

## 489 A Implementation Details

### 490 A.1 Deep Active Learning Decomposition

491 For any uncertainty sampling algorithm, picking the top- $B$  most uncertain examples can be easily  
492 decomposed into an iterative procedure that picks the next most uncertain example. Next, for diversity  
493 based deep active learning algorithms, one usually rely on a greedy iterative procedure to collect  
494 a batch, e.g. K-means++ for BADGE [Ash et al., 2019] and greedy K-centers for Coreset [Sener  
495 and Savarese, 2017]. Lastly, deep active learning algorithms such as Cluster-Margin [Citovsky  
496 et al., 2021] and GALAXY [Zhang et al., 2022] have already proposed their algorithms as iterative  
497 procedures that select unlabeled examples sequentially.

### 498 A.2 Implementation of Modified Submodular

499 Instead of requiring access to a balanced holdout set [Kothawade et al., 2021], we construct the  
500 balanced set using training examples. We use the Submodular Mutual Information function FLQMI  
501 as suggested by Table 1 of [Kothawade et al., 2021]. The proposed greedy submodular optimization is  
502 itself an iterative procedure that selects one example at a time. While SIMILAR usually performs  
503 well, our modification that discards the holdout set is unfortunately ineffective in our experiments.  
504 This is primarily due to the lack of the holdout examples, which may often happen in practical  
505 scenarios.

### 506 A.3 Stanford Car Multi-label Dataset

507 We transform the original labels into 10 binary classes of

- 508 1. If the brand is "Audi".
- 509 2. If the brand is "BMW".
- 510 3. If the brand is "Chevrolet".
- 511 4. If the brand is "Dodge".
- 512 5. If the brand is "Ford".
- 513 6. If the car type is "Convertible".
- 514 7. If the car type is "Coupe".
- 515 8. If the car type is "SUV".
- 516 9. If the car type is "Van".
- 517 10. If the car is made in or before 2009.

### 518 A.4 Negative Weighting for Common Classes

519 For multi-label classifications, for some classes, there could be more positive associations (label of 1s)  
520 than negative associations (label of 0s). Therefore, in those classes, the rarer labels are negative. In  
521 class diverse reward  $\langle v_{div}^t, y \rangle$  in Section 4.1, we implement an additional weighting of  $\mathbb{1}_{rare}^t * v_{div}^t$ ,  
522 where  $*$  denotes an elementwise multiplication. Here, each element  $\mathbb{1}_{rare,i}^t \in \{1, -1\}$  takes value  
523  $-1$  when  $COUNT^t(i)$  is larger than half the size of labeled set. This negative weighting can be seen  
524 as upsampling negative class associations when positive associations are the majority.

### 525 A.5 Model Training

526 All of our experiments are conducted using the ResNet-18 architecture [He et al., 2016] pretrained on  
527 ImageNet. We use the Adam optimizer [Kingma and Ba, 2014] with learning rate of  $1e-4$  and weight  
528 decay of  $5e-5$ .

### 529 A.6 Baseline Algorithms

530 In the original GALAXY work by [Zhang et al., 2022], their algorithm construct  $K$  one-vs-rest  
531 linear graphs, one for each class. GALAXY requires finding the shortest shortest path among all

532  $K$  graphs, an operation whose computation scales linearly in  $K$ . When  $K$  is large, this becomes  
 533 computationally prohibitive to run. Therefore, we instead include  $K$  separate GALAXY algorithms,  
 534 each only bisecting on one of the one-vs-rest graphs. This is equivalent with running  $K$  GALAXY  
 535 algorithms, one for each binary classification task between class  $i \in [K]$  and the rest. As a baseline,  
 536 we interleave these algorithms uniformly at random.

537 For Uncertainty sampling in multi-label settings, we simply have  $K$  individual uncertainty sampling  
 538 algorithms, where the  $i$ -th algorithm samples the most uncertain example based only on the binary  
 539 classification task of class  $i$ .

## 540 B Proof of Theorem 5.2

541 Our proof follows a similar procedure from regret analysis for Thompson Sampling of the stochastic  
 542 multi-armed bandit problem [Lattimore and Szepesvári, 2020]. Let  $\alpha^t := \{\alpha^{t,j}\}_{j=1}^B$  and  $y^t :=$   
 543  $\{y^{t,j}\}_{j=1}^B$  denote the actions and observations from the  $i$ -th round. We define the history up to  $t$  as  
 544  $H_t = \{\alpha^1, y^1, \alpha^2, y^2, \dots, \alpha^{t-1}, y^{t-1}\}$ . Moreover, for each  $i \in [M]$ , we define  $H_{t,i} = \{y^{t',j} \in H_t :$   
 545  $\alpha^{t',j} = i\}$  as the history of all observations made by choosing the  $i$ -th arm (algorithm).

546 Now we analyze reward estimates at each round  $t$ . When given history  $H_t$  and arm  $i \in [M]$ , each  
 547 observation  $y \in H_{t,i}$  is an unbiased estimate of  $\theta^i$  as  $y \sim \mathbb{P}_{\theta^i}$ . Therefore, for any fixed  $v^t$ ,  $\langle v^t, y \rangle$  is  
 548 an unbiased estimate of the expected reward  $\langle v^t, \theta^i \rangle$ , which we denote by  $\mu^{t,i}$ .

549 For each arm  $i$ , we can then obtain empirical reward estimate  $\bar{\mu}^{t,i}$  of the true expected reward  $\mu^{t,i}$   
 550 by  $\bar{\mu}^{t,i} := \frac{1}{|H_{t,i}|} \sum_{y \in H_{t,i}} \langle v^t, y \rangle$  where  $\bar{\mu}^{t,i} = 0$  if  $|H_{t,i}| = 0$ . Since expected rewards and reward  
 551 estimates are bounded by  $[-1, 1]$ , by standard sub-Gaussian tail bounds, we can then construct  
 552 confidence interval,

$$\mathbb{P}(\forall i \in [M], t \in [T], |\bar{\mu}^{t,i} - \mu^{t,i}| \leq d^{t,i}) \geq 1 - \frac{1}{T}$$

553 where  $d^{t,i} := \sqrt{\frac{8 \log(MT^2)}{1V|H_{t,i}|}}$ . Additionally, we define upper confidence bound as  $U^{t,i} =$   
 554  $\text{clip}_{[-1,1]}(\bar{\mu}^{t,i} + d^{t,i})$ .

555 At each iteration  $t$ , we have the posterior distribution  $\mathbb{P}(\Theta = \cdot | H_t)$  of the ground truth  $\Theta =$   
 556  $\{\theta^i\}_{i=1}^M$ .  $\hat{\Theta} = \{\hat{\theta}^i\}_{i=1}^M$  is sampled from this posterior. Consider  $i_*^t = \arg \max_{i \in [M]} \langle v^t, \theta^i \rangle$  and  
 557  $\alpha^{t,j} = \arg \max_{i \in [M]} \langle v^t, \hat{\theta}^i \rangle$ . The distribution of  $i_*^t$  is determined by the posterior  $\mathbb{P}(\Theta = \cdot | H_t)$ . The  
 558 distribution of  $\alpha^{t,j}$  is determined by the distribution of  $\hat{\Theta}$ , which is also  $\mathbb{P}(\Theta = \cdot | H_t)$ . Therefore,  $i_*^t$   
 559 and  $\alpha^{t,j}$  are identically distributed. Furthermore, since the upper confidence bounds are deterministic  
 560 functions of  $i$  when given  $H_t$ , we then have  $\mathbb{E}[U^{t,\alpha^{t,j}} | H_t] = \mathbb{E}[U^{t,i_*^t} | H_t]$ .

561 As a result, we upper bound the Bayesian regret by

$$\begin{aligned} BR(\text{TAILOR}) &= \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B \mu^{t,i_*^t} - \mu^{t,\alpha^{t,j}} \right] \\ &= \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B (\mu^{t,i_*^t} - U^{t,i_*^t}) + (U^{t,\alpha^{t,j}} - \mu^{t,\alpha^{t,j}}) \right]. \end{aligned}$$

562 Now, note that since  $\bar{\mu}^{t,i} \in [-1, 1]$  we have  $\text{clip}_{[-1,1]}(\bar{\mu}^{t,i} + d^{t,i}) = \text{clip}_{[-\infty,1]}(\bar{\mu}^{t,i} + d^{t,i})$ ,  
 563 where only the upper clip takes effect. Based on the sub-Gaussian confidence intervals  
 564  $\mathbb{P}(\forall i \in [M], t \in [T], |\bar{\mu}^{t,i} - \mu^{t,i}| \leq d^{t,i}) \geq 1 - \frac{1}{T}$ , we can derive the following two confidence

565 bounds:

$$\begin{aligned}
\mathbb{P}(\forall i \in [M], t \in [T], \mu^{t,i} > \overline{U}^{t,i}) &= \mathbb{P}(\forall i \in [M], t \in [T], \mu^{t,i} > \text{clip}_{[-1,1]}(\bar{\mu}^{t,i} + d^{t,i})) \\
&= \mathbb{P}(\forall i \in [M], t \in [T], \mu^{t,i} > \bar{\mu}^{t,i} + d^{t,i}), \text{ since } \mu^{t,i} \leq 1 \\
&= \mathbb{P}(\forall i \in [M], t \in [T], \mu^{t,i} - \bar{\mu}^{t,i} > d^{t,i}) \leq \frac{1}{2T} \\
\mathbb{P}(\forall i \in [M], t \in [T], U^{t,i} - \mu^{t,i} > 2d^{t,i}) &= \mathbb{P}(\forall i \in [M], t \in [T], \text{clip}_{[-1,1]}(\bar{\mu}^{t,i} + d^{t,i}) - \mu^{t,i} > 2d^{t,i}) \\
&\leq \mathbb{P}(\forall i \in [M], t \in [T], \bar{\mu}^{t,i} + d^{t,i} - \mu^{t,i} > 2d^{t,i}) \\
&= \mathbb{P}(\forall i \in [M], t \in [T], \bar{\mu}^{t,i} - \mu^{t,i} > d^{t,i}) \leq \frac{1}{2T}.
\end{aligned}$$

566 Now with the decomposition,

$$\begin{aligned}
BR(\text{TAILOR}) &= \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B \mu^{t,i_*^t} - \mu^{t,\alpha^{t,j}} \right] \\
&= \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B \mu^{t,i_*^t} - U^{t,i_*^t} \right] + \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B U^{t,\alpha^{t,j}} - \mu^{t,\alpha^{t,j}} \right]
\end{aligned}$$

567 we can bound the two expectations individually.

568 First, to bound  $\mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B \mu^{t,i_*^t} - U^{t,i_*^t} \right]$ , we note that  $\mu^{t,i_*^t} - U^{t,i_*^t}$  is negative with high  
569 probability. Also, the maximum value this can take is bounded by 2 as  $\mu^{t,i}, U^{t,i} \in [-1, 1]$ . Therefore,  
570 we have

$$\mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B \mu^{t,i_*^t} - U^{t,i_*^t} \right] \leq \left( \sum_{t=1}^T \sum_{j=1}^B 0 \cdot \mathbb{P}(\mu^{t,i_*^t} \leq U^{t,i_*^t}) + 2 \cdot \mathbb{P}(\mu^{t,i_*^t} > U^{t,i_*^t}) \right) \leq 2TB \cdot \frac{1}{2T} = B.$$

571 Next, to bound  $\mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B U^{t,\alpha^{t,j}} - \mu^{t,\alpha^{t,j}} \right]$  we decompose it similar to the above:

$$\begin{aligned}
\mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^B U^{t,\alpha^{t,j}} - \mu^{t,\alpha^{t,j}} \right] &\leq \left( \sum_{t=1}^T \sum_{j=1}^B 2\mathbb{P}(U^{t,\alpha^{t,j}} - \mu^{t,\alpha^{t,j}} > 2d^{t,i}) \right) + \left( \sum_{t=1}^T \sum_{j=1}^B 2d^{t,i} \right) \\
&\leq B + \left( \sum_{t=1}^T \sum_{j=1}^B \sqrt{\frac{32 \log(MT^2)}{1 \vee |H_{t,\alpha^{t,j}}|}} \right)
\end{aligned}$$

572 where recall that  $|H_{t,i}|$  is the number of samples collected using algorithm  $i$  in rounds  $\leq t$ .

573 To bound the summation, we utilize the fact that  $\frac{1}{1 \vee |H_{t,i}|} \leq \frac{B}{k}$  for each  $k \in [|H_{t,i}|, |H_{t+1,i}|]$ , since  
574  $|H_{t+1,i}| - |H_{t,i}| \leq B$ . As a result, we get

$$\begin{aligned}
&\sum_{t=1}^T \sum_{j=1}^B \sqrt{\frac{32 \log(MT^2)}{1 \vee |H_{t,\alpha^{t,j}}|}} \\
&\leq \sum_{t=1}^T \sum_{i=1}^M \sum_{k=1}^{|H_{T,i}|} \sqrt{\frac{32 \log(MT^2) \cdot B}{k}} \\
&\leq O(\sqrt{B(\log T + \log M)} \sum_{i=1}^M \sqrt{|H_{T,i}|}) \\
&\leq O(\sqrt{B(\log T + \log M)}) \cdot O(\sqrt{BMT}) = O(B\sqrt{MT(\log T + \log M)})
\end{aligned}$$

575 where last two inequalities follow from simple algebra and the fact that  $\sum_{i=1}^M |H_{T,i}| = TB$ .

576 Finally, to combine all of the bounds above, we get  $BR(\text{TAILOR}) \leq B + B +$   
577  $O(B\sqrt{MT(\log T + \log M)}) = O(B\sqrt{MT(\log T + \log M)})$ .

578 **C Time Complexity**

579 Let  $N_{train}$  denote the total neural network training. The time complexity of collecting each batch  
580 for each active learning algorithm  $\mathcal{A}_i$  can be separated into  $P_i$  and  $Q_i$ , which are the computa-  
581 tion complexity for preprocessing and selection of each example respectively. As examples of  
582 preprocessing, BADGE [Ash et al., 2019] computes gradient embeddings, SIMILAR [Kothawade  
583 et al., 2021] further also compute similarity kernels, GALAXY [Zhang et al., 2022] constructs  
584 linear graphs, etc. The selection complexities are the complexities of each iteration of K-means++  
585 in BADGE, greedy submodular optimization in SIMILAR, and shortest shortest path computa-  
586 tion in GALAXY. Therefore, for any individual algorithm  $\mathcal{A}_i$ , the computation complexity is then  
587  $O(N_{train} + TP_i + TBQ_i)$  where  $T$  is the total number of rounds and  $B$  is the batch size. When  
588 running TAILOR, as we do not know which algorithms are selected, we provide a worst case upper  
589 bound of  $O(N_{train} + T \cdot (\sum_{i=1}^M P_i) + TB \cdot \max_{i \in [M]} Q_i)$ , where the preprocessing is done for every  
590 candidate algorithm. In practice, some of the preprocessing operations such as gradient embedding  
591 computation could be shared among multiple algorithms, thus only need to be computed once. While  
592 the computation of rewards and Thompson sampling updates incur some extra complexity, they  
593 are usually dominated in practice by the complexity of neural network training and running each  
594 candidate algorithm.



609 **E Full Results**

610 All of the results below are averaged from four individual trials except for Imagenet, which is the  
 611 result of a single trial.

612 **E.1 Multi-label Classification**

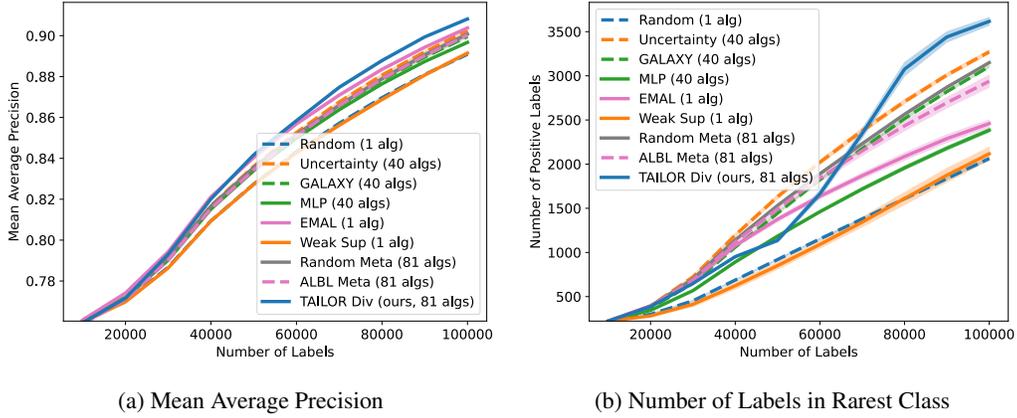


Figure 7: CeleBA

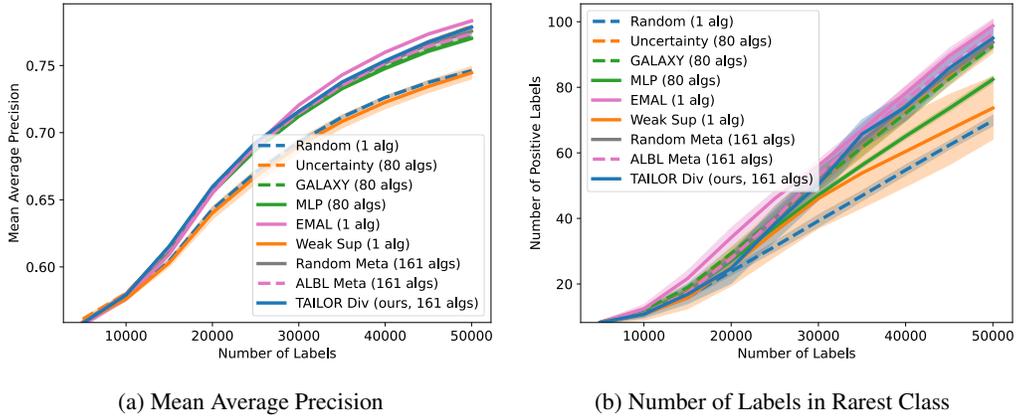


Figure 8: COCO

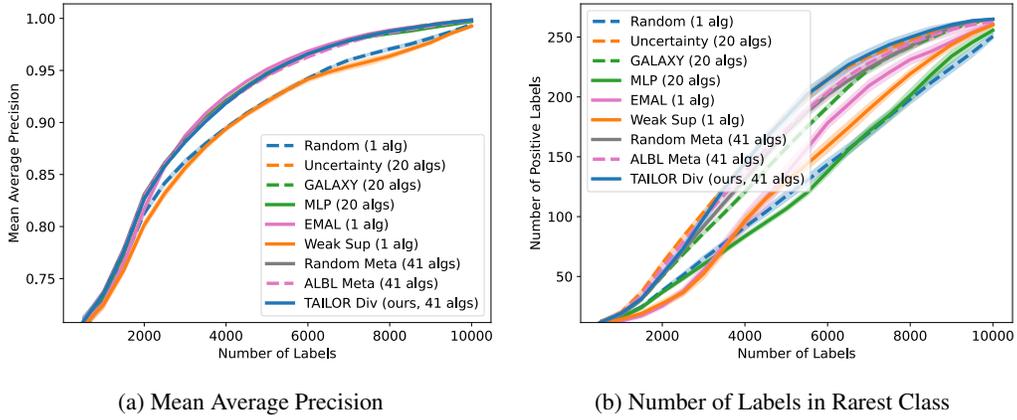


Figure 9: VOC

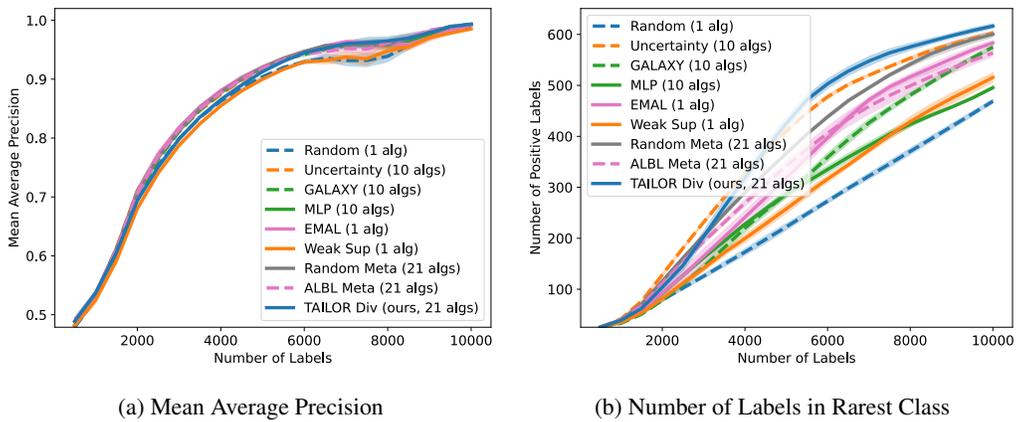


Figure 10: Stanford Car

613 E.2 Multi-class Classification

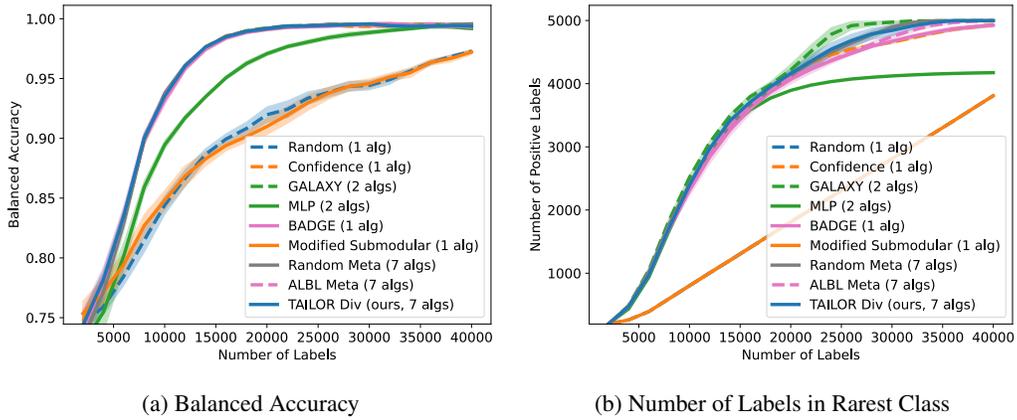


Figure 11: CIFAR-10, 2 classes

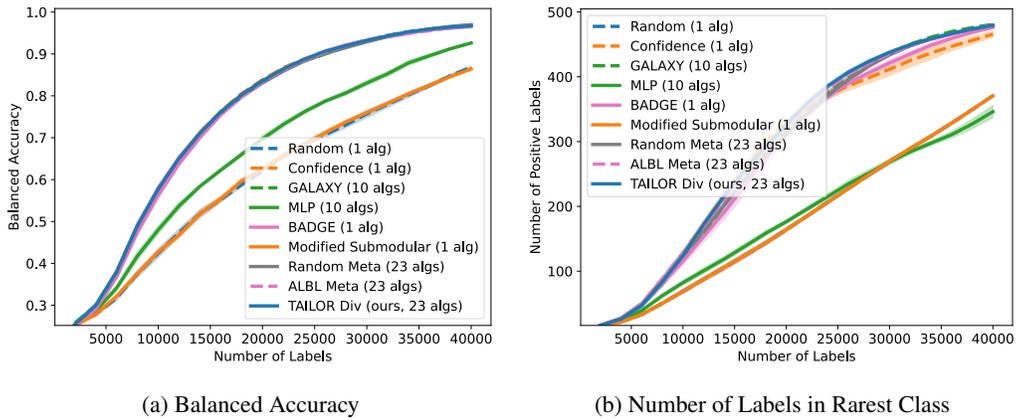


Figure 12: CIFAR-100, 10 classes

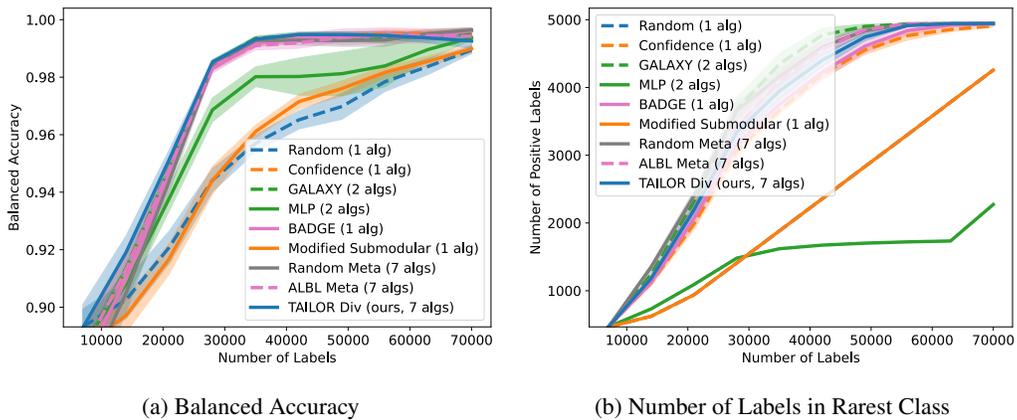
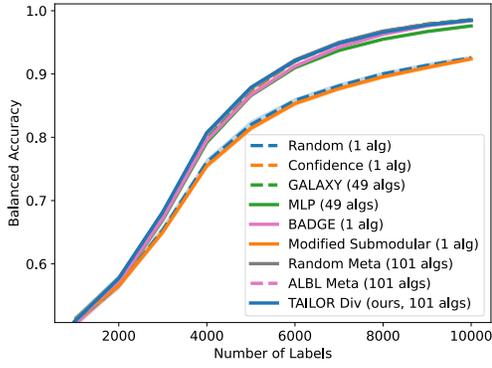
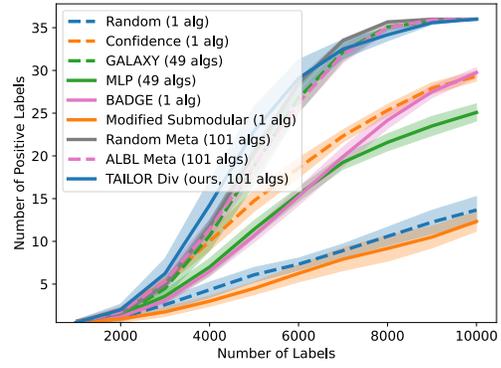


Figure 13: SVHN, 2 classes

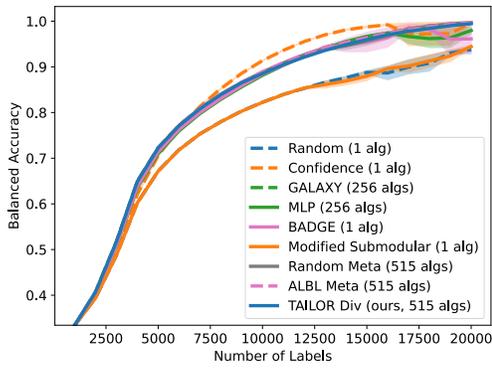


(a) Balanced Accuracy

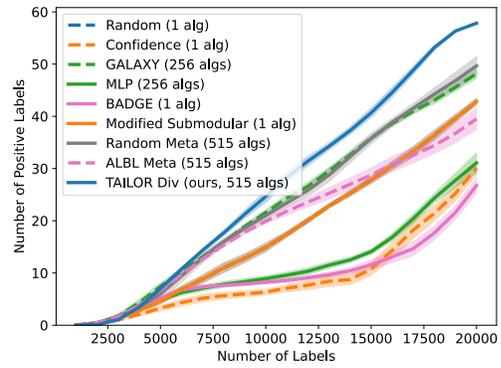


(b) Number of Labels in Rarest Class

Figure 14: Kuzushiji-49



(a) Balanced Accuracy



(b) Number of Labels in Rarest Class

Figure 15: Caltech256

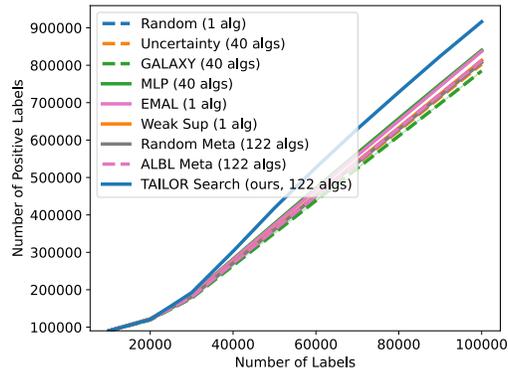


Figure 16: CelebA, Total Number of Positive Labels

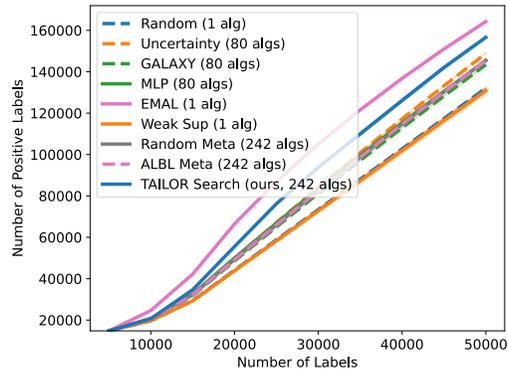


Figure 17: COCO, Total Number of Positive Labels

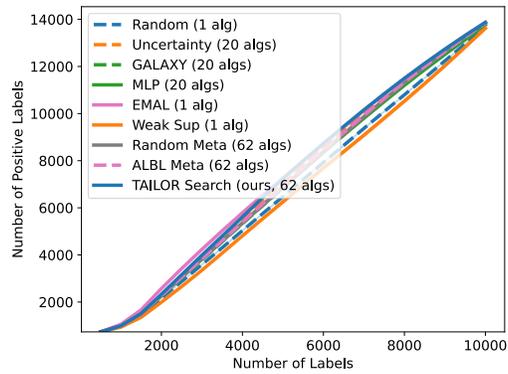


Figure 18: VOC, Total Number of Positive Labels

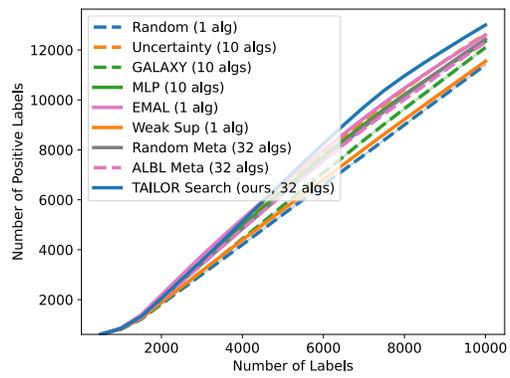


Figure 19: Stanford Car, Total Number of Positive Labels

615 **F What Algorithms Does TAILOR Choose?**

616 In the following two figures, we can see TAILOR chooses a non-uniform set of algorithms to focus on  
 617 for each dataset. On CelebA, TAILOR out-perform the best baseline, EMAL sampling, by a significant  
 618 margin. As we can see, TAILOR rely on selecting a *combination* of other candidate algorithms instead  
 619 of only selecting EMAL.

620 On the other hand, for the Stanford car dataset, we see TAILOR 's selection mostly align with the  
 baselines that perform well especially in the later phase.

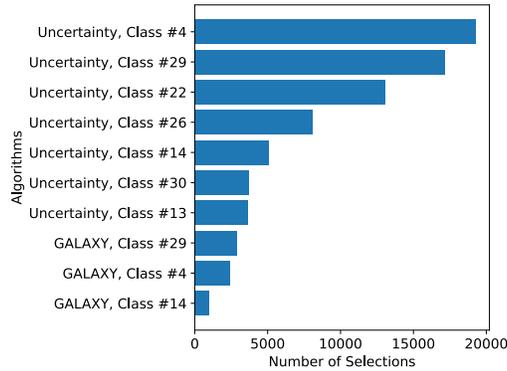


Figure 20: TAILOR Top-10 Most Selected Candidate Algorithms on CelebA Dataset

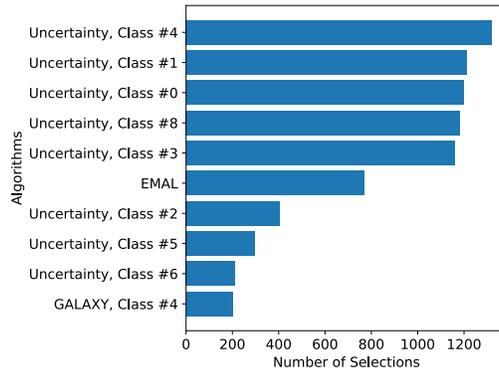


Figure 21: TAILOR Top-10 Most Selected Candidate Algorithms on Stanford Car Dataset

621

622 In the following figures, we plot the number of times the most frequent candidate algorithm is chosen.  
 623 As can be shown, TAILOR chooses candidate algorithm much more aggressively than other meta  
 624 algorithms in eight out of the ten settings.

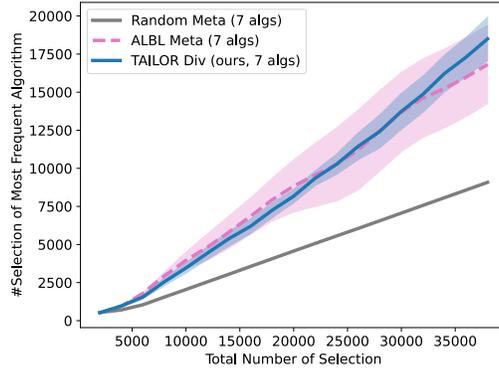


Figure 22: CIFAR-10, 2 Classes, Number of Pulls of The Most Frequent Selection

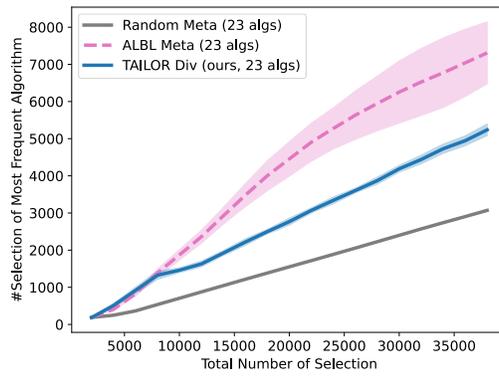


Figure 23: CIFAR-100, 10 Classes, Number of Pulls of The Most Frequent Selection

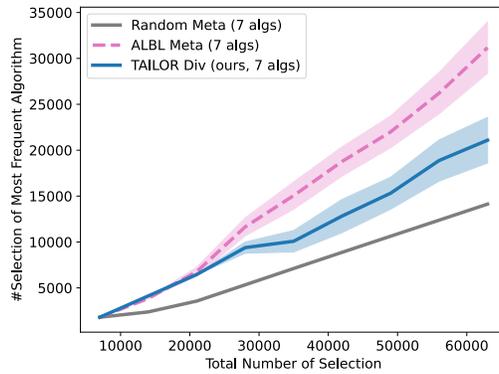


Figure 24: SVHN, 2 Classes, Number of Pulls of The Most Frequent Selection

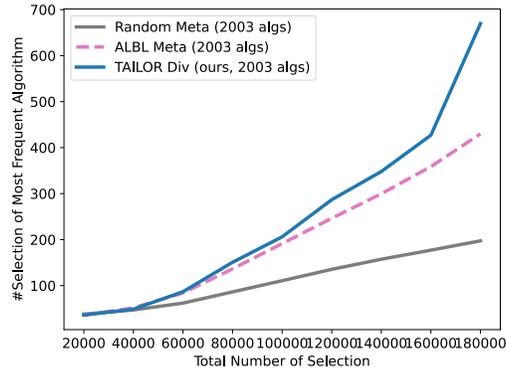


Figure 25: ImageNet-1k, Number of Pulls of The Most Frequent Selection

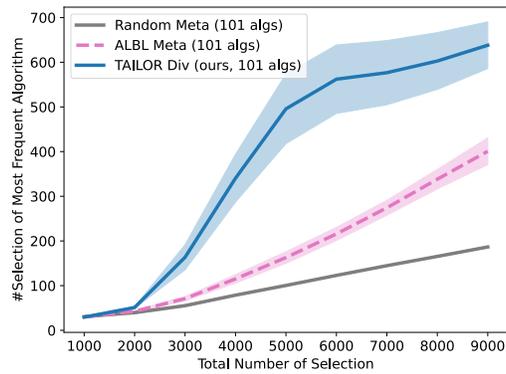


Figure 26: Kuzushiji-49, Number of Pulls of The Most Frequent Selection

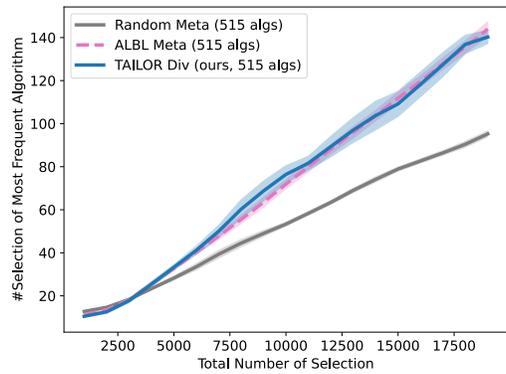


Figure 27: Caltech256, Number of Pulls of The Most Frequent Selection

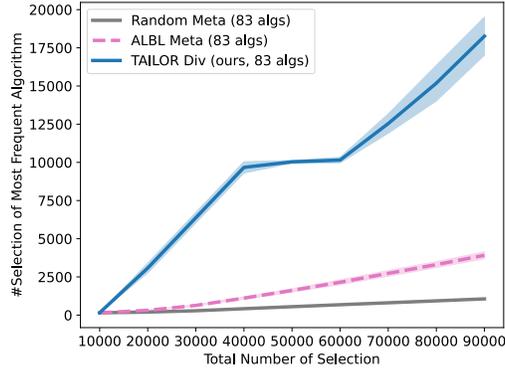


Figure 28: CelebA, Number of Pulls of The Most Frequent Selection

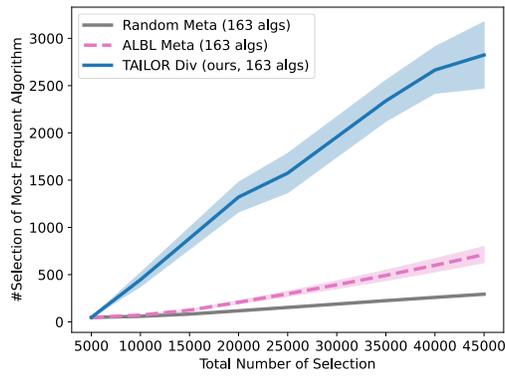


Figure 29: COCO, Number of Pulls of The Most Frequent Selection

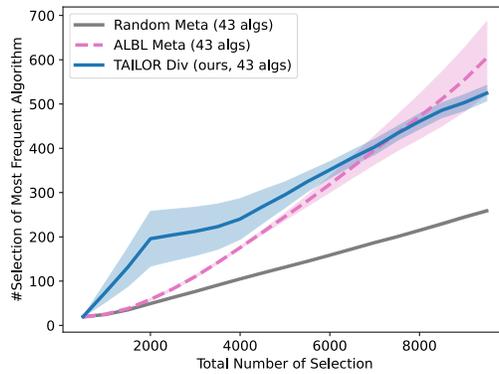


Figure 30: VOC, Number of Pulls of The Most Frequent Selection

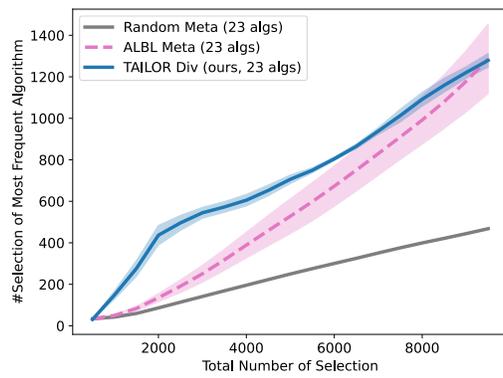


Figure 31: Stanford Car, Number of Pulls of The Most Frequent Selection