

Selecting Near-Optimal Learners via Incremental Data Allocation

Ashish Sabharwal

Allen Institute for AI
Seattle, WA, USA
AshishS@allenai.org

Horst Samulowitz and Gerald Tesaro

IBM Watson Research Center
Yorktown Heights, NY, USA
samulowitz, gtesauro@us.ibm.com

Abstract

We study a novel machine learning (ML) problem setting of sequentially allocating small subsets of training data amongst a large set of classifiers. The goal is to select a classifier that will give near-optimal accuracy when trained on all data, while also minimizing the cost of misallocated samples. This is motivated by large modern datasets and ML toolkits with many combinations of learning algorithms and hyper-parameters. Inspired by the principle of “optimism under uncertainty,” we propose an innovative strategy, Data Allocation using Upper Bounds (DAUB), which robustly achieves these objectives across a variety of real-world datasets.

We further develop substantial theoretical support for DAUB in an idealized setting where the expected accuracy of a classifier trained on n samples can be known exactly. Under these conditions we establish a rigorous sub-linear bound on the regret of the approach (in terms of misallocated data), as well as a rigorous bound on suboptimality of the selected classifier. Our accuracy estimates using real-world datasets only entail mild violations of the theoretical scenario, suggesting that the practical behavior of DAUB is likely to approach the idealized behavior.

Introduction

The goal of our work is to develop novel practical methods to enhance tractability of Data Science practice in the era of Big Data. Consider, for example, the following very common scenario: A Data Science practitioner is given a data set comprising a training set, a validation set, and a collection of classifiers in an ML toolkit, each of which may have numerous possible hyper-parameterizations. The practitioner would like to determine which classifier/parameter combination (hereafter referred to as “learner”) would yield the highest validation accuracy, after training on all examples in the training set. However, the practitioner may have quite limited domain knowledge of salient characteristics of the data, or indeed of many of the algorithms in the toolkit.

In such a scenario, the practitioner may inevitably resort to the traditional approach to finding the best learner (cf. Caruana and Niculescu-Mizil 2006), namely, brute-force training of all learners on the full training set, and selecting the one with best validation accuracy. Such an

approach is acceptable if the computational cost of training all learners is not an issue. However, in the era of Big Data, this is becoming increasingly infeasible. Web-scale datasets are proliferating from sources such as Twitter, TREC, SNAP, ImageNet, and the UCI repository, particularly in domains such as vision and NLP. ImageNet datasets can exceed 100 gigabytes, and the recent “YouTube-Sports-1M” video collection exceeds 40 terabytes. Moreover, the diverse set of learners available in today’s ML packages (Hall et al. 2009; Pedregosa and others 2011; Schaul et al. 2010; McCallum 2002) are continually expanding, and many of the most successful recent algorithms entail very heavy training costs (e.g., Deep Learning neural nets with Dropout).

The above factors motivate a search for techniques to reduce training cost while still reliably finding a near-optimal learner. One could consider training each learner on a small subset of the training examples, and choose the best performing one. This entails less computation, but could result in significant loss of learner accuracy, since performance on a small subset can be a misleading predictor of performance on the full dataset. As an alternative, the small-subset results could be projected forward using parameterized accuracy models to predict full training set accuracy. Creating such models is, however, a daunting task (Guerra, Prudencio, and Ludermir 2008), potentially needing prior knowledge about learners and domain, characteristic features of the data, etc.

In this paper, we develop a novel formulation of what it means to solve the above dual-objective problem, and we present a novel solution approach, inspired by multi-armed bandit literature (Auer, Cesa-Bianchi, and Fischer 2002; Thompson 1933; Scott 2010; Agrawal and Goyal 2012). Our method develops model-free, cost-sensitive strategies for sequentially allocating small batches of training data to selected learners, wherein “cost” reflects misallocated samples that were used to train other learners that were ultimately not selected. We express the cost in terms of the *regret* of the approach, comparing the algorithm’s cost with that of an oracle which only allocates data to the best learner.

Our main contributions are as follows. First, we give a precise definition of a new ML problem setting, called the *Cost-Sensitive Training Data Allocation Problem*. Second, we present a simple, knowledge-free, easy-to-use and practical new algorithm for this setting, called DAUB (Data Allocation with Upper Bounds). Third, we give empiri-

cal demonstrations that DAUB achieves significant savings in training time while reliably achieving optimal or near-optimal learner accuracy over multiple real-world datasets. Fourth, we provide theoretical support for DAUB in an idealization of the real-world setting, wherein DAUB can work with noiseless accuracy estimates when training on n samples, in lieu of actual noisy estimates. The real-world behavior of DAUB will progressively approach the idealized behavior as n becomes large. In this setting, we establish a bound on accuracy of learners selected by DAUB, a sub-linear bound on the data misallocated by DAUB, and an associated bound on the computational training cost (regret).

Related work on traditional bandit strategies mentioned above, such as the celebrated UCB1 (Auer, Cesa-Bianchi, and Fischer 2002) and Thompson sampling (Thompson 1933; Agrawal and Goyal 2012), presume that additional trials of a given arm yield stationary payoffs. Whereas in our scenario, additional data allocations to a learner yield increasing values of its accuracy. There are also existing methods to optimize a single arbitrary function while minimizing the number of evaluations (cf. Munos 2014). These also do not fit our setting: we are dealing with multiple unknown but well-behaved functions, and wish to rank them on estimated accuracy after training on the full dataset, based on their upper-bounds from much fewer samples.

Somewhat related is algorithm portfolio selection (Rice 1976) which seeks the most suitable algorithm (e.g., learner) for a given problem instance, based on knowledge from other instances and features characterizing the current instance. Note, however, that most selection algorithms use parameterized accuracy models which are fit to data (e.g., Hutter et al. 2014). Also related is work on hyper-parameter optimization, where one searches for novel configurations of algorithms to improve performance (Snoek, Larochelle, and Adams 2012; Bergstra, Yamins, and Cox 2013; Bergstra and Bengio 2012; Bergstra et al. 2011) or a combination of both (Feurer, Springenber, and Hutter 2015). An example is Auto-Weka (Thornton et al. 2013), which combines selection and parameter configuration based on Bayesian optimization (cf. Brochu, Cora, and de Freitas 2009). Predicting generalization error on unseen data has in fact been recognized as a major ML challenge (Guyon et al. 2006).

A recent non-frequentist approach (Hoffman, Shahriari, and de Freitas 2014) takes a Bayesian view of multi-armed bandits, applicable especially when the number of arms exceeds the number of allowed evaluations, and applies it also to automatic selection of ML algorithms. Like some prior methods, it evaluates algorithms on a small fixed percentage (e.g., 10%) of the full dataset. Unlike the above approaches, we do not assume that training (and evaluation) on a small fixed fraction of the data reliably ranks full-training results.

Finally, Domhan, Springenberg, and Hutter (2015) recently proposed extrapolating learning curves to enable early termination of non-promising learners. Their method is designed specifically for neural networks and does not apply directly to many classifiers (SVMs, trees, etc.) that train non-iteratively from a single pass through the dataset. They also do not focus on a theoretical justification and fit accuracy estimates to a library of hand-designed learning curves.

Cost-Sensitive Training Data Allocation

We begin by formally defining the problem of cost-sensitive training data allocation. As before, we use *learner* to refer to a classifier along with a hyper-parameter setting for it. Let $\mathcal{C} = C_1, C_2, \dots, C_M$ be a set of M learners which can be trained on subsets of a training set T_r and evaluated on a validation set T_v . Let $|T_r| = N$. For $k \in \mathbb{N}$, let $[k]$ denote the set $\{1, 2, \dots, k\}$.

For $i \in [M]$, let $c_i : [N] \rightarrow \mathbb{R}^+$ be a *cost function* denoting expected computational cost of training learner C_i when n training examples are drawn uniformly at random from T_r .¹ We make two common assumptions about the training process, namely, that it looks at all training data and its complexity grows at least linearly. Formally, $c_i(n) \geq n$ and $c_i(m) \geq \frac{m}{n} c_i(n)$ for $m > n$.

For $i \in [M]$, let $f_i : [N] \rightarrow [0, 1]$ be an *accuracy function* where $f_i(n)$ denotes expected accuracy of C_i on T_v when trained on n training examples chosen at random from T_r . The corresponding *error function*, $e_i(n)$, is defined as $1 - f_i(n)$. Note that our tool also supports accuracy functions not tied to a fixed validation set T_v (e.g., cross-validation) and other measures such as precision, recall, and F1-score; our analysis applies equally well to these measures.

We denote a training data *allocation* of n training samples to learner i by a pair $a = (i, n)$. Let $S = ((i^{(1)}, n^{(1)}), (i^{(2)}, n^{(2)}), \dots, (i^{(s)}, n^{(s)}))$ be a *sequence of allocations* to learners in \mathcal{C} . We will use S_i to denote the *induced subsequence* containing all training data allocations to learner C_i , i.e., the subsequence of S induced by all pairs $(i^{(k)}, n^{(k)})$ such that $i^{(k)} = i$. In our context, if allocations $(i, n^{(k)})$ and $(i, n^{(\ell)})$ are in S_i with $k < \ell$, then $n^{(k)} < n^{(\ell)}$.

Evaluating $f_i(n)$ amounts to training learner C_i on n examples from T_r and evaluating its accuracy. This, in expectation, incurs a computational cost of $c_i(n)$. In general, the expected *training complexity* or cost associated with \mathcal{C} under the data allocation sequence S is $\text{cost}(S) = \sum_{(i,n) \in S} c_i(n)$.

Our goal is to search for an $\tilde{i} \in [M]$ such that $f_{\tilde{i}}(N)$ is maximized, while also ensuring that overall training cost is not too large relative to $c_{\tilde{i}}(N)$. This bi-objective criterion is not easy to achieve. E.g., a brute-force evaluation, corresponding to $S = ((1, N), (2, N), \dots, (M, N))$ and $\tilde{i} = \arg \max_{i \in [M]} f_i(N)$, obtains the optimal \tilde{i} but incurs maximum training cost of $c_i(N)$ for all suboptimal learners. On the other hand, a low-cost heuristic $S = ((1, n), (2, n), \dots, (m, n), (\tilde{i}, N))$ for some $n \ll N$ and $\tilde{i} = \arg \max_i f_i(n)$, incurs a small training overhead of only $c_i(n)$ for each suboptimal C_i , but may choose an arbitrarily suboptimal \tilde{i} .

We seek an in-between solution, ideally with the best of both worlds: a bounded optimality gap on $C_{\tilde{i}}$'s accuracy, and a bounded regret in terms of data misallocated to sufficiently suboptimal learners. Informally speaking, we will ensure that learners with performance at least Δ worse than

¹While we define the core concepts in terms of expected values suitable for a formal definition and idealized analysis, the actual DAUB algorithm will operate on observed values of c_i on particular subsets of n training examples chosen at runtime.

optimal are allocated only $o(N)$ training examples, i.e., an asymptotically vanishing fraction of T_r . Under certain conditions, this will ensure that the training cost regret is sub-linear. We next formally define the notions of suboptimality and regret in this context.

Definition 1. Let \mathcal{C} be a collection of M learners with accuracy functions f_i , $\Delta \in (0, 1]$, $n \in \mathbb{N}^+$, and $i^* = \arg \max_{i \in [M]} f_i(n)$. A learner $C_j \in \mathcal{C}$ is called (n, Δ) -suboptimal for \mathcal{C} if $f_{i^*}(n) - f_j(n) \geq \Delta$, and (n, Δ) -optimal otherwise.

Definition 2. Let S be a data allocation sequence for a collection \mathcal{C} of learners with accuracy functions f_i , $\Delta \in (0, 1]$, and $n \in \mathbb{N}^+$. The (n, Δ) -regret of S for \mathcal{C} is defined as $\sum_{i: C_i \text{ is } (n, \Delta)\text{-suboptimal}} \text{cost}(S_i)$.

The regret of S is thus the cumulative cost of training all (n, Δ) -suboptimal learners when using S .

Definition 3 (COST-SENSITIVE TRAINING DATA ALLOCATION PROBLEM). Let $\mathcal{C} = \{C_1, \dots, C_M\}$ be a set of learners, T_r be a training set for \mathcal{C} containing N examples, T_v be a validation set, c_i and f_i , for $i \in [M]$, be the training cost and accuracy functions, resp., for learner C_i , and $\Delta \in (0, 1]$ be a constant. The *Cost-Sensitive Training Data Allocation Problem* is to compute a training data allocation sequence S for \mathcal{C} and T_r as well as a value $\tilde{i} \in [M]$ such that:

1. S contains (\tilde{i}, N) ,
2. $C_{\tilde{i}}$ is (N, Δ) -optimal,
3. $\text{cost}(S_{\tilde{i}}) \leq d \cdot c_{\tilde{i}}(N)$ for some fixed constant d , and
4. (N, Δ) -regret of S is $o(M \cdot \text{cost}(S_{\tilde{i}}))$ in N .

A solution to this problem thus identifies an (N, Δ) -optimal learner $C_{\tilde{i}}$, trained on all of T_r , incurring on $C_{\tilde{i}}$ no more than a constant factor overhead relative to the minimum training cost of $c_{\tilde{i}}(N)$, and with a guarantee that any (N, Δ) -suboptimal learner C_i incurred a vanishingly small training cost compared to training $C_{\tilde{i}}$ (specifically, $\text{cost}(S_i)/\text{cost}(S_{\tilde{i}}) \rightarrow 0$ as $N \rightarrow \infty$).

The DAUB Algorithm

Algorithm 1 describes our *Data Allocation using Upper Bounds* strategy. The basic idea is to project an optimistic upper bound on full-training accuracy $f_i(N)$ of learner i using recent evaluations $f_i(n)$. The learner with highest upper bound is then selected to receive additional samples. Our implementation of DAUB uses monotone regression to estimate upper bounds on $f_i(N)$ as detailed below, since observed accuracies are noisy and may occasionally violate known monotonicity of learning curves. Whereas in the noise-free setting, a straight line through the two most recent values of $f_i(n)$ provides a strict upper bound on $f_i(N)$.

As a bootstrapping step, DAUB first allocates b, br , and $br^2 \leq N$ training examples to each learner C_i , trains them, and records their training and validation accuracy in arrays f_i^T and f_i^V , resp. If f_i^V at the current point is smaller than at the previous point, DAUB uses a simple *monotone regression* method, making the two values meet in the middle.

```

Input : Learners  $\mathcal{C} = \{C_1, \dots, C_M\}$ , training examples  $T_r$ ,
          $N = |T_r|$ , validation set  $T_v$ 
Output : Learner  $C_{\tilde{i}}$  trained on  $T_r$ , data allocation sequence  $S$ 
Params : Geom. ratio  $r > 1$ , granularity  $b \in \mathbb{N}^+$  s.t.  $br^2 \leq N$ 

DAUB ( $\mathcal{C}, T_r, T_v, r, b$ )
begin
   $S \leftarrow$  empty sequence
  for  $i \in 1..M$  do
    for  $k \in 0..2$  do
       $\text{append}(i, br^k)$  to  $S$ ;  $\text{TrainLearner}(i, br^k)$ 
     $n_i \leftarrow br^2$ ;  $u_i \leftarrow \text{UpdateBound}(i, n_i)$ 
  while  $(\max_i n_i) < N$  do
     $j \leftarrow \arg \max_{i \in [M]} u_i$  (break ties arbitrarily)
     $n_j \leftarrow \min\{\lceil rn_j \rceil, N\}$ 
     $\text{append}(j, n_j)$  to  $S$ ;  $\text{TrainLearner}(j, n_j)$ 
     $u_j \leftarrow \text{UpdateBound}(j)$ 
  select  $\tilde{i}$  such that  $n_{\tilde{i}} = N$ 
  return  $C_{\tilde{i}}, S$ 
end

TrainLearner ( $i \in [M], n \in [N]$ )
begin
   $T \leftarrow n$  examples sampled from  $T_r$ ;  $\text{Train } C_i$  on  $T$ 
   $f_i^T[n] \leftarrow$  training accuracy of  $C_i$  on  $T$ 
   $f_i^V[n] \leftarrow$  validation accuracy of  $C_i$  on  $T_v$ 
  if  $n/r \geq b$  then
     $\delta \leftarrow (f_i^V[n] - f_i^V[n/r])$ 
    if  $\delta < 0$  then  $f_i^V[n/r] \leftarrow f_i^V[n/r] - \delta/2$ ;  $f_i^V[n] \leftarrow f_i^V[n] + \delta/2$ 
  end

UpdateBound ( $i \in [M], n \in [N]$ )
begin
   $f_i^V[n] \leftarrow \text{LinearRegrSlope}(f_i^V[n/r^2], f_i^V[n/r], f_i^V[n])$ 
   $ub_i^V[n] \leftarrow f_i^V[n] + (N - n)f_i^V[n]$ 
  return  $\min\{f_i^T[n], ub_i^V[n]\}$ 
end

```

Algorithm 1: Data Allocation using Upper Bounds

After bootstrapping, in each iteration, it identifies a learner C_j that has the most promising upper bound estimate (computed as discussed next) on the unknown projected expected accuracy $f_j(N)$ and allocates r times more examples (up to N) to it than what C_j was allocated previously. For computing the upper bound estimate, DAUB uses two sources. First, assuming training and validation data come from the same distribution, $f_i^T[n_i]$ provides such an estimate. Further, as will be justified in the analysis of the idealized scenario called DAUB*, $ub_i^V[n_i] = f_i^V[n_i] + (N - n_i)f_i^V[n_i]$ also provides such an estimate under certain conditions, where $f_i^V[n_i]$ is the estimated derivative computed as the slope of the linear regression best fit line through $f_i^V[n]$ for $n \in \{n_i/r^2, n_i/r, n_i\}$. Once some learner $C_{\tilde{i}}$ is allocated all N training examples, DAUB halts and outputs $C_{\tilde{i}}$ along with the allocation sequence it used.

Theoretical Support for DAUB

To help understand the behavior of DAUB, we consider an idealized variant, DAUB*, that operates precisely like DAUB but *has access to the true expected accuracy and cost*

functions, $f_i(n)$ and $c_i(n)$, not just their observed estimates. As n grows, learning variance (across random batches of size n) decreases, observed estimates of $f_i(n)$ and $c_i(n)$ converge to these ideal values, and the behavior of DAUB* thus approaches that of DAUB*.

Let $f^* = \max_{i \in [M]} f_i(N)$ be the (unknown) target accuracy and C_{i^*} be the corresponding (unknown) optimal learner. For each C_i , let $u_i : [N] \rightarrow [0, 1]$ be an arbitrary projected upper bound estimate that DAUB* uses for $f_i(N)$ when it has allocated $n < N$ training examples to C_i . We will assume w.l.o.g. that u_i is non-increasing at the points where it is evaluated by DAUB*.² For the initial part of the analysis, we will think of u_i as a black-box function, ignoring how it is computed. Let $u^{\min}(N) = \min_{i \in [N]} \{u_i(N)\}$. It may be verified that once $u_j(n)$ drops below $u^{\min}(N)$, DAUB* will stop allocating more samples to C_j . While this gives insight into the behavior of DAUB*, for the analysis we will use a slightly weaker form n_i^* that depends on the target accuracy f^* rather than $u^{\min}(N)$.

Definition 4. $u_i : [N] \rightarrow [0, 1]$ is a *valid projected upper bound function* if $u_i(n) \geq f_i(N)$ for all $n \in [N]$.

Definition 5. Define $n_i^* \in \mathbb{N}$ as N if $u_i(N) \geq f^*$ and as $\min\{\ell \mid u_i(\ell) < f^*\}$ otherwise.

A key observation is that when using u_j as the only source of information about C_j , one must allocate at least n_j^* examples to C_j before acquiring enough information to conclude that C_j is suboptimal. Note that n_j^* depends on the interaction between u_j and f^* , and is thus unknown. Interestingly, we can show that DAUB*($\mathcal{C}, T_r, T_v, r, b$) allocates to C_j at most a constant factor more examples, specifically fewer than rn_j^* in each step and $\frac{r^2}{r-1}n_j^*$ in total, if it has access to valid projected upper bound functions for C_j and C_{i^*} (cf. Lemma 1 in Appendix). In other words, DAUB*'s allocation is essentially optimal w.r.t. u_j .

Remark 1. A careful selection of the learner in each round is critical for allocation optimality w.r.t. u_j . Consider a simpler alternative: In round k , train all currently active learners on $n = br^k$ examples, compute all $f_i(n)$ and $u_i(n)$, and permanently drop C_j from consideration if $u_j(n) < f_i(n)$ for some C_i . This will not guarantee allocation optimality; any permanent decisions to drop a classifier must necessarily be conservative to be correct. By instead only temporarily suspending suboptimal looking learners, DAUB* guarantees a much stronger property: C_j receives no more allocation as soon as $u_j(n)$ drops below the (unknown) target f^* .

The following observation connects data allocation to training cost: if DAUB* allocates at most n training examples to a learner C_j in each step, then its overall cost for C_j is at most $\frac{r}{r-1}c_j(n)$ (cf. Lemma 2 in Appendix). Combining this with Lemma 1, we immediately obtain the following result regarding DAUB*'s regret:³

²Since DAUB* evaluates u_i for increasing values of n , it is easy to enforce monotonicity.

³All proofs are deferred to the Appendix.

Theorem 1. Let $\mathcal{C}, T_r, T_v, N, M, c_i$ and f_i for $i \in [M]$ be as in Definition 3. Let $r > 1, b \in \mathbb{N}^+$, and S be the allocation sequence produced by DAUB*($\mathcal{C}, T_r, T_v, r, b$). If the projected upper bound functions u_j and u_{i^*} used by DAUB* are valid, then $\text{cost}(S_j) \leq \frac{r}{r-1}c_j(rn_j^*)$.

In the remainder of the analysis, we will (a) study the validity of the actual projected upper bound functions used by DAUB* and (b) explore conditions under which (N, Δ) -suboptimality of C_j guarantees that n_j^* is a vanishingly small fraction of N , implying that DAUB* incurs a vanishingly small training cost on any (N, Δ) -suboptimal learner.

Obtaining Valid Projected Upper Bounds. If f_i for $i \in [M]$ were arbitrary functions, it would clearly be impossible to upper bound $f_i(N)$ by looking only at estimates of $f_i(n)$ for $n < N$. Fortunately, each f_i is the expected accuracy of a learner and is thus expected to behave in a certain way. In order to bound DAUB*'s regret, we make two assumptions on the behavior of f_i . First, f_i is non-decreasing, i.e., more training data does not hurt validation accuracy. Second, f_i has a *diminishing returns* property, namely, as n grows, the additional validation accuracy benefit of including more training examples diminishes. Formally:

Definition 6. $f : \mathbb{N} \rightarrow [0, 1]$ is *well-behaved* if it is non-decreasing and its discrete derivative, f' , is non-increasing.

These assumptions on expected accuracy are well-supported from the PAC theory perspective. Let $ub_i(n)$ be the projected upper bound function used by DAUB* for C_i , namely the minimum of the *training accuracy* $f_i^T(n)$ of C_i at n and the *validation accuracy* based expression $f_i(n) + (N - n)f_i'(n)$. For DAUB*, we treat $f_i'(n)$ as the one-sided discrete derivative defined as $(f_i(n) - f_i(n - s))/s$ for some parameter $s \in \mathbb{N}^+$. We assume the training and validation sets, T_r and T_v , come from the same distribution, which means $f_i^T(n)$ itself is a valid projected upper bound. Further, we can show that if $f_i(n)$ is well-behaved, then $ub_i(n)$ is a *valid projected upper bound function* (cf. Lemma 3 in Appendix).

Thus, instead of relying on a parameterized functional form to model $f_i(n)$, DAUB* evaluates $f_i(n)$ for certain values of n and computes an expression that is guaranteed to be a valid upper bound on $f_i(N)$ if f_i is well-behaved.

Bounding Regret. We now fix ub_i as the projected upper bound functions and explore how (N, Δ) -suboptimality and the well-behaved nature of f_i together limit how large n_i^* is.

Definition 7. For $\Delta \in (0, 1]$ and a well-behaved accuracy function f_i , define $n_i^\Delta \in \mathbb{N}$ as N if $f_i(N) > \Delta/N$ and as $\min\{\ell \mid f_i'(\ell) \leq \Delta/N\}$ otherwise.

Using first order Taylor expansion, we can prove that $ub_i(n_i^\Delta) < f^*$, implying $n_j^* \leq n_j^\Delta$ (cf. Lemma 4 in Appendix). Combining this with Theorem 1, we obtain:

Theorem 2. Let $\mathcal{C}, T_r, T_v, N, M, c_i$ and f_i for $i \in [M]$ be as in Definition 3. Let $r > 1, b \in \mathbb{N}^+, \Delta \in (0, 1], C_j \in \mathcal{C}$ be an (N, Δ) -suboptimal learner, and S be the allocation sequence produced by DAUB*($\mathcal{C}, T_r, T_v, r, b$). If f_j and f_{i^*} are well-behaved, then $\text{cost}(S_j) \leq \frac{r}{r-1}c_j(rn_j^\Delta)$.

The final piece of the analysis is an asymptotic bound on n_i^Δ . To this end, we observe that the derivative of any bounded, well-behaved, discrete function of \mathbb{N} behaves asymptotically as $o(1/n)$ (cf. Proposition 1 in Appendix). Applying this to f_j , we can prove that if $f'_j(N) \leq \Delta/N$, then n_j^Δ is $o(N)$ in N (cf. Lemma 5 in Appendix).

This leads to our main result regarding DAUB*'s regret:

Theorem 3 (Sub-Linear Regret). *Let $\mathcal{C}, T_r, T_v, N, M, c_i$ and f_i for $i \in [M]$ be as in Definition 3. Let $r > 1, b \in \mathbb{N}^+$, and $\Delta \in (0, 1]$. Let $J = \{j \mid C_j \in \mathcal{C} \text{ is } (N, \Delta)\text{-suboptimal}\}$. For all $j \in J$, suppose f_j is well-behaved and $f'_j(N) \leq \Delta/N$. If DAUB*($\mathcal{C}, T_r, T_v, r, b$) outputs S as the training data allocation sequence along with a selected learner $C_{\bar{i}}$ trained on all of T_r , then:*

1. $\text{cost}(S_{\bar{i}}) \leq \frac{r}{r-1} c_{\bar{i}}(N)$;
2. (N, Δ) -regret of S is $o(\sum_{j \in J} c_j(N))$ in N ; and
3. If $c_j(n) = O(c_{\bar{i}}(n))$ for all $j \in J$, then the (N, Δ) -regret of S is $o(M \cdot \text{cost}(S_{\bar{i}}))$ in N .

Thus, DAUB* successfully solves the cost-sensitive training data allocation problem whenever for $j \in J$, f_j is well-behaved and $c_j(n) = O(c_{\bar{i}}(n))$, i.e., training any suboptimal learner is asymptotically not any costlier than training an optimal learner. While more refined versions of this result can be generated, the necessity of an assumption on the cost function is clear: if a suboptimal learner C_j was arbitrarily costlier to train than optimal learners, then, in order to guarantee near-optimality, one must incur a significant misallocation cost training C_j on some reasonable subset of T_r in order to ascertain that C_j is in fact suboptimal.

Tightness of Bounds. The cost bound on misallocated data in Theorem 2 in terms of n_i^Δ is in fact tight (up to a constant factor) in the worst case, unless further assumptions are made about the accuracy functions. In particular, every algorithm that guarantees (N, Δ) -optimality without further assumptions must, in the worst case, incur a cost of the order of $c_j(n_j^\Delta)$ for every suboptimal $C_j \in \mathcal{C}$ (cf. Theorem 5 in Appendix for a formal statement):

Theorem 4 (Lower Bound, informal statement). *Let $\Delta \in (0, 1]$ and \mathcal{A} be a training data allocation algorithm that always outputs an (N, Δ) -optimal learner. Then there exists an (N, Δ) -suboptimal learner C_j that would force \mathcal{A} to incur a misallocated training cost larger than $c_j(n_j^\Delta)/2$.*

Experiments

Our experiments make use of 41 classifiers covering a wide range of algorithms (SVMs, Decision Trees, Neural Networks, Logistic Regression, etc.) as implemented in WEKA (Hall et al. 2009). All experiments were conducted on AMD Opteron 6134 machines with 32 cores and 64 GB memory, running Scientific Linux 6.1.⁴

We first evaluate DAUB on one real-world binary classification dataset, ‘‘Higgs boson’’ (Baldi, Sadowski, and Whiteson 2014) and one artificial dataset, ‘‘Parity with distractors,’’

⁴Code and data, including full parameterization for each classifier, are available from the authors.

	HIGGS			PARITY		
	Full	no f^T	DAUB	Full	no f^T	DAUB
Iterations	41	105	56	41	26	19
Allocation	1,578k	590k	372k	860k	171k	156k
Time (sec)	49,905	17,186	2,001	5,939	617	397

Table 1: Comparison of full training, DAUB without training accuracy bound f^T , and DAUB.

to examine robustness of DAUB’s strategy across two extremely different types of data. In the latter task the class label is the parity of a (hidden) subset of binary features—the remaining features serve as distractors, with no influence on the class label. We generated 65,535 distinct examples based on 5-bit parity with 11 distractors, and randomly selected 21,500 samples each for T_r and T_v . For the Higgs and other real-world datasets, we first randomly split the data with a 70/30 ratio and selected 38,500 samples for T_r from the 70% split and use the 30% as T_v . We coarsely optimized the DAUB parameters at $b = 500$ and $r = 1.5$ based on the Higgs data, and kept those values fixed for all datasets. This yielded 11 possible allocation sizes: 500, 1000, 1500, 2500, 4000, 5000, 7500, 11500, 17500, 25500, 38500⁵.

Results for HIGGS and PARITY are as follows. The accuracy loss of the ultimate classifiers selected by DAUB turned out to be quite small: DAUB selected the top classifier for HIGGS (i.e. 0.0% loss) and one of the top three classifiers for PARITY (0.3% loss). In terms of complexity reduction, Table 1 shows clear gains over ‘‘full’’ training of all classifiers on the full T_r , in both total allocated samples as well total CPU training time, for both standard DAUB as well as a variant which does not use training set accuracy f^T as an upper bound. Both variants reduce the allocated samples by $\sim 2x$ - $4x$ for HIGGS, and by $\sim 5x$ for PARITY. The impact on CPU runtime is more pronounced, as many suboptimal classifiers with supra-linear runtimes receive very small amounts of training data. As the table shows, standard DAUB reduces total training time by a factor of $\sim 25x$ for HIGGS, and $\sim 15x$ for PARITY.

Figures 1 and 2 provide additional insight into DAUB’s behavior. Figure 1 shows how validation accuracy progresses with increasing training data allocation to several classifiers on the HIGGS dataset. The plots for the most part conform to our ideal-case assumptions of increasing accuracy with diminishing slope, barring a few monotonicity glitches⁶ due to stochastic sampling noise. We note that, while there is one optimal classifier C^* (a parameterization of a Rotation Forest) with best validation accuracy after training on all of T_r , there are several other classifiers that outperformed C^* in early training. For instance, LADTree is better than C^* until 5,000 examples but then flattens out.

Figure 2 gives perspective on how DAUB distributes data allocations among the 41 classifiers when run on the HIGGS

⁵Unfortunately some of our classifiers crash and/or run out of memory above 38500 samples.

⁶In our experience, most of these glitches pertain to weak classifiers and thus would not significantly affect DAUB, since DAUB mostly focuses its effort on the strongest classifiers.

Dataset	Application Area	Full Training		DAUB				
		Allocation	Time (s)	Iterations	Allocation	Time (s)	Speedup	Loss
Buzz	social media	1,578k	56,519	57	302k	5,872	10x	0.0%
Cover Type	forestry	1,578k	43,578	13	160k	3,848	11x	1.1%
HIGGS	signal processing	1,578k	49,905	56	372k	2,001	25x	0.0%
Million Songs	music	1,578k	115,911	53	333k	17,208	7x	0.6%
SUSY	high-energy physics	1,578k	26,438	31	214k	837	31x	0.9%
Vehicle Sensing	vehicle management	1,578k	68,139	50	296k	5,603	12x	0.0%

Table 2: Comparison of full training and DAUB across six benchmarks.




Figure 1: Higgs validation accuracy curves of several classifiers that initially outperform the classifier (Rotation Forest) that is ultimately the best when given all training data.

dataset. The classifiers here are sorted by decreasing validation accuracy $f_i(N)$. While DAUB manages to select C^* in this case, what’s equally critical is the distribution of allocated training data. The figure shows that DAUB allocates most of the training data to the top eight classifiers. Most classifiers receive 2500 or fewer samples, and only four classifiers receive more than 10k samples, with all of them within 1.5% of the optimal performance.

Finally, in Table 2 we report results of DAUB on Higgs plus five other real-world benchmarks as indicated: Buzz (Kawala et al. 2013); Covertypes (Blackard and Dean 2000); Million Song Dataset (Bertin-Mahieux et al. 2011); SUSY (Baldi, Sadowski, and Whiteson 2014); and Vehicle-SensIT (Duarte and Hu 2004). These experiments use exactly the same parameter settings as for HIGGS and PARITY. As before, the table shows a comparison in terms of allocated training samples and runtime. In addition it displays the incurred accuracy loss of DAUB’s final selected classifier. The highest loss is $\sim 1\%$, well within an acceptable range. The average incurred loss across all six benchmarks is 0.4% and the average speedup is 16x. Our empirical findings thus show that in practice DAUB can consistently select near-optimal classifiers at a substantial reduced computational cost when compared to full training of all classifiers.




Figure 2: Number of Higgs examples allocated by DAUB to various classifiers, sorted by decreasing validation performance at $N = 38500$.

Conclusion

We reiterate the potential practical impact of our original *Cost-Sensitive Training Data Allocation* problem formulation, and our proposed DAUB algorithm for solving this problem. In our experience, DAUB has been quite easy to use, easy to code and tune, and is highly practical in robustly finding near-optimal learners with greatly reduced CPU time across datasets drawn from a variety of real-world domains. Moreover, it does not require built-in knowledge of learners or properties of datasets, making it ideally suited for practitioners without domain knowledge of the learning algorithms or data characteristics. Furthermore, all intermediate results can be used to interactively inform the practitioner of relevant information such as progress (e.g., updated learning curves) and decisions taken (e.g., allocated data). Such a tool was introduced by Biem et al. (2015) and a snapshot of it is depicted in the Appendix.

Our theoretical work on the idealized DAUB* scenario also reveals novel insights and provides important support for the real-world behavior of DAUB with noisy accuracy estimates. As dataset sizes scale, we expect that DAUB will better and better approach the idealized behavior of DAUB*, which offers strong bounds on both learner sub-optimality as well as regret due to misallocated samples.

There are many opportunities for further advances in both

the theoretical and practical aspects of this work. It should be possible to develop more accurate bound estimators given noisy accuracy estimates, e.g., using monotone spline regression. Likewise, it may be possible to extend the theory to encompass noisy accuracy estimates, for example, by making use of PAC lower bounds on generalization error to establish upper bounds on learner accuracy. DAUB could be further combined in an interesting way with methods (Ali, Caruana, and Kapoor 2014) to optimally split data between training and validation sets.

References

- Agrawal, S., and Goyal, N. 2012. Analysis of Thompson sampling for the multi-armed bandit problem. In *COLT-2012*, 39.1–39.26.
- Ali, A.; Caruana, R.; and Kapoor, A. 2014. Active learning with model selection. In *Proc. of AAAI-2014*.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47(2-3):235–256.
- Baldi, P.; Sadowski, P.; and Whiteson, D. 2014. Searching for exotic particles in high-energy physics with deep learning. *Nature Communications* 5.
- Bergstra, J., and Bengio, Y. 2012. Random search for hyperparameter optimization. *Journal of Machine Learning Research* 13:281–305.
- Bergstra, J.; Bardenet, R.; Bengio, Y.; and Kégl, B. 2011. Algorithms for hyper-parameter optimization. In *NIPS*, 2546–2554.
- Bergstra, J.; Yamini, D.; and Cox, D. D. 2013. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *ICML-2013*.
- Bertin-Mahieux, T.; Ellis, D. P.; Whitman, B.; and Lamere, P. 2011. The million song dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)*.
- Biem, A.; Butrico, M. A.; Feblowitz, M. D.; Klinger, T.; Malitsky, Y.; Ng, K.; Perer, A.; Reddy, C.; Riabov, A. V.; Samulowitz, H.; Sow, D.; Tesauro, G.; and Turaga, D. 2015. Towards cognitive automation of data science. In *Proc. of AAAI-2015, Demonstrations Track*.
- Blackard, J. A., and Dean, D. J. 2000. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture* 24(3):131–151.
- Brochu, E.; Cora, V. M.; and de Freitas, N. 2009. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report UBC TR-2009-23, Department of Computer Science, University of British Columbia.
- Caruana, R., and Niculescu-Mizil, A. 2006. An empirical comparison of supervised learning algorithms. In *ICML-2006*, 161–168.
- Domhan, T.; Springenberg, J. T.; and Hutter, F. 2015. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *IJCAI-2015*.
- Duarte, M., and Hu, Y. H. 2004. Vehicle classification in distributed sensor networks. In *Journal of Parallel and Distributed Computing*.
- Feurer, M.; Springenberg, J.; and Hutter, F. 2015. Initializing bayesian hyperparameter optimization via meta-learning. In *Proc. of AAAI-2015*.
- Guerra, S. B.; Prudencio, R. B. C.; and Ludermir, T. B. 2008. Predicting the performance of learning algorithms using support vector machines as meta-regressors. In *ICANN*.
- Guyon, I.; Alamdari, A. R. S. A.; Dror, G.; and Buhmann, J. M. 2006. Performance prediction challenge. In *IJCNN-2006*, 1649–1656.
- Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. H. 2009. The WEKA Data Mining Software: An update. *SIGKDD Explorations* 11(1).
- Hoffman, M. D.; Shahriari, B.; and de Freitas, N. 2014. On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In *AISTATS*, 365–374.
- Hutter, F.; Xu, L.; Hoos, H. H.; and Leyton-Brown, K. 2014. Algorithm runtime prediction: Methods & evaluation. *Artif. Intell.* 206:79–111.
- Kawala, F.; Douzal-Chouakria, A.; Gaussier, E.; and Dimert, E. 2013. Prédiction d’activité dans les réseaux sociaux en ligne. In *Conférence sur les Modèles et l’Analyse des Réseaux Approches Mathématiques et Informatique (MARAMI)*.
- McCallum, A. 2002. MALLET: A machine learning for language toolkit. <http://mallet.cs.umass.edu>.
- Munos, R. 2014. From bandits to Monte-Carlo Tree Search: The optimistic principle applied to optimization and planning. *Foundations and Trends in Machine Learning* 7(1):1–130.
- Pedregosa, F., et al. 2011. Scikit-learn: Machine learning in Python. *JMLR* 12:2825–2830.
- Rice, J. 1976. The algorithm selection problem. *Advances in Computers* 15:65–118.
- Schaul, T.; Bayer, J.; Wierstra, D.; Sun, Y.; Felder, M.; Sehnke, F.; Rückstieß, T.; and Schmidhuber, J. 2010. PyBrain. *JMLR*.
- Scott, S. L. 2010. A modern Bayesian look at the multi-armed bandit. *Appl. Stochastic Models Bus. Ind.* 26:639–658.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems* 25.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25:285–294.
- Thornton, C.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2013. Auto-WEKA: Combined selection and hyperparameter optimization of classification algorithms. In *Proc. of KDD-2013*, 847–855.

Appendix: Proof Details

Lemma 1. *Let $\mathcal{C}, T_r, T_v, N, M$, and f_i for $i \in [M]$ be as in Definition 3. Let $r > 1$ and $b \in \mathbb{N}^+$. If the projected upper bound functions u_j and u_{i^*} used by $DAUB^*(\mathcal{C}, T_r, T_v, r, b)$ are valid, then it allocates to C_j fewer than rn_j^* examples in each step and $\frac{r^2}{r-1}n_j^*$ examples in total.*

Proof of Lemma 1. Suppose, for the sake of contradiction, that $DAUB^*$ allocates at least rn_j^* examples to learner C_j at some point in its execution. Since $r > 1$ and all allocation sizes are at most N , $n_j^* < N$. Further, since n_j , the number of examples allocated to C_j , is always incremented geometrically by a factor of at most r , at some previous point in the algorithm, we must have $n_j \in \{n_j^*, \dots, rn_j^* - 1\}$. Since the projected upper bound function u_j is non-increasing, at that

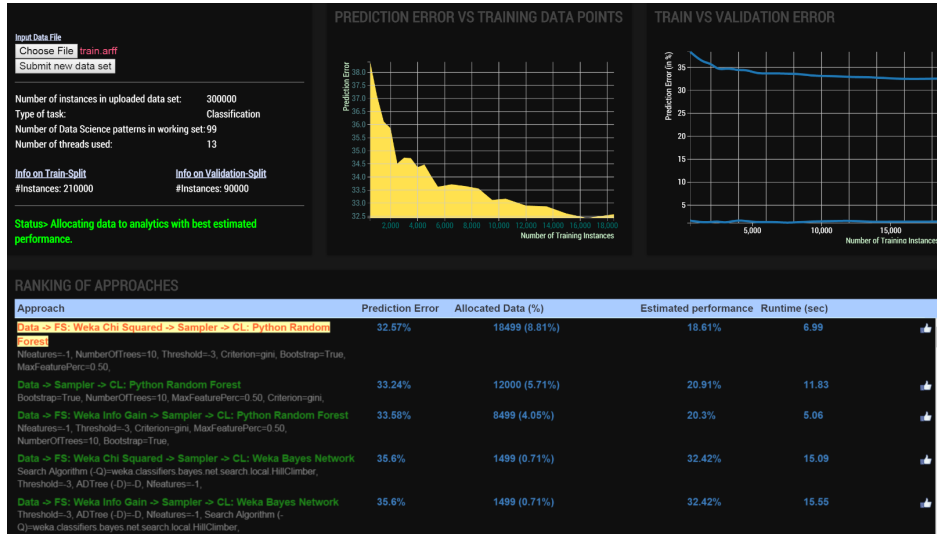


Figure 3: Implementation of a practitioners tool based on DAUB. The interactive dashboard shows the performance of various learners in terms of loss on the validation set based on training on currently allocated training data points. For each learner there are continuously updated plots that show the learning curve and a comparison of training and validation performance.

previous point in the algorithm, $u_j(n_j) \leq u_j(n_j^*) < f^*$ by the definition of n_j^* . On the other hand, since the projected upper bound function u_{i^*} is valid, the projected upper bound for C_{i^*} would always be at least f^* .

Therefore, the algorithm, when choosing which learner to allocate the next set of examples to, will, from this point onward, always prefer C_{i^*} (and possibly another learner appearing to be even better) over C_j , implying that n_j will never exceed its current value. This contradicts the assumption that DAUB* allocates at least rn_j^* examples to C_j at some point during its execution.

For the bound on the total number of examples allocated to C_j , let $k = \lfloor \log_r \frac{rn_j^*}{b} \rfloor$. Since DAUB* allocates fewer than rn_j^* examples to C_j in any single step and the allocation sizes start at b and increase by a factor of r in each step, C_j must have received precisely $b + br + br^2 + \dots + br^k$ examples in total. This is smaller than $\frac{r}{r-1} br^k \leq \frac{r^2}{r-1} n_j^*$. \square

Lemma 2. Let $\mathcal{C}, T_r, T_v, N, M$, and c_i for $i \in [M]$ be as in Definition 3. Let $r > 1, b \in \mathbb{N}^+$, and S be the training data allocation sequence produced by $DAUB^*(\mathcal{C}, T_r, T_v, r, b)$. If $DAUB^*$ allocates at most n training examples to a learner $C_j \in \mathcal{C}$ in each step, then $cost(S_j) \leq \frac{r}{r-1} c_j(n)$.

Proof of Lemma 2. As in the proof of Lemma 1, let $k = \lfloor \log_r(n/b) \rfloor$ and observe that the data allocation subsequence S_j for learner C_j must have been $(b, br, br^2, \dots, br^k)$. The corresponding training cost for C_j is $cost(S_j) = c_j(b) + c_j(br) + c_j(br^2) + \dots + c_j(br^k)$. By the assumption that c_j grows at least linearly: $c_j(br^\ell) \leq \frac{c_j(br^k)}{r^{k-\ell}}$ for $\ell \leq k$. It follows that:

$$\begin{aligned} cost(S_j) &\leq c_j(br^k) \cdot (r^{-k} + r^{-k+1} + r^{-k+2} + \dots + 1) \\ &< \frac{r}{r-1} c_j(br^k) \leq \frac{r}{r-1} c_j(n) \end{aligned}$$

This finishes the proof. \square

Lemma 3. If $f_i(n)$ is well-behaved, then $ub_i(n)$ is a valid projected upper bound function.

Proof of Lemma 3. Recall that $ub_i(n)$ is the minimum of $f_i^T(n)$ and $f_i(n) + (N-n)f_i'(n)$. Since $f_i^T(n)$ is a non-increasing function and is already argued to be an upper bound on $f_i(N)$, it suffices to show that $g(n) = f_i(n) + (N-n)f_i'(n)$ is also a non-increasing function of n and $g(n) \geq f_i(N)$.

$$\begin{aligned} g(n+1) &= f_i(n+1) + (N-n-1)f_i'(n+1) \\ &\leq (f_i(n) + f_i'(n)) + (N-n-1)f_i'(n+1) \\ &\leq (f_i(n) + f_i'(n)) + (N-n-1)f_i'(n) \\ &\quad \text{because } f_i' \text{ is non-increasing} \\ &= f_i(n) + (N-n)f_i'(n) = g(n) \end{aligned}$$

The first inequality follows from the assumptions on the behavior of f_i w.r.t. n and its first-order Taylor expansion. Specifically, recall that $f_i'(n)$ is defined as $(f_i(n) - f_i(n-s))/s$ for some (implicit) parameter s . For concreteness, let's refer to that implicit function as $h_i(n, s)$. Let $h_i'(n, s)$ denote the discrete derivative of $h_i(n, s)$ w.r.t. n . The non-increasing nature of $f_i'(n)$ w.r.t. n implies $h_i'(n, s)$ for any fixed n , is a non-decreasing function of s . In particular, $f_i'(n) = h_i'(n, s) \geq h_i'(n, 1)$ for any $s \geq 1$. It follows that $f_i(n+1) = f_i(n) + h_i'(n+1, 1) \leq f_i(n) + h_i'(n+1, s) = f_i(n) + f_i'(n+1) \leq f_i(n) + f_i'(n)$, as desired.

Thus, $g(n)$ is non-increasing. Since $g(N) = f(N)$ by definition, we have $g(n) \geq f(N)$, finishing the proof. \square

Lemma 4. For an (N, Δ) -suboptimal learner C_j with well-behaved accuracy function $f_j, n_j^* \leq n_j^\Delta$.

Proof of Lemma 4. If $f'_j(N) > \Delta/N$, then $n_j^\Delta = N$ and the statement of the lemma holds trivially. Otherwise, by the definition of n_j^Δ , $f'_j(n_j^\Delta) \leq \Delta/N$. In order to prove $n_j^* \leq n_j^\Delta$, we must show that $ub_j(n_j^\Delta) < f^*$. We do this by using first-order Taylor expansion:

$$\begin{aligned} ub_j(n_j^\Delta) &\leq f_j(n_j^\Delta) + (N - n_j^\Delta) f'_j(n_j^\Delta) \\ &\leq f_j(n_j^\Delta) + (N - n_j^\Delta) \frac{\Delta}{N} \quad \text{by above observation} \\ &< f_j(n_j^\Delta) + \Delta \\ &\leq f_j(N) + \Delta \quad \text{since } n_j^\Delta \leq N; f_j \text{ is non-decreasing} \\ &\leq f_{i^*}(N) \quad \text{by } (N, \Delta)\text{-suboptimality of } C_j \end{aligned}$$

Hence, $ub_j(n_j^\Delta) < f_{i^*}(N) = f^*$. \square

Proposition 1. *If $g : \mathbb{N} \rightarrow [m_1, m_2]$ is a well-behaved function for $m_1, m_2 \in \mathbb{R}$, then its discrete derivative $g'(n)$ decreases asymptotically as $o(1/n)$.*

Proof of Proposition 1. Recall the definition of the discrete derivative from Section and assume for simplicity of exposition that the parameter s is 1, namely, $g'(n) = g(n) - g(n-1)$. The argument can be easily extended to $s > 1$. Applying the definition repeatedly, we get $g(n) = g(1) + (g'(2) + \dots + g'(n)) \geq m_1 + g'(2) + \dots + g'(n)$. If $g'(n)$ was $\Omega(1/n)$, then there would exist n_0 and c such that for all $n \geq n_0$, $g'(n) \geq c/n$. This would mean $g(n) \geq m_1 + \sum_{n \geq n_0} c/n$. This summation, however, diverges to infinity while $g(n)$ is bounded above by m_2 . It follows that $g'(n)$ could not have been $\Omega(1/n)$ to start with. It must thus decrease asymptotically strictly faster than $1/n$, that is, be $o(1/n)$. \square

Lemma 5. *For an (N, Δ) -suboptimal learner C_j with a well-behaved accuracy function f_j satisfying $f'_j(N) \leq \Delta/N$, we have that n_j^Δ is $o(N)$ in N .*

Proof of Lemma 5. From Proposition 1, $f'_j(N) = o(1/N)$, implying $\Delta/f'_j(N) = \omega(N\Delta)$. This means that a value N' that is $o(N/\Delta)$ and suffices to ensure $\Delta/f'_j(N') \geq N$ for all large enough N , that is, $f'_j(N') \leq \Delta/N$. Since n_j^Δ is, by definition, no larger than N' , n_j^Δ must also be $o(N/\Delta)$. \square

Proof of Theorem 3. Since DAUB* never allocates more than N training examples in a single step to C_i for any $i \in [M]$, it follows from Lemma 2 that $cost(S_i) \leq \frac{r}{r-1} c_i(N)$. In particular, $cost(S_i) \leq \frac{r}{r-1} c_i(N)$.

The (N, Δ) -regret of DAUB*, by definition, is $\sum_{j \in J} cost(S_j)$. By Theorem 2, this is at most $\frac{r}{r-1} \sum_{j \in J} c_j(rn_j^\Delta)$. Since the cost function c_j is assumed to increase at least linearly, this quantity is at most $\frac{r}{r-1} \sum_{j \in J} \frac{rn_j^\Delta}{N} c_j(N)$. From Lemma 5, we have that $n_j^\Delta = o(N/\Delta)$ and hence $\frac{rn_j^\Delta}{N} = o(1)$ in N . Plugging this in and dropping the constants r and Δ from the asymptotics, we obtain that the regret is $o(\sum_{j \in J} c_j(N))$.

Finally, if $c_j(n) = O(c_i(n))$, then $\sum_{j \in J} c_j(N) = O(\sum_{j \in J} c_i(N))$, which is simply $O(M \cdot c_i(N))$. Since

$cost(S_i) \geq c_i(N)$, this quantity is also $O(M \cdot cost(S_i))$ in N . It follows from the above result that the (N, Δ) -regret of DAUB* is $o(M \cdot cost(S_i))$ in N , as claimed. \square

Theorem 5 (Lower Bound). *Let $\Delta \in (0, 1]$ and \mathcal{A} be a training data allocation algorithm that, when executed on a training set of size N , is guaranteed to always output an (N, Δ) -optimal learner. Let $\mathcal{C}, T_r, T_v, N, M, c_i$ and f_i for $i \in [M]$ be as in Definition 3. Let $\gamma = (\sqrt{5} - 1)/2$ and $C_j \in \mathcal{C}$ be an (N, Δ) -suboptimal learner. Then there exists a choice of $f_j(n)$ such that f_j is well-behaved, $f'_j(N) \leq \Delta/N$, and $\mathcal{A}(\mathcal{C}, T_r, T_v)$ allocates to C_j more than γn_j^Δ examples, thus incurring a misallocated training cost on C_j larger than $\gamma c_j(n_j^\Delta)$.*

Proof of Theorem 5. We will argue that, under certain circumstances, \mathcal{A} must allocate at least γN_j^Δ examples to C_j in order to guarantee (N, Δ) optimality.

To prove the desired result by contradiction, we consider the optimal learner C_{i^*} and will construct specific accuracy functions $f_j(n)$ and $f_{i^*}(n)$ such that they are identical for all $n \leq \gamma N_j^\Delta$, have derivative no more than Δ/n at N_j^Δ , but differ by at least Δ when evaluated at $n = N$. This would imply that \mathcal{A} simply cannot distinguish between the accuracy of C_j and C_{i^*} by evaluating them on at most γN_j^Δ examples and thus cannot guarantee (N, Δ) -optimality of the learner it outputs.

Suppose we can preserve the properties of f_j required by the theorem, including $f'_j(N_j^\Delta) \leq \Delta/N$, but can enforce that $f'_j(\gamma N_j^\Delta) \geq d\Delta/N$ for some $d > 1$ whose value we will determine shortly. Further, let us alter f_j such that it remains unchanged for $n \leq \gamma N_j^\Delta$ and is altered for larger n such that $f'_j(n) = d\Delta/N$ for $\gamma N_j^\Delta < n \leq N_j^\Delta$ and $f'_j(n) = \Delta/N$ for $n > N_j^\Delta$. Recalling that $N_j^\Delta \leq N$, this modified f_j then satisfies:

$$\begin{aligned} f_j(N) - f_j(\gamma N_j^\Delta) &= (N_j^\Delta - \gamma N_j^\Delta) \frac{d\Delta}{N} + (N - N_j^\Delta) \frac{\Delta}{N} \\ &= (1 - \gamma - 1/d) N_j^\Delta \frac{d\Delta}{N} + \Delta \end{aligned}$$

which is at least Δ as long as $1/d \leq 1 - \gamma$. Now define f_{i^*} s.t. $f_{i^*}(n) = f_j(n)$ for $n \leq \gamma N_j^\Delta$ and $f_{i^*}(n) = f_j(\gamma N_j^\Delta)$ for $n > \gamma N_j^\Delta$. Thus, f_j and f_{i^*} are identical for $n \leq \gamma N_j^\Delta$ but $f_{i^*}(N) - f_j(N) \geq \Delta$, as desired.

It remains to show that we can choose a valid f_j satisfying the properties required by the theorem but such that $f'_j(\gamma N_j^\Delta) \geq d\Delta/N$. To achieve this, consider $f_j(n) = 1 - c\Delta/n$ for any $c > 0$. Then f_j is non-decreasing, $f'_j(n) = c\Delta/n^2$ is non-increasing, $n_j^\Delta = \sqrt{cN}$, $f'_j(n_j^\Delta) = \Delta/N$, and $f'_j(\gamma n_j^\Delta) = f'_j(\gamma \sqrt{cN}) = \Delta/(\gamma^2 N)$. Notice that $\gamma = (\sqrt{5} - 1)/2$ satisfies the condition $\gamma^2 \leq 1 - \gamma$; it is in fact the largest value that satisfies the condition. Hence, we can choose $d = 1/\gamma^2$ in order to ensure $1/d \leq 1 - \gamma$, which in turn ensures that the altered $f_j(N)$ and $f_j(\gamma N)$, as constructed above, differ by at least Δ . This finishes the construction for the lower bound. \square