(x'_n) \leq \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \cdot 1 - \frac{1}{N} \sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \cdot 0 = \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP}$$

By natural constraint in Problem B.2, $\sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} = 1$. Thus, we have

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{a}}_k(x_n) - \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{a}}_k(x_n) \leq \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} = \frac{1}{N}$$

On the other hand, $\mathbf{Z}^{*,LP}$ is the optimal solution to Problem B.2 and \mathbf{Z}^* is a feasible solution. Thus we have

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^* \hat{\mathbf{a}}_k(x_n) \leq \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{a}}_k(x_n)$$

Combining the two inequalities leads to

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{a}}_k(x_n) \geq \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^* \hat{\mathbf{a}}_k(x_n) - \frac{1}{N}$$

Step 3: Finally, we are ready to show the budget constraint is satisfied. By Lemma B.5, $\tilde{\mathbf{Z}}_{n,\cdot}^{p^*} = \mathbf{Z}_{n,\cdot}^{*,LP}$, $\forall n \neq n'$, we have

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{c}}_k - \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{c}}_k = \frac{1}{N} \sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k$$

Denote $s^{p^*}(x_{n'})$ by k_1 . By construction, we have

$$\sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k = \hat{\mathbf{c}}_{k_1} - \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k$$

Let S be the set of any k such that $\mathbf{Z}_{n',k}^{*,LP} \neq 0$. Then we can further write

$$\sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k = \hat{\mathbf{c}}_{k_1} - \sum_{k \in S} \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k$$

Note that $k \in S$ implies $k \in \arg \max_k \hat{\mathbf{a}}_k(x_{n'}) - \hat{\mathbf{c}}_k p^*$ (Suppose not. Then there exists some k' , such that $\hat{\mathbf{a}}_{k'}(x_{n'}) - \hat{\mathbf{c}}_{k'} p^* > \hat{\mathbf{a}}_k(x_{n'}) - \hat{\mathbf{c}}_k p^*$. Multiplying both sides by $-\frac{1}{N}$ and then adding \mathbf{q}_n^* gives $-\frac{1}{N} \hat{\mathbf{a}}_{k'}(x_{n'}) + \frac{1}{N} \hat{\mathbf{c}}_{k'} p^* + \mathbf{q}_n^* < -\frac{1}{N} \hat{\mathbf{a}}_k(x_{n'}) + \frac{1}{N} \hat{\mathbf{c}}_k p^* + \mathbf{q}_n^*$. By complementary slackness of Problem B.2, $\mathbf{Z}_{n,k}^{*,LP} (-\frac{1}{N} \hat{\mathbf{a}}_k(x_{n'}) + \frac{1}{N} \hat{\mathbf{c}}_k p^* + \mathbf{q}_n^*) = 0$. $k \in S$ implies $\mathbf{Z}_{n,k}^{*,LP} \neq 0$ and thus $-\frac{1}{N} \hat{\mathbf{a}}_k(x_{n'}) + \frac{1}{N} \hat{\mathbf{c}}_k p^* + \mathbf{q}_n^* = 0$. Thus, $-\frac{1}{N} \hat{\mathbf{a}}_{k'}(x_{n'}) + \frac{1}{N} \hat{\mathbf{c}}_{k'} p^* + \mathbf{q}_n^* < 0$, which contradicts with the feasibility constraint in the dual problem.). Recall that k_1 is determined by $\arg \max_k \hat{\mathbf{a}}_k(x_{n'}) - \hat{\mathbf{c}}_k p^*$ and we break ties by picking k with smallest cost. Thus, for any $k \in S$, $\hat{\mathbf{c}}_k \geq \hat{\mathbf{c}}_{k_1}$. Therefore,

$$\sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k \leq \hat{\mathbf{c}}_{k_1} - \sum_{k \in S} \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_{k_1} = (1 - \sum_{k \in S} \mathbf{Z}_{n',k}^{*,LP}) \hat{\mathbf{c}}_{k_1}$$

By feasibility constraint in Problem B.2, $\sum_{k \in S} \mathbf{Z}_{n',k}^{*,LP} = \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} = 1$. Thus, the above inequality becomes $\sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k \leq 0$. Thus,

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{c}}_k - \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{c}}_k = \frac{1}{N} \sum_{k=1}^K \tilde{\mathbf{Z}}_{n',k}^{p^*} \hat{\mathbf{c}}_k - \frac{1}{N} \sum_{k=1}^K \mathbf{Z}_{n',k}^{*,LP} \hat{\mathbf{c}}_k \leq 0$$

$\mathbf{Z}^{*,LP}$ is a feasible solution to Problem B.2, so it must satisfy the budget constraint and thus $\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{c}}_k \leq b$. Hence, we must have

$$\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \tilde{\mathbf{Z}}_{n,k}^{p^*} \hat{\mathbf{c}}_k \leq \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{Z}_{n,k}^{*,LP} \hat{\mathbf{c}}_k \leq b \leq b$$

i.e., $\tilde{\mathbf{Z}}^{p^*}$ satisfies the budget constraint in Problem B.1.

Finally, combining step 2 and step 3 finishes the proof. \square

B.3. Proof of Theorem 4.3

Proof. Let us first establish a few lemmas consisting of the main components of the proof.

Lemma B.6. Suppose $\delta \geq \frac{\|\mathbf{c}\|_\infty}{b} \left[\sqrt{\frac{\log 4 - \log \epsilon}{N}} + \sqrt{\frac{\log 4 - \log \epsilon}{N^{Tr}}} \right]$. Then with probability at least $1 - \epsilon$, $s^{\hat{b}}$ is a feasible solution to Problem 4.1.

Proof. We first note that Problem 4.2 is a linear programming, and its dual problem is

$$\begin{aligned} \max_{\mathbf{z} \in \mathbb{R}^{N \times K}}: & \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \mathbf{z}_{n,k} \hat{\mathbf{a}}_k(x_n^{Tr}) \\ \text{s.t.} & \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \sum_{k=1}^K \mathbf{z}_{n,k} \hat{\mathbf{c}}_k \leq (1 - \delta) \hat{b} \\ & \sum_{k=1}^K \mathbf{z}_{n,k} = 1, \mathbf{z}_{n,k} \in [0, 1], \forall n, k \end{aligned} \quad (\text{B.4})$$

Note that this is in the same form of Problem B.2 except that the data become $\{x_n^{Tr}\}_{n=1}^{N^{Tr}}$ instead of $\{x_n\}_{n=1}^N$. Using a similar argument in the proof for Theorem 4.2, $s^{\hat{b}}(x_n^{Tr})$ is a feasible solution to

$$\begin{aligned} \max & \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{\hat{b}}}(x_n^{Tr}) \\ \text{s.t.} & \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^{\hat{b}}]}(x_n^{Tr}, \mathbf{c}) \leq (1 - \delta)b, \end{aligned}$$

and thus we have

$$\frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^{\hat{b}}]}(x_n^{Tr}, \mathbf{c}) \leq (1 - \delta)b$$

Note that training data x_n^{Tr} are i.i.d samples from the true distribution and $0 \leq \eta^{[s]}(x_n^{Tr}, \mathbf{c}) \leq \|\mathbf{c}\|_\infty$. Thus, by Hoeffding's inequality, with probability $1 - \epsilon$, we have

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^{\hat{b}}]}(x_n^{Tr}, \mathbf{c}) - \mathbb{E} \left[\eta^{[s^{\hat{b}}]}(x, \mathbf{c}) \right] \right\| \leq \|\mathbf{c}\|_\infty \sqrt{\frac{\log 2 - \log \epsilon}{2N^{Tr}}}$$

The data stream x_n is also from the same distribution, and thus we also have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N \eta^{[s^{\hat{b}}]}(x_n, \mathbf{c}) - \mathbb{E} \left[\eta^{[s^{\hat{b}}]}(x, \mathbf{c}) \right] \right\| \leq \|\mathbf{c}\|_\infty \sqrt{\frac{\log 2 - \log \epsilon}{2N}}$$

Applying union bound, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^{\hat{b}}]}(x_n^{Tr}, \mathbf{c}) - \mathbb{E} \left[\eta^{[s^{\hat{b}}]}(x, \mathbf{c}) \right] \right\| \leq \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}}$$

and

$$\left\| \frac{1}{N} \sum_{n=1}^N \eta^{[s^{\hat{b}}]}(x_n, \mathbf{c}) - \mathbb{E} \left[\eta^{[s^{\hat{b}}]}(x, \mathbf{c}) \right] \right\| \leq \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N}}$$

Using triangle inequality, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^p]}(x_n^{Tr}, \mathbf{c}) - \frac{1}{N} \sum_{n=1}^N \eta^{[s^p]}(x_n, \mathbf{c}) \right\| \leq \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}}.$$

Thus we have

$$\begin{aligned} \frac{1}{N} \sum_{n=1}^N \eta^{[s^p]}(x_n, \mathbf{c}) &\leq \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^p]}(x_n^{Tr}, \mathbf{c}) + \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}} \\ &\leq (1 - \delta)b + \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \|\mathbf{c}\|_\infty \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}} \leq b \end{aligned}$$

where the last inequality is due to the assumption on δ . That is to say, with probability $1 - \epsilon$, $s^{\hat{p}}$ is a feasible solution to Problem 4.1, which completes the proof. \square

Lemma B.7. Construct the set $\Omega_M \triangleq \{0, \frac{1}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}, \frac{2}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}, \dots, \frac{1}{\min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}\}$ and

$$\begin{aligned} \hat{p}(\Omega_M) &\triangleq \arg \max_{p \in \Omega_M} \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^p}(x_n^{Tr}) \\ &\text{s.t. } \frac{1}{N^{Tr}} \eta^{[s^p]}(x_n, \mathbf{c}) \leq (1 - \delta)b. \end{aligned}$$

Then with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \right\| \leq O\left(\sqrt{\frac{\log N + \log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log N^{Tr} + \log 8 - \log \epsilon}{2N^{Tr}}}\right).$$

Proof. Note that $\Omega_M \subseteq \mathbb{R}$. Consider an element $p \in \mathbb{R}$.

(i) $p \geq \frac{1}{\min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}$: This effectively means the API with the smallest cost is always selected. In other words, we always have

$$bs = \arg \max \hat{\mathbf{a}}_k(x) - p\hat{\mathbf{c}}_k$$

To see this, simply note that for any other k_1 , we have

$$\begin{aligned} \hat{\mathbf{a}}_{bs}(x) - p\hat{\mathbf{c}}_{bs} - (\hat{\mathbf{a}}_{k_1}(x) - p\hat{\mathbf{c}}_{k_1}) &= \hat{\mathbf{a}}_{bs}(x) - \hat{\mathbf{a}}_{k_1}(x) + p(\hat{\mathbf{c}}_{k_1} - \hat{\mathbf{c}}_{bs}) = \hat{\mathbf{a}}_{bs}(x) - \hat{\mathbf{a}}_{k_1}(x) + p\hat{\mathbf{c}}_{k_1} \\ &\geq 0 - 1 + p\hat{\mathbf{c}}_{k_1} \geq -1 + \hat{\mathbf{c}}_{k_1} \cdot \frac{1}{\min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} \geq 0 \end{aligned}$$

Thus, for such p , the objective value is the same as that for $\frac{1}{\min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} \in \Omega_M$.

(ii): $0 \leq p \leq \frac{1}{\min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}$: By construction of Ω_M , there exists some m , such that $\frac{m}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} \leq p \leq \frac{m+1}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}$. Let $p_j \triangleq \frac{j}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}$ for ease of notations. Clearly, we have $p_m \in \Omega_M$.

Now let us partition the space of $\hat{\mathbf{a}}(x)$ into M regions, denoted by A_1, A_2, \dots, A_M . Abusing the notation a little bit, let $\phi(p, x) \triangleq \arg \max \hat{\mathbf{a}}_k(x) - p\hat{\mathbf{c}}_k$. A_1 is the set of all $\hat{\mathbf{a}}(x)$ such that $\phi(p, x)$ is a constant. A_2 is the set of all $\hat{\mathbf{a}}(x)$ such that $\phi(p, x)$ is a constant for p larger than p_1 . Generally, A_j is the set of all $\hat{\mathbf{a}}(x)$ such that $\phi(p, x)$ is a constant for p larger than p_{j-1} subtracting A_{j-1} . Formally,

$$A_j = \begin{cases} \{\hat{\mathbf{a}}(x) : \phi(p, x) \text{ is a constant}\}, & j = 1 \\ \{\hat{\mathbf{a}}(x) : \phi(p, x) \text{ is a constant if } p \geq p_{j-1}\} - A_{j-1}, & j > 1 \end{cases}$$

One can easily verify that $\{A_j\}$ form a partition of the space of the estimated accuracy, and further more, $\|A_j\| \leq \frac{\|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k}$. By the assumption of the distribution, there exists some constant u , such that $Pr(A) \leq u\|A\|$, for any A in the probability space. Thus, we must have

$$\Pr[\hat{\mathbf{a}}(x) \in A_j] \leq \|A_j\|u = \frac{u\|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k^2}$$

Now note that, when $p_m = \frac{m}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} \leq p \leq \frac{m+1}{(M-1) \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} = p_{m+1}$, only elements in A_m may affect the reward. More precisely, we have

$$\frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p_m}}(x_n) = \frac{1}{N} \sum_{x_n \in A_m} r^{s^p}(x_n) - \frac{1}{N} \sum_{x_n \in A_m} r^{s^{p_m}}(x_n)$$

Note that each estimated accuracy is an i.i.d sample from the true distribution, and its value is from $[0, 1]$, by Hoeffding's inequality, with probability $1 - \epsilon$, we have

$$\left\| \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{x_n \in A_j} - \Pr[x_n \in A_j] \right\| \leq \sqrt{\frac{\log 2 - \log \epsilon}{2N}}$$

Applying the union bound, we have for any j , with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{x_n \in A_j} - \Pr[x_n \in A_j] \right\| \leq \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}}$$

Therefore, we have with probability $1 - \epsilon$,

$$\begin{aligned} \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p_m}}(x_n) &= \frac{1}{N} \sum_{x_n \in A_m} r^{s^p}(x_n) - \frac{1}{N} \sum_{x_n \in A_m} r^{s^{p_m}}(x_n) \\ &\geq \sum_{x_n \in A_m} 0 - \frac{1}{N} \sum_{x_n \in A_m} 1 = \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{x_n \in A_m} \\ &\geq \Pr[x_n \in A_m] - \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \\ &\geq -\sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \end{aligned}$$

and similarly

$$\begin{aligned} \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p_m}}(x_n) &= \frac{1}{N} \sum_{x_n \in A_m} r^{s^p}(x_n) - \frac{1}{N} \sum_{x_n \in A_m} r^{s^{p_m}}(x_n) \\ &\leq \sum_{x_n \in A_m} 1 - \frac{1}{N} \sum_{x_n \in A_m} 0 = \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{x_n \in A_m} \\ &\leq \Pr[x_n \in A_m] + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \\ &\leq \frac{u\|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \end{aligned}$$

That is to say,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p_m}}(x_n) \right\| \leq \frac{u\|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \quad (\text{B.5})$$

Similarly, for the training dataset, we can also get, with probability $1 - \epsilon$,

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^p}(x_n^{Tr}) - \frac{1}{N^{Tr}} \sum_{n=1}^N r^{s^{p_m}}(x_n^{Tr}) \right\| \leq \frac{u \|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}}$$

Combining case (i) and case (ii), we have just shown that for any $p \in \mathbb{R}$, there exists another $p' \in \Omega_M$, such that

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^p}(x_n^{Tr}) - \frac{1}{N^{Tr}} \sum_{n=1}^N r^{s^{p'}}(x_n^{Tr}) \right\| \leq \frac{u \|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}} \quad (\text{B.6})$$

Thus, applying Lemma B.2, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{\hat{p}}}(x_n^{Tr}) - \frac{1}{N^{Tr}} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n^{Tr}) \right\| \leq \frac{u \|\mathbf{c}\|_1}{M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + \sqrt{\frac{\log M + \log 2 - \log \epsilon}{2N}}$$

Now by Lemma B.3, for each fixed j , we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{p_j}}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{p_j}}(x_n^{Tr}) \right\| \leq \left[\sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}} \right]$$

Applying union bound, with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{p_j}}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{p_j}}(x_n^{Tr}) \right\| \leq \left[\sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N^{Tr}}} \right]$$

for all j . Specifically, we have

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{\hat{p}(\Omega_M)}}(x_n^{Tr}) \right\| \leq \left[\sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N^{Tr}}} \right] \quad (\text{B.7})$$

and

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{p'}}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^{p'}}(x_n^{Tr}) \right\| \leq \left[\sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log M + \log 4 - \log \epsilon}{2N^{Tr}}} \right] \quad (\text{B.8})$$

Now combining equations B.5, B.6, B.7, and B.8 with triangle inequality, we have with probability $1 - \epsilon$,

$$\begin{aligned} & \left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \right\| \\ & \leq \frac{u \|\mathbf{c}\|_1}{4M \min_{\mathbf{c}_k \neq 0} \mathbf{c}_k} + 4 \sqrt{\frac{\log M + \log 8 - \log \epsilon}{2N}} + 2 \sqrt{\frac{\log M + \log 8 - \log \epsilon}{2N^{Tr}}} \end{aligned}$$

Setting $M = \min\{N^{Tr}, N\}$, we have

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \right\| \leq O\left(\sqrt{\frac{\log N + \log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log N^{Tr} + \log 8 - \log \epsilon}{2N^{Tr}}}\right)$$

which completes the proof. \square

Lemma B.8. *Let*

$$p(\Omega_M) \triangleq \arg \max_{p \in \Omega_M} \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n)$$

$$s.t. \frac{1}{N} \eta^{[s^p]}(x_n, \mathbf{c}) \leq (1 - \delta)b - \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}} \right]$$

Then with probability $1 - \epsilon$,

$$\frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \geq \frac{1}{N} \sum_{n=1}^N r^{s^p(\Omega_M)} - \sqrt{\frac{\log 8 - \log \epsilon}{2N}} - \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}}$$

and

$$\frac{1}{N} \sum_{n=1}^N \eta^{[s^{\hat{p}(\Omega_M)}]}(x_n, \mathbf{c}) \leq (1 - \delta)b + 2\|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}} \right].$$

Proof. By Lemma B.3, with probability $1 - \epsilon$, we have

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^p}(x_n^{Tr}) \right\| \leq \sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}}.$$

and similarly, with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N \eta^{[s^p]}(x_n, \mathbf{c}) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^p]}(x_n^{Tr}, \mathbf{c}) \right\| \leq \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}} \right].$$

which is the same as

$$\left\| \frac{1}{N} \sum_{n=1}^N \eta^{[s^p]}(x_n, \mathbf{c}) - (1 - \delta)b - \left[\frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^p]}(x_n^{Tr}, \mathbf{c}) - (1 - \delta)b \right] \right\| \leq \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 4 - \log \epsilon}{2N}} + \sqrt{\frac{\log 4 - \log \epsilon}{2N^{Tr}}} \right].$$

Now applying union bound, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) - \frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} r^{s^p}(x_n^{Tr}) \right\| \leq \sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}}.$$

and

$$\left\| \frac{1}{N} \sum_{n=1}^N \eta^{[s^p]}(x_n, \mathbf{c}) - (1 - \delta)b - \left[\frac{1}{N^{Tr}} \sum_{n=1}^{N^{Tr}} \eta^{[s^p]}(x_n^{Tr}, \mathbf{c}) - (1 - \delta)b \right] \right\| \leq \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}} \right].$$

both hold. By Lemma B.4, we can conclude that

$$\frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \geq \frac{1}{N} \sum_{n=1}^N r^{s^p(\Omega_M)} - \sqrt{\frac{\log 8 - \log \epsilon}{2N}} - \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}}.$$

and

$$\frac{1}{N} \sum_{n=1}^N \eta^{[s^{\hat{p}(\Omega_M)}]}(x_n, \mathbf{c}) - (1 - \delta)b \leq 2\|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2N^{Tr}}} \right].$$

with probability $1 - \epsilon$, which completes the proof. \square

Lemma B.9. For $\delta = \Omega(\sqrt{\frac{\log 4 - \log \epsilon}{N}} + \sqrt{\frac{\log 4 - \log \epsilon}{NTr}})$, we have with probability $1 - \epsilon$,

$$\frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \geq -O\left(\sqrt{\frac{\log N - \log \epsilon}{2N}} + \sqrt{\frac{\log N - \log \epsilon}{2NTr}}\right).$$

Proof. Let $\tilde{p}(\Delta)$ be the optimal solution to the following problem

$$\max_{p \in \mathbb{R}} \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) \text{ s.t. } \frac{1}{N} \eta^{[s^p]}(x_n, \mathbf{c}) \leq b - \Delta$$

On one hand, p^* apparently is a feasible solution to the above problem with $\Delta = 0$, so we must have

$$\frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \leq \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(0)}}(x_n) \quad (\text{B.9})$$

Let $\Delta' = \delta + \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2NTr}} \right]$. Then $\tilde{p}(\Delta')$ corresponds to the following problem

$$\begin{aligned} & \max_{p \in \mathbb{R}} \frac{1}{N} \sum_{n=1}^N r^{s^p}(x_n) \\ & \text{s.t. } \frac{1}{N} \eta^{[s^p]}(x_n, \mathbf{c}) \leq (1 - \delta)b - \|\mathbf{c}\|_\infty \left[\sqrt{\frac{\log 8 - \log \epsilon}{2N}} + \sqrt{\frac{\log 8 - \log \epsilon}{2NTr}} \right] \end{aligned}$$

Then using the same argument in the proof of Lemma B.7, we have with probability $1 - \epsilon$,

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(\Delta')}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) \right\| \leq \sqrt{\frac{\log N + \log 8 + \log \epsilon}{2N}} + \sqrt{\frac{\log N + \log 8 - \log \epsilon}{2NTr}}. \quad (\text{B.10})$$

Furthermore, it is clear that $\tilde{p}(\Delta)$ is decreasingly-monotone with respect to Δ . In fact, removing the budget by Δ_1 , at most $\frac{\Delta_1}{\min_{\mathbf{c}_k > \mathbf{c}_j} c_k - c_j}$ data's APIs need to be changed, and thus incurs at most $\frac{\Delta_1}{\min_{\mathbf{c}_k > \mathbf{c}_j} c_k - c_j}$ accuracy decrease. That is to say, we must have

$$\left\| \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(\Delta')}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(0)}}(x_n) \right\| \leq \frac{\Delta_1}{\min_{\mathbf{c}_k > \mathbf{c}_j} c_k - c_j} \quad (\text{B.11})$$

Now combining equations B.9, B.10, B.11, we can obtain

$$\begin{aligned} & \frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \\ &= \frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(\Delta')}}(x_n) \\ &+ \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(\Delta')}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(0)}}(x_n) + \frac{1}{N} \sum_{n=1}^N r^{s^{\tilde{p}(0)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \\ &\geq -\sqrt{\frac{\log N + \log 8 + \log \epsilon}{2N}} - \sqrt{\frac{\log NTr + \log 8 - \log \epsilon}{2NTr}} - \frac{\Delta_1}{\min_{\mathbf{c}_k > \mathbf{c}_j} c_k - c_j} - 0 \end{aligned}$$

When $\delta = \Omega\left(\sqrt{\frac{\log 4 - \log \epsilon}{N}} + \sqrt{\frac{\log 4 - \log \epsilon}{NTr}}\right)$, we have

$$\begin{aligned} & \frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \\ &\geq -O\left(\sqrt{\frac{\log N - \log \epsilon}{2N}} + \sqrt{\frac{\log NTr - \log \epsilon}{2NTr}}\right) \end{aligned}$$

which completes the proof. \square

Now we are ready to prove the main theorem. We start by showing the bound on the reward. Suppose $\delta = \Theta\left(\sqrt{\frac{\log N - \log \epsilon}{N}} + \sqrt{\frac{\log N^{Tr} - \log \epsilon}{N^{Tr}}}\right)$. By union bound, with probability $1 - \epsilon$, Lemma B.7, Lemma B.8, and Lemma B.9 all hold, and we have

$$\begin{aligned}
 & \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \\
 &= \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) + \frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}(\Omega_M)}}(x_n) \\
 &+ \frac{1}{N} \sum_{n=1}^N r^{s^{p(\Omega_M)}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \\
 &\geq -O\left(\sqrt{\frac{\log N - \log \epsilon}{N}} + \sqrt{\frac{\log N^{Tr} - \log \epsilon}{N^{Tr}}}\right)
 \end{aligned}$$

Now note that by Theorem 4.2, we have with probability 1,

$$\frac{1}{N} \sum_{n=1}^N r^{s^{p^*}}(x_n) \geq \frac{1}{N} \sum_{n=1}^N r^{s^*}(x_n) - \frac{1}{N}$$

Combing the above two inequalities, we have

$$\frac{1}{N} \sum_{n=1}^N r^{s^{\hat{p}}}(x_n) - \frac{1}{N} \sum_{n=1}^N r^{s^*}(x_n) \geq -O\left(\sqrt{\frac{\log N - \log \epsilon}{N}} + \sqrt{\frac{\log N^{Tr} - \log \epsilon}{N^{Tr}}}\right)$$

Next we consider the feasibility requirement. By Lemma B.6, with probability $1 - \epsilon$, $s^{\hat{p}}$ is a feasible solution to Problem 4.1. That is to say, $s^{\hat{p}}$ with probability $1 - \epsilon$ is a feasible solution. Applying union bound completes the proof. \square

C. Experimental Details

We provide missing experimental details in this part.

Experimental setup All experiments were run on a machine with 8 Intel Xeon Platinum 2.5 GHz cores, 32 GB RAM, and 500GB disk with Ubuntu 16.04 LTS as the OS. Our code is implemented in Python 3.7. Each experiments, except the case study, were run for five times to mitigate the randomness introduced by training-testing splitting.

ML tasks and services Recall that We focus on three multi-label classification tasks, multi-label image classification (*MIC*), scene text recognition (*STR*), and named entity recognition (*NER*).

MIC is a computer vision task, where the goal is to assign a set of labels to a given image. For *MIC*, we use 3 different commercial ML cloud services, Google Vision (*Goo*), Microsoft Vision (*Mic*), and Everyapixel(*Eve*). We also use a single shot detector model (SSD) pretrained on OpenImageV4 (Kuznetsova et al., 2020), which is freely available from GitHub (SSD). All of those APIs produce labels from a large (and unknown) set, but the datasets we consider have bounded number of labels. For example, there are only 80 distinct labels in COCO dataset. Thus, we remove the predicted labels which are not in the full label set. For example, if Google API gives label $\{person, car, man\}$ for an image in COCO, but *man* is not in the full label set of COCO, then we will use $\{person, car\}$ as the label set produced from Google.

STR is a computer vision task, where the goal is to predict all texts in a natural scene image. In the context of multi-label classification, we view each predicted word as a label, and all possible words as the label set. For *STR*, the ML services used in the experiments are Google Vision (*Goo*), iFLYTEK API (*Ifi*), and Tencent API (*Ten*). We also use PP-OCR (*Pad*), an open source model from GitHub.

NER is a natural language processing task where the goal is to extract all possible entities from a given text. For example, for the sentence *ICML was held in Long Beach in 2019*, *ICML* should be extracted as an organization, and *Long Beach* should be identified as a location. In this paper, we consider three common types of entities, *person*, *location*, and *organization*. For any given text, each possible entity is viewed as a label, and the label set is the number of unique entities in the entire dataset. For *NER*, we use three common APIs: Amazon Comprehend (*Ama*), Google NLP (*GoN*), and IBM natural language API (*IBM*). a multi-task convolutional neural network model(*Spa*) from GitHub is also used.

Datasets. The experiments were conducted on 9 datasets. For *MIC*, we use three popular datasets including PASCAL (Everingham et al., 2015), MIR (Huiskes & Lew, 2008) and COCO (Lin et al., 2014). PASCAL is a standard object recognition dataset with 20 distinct labels, and COCO is another one with 80 unique labels. PASCAL’s label set contains 20 common objects: *person, bird, cat, cow, dog, horse, sheep, aeroplane, bicycle, boat, bus, car, motorbike, train, bottle, chair, dining table, potted plant, sofa, tv/monitor*. The 80 objects in COCO include: *person, bicycle, car, motorcycle, airplane, bus, train, truck, boat, traffic light, fire hydrant, stop sign, parking meter, bench, bird, cat, dog, horse, sheep, cow, elephant, bear, zebra, giraffe, backpack, umbrella, handbag, tie, suitcase, frisbee, skis, snowboard, sports ball, kite, baseball bat, baseball glove, skateboard, surfboard, tennis racket, bottle, wine glass, cup, fork, knife, spoon, bowl, banana, apple, sandwich, orange, broccoli, carrot, hot dog, pizza, donut, cake, chair, couch, potted plant, bed, dining table, toilet, tv, laptop, mouse, remote, keyboard, cell phone, microwave, oven, toaster, sink, refrigerator, book, clock, vase, scissors, teddy bear, hair drier, toothbrush*. For those two datasets, we use their original associated labels as the label set. MIR is a dataset designed for image retrieval. There are originally 25 labels: *animals, baby, bird, car, clouds, dog, female, flower, food, indoor, lake, male, night, people, plant_life, portrait, river, sea, sky, structures, sunset, transport, tree, water*. We remove the label *night* since it is not in the label set of any of the APIs or the GitHub model. On average, there are 1.44 labels per image for PASCAL, 3.71 labels per image for MIR, and 2.91 labels per image for COCO. The dataset statistic is summarized in Table 2. Most of the datasets are open and under Creative Commons license (e.g., the dataset COCO (Lin et al., 2014)). The details can be found in their corresponding paper and repository. As those datasets are actually open, they do not require an in-person consent from the authors/developers. The datasets themselves may contain personal information (e.g., there are personal images in COCO). Though, they have been render anonymous. For the purpose of deciding which API to call, we also do not use personally identifiable information.

For *STR*, we use three large scale Chinese text recognition datasets, MTWI (He et al., 2018), ReCTS (Zhang et al., 2019) and LSVT (Sun et al., 2019). The label set contains all possible Chinese characters as well as digits (0-9). MTWI contains images from the internet mainly targeting at advertisements. Thus, most of its images have dense texts. ReCTS includes photos taken on sign boards and thus has relatively fewer words. The images from LSVT are typically street view images

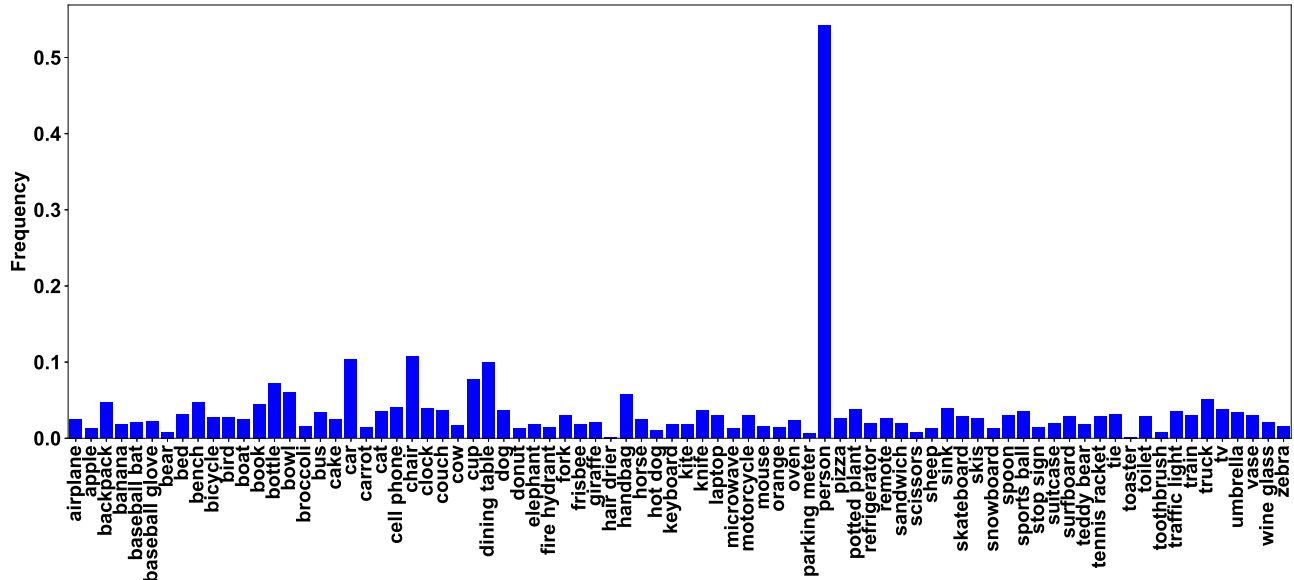


Figure 6. Label Distribution on COCO.

and hence have medium number of words. All images in MTWI and ReCTS are fully annotated and used in our experiments. LSVT contains both fully and partially annotated images, and we only use the subset with full annotations.

The other datasets, CONLL (Sang & Meulder, 2003), ZHNER (ZHN) and GMB (Bos, 2013), are used for *NER* task. CONLL contains English sentences from newspapers, and texts from GMB are also English and from a wider range of sources. On the other hand, ZHNER is a Chinese text dataset. We consider four common types of entities: organization, person, and location. In this paper, we focus on three common types of entities that all datasets contain: *persons*, *locations*, and *organizations*. Each sentence from those datasets is extracted as a data point, and the associated label set is simply all entities in this sentence. An entity is considered correctly extracted if and only if it is labeled as an entity and its entity type is correct.

GitHub model cost We evaluate the inference time of all GitHub models on an Amazon EC2 p2.x instance, which is \$0.90 per hour. For multi-label image classification, the GitHub model (SSD) takes 6s to classify each image, resulting in an equivalent cost of \$0.0015 per image. For the named entity task, the GitHub model (Spa) can extract the entities from a sentence in 0.015s, leading to \$ 0.00000375 per sentence. The GitHub model (Pad) with the mobile version 3.0 text detector and recognizer requires 1.5 on average to extract text from an image, causing a cost of \$ 0.000375 per image. Compared to the commercial APIs, this cost is much cheaper.

Case study on COCO Now we provide more details about the case study on the multi-label image classification dataset, COCO. There are in total 123,287 images containing labels from 80 different categories in COCO. Figure 6 gives the label distribution. First note that the label distribution is quite skewed. overall, the label person is the most frequent: more than 50% of the images contain the person label. Among others, car, chair, and dining tables are also quite common labels in this dataset with more than 10% occurrence. On the other hand, there are also quite some rare labels. For example, half driver and toaster appear in less than 1% of the images. Such imbalance between different labels imposes a high data and computational complexity to directly apply previous approach that learns a decision rule per label, and thus verifies the necessity of the proposed framework, FrugalMCT.

To further understand when and why FrugalMCT gives a better performance than single API, we present the precision and recall per class for each API, majority vote, and FrugalMCT in Figure 7 and Figure 8. We first note that there is no API universally better than other APIs for each label. For example, GitHub and Microsoft APIs can hardly correctly predict the label “toaster”, but Everyapixel and Google APIs have a relatively high accuracy on label “toaster”. On the other hand, Everyapixel has a low accuracy on label “kite” and “knite”, while Microsoft, Google, GitHub APIs can usually predict those

Table 6. Comparison of ensemble methods as well as cost-aware approaches. For FrugalMCT and FrugalML, we pick their corresponding strategies that minimize the cost while ensures that the accuracy reaches the highest possible.

| | best single API | | FrugalML | | FrugalMCT | | majority vote | | weighted maj vote | |
|--------|-----------------|------|----------|------|-----------|------|---------------|--------|-------------------|--------|
| | acc | cost | acc | cost | acc | cost | acc | cost | acc | cost |
| PASCAL | 74.8 | 10 | 76.9 | 11 | 78.5 | 8 | 77.8 | 31.01 | 77.8 | 31.01 |
| MIR | 41.2 | 10 | 43.8 | 8 | 49.2 | 14 | 41.4 | 31.01 | 48.7 | 31.01 |
| COCO | 47.5 | 10 | 49.3 | 8 | 54 | 12 | 50.1 | 31.01 | 52.8 | 31.01 |
| MTWI | 67.9 | 210 | 68.1 | 213 | 71.1 | 208 | 75.4 | 275.01 | 75.4 | 275.01 |
| ReCTS | 61.3 | 210 | 63.4 | 213 | 64.7 | 208 | 70.2 | 275.01 | 70.2 | 275.01 |
| LSVT | 53.8 | 210 | 56.2 | 213 | 57.2 | 208 | 62.8 | 275.01 | 62.8 | 275.01 |
| CONLL | 52.6 | 3 | 55.7 | 32 | 56.8 | 36.8 | 58.5 | 43.01 | 58.5 | 43.01 |
| ZHNER | 61.3 | 30 | 67.4 | 31.2 | 71.8 | 36.8 | 66 | 43.01 | 66 | 43.01 |
| GMB | 50.1 | 30 | 52.6 | 30.1 | 53.1 | 20.5 | 51.3 | 43.01 | 51.5 | 43.01 |

labels with higher accuracy. This implies that combining different APIs may produce an accuracy better than any single one of them. There are also some easy labels on which all APIs give a high accuracy. For example, on the label “zebra”, all APIs give a 90% precision and recall. This actually suggests that it is not always necessary to use all API. For example, if GitHub predicts an image has the label “zebra”, and we know there is no other labels in this image, then probably there is no need to call any other APIs.

Another interesting observation is that FrugalMCT improves the precision and recall for almost every label compared to any single API. This is primarily because FrugalMCT appropriately utilizes the predicted label information from GitHub model to infer which API is better on certain input, and combine its performance with the base API aptly. Yet, the precision and recall difference can be quite different for different APIs. For example, as shown in Figure 8(c), the recall for “airplane” is much higher than its precision, but banana’s precision is much higher than its recall. For applications that have specific precision and recall requirements, we may adopt different accuracy metrics in FrugalMCT. Another interesting observation is that the precision and recall for some labels is extremely. For example, “hair drier” cannot be predicted by FrugalMCT, which is due to that no API actually predicts this label correctly. How to extend FrugalMCT to recognize unseen labels remains an open question.

Table 7. Accuracy predictor performance. RMSE and PCC stand for root mean square error and Pearsons correlation coefficient.

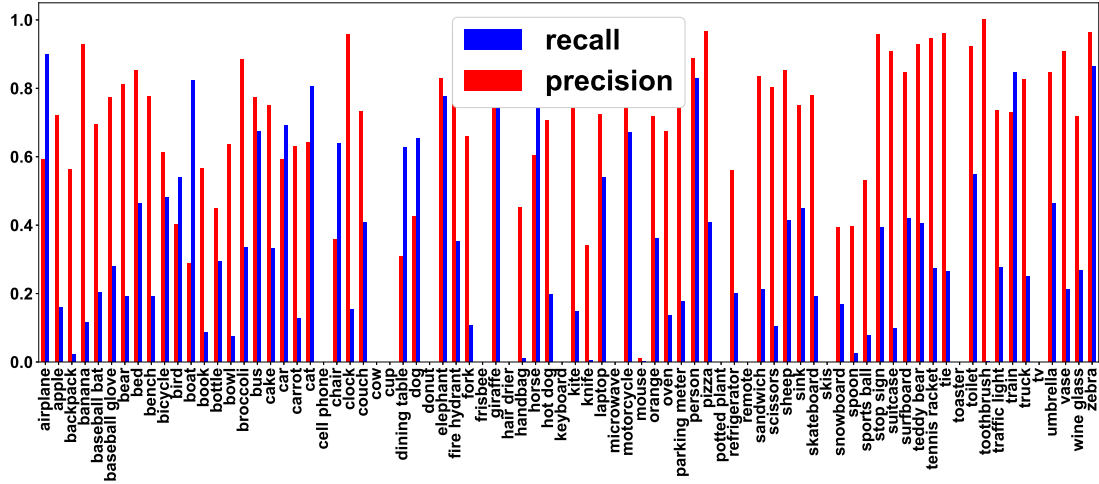
| Data | RMSE | | PCC | |
|--------|-----------|------|-----------|--------|
| | FrugalMCT | DAP | FrugalMCT | DAP |
| PASCAL | 0.28 | 0.35 | 0.55 | 0.012 |
| MIR | 0.22 | 0.31 | 0.55 | -0.013 |
| COCO | 0.24 | 0.31 | 0.63 | 0.001 |
| MTWI | 0.17 | 0.21 | 0.57 | 0.004 |
| ReCTS | 0.22 | 0.27 | 0.57 | 0.001 |
| LSVT | 0.19 | 0.24 | 0.61 | -0.003 |
| CONLL | 0.29 | 0.41 | 0.72 | -0.003 |
| ZHNER | 0.31 | 0.36 | 0.48 | -0.005 |
| GMB | 0.28 | 0.40 | 0.69 | -0.006 |

Ensemble method comparison For comparison, we compare FrugalMCT against FrugalML as well as two ensemble methods, majority vote and weighted majority vote. In majority vote, for each label, we accept it if at least half of the APIs predict it. In weighted majority vote, we assign each API’s accuracy as its weight. Next, for each label, we compute a label score, which is equal to the sum of each API’s confidence score on this label weighted by its corresponding weight. If an API does not predict a label, then its confidence score is viewed as 0. Finally, we only accept the label if its label score is larger than a threshold. We pick a threshold that maximizes the overall accuracy by grid search.

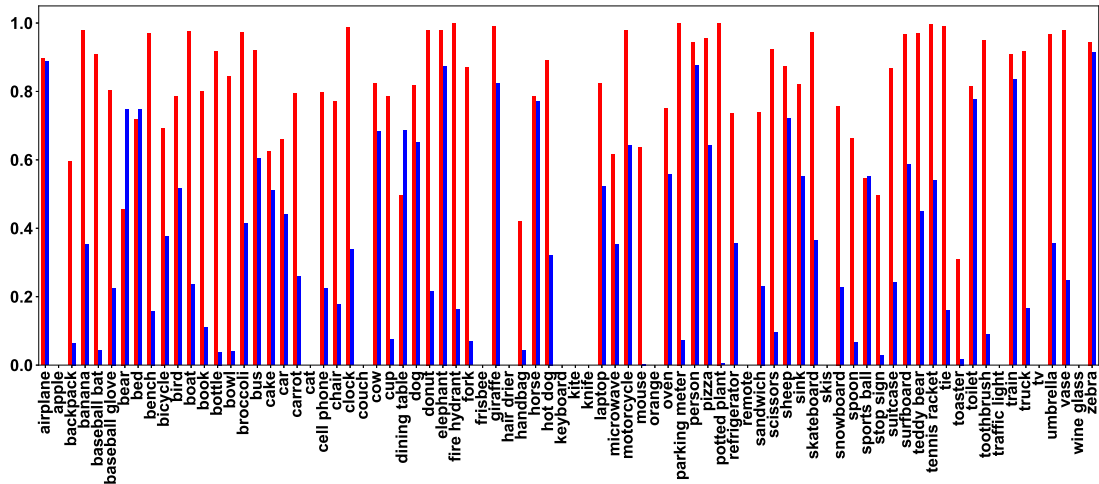
The results are summarized in Table 6. Overall, we observe that FrugalMCT and ensemble methods have similar performance across different tasks and datasets, but with a much lower cost. In fact, for datasets including COCO and ZHNER, FrugalMCT can achieve an accuracy even higher than ensemble methods.

Accuracy predictor performance Note that FrugalMCT’s performance highly depends on its accuracy predictors’ performance. To obtain a quantitative sense of the accuracy predictors, we evaluate the accuracy predictors’ performance in Table 7. RMSE measures the standard deviation of the difference between accuracy predictor’s output and the corresponding true accuracy. PCC stands for Pearson correlation coefficient, which roughly measures the linear correlation between the true accuracy and the predicted value from the accuracy predictors. Overall, FrugalMCT’s random forest predictors enjoy a much smaller RMSE and higher PCC than DAP (the dummy accuracy predictors), which matches the fact that FrugalMCT gives a higher end to end performance than using the DAP.

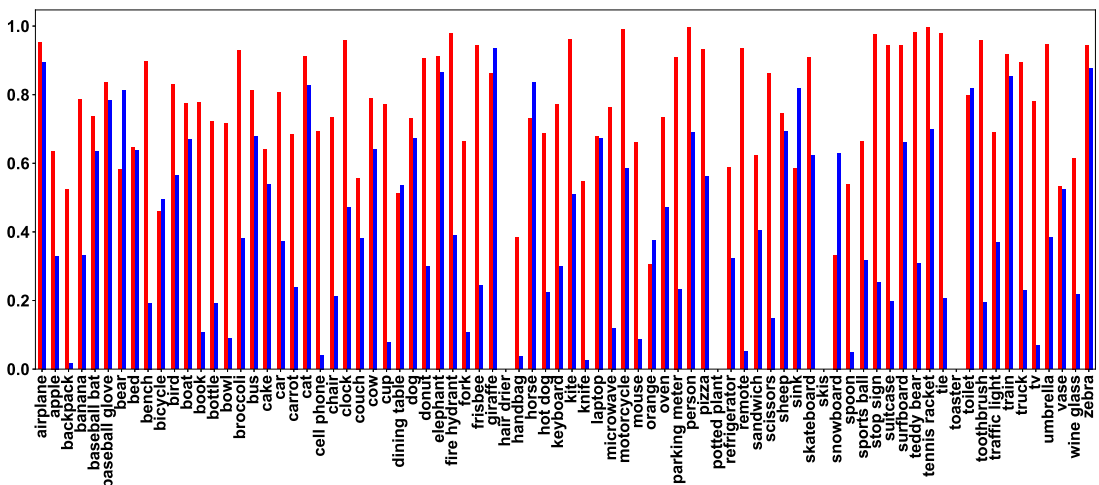
Efficient Online ML API Selection for Multi-Label Classification Tasks



(a) GitHub



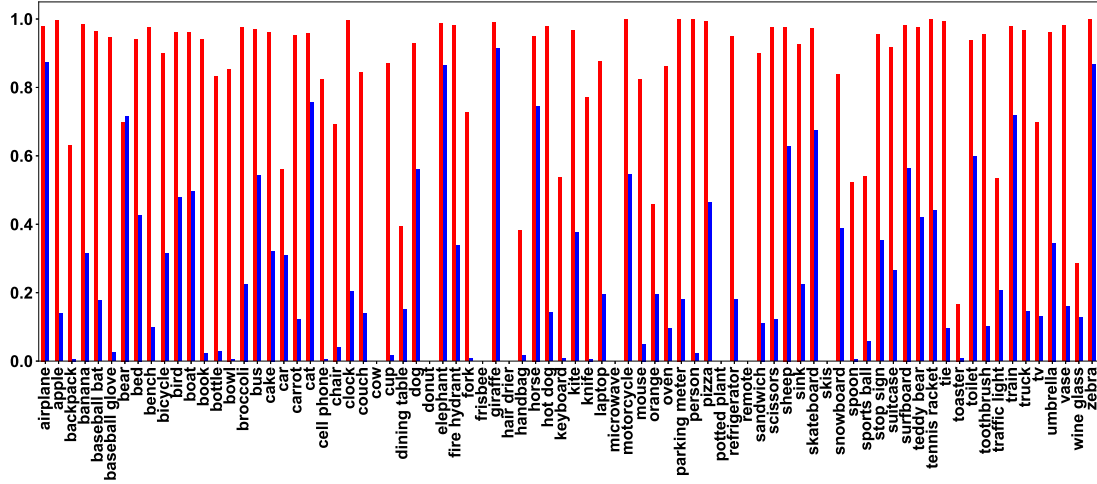
(b) Everypixel



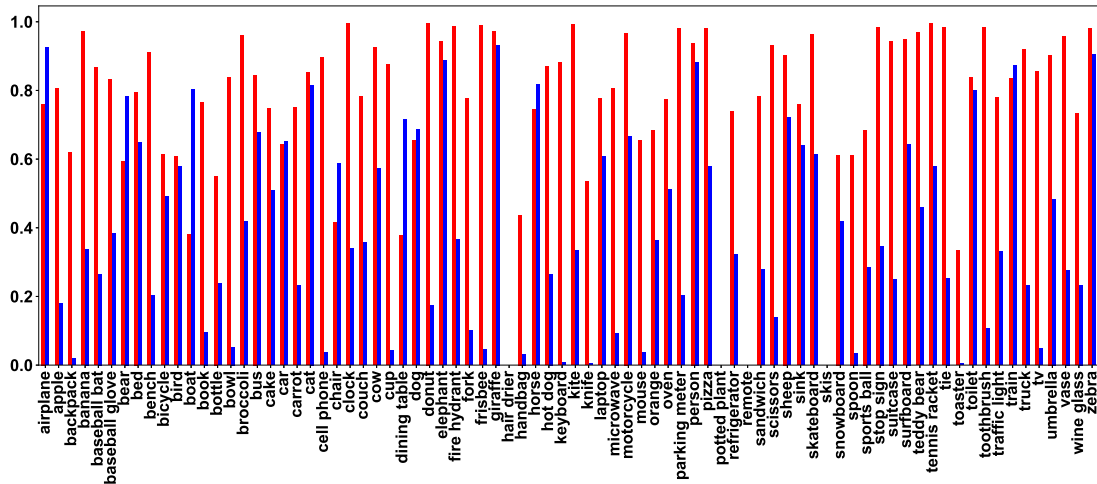
(c) Microsoft

Figure 7. The per class precision and recall of different APIs .

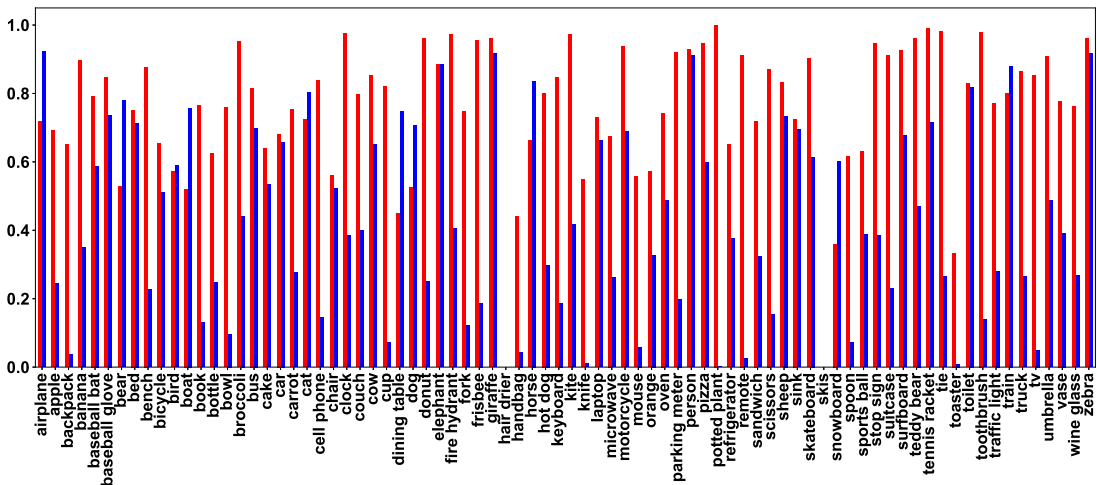
Efficient Online ML API Selection for Multi-Label Classification Tasks



(a) Google



(b) Majority Vote



(c) FrugalMCT

Figure 8. The per class precision and recall of different APIs .