

Accelerating Extreme Classification via Adaptive Feature Agglomeration

Ankit Jalan and Purushottam Kar

Department of CSE, IIT Kanpur, INDIA
aankitjalan@gmail.com, purushot@cse.iitk.ac.in

Abstract

Extreme classification seeks to assign each data point, the most relevant labels from a universe of a million or more labels. This task is faced with the dual challenge of high precision and scalability, with millisecond level prediction times being a benchmark. We propose DEFRAG, an adaptive feature agglomeration technique to accelerate extreme classification algorithms. Despite past works on feature clustering and selection, DEFRAG distinguishes itself in being able to scale to millions of features, and is especially beneficial when feature sets are sparse, which is typical of recommendation and multi-label datasets. The method comes with provable performance guarantees and performs efficient task-driven agglomeration to reduce feature dimensionalities by an order of magnitude or more. Experiments show that DEFRAG can not only reduce training and prediction times of several leading extreme classification algorithms by as much as 40%, but also be used for feature reconstruction to address the problem of missing features, as well as offer superior coverage on rare labels.

1 Introduction

The task of taking assigning data points, one or more labels from a vast universe of millions of labels is often referred to as the *extreme classification* problem. Although reminiscent of the classical multi-label learning problem, the emphasis on addressing extremely large label spaces distinguishes extreme classification. Recent advances in extreme classification have allowed problems such as ranking, recommendation and retrieval to be viewed and formulated as multi-label problems, indeed with millions of labels.

This focus on extremely large label sets has given us state-of-the-art methods for product recommendation [Jain *et al.*, 2016], search advertising [Prabhu *et al.*, 2018], and video recommendation [Weston *et al.*, 2013], as well as led to advances in our understanding of scalable optimization [Prabhu *et al.*, 2018], and distributed and parallel processing [Yen *et al.*, 2017; Babbar and Schölkopf, 2017]. Recent advances have utilized a variety of techniques – label embeddings, random forests, binary relevance, which we review in §3.

Nevertheless, extreme classification algorithms continue to face several challenges that we enumerate below.

1) **Precision:** data points often have only 5-6 or fewer labels relevant to them (e.g. very few products, of the possibly millions on sale on an online marketplace, would interest any given customer). It is challenging to accurately identify these 5-6 relevant labels among the millions of irrelevant ones.

2) **Prediction:** given their use in (live) recommendation systems, extremely rapid predictions are expected, typically within milliseconds. This often restricts the algorithmic techniques that can be used, to computationally frugal ones.

3) **Processing:** extreme classification datasets contain not only millions of labels, but also millions of data points, each represented as a million-dimensional vector itself. It is challenging to offer scalable training on such large datasets.

4) **Parity:** huge label sets often exhibit power-law behavior with most labels being *rare* i.e. relevant to very few data points. This makes it tempting for algorithms to focus only on popular labels, neglecting the vast majority of rare ones. However, this is detrimental for recommendation outcomes.

In this work, we develop the DEFRAG method and variants to address these specific challenges for a large family of algorithms. Our contributions are summarized below.

Our Contributions.

1) We propose the DEFRAG algorithm that accelerates extreme classification algorithms by performing efficient feature agglomeration on datasets with millions of features and data points. DEFRAG performs agglomeration by constructing a balanced hierarchy which novel, and offers faster and better agglomerates than traditional clustering methods.

2) We show that DEFRAG *provably* preserves the performance of a large family of extreme classification algorithms. This is corroborated experimentally where using DEFRAG significantly reduces training and prediction times of algorithms but with no significant reduction in precision levels.

3) We exploit DEFRAG’s agglomerates in a novel manner to develop the REFRAG algorithm to address the parity problem by performing efficient label re-ranking. This vastly improves the coverage of existing algorithms by accurately predicting extremely rare labels.

4) We develop the FIAT algorithm to perform scalable feature imputation which preserves prediction accuracy even when a large fraction of data features are removed.

5) We perform extensive experimentation on large-scale datasets to establish that DEFRAG not only offers significant reductions in training and prediction times, but that it does so with little or no reduction in precision.

2 Problem Formulation and Notation

The training data will be provided as n labeled data points $(\mathbf{x}^i, \mathbf{y}^i), i = 1, \dots, n$ where $\mathbf{x}^i \in \mathbb{R}^d$ is the feature vector and $\mathbf{y}^i \in \{0, 1\}^L$ is the label vector. There may be several (upto L) labels associated with each data point. Extreme classification datasets exhibit extreme *sparsity* in feature and label vectors. Let \hat{d} denote the average number of non-zero features per data point and \hat{L} denote the average number of active labels per data point. §6 shows that $\hat{d} \ll d$ and $\hat{L} \ll L$. We will denote the feature matrix using $X = [\mathbf{x}^1, \dots, \mathbf{x}^n] \in \mathbb{R}^{d \times n}$ and the label matrix using $Y = [\mathbf{y}^1, \dots, \mathbf{y}^n] \in \{0, 1\}^{L \times n}$.

Notation. Let $\mathcal{F} = \{F_1, \dots, F_K\}$ denote any K -partition of the feature set $[d]$ i.e. $F_i \cap F_j = \emptyset$ if $i \neq j$ and $\bigcup_{k=1}^K F_k = [d]$. Let $d_k := |F_k|$ denote the size of the k^{th} cluster. For any vector $\mathbf{z} \in \mathbb{R}^d$, let z_j denote its j^{th} coordinate. For any set $F_k \in \mathcal{F}$, let $\mathbf{z}_{F_k} := [z_j]_{j \in F_k}^T \in \mathbb{R}^{d_k}$ denote the (shorter) vector containing only coordinates from the set F_k .

Feature Agglomeration. Feature agglomeration involves creating clusters of features and then summing up features within a cluster. If \mathcal{F} is a partition of the features $[d]$, then corresponding to every cluster $F_k \in \mathcal{F}$, we create a single “super”-feature. Thus, given a vector $\mathbf{z} \in \mathbb{R}^d$, we can create an agglomerated vector $\tilde{\mathbf{z}}^{[\mathcal{F}]} \in \mathbb{R}^K$ (abbreviated to just $\tilde{\mathbf{z}}$ for sake of notational simplicity) with just K features using the clustering \mathcal{F} . The k^{th} dimension of $\tilde{\mathbf{z}}$ will be $\tilde{z}_k = \sum_{j \in F_k} z_j$ for $k = 1, \dots, K$. The DEFRAG algorithm will automatically learn relevant feature clusters \mathcal{F} .

3 Related Works

We discuss relevant works in extreme classification and scalable clustering and feature agglomeration techniques here.

Binary Relevance. Also known as *one-vs-all* methods, these techniques, for example DiSMEC [Babbar and Schölkopf, 2017], PPDSParse [Yen *et al.*, 2017], and ProXML [Babbar and Schölkopf, 2019], learn L binary classifiers: for each label $l \in [L]$, a binary classifier is learnt to distinguish data points that contain label l from those that do not. Binary relevance methods offer some of the highest precision values among extreme classification algorithms [Prabhu *et al.*, 2018]. However, despite advances in parallel training and active set methods, they still incur training and prediction times that are prohibitive for most applications.

Label/Feature Embedding. These techniques project feature and/or label vectors onto a low dimensional space i.e. $\mathbf{x}^i \mapsto \hat{\mathbf{x}}^i, \mathbf{y}^i \mapsto \hat{\mathbf{y}}^i$ where $\hat{\mathbf{x}}^i, \hat{\mathbf{y}}^i \in \mathbb{R}^p, p \ll \min\{d, L\}$ using random or learnt projections. Prediction and training is performed in the low dimensional space \mathbb{R}^p for speed. These methods SLEEC [Bhatia *et al.*, 2015], AnnexML [Tagami, 2017] and LEML [Yu *et al.*, 2014] offer strong theoretical guarantees, but are usually forced to choose a moderate value




Figure 1: An illustration of the feature agglomeration process.

of p to maintain scalability. This often results in low precision values and causes these methods to struggle on rare labels.

Data Partitioning. These techniques learn a decision tree over the data points which are hierarchically clustered into several leaves, with the hope that the similar data points, i.e. those with similar label vectors, end up in the same leaf. A simple classifier (usually constant) performs label prediction at a leaf. These methods PfastreXML [Jain *et al.*, 2016], FastXML [Prabhu and Varma, 2014] and CRAFTML [Siblini *et al.*, 2018] offer fast prediction times due to prediction being logarithmic in number of leaves in a balanced tree.

Label Partitioning. These methods instead learn to organize labels into (overlapping) clusters, using hierarchical partitioning techniques. Prediction is done by taking a data point to one or more of the leaves of the tree and using a simple method such as 1-vs-all among labels present at that leaf. These methods PLT [Jasinska *et al.*, 2016], Parabel [Prabhu *et al.*, 2018], LPSR [Weston *et al.*, 2013] offer fast prediction times due to the tree structure, as well as high precision by using a 1-vs-all classifier at the leaves, as Parabel does.

Large Scale (Feature) Clustering. Clustering, as well as feature clustering and agglomeration, are well-studied topics. Past works include techniques for scalable balanced k-means using alternating minimization techniques SCBC [Banerjee and Ghosh, 2006] and BCLS [Liu *et al.*, 2017], scalable spectral clustering using landmarking LSC [Chen and Cai, 2011], and scalable information-theoretic clustering ITDC [Dhillon *et al.*, 2003]. We do compare DEFRAG against all these algorithms. These algorithms were chosen since they were able to scale to at least the smallest datasets in our experiments.

4 Adaptive Extreme Feature Agglomeration

We now describe the DEFRAG method, discuss its key advantages and then develop the REFRAG method for rare label prediction and the FIAT method for feature imputation. Recall from §2 that given a K -partition \mathcal{F} of the features $[d]$, feature agglomeration takes each cluster $F_k \in \mathcal{F}$ and agglomerates all features $j \in F_k$ by summing up their feature values.

DEFRAG: aDaptive Extreme FeatuRe AGglomeration. Given a dataset with d -dimensional features, DEFRAG first clusters these features into balanced clusters, with each cluster containing, say no more than d_0 features. Suppose this process results in K clusters. DEFRAG then uses feature agglomeration (see §2 and Figure 1) to obtain K -dimensional

Algorithm 1 DEFRAG: Make-Tree

Input: Feature set $S \subseteq [d]$, representative vectors $\mathbf{z}^i \in \mathbb{R}^p$ for each feature $i \in S$, maximum leaf size d_0

Output: A tree with each leaf having upto d_0 features

```
1: if  $|S| \leq d_0$  then
2:    $\mathbf{n} \leftarrow \text{Make-Leaf}(S)$  //No need to split this node
3: else
4:    $\mathbf{n} \leftarrow \text{Make-Internal-Node}()$ 
5:    $\{S_+, S_-\} \leftarrow \text{Balanced-Split}(S, \{\mathbf{z}^i, i \in S\})$ 
      // Balanced spherical k-means or nDCG split
6:    $\mathbf{n}_+ \leftarrow \text{Make-Tree}(S_+, \{\mathbf{z}^i, i \in S_+\}, d_0)$ 
7:    $\mathbf{n}_- \leftarrow \text{Make-Tree}(S_-, \{\mathbf{z}^i, i \in S_-\}, d_0)$ 
8:    $\mathbf{n}.\text{Left-Child} \leftarrow \mathbf{n}_+$ 
9:    $\mathbf{n}.\text{Right-Child} \leftarrow \mathbf{n}_-$ 
10: end if
11: return Root node of this tree  $\mathbf{n}$ 
```

features for all data points in the dataset which are then used for training and testing.

DEFRAG first creates a representative vector for each feature $j \in [d]$ and then performs *hierarchical clustering* on them (see Algorithm 1) to obtain feature clusters. At each internal node of the hierarchy, features at that node are split into two children nodes of equal sizes by solving either a balanced spherical 2-means problem or else by minimizing a *ranking loss* like nDCG [Prabhu and Varma, 2014] which we call DEFRAG-N (see Appendix A for details). This process is continued till we are left with less than d_0 features at a node, in which case the node is made a leaf. We now discuss two methods to construct these representative vectors.

DEFRAG-X This variant clusters together *co-occurrent* features e.g. $j, j' \in [d]$ where data points with a non-zero (or high) value for feature j also have a non-zero (or high) value for feature j' . DEFRAG-X represents each feature $j \in [d]$ as an n -dimensional vector $\mathbf{p}^j = [\mathbf{x}_j^1, \dots, \mathbf{x}_j^n]^\top \in \mathbb{R}^n$, essentially as the list of values that feature takes in all data points.

DEFRAG-XY This variant clusters together *co-predictive* features e.g. $j, j' \in [d]$ where data points where feature j is non-zero have similar labels as data points where feature j' is non-zero. To do so, DEFRAG-XY represents each feature $j \in [d]$ as an L -dimensional vector $\mathbf{q}^j = \sum_{i=1}^n \mathbf{x}_j^i \mathbf{y}^i \in \mathbb{R}^L$, essentially as a weighted aggregate of the label vectors of data points where the feature j is non-zero.

Suited for sparse, high-dim. data. DEFRAG is superior to classical dimensionality reduction techniques like PCA/random projection for high-dimensional, sparse data.

1) Applying feature agglomeration to a vector simply involves summing up the coordinates of that vector and is much cheaper than performing PCA or a random projection.

2) PCA/random projection densify vectors and so methods such as LEML and SLEEC are compelled to use a small embedding dimension (≈ 500) for sake of scalability which leads to information loss. Feature agglomeration, however, does not densify vectors: if a vector $\mathbf{x} \in \mathbb{R}^d$ has only s non-zero coordinates, then for any feature K -clustering \mathcal{F} , the vector $\tilde{\mathbf{x}}^{[\mathcal{F}]} \in \mathbb{R}^K$ cannot have more than s non-zero coordinates. This allows DEFRAG to operate with relatively large

values of K (e.g. $K = d/8$ is default in our experiments as we set $d_0 = 8$) without worrying about memory or time issues. Thus, DEFRAG can offer *mildly* agglomerated vectors which preserve much of the information of the original vector, yet offer speedups due to the reduced dimensionality.

3) Feature agglomeration has an implicit *weight-tying* effect since once we learn a model over the agglomerated features, all features belonging to a given cluster $F_k \in \mathcal{F}$ effectively receive the same model weight. This reduces the capacity of the model and can improve generalization error.

Provably bounded distortion, Performance preservation.

We show in §5 that if we obtain a feature clustering \mathcal{F} with small clustering error, then feature agglomeration using \mathcal{F} *provably* preserves the performance of *all* linear models. Specifically, for every model $\mathbf{w} \in \mathbb{R}^d$ over the original vectors, there must exist a model $\tilde{\mathbf{w}} \in \mathbb{R}^K$ over the agglomerated vectors such that for *any vector* $\mathbf{x} \in \mathbb{R}^d$, we have $\mathbf{w}^\top \mathbf{x} \approx \tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^{[\mathcal{F}]}$. This ensures that similar 1-vs-all models can be learnt over $\tilde{\mathbf{x}}^{[\mathcal{F}]}$, as well as similar trees and label partitions can be built. We note that *all* algorithms discussed in §3 ultimately use just linear models as components (e.g. binary relevance methods learn L linear classifiers, embedding methods learn linear projections, data and label partitioning methods learn linear models to split internal nodes).

Task adaptivity. DEFRAG-XY takes into account labels in its feature representation which makes it task-adaptive as compared to dimensionality reduction or clustering methods like k-means, PCA which do not take consider labels. Indeed, we will see that on many datasets, DEFRAG-XY outperforms DEFRAG-X which also does not take labels into account.

Novelty, Speed and Scalability. Hierarchical *feature* agglomeration is novel in the context of extreme classification although hierarchical *data* partitioning (Parabel) and hierarchical *label* partitioning (PfasteXML) have been successfully attempted before. The representative vectors created by DEFRAG are themselves sparse and hierarchical feature agglomeration offers speedy feature clustering. DEFRAG's overhead on the training process is thus, very small.

Time Complexity. Let $\text{nnz}(X) = n \cdot \hat{d}$ be the number of non-zero elements in the feature matrix X . Computing the feature representations $\mathbf{p}^j, j \in [d]$ takes $\mathcal{O}(\text{nnz}(X))$ time. The total time taken to perform balanced spherical 2-means clustering for all nodes at a certain level in the tree is $\mathcal{O}(\text{nnz}(X))$ as well. Since DEFRAG performs balanced splits, there can be at most $\mathcal{O}(\log d)$ levels in the tree, thus giving us a total time complexity of $\mathcal{O}(\text{nnz}(X) \log d)$.

FIAT: Feature Imputation via Agglomeration. Co-occurrence based feature imputation has been popularly used to overcome the problem of missing features. However, this becomes prohibitive for extreme classification settings since the standard co-occurrence matrix $C = XX^\top$ is too dense to store and operate. We exploit the feature clusters offered by DEFRAG to create a scalable co-occurrence based feature imputation algorithm FIAT. For any feature cluster $F_k \in \mathcal{F}$ let $X_{F_k} \in \mathbb{R}^{d \times n}$ denote the matrix with only those rows that belong to the cluster F_k . Given this, we compute a pseudo co-occurrence matrix $C^{\mathcal{F}} = \sum_{k=1}^K X_{F_k} X_{F_k}^\top \in \mathbb{R}^d$.

Note that $C^{\mathcal{F}}$ has a block-diagonal structure and has only upto $\frac{d}{K}$ non-zero entries where K is the number of clusters. Thus, it is much cheaper to store and operate. Given a feature vector $\mathbf{x} \in \mathbb{R}^d$ that we suspect has missing features, we perform feature imputation on it by simply calculating $C^{\mathcal{F}}\mathbf{x}$.

REFRAG: REranking via FeatuRE AGglomeration. The presence of a vast majority of rare labels that occur in very few data points can cause algorithms to neglect rare labels in favor of popular ones [Wei and Li, 2018]. To address this, we propose an efficient reranking solution based on the pseudo co-occurrence matrix $C^{\mathcal{F}}$ described earlier. First compute the matrix product $C^{\mathcal{F}}XY^{\top} \in \mathbb{R}^{d \times L}$. The l^{th} column of this matrix $l \in [L]$ can be interpreted as giving us a *prototype data point* $\xi^l \in \mathbb{R}^d$ for the label l .

These prototypes can be used to get the affinity score of a test data point \mathbf{x}^t to a label $l \in [L]$ as $e^{-\frac{\gamma}{2} \cdot \|\mathbf{x}^t - \xi^l\|_2^2}$. Once a base classification algorithm such as Parabel or DiSMEC has given scores for the test point \mathbf{x}^t with respect to various labels, instead of predicting the labels with the highest scores right-away, we combine the classifier scores with these affinity scores and then make the predictions. We note that a similar approach was proposed by [Jain *et al.*, 2016] who did achieve enhanced performance on rare labels.

However, whereas their method requires an optimization problem to be solved to obtain the prototypes, we have a closed form expression for prototypes in our model given the efficiently computable pseudo co-occurrence matrix $C^{\mathcal{F}}$.

Due to lack of space, further algorithmic details as well as proofs of theorems in §5 are presented in the full version of the paper available at the URL given below.

5 Performance Guarantees

In this section we establish that DEFRAG provably preserves the performance of extreme classification algorithms. For any vector $\mathbf{v} \in \mathbb{R}^p$ we will utilize the orthogonal decomposition $\mathbf{v} = \mathbf{v}^{\parallel} + \mathbf{v}^{\perp}$ where \mathbf{v}^{\parallel} is the component of \mathbf{v} along the all-ones vector $\mathbf{1}_p = (1, \dots, 1) \in \mathbb{R}^p$ and \mathbf{v}^{\perp} is the component orthogonal to it i.e. $\mathbf{1}_p^{\top} \mathbf{v}^{\perp} = 0$. At the core of our results is the following lemma. Given a real valued matrix $Z \in \mathbb{R}^{d \times p}$ for some $p > 0$ and a K -partition \mathcal{F} of the feature set $[d]$, we will let $Z_k \in \mathbb{R}^{d_k \times p}$ denote the matrix formed out of the rows of the matrix that correspond to the partition F_k .

Lemma 1. *Given any matrix $Z \in \mathbb{R}^{d \times p}$ and any K -partition $\mathcal{F} = \{F_1, \dots, F_K\}$ of $[d]$, suppose there exist vectors $\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^K \in \mathbb{R}^p$ such that $Z_k = \mathbf{1}_{d_k}(\boldsymbol{\mu}^k)^{\top} + \Delta_k$ where $\mathbf{1}_{d_k} := (1, \dots, 1)^{\top} \in \mathbb{R}^{d_k}$, then for every $\mathbf{w} \in \mathbb{R}^d$ and every $k \in [K]$, there must exist a real value $c_{\mathbf{w},k} \in \mathbb{R}$ such that*

$$(\mathbf{w}_{F_k} - c_{\mathbf{w},k} \cdot \mathbf{1}_{d_k})^{\top} Z_k Z_k^{\top} (\mathbf{w}_{F_k} - c_{\mathbf{w},k} \cdot \mathbf{1}_{d_k}) \leq \|\Delta_k^{\top} \mathbf{w}_{F_k}^{\perp}\|_2^2.$$

Lemma 1 will be used to show below that, if a group of features is “well-clustered”, then it is possible to tie together weights corresponding to those features in *every* linear model.

Theorem 2. *Upon executing DEFRAG-X with a feature matrix $X = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ and label matrix $Y = [\mathbf{y}^1, \dots, \mathbf{y}^n]$, suppose we obtain a feature K -partition $\mathcal{F} = [F_1, \dots, F_K]$ with err_k denoting the Euclidean clustering error within the*

k^{th} cluster; then for any loss function $\ell(\cdot)$ that is L -Lipschitz and for every linear model $\mathbf{w} \in \mathbb{R}^d$, there must exist a model $\tilde{\mathbf{w}} \in \mathbb{R}^K$ such that for all subsets of data points $S \subseteq [n]$,

$$\sqrt{\sum_{i \in S} (\ell(\mathbf{w}^{\top} \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^{\top} \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2} \leq L \cdot \sum_{k=1}^K \|\mathbf{w}_{F_k}^{\perp}\|_2 \cdot err_k.$$

To simplify this result, let $w_0 = \max_{k \in [K]} \|\mathbf{w}_{F_k}^{\perp}\|_2^2 \leq \max_{k \in [K]} \|\mathbf{w}_{F_k}\|_2^2$ (since cluster sizes d_k are typically small, w_0 is small too) and use the fact that $err_k \geq 0$ to get

$$\sqrt{\sum_{i \in S} (\ell(\mathbf{w}^{\top} \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^{\top} \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2} \leq L \cdot w_0 \cdot \sum_{k=1}^K err_k.$$

A few points are notable about the above results.

Uniform Model Preservation. Theorem 2 guarantees that if the clustering error is small (and DEFRAG does minimize clustering error), then for *every* possible linear model $\mathbf{w} \in \mathbb{R}^d$ over the original features \mathbf{x}^i , we can learn a model $\tilde{\mathbf{w}} \in \mathbb{R}^d$ over the agglomerated features $\tilde{\mathbf{x}}^i$ such that both models behave similarly with respect to any Lipschitz loss function. It is notable that Theorem 2 holds *simultaneously* for all linear models \mathbf{w} , thus making it algorithm agnostic.

Classifier Preservation. Most leading algorithms (Parabel, DiSMEC, PfastreXML, PPDSparse, SLEEC) construct classifiers by learning several linear models using hinge loss or exponential loss which are Lipschitz. By preserving the performance of all such individual linear models, DEFRAG preserves the overall performance of these algorithms too. Note that Theorem 2 holds uniformly over *all* subsets $S \subseteq [n]$ of data points, which is useful since these algorithms often learn several linear models on various subsets of the data.

Graceful Adaptivity. Suppose for a model \mathbf{w} , the weights within a cluster F_k are similar i.e. $\mathbf{w}_{F_k} \approx w \cdot \mathbf{1}_{d_k}$, $w \in \mathbb{R}$. Then this implies $\mathbf{w}_{F_k}^{\perp} \approx \mathbf{0}$ and the contribution of this cluster to the total error will be very small. This indicates that if some of the original weights are anyway tied together, DEFRAG automatically offers extremely accurate reconstructions.

In Appendix B.1, we show that DEFRAG-XY preserves the performance of label clustering methods such as Parabel.

6 Experimental Results

We studied the effects of using DEFRAG variants with several extreme classification algorithms, as well as compared DEFRAG with other clustering algorithms. Our implementation of DEFRAG is available at the URL given below.

Code Link: <https://github.com/purushottamkar/defrag/>

Datasets and Implementations. All datasets, train-test splits, and implementations of extreme classification algorithms were sourced from the Extreme Classification Repository [Bhatia *et al.*, 2019] (see Table 4 in the appendix). Implementations of clustering algorithms were sourced from the authors whenever possible. For SCBC, LSC and ITDC, public implementations were not available and scalable implementations were created in the Python language.

Method	LMI	Bal.	Ent.	Time (min)	P1(%)	P3(%)	P5(%)
EURLex-4K							
ITDC	0.47	Inf	0.87	0.7	73.85	59.63	49.19
SCBC	0.39	321	0.75	0.55	71.96	59.50	49.41
LSC	0.60	130	0.93	5.7	71.44	58.68	48.71
BCLS	0.50	Inf	0.91	2.5	74.14	60.96	50.53
DEFRAG-X	0.37	1.11	0.99	0.03	78.97	65.68	54.46
Wiki10-31K							
ITDC	0.52	Inf	0.88	23	82.03	69.72	60.07
DEFRAG-G	0.46	1.08	0.99	0.48*	82.72	69.62	60.23
DEFRAG-X	0.36	1.08	0.99	0.45	84.99	73.47	63.91

Table 1: A comparison of DEFRAG with other clustering algorithms on clustering quality (see Appendix C for definitions of clustering metrics), as measured by loss of mutual information (LMI), balance factor, normalized entropy, clustering time, and classification performance when the Parabel algorithm was executed upon agglomerated features given by the clustering algorithms. BCLS, LSC and SCBC could not scale to Wiki10. A balance factor of Inf indicates the presence of an empty cluster. DEFRAG not only outperforms other clustering algorithms in terms of clustering quality and classification accuracy, but offers clustering times that can be an order of magnitude smaller. DEFRAG-G denotes the DEFRAG-X algorithm executed on word features learnt by the GloVe algorithm [Pennington *et al.*, 2014]. DEFRAG could not be outperformed by carefully crafted word vector representations like GloVe either.

*The clustering time for DEFRAG-G does not include time taken to extract (dense) GloVe word features from raw text.

Hyperparameters. If available, hyperparameter settings recommended by authors were used for all methods. If unavailable, a fine grid search was performed over a reasonable range to offer adequately tuned hyperparameters to the methods. DEFRAG had its only hyperparameter, the max size of a feature cluster d_0 (see Algorithm 1), fixed to 8.

Comparison with other clustering methods. Table 1 compares DEFRAG with other clustering algorithms on clustering quality, execution time and classification performance (see Appendix C for definitions of clustering metrics). Features were agglomerated according to feature clusters given by all algorithms and Parabel was executed on them. DEFRAG handily outperforms all other methods.

Dataset-wise and Method-wise performance Table 3 presents the outcome of using DEFRAG with several leading algorithms on 8 datasets. On Wiki10 and Delicious, DEFRAG-XY+Parabel offers the best overall performance across all methods. More generally, the table shows 21 instances, across the 8 datasets, of how DEFRAG performs with various algorithms. In 3 of these instances, DEFRAG outperforms the base method (EURLex-PPDSparse, Wiki10-Parabel and Delicious-Parabel), in 7 others, DEFRAG lags by less than 2.5%, in 8 others, it lags by less than 5%. Only in 3 cases is the lag $> 5\%$. DEFRAG variants do seem to work best with the Parabel method.

Trade-offs offered by DEFRAG. It is easy to see that if we create a small number of clusters K , by setting d_0 to be a large value, then the agglomerated vectors will be lower dimensional and as such, offer faster training/prediction and smaller model sizes. However this may cause a dip in predic-




Figure 2: Trade-offs offered by DEFRAG variants. The maximum size of clusters in DEFRAG variants d_0 was changed among the values 32, 16, 8 (default), and 4. The plots show how this affects prediction accuracy, training/test time and model sizes. The black line shows Parabel’s default performance with the black triangle marking its model size, test time etc. Aggressive clustering (e.g. $d_0 = 32$) gives faster training/prediction times and smaller model sizes but also some drop in accuracy. DEFRAG-X3 and DEFRAG-XY3 refer to an ensemble of 3 independent realizations of DEFRAG (see Appendix A.3). Appendix C.2 presents several more trade-off plots.

Method	P1	P3	P5	N1	N3	N5
Wiki-10K						
PfastReXML	19.02	18.34	18.43	19.02	18.49	18.52
REFRAG	20.56	19.51	19.26	20.56	19.74	19.54
Delicious-200K						
PfastReXML	3.15	3.87	4.43	3.15	3.68	4.06
REFRAG	7.34	8.05	8.66	7.34	7.86	8.27

Table 2: REFRAG with propensity scored metrics. N1,3,5 refer to propensity weighted nDGG@k. ^{||} Values from [Bhatia *et al.*, 2019].

tion accuracy. Figure 2 shows that DEFRAG variants offer attractive trade-offs in this respect.

Rare-label prediction with REFRAG. Table 2 shows that REFRAG offers much better propensity-weighted metrics [Jain *et al.*, 2016] (which down-weigh popular and emphasize rare labels) than PfastreXML which also attempts label reranking. Figure 3 shows that REFRAG achieves much better coverage (3.85) than Parabel (1.23) on Delicious and in general, predicts rare labels far more accurately. Figure 3 also shows that FIAT offers resilience to feature erasures.

Acknowledgments

The authors thank the reviewers for comments on improving the presentation of the paper. The authors are also thankful to the lab-team members at the CSE department, IIT Kanpur esp. M. Bagga, S. Malhotra, N. Yadav, B. K. Mishra for their support in running the experiments. P. K. thanks Microsoft Research India and Tower Research for research grants.




Figure 3: REFRAG offers far superior coverage of rare labels than Parabel on Wiki10 and Delicious datasets. Please see Figure 4 and caption thereof in Appendix C.2 for definitions of coverage and other details. In the last two plots, a fraction of features was randomly erased from the feature vectors of all test data points. The FIAT algorithm is more robust to such erasures than the default Parabel algorithm. The gap between FIAT and Parabel widens as erasures become more common. Please see Figure 5 in Appendix C.2 for more details and plots.

Method	P1(%)	P3(%)	P5(%)	Total Time (hr)	Train Time (hr)	Test Time (ms)	Model Size (GB)
EURLex-4K							
PfastReXML	70.41	59.22	50.56	0.07	0.07	1.26	0.26
DEFRAG-X	68.50	56.57	47.78	0.05	0.05	1.55	0.19
SLEEC	72.96	56.03	45.49	0.06	0.06	1.62	0.70
DEFRAG-X	67.89	51.55	42.04	0.03	0.03	1.05	0.31
Dismec	82.85	70.37	58.69	0.04	0.04	1.1	0.08
DEFRAG-X	79.12	66.39	54.97	0.02	0.02	0.5	0.02
PPDSparse	72.90	57.1	45.8	0.010	0.010	0.03	0.01
DEFRAG-X	71.40	57.9	47.4	0.006	0.006	0.05	0.01
Parabel	82.28	68.81	57.58	0.010	0.010	0.73	0.03
DEFRAG-X	78.97	65.68	54.46	0.009	0.008	0.59	0.01
DEFRAG-XY	79.23	65.77	54.65	0.009	0.008	0.59	0.01
ProXML [‡]	83.40	70.90	59.10	-	-	-	-
Wiki10-31K							
PfastReXML	75.67	64.55	57.35	0.27	0.27	11.94	1.12
DEFRAG-X	69.79	58.54	52.52	0.14	0.13	11.74	0.80
SLEEC	84.28	72.05	61.80	0.38	0.38	6.00	3.9
DEFRAG-X	83.87	70.35	59.76	0.17	0.17	3.50	2.0
Dismec	84.12	74.71	65.94	1.48	1.48	42	7.1
DEFRAG-X	82.30	72.14	63.78	0.66	0.65	15	1.5
PPDSparse	74.68	60.03	49.12	0.59	0.59	2.2	0.04
DEFRAG-X	63.55	50.42	41.20	0.50	0.50	3.7	0.03
Parabel	84.19	72.46	63.37	0.13	0.13	2.04	0.18
DEFRAG-X	84.99	73.47	63.91	0.10	0.09	1.46	0.14
DEFRAG-XY	85.08	73.76	64.06	0.11	0.09	1.47	0.14
Amazon-670K							
PfastReXML	36.90	34.22	32.10	5.70	5.70	6.10	10.98
DEFRAG-X	32.67	30.27	28.40	2.70	2.69	7.21	9.40
SLEEC	32.48	28.87	26.31	2.22	2.22	1.43	8.0
DEFRAG-X	31.40	28.04	25.69	1.63	1.63	1.62	4.2
Parabel	44.92	39.77	35.98	0.24	0.24	0.81	1.94
DEFRAG-X	42.71	37.71	33.93	0.23	0.21	0.76	1.68
DEFRAG-XY	42.62	37.72	33.94	0.23	0.21	0.77	1.66
DiSMEC [§]	44.70	39.70	36.10	-	-	-	-
ProXML [‡]	43.50	38.70	35.30	-	-	-	-
AmazonCat-13K							
PfastReXML	85.56	75.19	62.84	11.66	11.66	0.54	19.02
DEFRAG-X	84.71	73.48	61.19	7.01	7.00	0.53	16.17
Dismec [†]	93.80	79.07	64.05	6.68	6.68	1.45	6.0
DEFRAG-X	89.39	74.90	60.67	4.19	4.18	0.89	1.1
Parabel	93.06	79.15	64.51	0.43	0.43	0.62	0.61
DEFRAG-X	91.70	77.25	62.79	0.44	0.40	0.57	0.39
DEFRAG-XY	92.36	78.20	63.55	0.43	0.40	0.58	0.38
Delicious-200K							
Parabel	46.86	40.08	36.70	5.33	5.33	2.22	6.36
DEFRAG-X	47.23	40.53	37.19	3.23	3.16	1.05	4.83
DEFRAG-XY	47.61	40.90	37.66	3.34	3.12	1.06	4.76
PfastReXML	41.72	37.83	35.58	-	-	-	-
DiSMEC [§]	45.50	38.70	35.50	-	-	-	-
WikiLSHTC-325K							
PfastReXML	58.47	37.70	27.57	11.23	11.23	2.66	14.20
DEFRAG-X	50.86	32.08	23.40	6.03	6.03	2.14	12.82
Parabel	65.04	43.24	32.05	0.58	0.58	0.91	3.09
DEFRAG-X	59.49	39.25	29.20	0.56	0.50	0.79	2.50
DEFRAG-XY	61.38	40.42	29.99	0.54	0.50	0.78	2.44
PPDSparse	64.08	41.26	30.12	-	-	-	-
DiSMEC [§]	64.40	42.50	31.50	-	-	-	-
ProXML [‡]	63.60	41.50	30.80	-	-	-	-
Wikipedia-500K*							
Parabel	68.70	49.57	38.64	5.13	5.13	3.11	5.68
DEFRAG-X	65.15	44.96	34.85	3.27	3.14	1.62	5.25
DEFRAG-XY	64.73	44.79	34.76	3.31	3.20	1.62	5.22
DiSMEC [§]	70.20	50.60	39.70	-	-	-	-
ProXML [‡]	69.00	49.10	38.80	-	-	-	-
Amazon-3M*							
Parabel	47.42	44.66	42.55	3.14	3.14	0.73	31.43
DEFRAG-X	45.68	42.85	40.76	2.93	2.83	0.66	25.34
DEFRAG-XY	45.11	42.36	40.30	2.87	2.80	0.63	25.22
PfastReXML	43.83	41.81	40.09	-	-	-	-

Table 3: DEFRAG’s performance when used with various extreme classification algorithms. “Total time” for DEFRAG includes clustering=train time. On EURLex, Wiki10 and Delicious, DEFRAG actually achieves better classification accuracy than the base classifier itself (bold items). DEFRAG allows expensive 1-vs-all methods like DiSMEC and PPDSparse to be executed in a scalable manner with training time reductions of upto 40% on AmazonCat and model size reductions of upto 20% on WikiLSHTC. Precision values on certain datasets are being reported for sake of easy comparison. These were not obtained in our experiments and are being sourced from original publications:

[§][Babbar and Schölkopf, 2017], [†][Babbar and Schölkopf, 2019], ^{||}[Bhatia *et al.*, 2019].

* For all but Wiki-500K and Amazon-3M, DEFRAG was executed in an ensemble of 3 independent realizations (see Appendix A.3).

[†] Except for DiSMEC on AmazonCat (which required 12 cores to execute scalably), all times are reported on a single core.

References

- [Babbar and Schölkopf, 2017] Rohit Babbar and Bernhard Schölkopf. DiSMEC - Distributed Sparse Machines for Extreme Multi-label Classification. In *10th ACM International Conference on Web Search and Data Mining (WSDM)*, 2017.
- [Babbar and Schölkopf, 2019] Rohit Babbar and Bernhard Schölkopf. Data Scarcity, Robustness and Extreme Multi-label Classification. In *30th European Conference on Machine Learning (ECML) - to appear*, 2019.
- [Banerjee and Ghosh, 2006] Arindam Banerjee and Joydeep Ghosh. Scalable Clustering Algorithms with Balancing Constraints. *Data Mining and Knowledge Discovery*, 13(3):365–395, 2006.
- [Bhatia *et al.*, 2015] Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain. Sparse Local Embeddings for Extreme Multi-label Classification. In *29th Annual Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [Bhatia *et al.*, 2019] Kush Bhatia, Kunal Dahiya, Himanshu Jain, Yashoteja Prabhu, and Manik Varma. The Extreme Classification Repository. <http://manikvarma.org/downloads/XC/XMLRepository.html>, 2019. Retrieved on February 15, 2019.
- [Chen and Cai, 2011] Xinlei Chen and Deng Cai. Large Scale Spectral Clustering with Landmark-Based Representation. In *25th AAAI Conference on Artificial Intelligence (AAAI)*, 2011.
- [Dhillon *et al.*, 2003] Inderjit S. Dhillon, Subramanyam Mallela, and Rahul Kumar. A Divisive Information-Theoretic Feature Clustering Algorithm for Text Classification. *Journal of Machine Learning Research*, 3:1265–1287, 2003.
- [Jain *et al.*, 2016] Himanshu Jain, Yashoteja Prabhu, and Manik Varma. Extreme Multi-label Loss Functions for Recommendation, Tagging, Ranking & Other Missing Label Applications. In *22nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2016.
- [Jasinska *et al.*, 2016] Kalina Jasinska, Krzysztof Dembczyński, Róbert Busa-Fekete, Karlson Pfannschmidt, Timo Klerx, and Eyke Hüllermeier. Extreme F-Measure Maximization using Sparse Probability Estimates. In *33rd International Conference on Machine Learning (ICML)*, 2016.
- [Liu *et al.*, 2017] Hanyang Liu, Junwei Han, Feiping Nie, and Xuelong Li. Balanced Clustering with Least Square Regression. In *31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global Vectors for Word Representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [Prabhu and Varma, 2014] Yashoteja Prabhu and Manik Varma. FastXML: A Fast, Accurate and Stable Tree-classifier for eXtreme Multi-label Learning. In *20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2014.
- [Prabhu *et al.*, 2018] Yashoteja Prabhu, Anil Kag, Shrutendra Harsola, Rahul Agrawal, and Manik Varma. Parabel: Partitioned Label Trees for Extreme Classification with Application to Dynamic Search Advertising. In *27th International World Wide Web Conference (WWW)*, 2018.
- [Siblini *et al.*, 2018] Wissam Siblini, Pascale Kuntz, and Frank Meyer. CRAFTML, an Efficient Clustering-based Random Forest for Extreme Multi-label Learning. In *35th International Conference on Machine Learning (ICML)*, 2018.
- [Tagami, 2017] Yukihiro Tagami. AnnexML: Approximate Nearest Neighbor Search for Extreme Multi-label Classification. In *23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2017.
- [Wei and Li, 2018] Tong Wei and Yu-Feng Li. Does Tail Label Help for Large-Scale Multi-Label Learning. In *27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [Weston *et al.*, 2013] Jason Weston, Ameesh Makadia, and Hector Yee. Label Partitioning For Sublinear Ranking. In *30th International Conference on Machine Learning (ICML)*, 2013.
- [Yen *et al.*, 2017] Ian E. H. Yen, Xiangru Huang, Wei Dai, Pradeep Ravikumar, Inderjit S. Dhillon, and Eric Xing. PPDSParse: A Parallel Primal-Dual Sparse Method for Extreme Classification. In *23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2017.
- [Yu *et al.*, 2014] Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit S. Dhillon. Large-scale Multi-label Learning with Missing Labels. In *31st International Conference on Machine Learning (ICML)*, 2014.

Algorithm 2 Balanced Spherical k-means

Input: Feature set $S \subseteq [d]$, representative vectors for each feature $\mathbf{z}^i \in \mathbb{R}^p, i \in S$

Output: A balanced split of the feature set S

- 1: Initialize centroids $\mathbf{c}^+, \mathbf{c}^-$ to two randomly selected representative vectors
- 2: **while** not converged **do**
- 3: Calculate scores $s_i = (\mathbf{c}^+ - \mathbf{c}^-)^\top \mathbf{z}^i$ for each $i \in S$
- 4: $S_1 \leftarrow$ the $\lfloor |S|/2 \rfloor$ features with highest scores
- 5: $S_2 \leftarrow$ the $\lfloor |S|/2 \rfloor$ features with lowest scores
- 6: Recompute centroids $\mathbf{c}^\pm = \frac{1}{|S_\pm|} \sum_{i \in S_\pm} \mathbf{z}^i$
- 7: **end while**
- 8: **return** $\{S_+, S_-\}$

Algorithm 3 Balanced nDCG Split

Input: Feature set $S \subseteq [d]$, representative vectors for each feature $\mathbf{z}^i \in \mathbb{R}_+^p, i \in S$

Output: A balanced split of the feature set S

- 1: Select two vectors $\mathbf{z}^+, \mathbf{z}^-$ randomly from representative vectors
- 2: Initialize centroids as $\mathbf{r}^\pm \leftarrow \text{rank}(\mathbf{z}^\pm)$
- 3: **while** not converged **do**
- 4: Calculate $s_i = \text{nDCG}(\mathbf{r}^+, \mathbf{z}^i) - \text{nDCG}(\mathbf{r}^-, \mathbf{z}^i)$ for $i \in S$
- 5: $S_1 \leftarrow$ the $\lfloor |S|/2 \rfloor$ features with highest scores
- 6: $S_2 \leftarrow$ the $\lfloor |S|/2 \rfloor$ features with lowest scores
- 7: Recompute centroids $\mathbf{r}^\pm = \text{rank} \left(\sum_{i \in S_\pm} I(\mathbf{z}^i) \cdot \mathbf{z}^i \right)$
- 8: **end while**
- 9: **return** $\{S_+, S_-\}$

A Algorithmic Details from §4

In this section we present details of implementations of algorithms that were omitted from the main text due to lack of space. We start off with implementation details of the DEFRAG-X and DEFRAG-XY methods.

A.1 DEFRAG Implementation Details

As mentioned in the main text, DEFRAG chooses to represent each feature $j \in [d]$ as a vector, either as an n -dimensional vector $\mathbf{p}^j = [\mathbf{x}_j^1, \dots, \mathbf{x}_j^n]^\top \in \mathbb{R}^n$ of the values that feature takes on in the n data points (used by DEFRAG-X), or as an L -dimensional vector $\mathbf{q}^j = \sum_{i=1}^n \mathbf{x}_j^i \mathbf{y}^i \in \mathbb{R}^L$, essentially as a weighted aggregate of the label vectors of data points where the feature j is non-zero.

Irrespective of the representation used, DEFRAG next performs hierarchical clustering on the representative vectors as described in Algorithm 1. Starting with the root node which contains all features, nodes are split evenly till the number of features falls below a set threshold at which point the node is made a leaf. For our experiments in Figure 2, we varied this threshold among the values 4, 8, 16, 32, thus obtaining respectively $d/4, d/8, d/16, d/32$ clusters, but for all other experiments in Figure 3 and Tables 1 and 3, we fixed this threshold to 8, thus obtaining $d/8$ clusters.

The clustering was done using one of two methods - balanced spherical k-means, or balanced nDCG splitting. The details of these implementations are given below. Given the extreme sparsity of both the data and label vectors, the feature representations we obtain via DEFRAG-X and DEFRAG-XY, i.e. $\mathbf{p}^j, \mathbf{q}^j$ are themselves very sparse and so case Euclidean notions of proximity such as dot products and norms may not be appropriate for clustering. Thus, we also develop a version which we call DEFRAG-N which minimizes an nDCG ranking loss at each node [Prabhu and Varma, 2014]. We found that balancing did not affect the performance of DEFRAG-N too much. This may be due to the fact that the minimizing the nDCG loss naturally produces rather balanced clusters, something that has been independently observed by [Jain *et al.*, 2016; Prabhu and Varma, 2014].

Balanced Spherical k-means. Algorithm 2 presents the node splitting routine using the balanced k-means algorithm. The algorithm essentially follows the traditional Lloyd’s algorithm except at the cluster assignment stage when, instead of just assigning each point to the nearest cluster, the algorithm performs a fair split.

This algorithm can be derived as a special case (for $k = 2$) of the refinement step in the constrained clustering routine proposed in [Banerjee and Ghosh, 2006]. It is notable that [Prabhu *et al.*, 2018] derive essentially the same algorithm for splitting labels into balanced cluster, but they derive their approach starting from a different graph flow-based approach to constrained clustering.

nDCG Splitting. Given that the vector representations of the features we use are sparse vectors, we employed a ranking-based splitting technique as well. The technique was adapted from the the work of [Prabhu and Varma, 2014] and is described here (their work applies the technique to cluster binary vectors whereas we apply it to cluster feature representative vectors which need not be binary). Note that the technique requires that the vector representations contain only non-negative values. This is indeed true in all our experiments since the features in our datasets are created out of bag-of-words representations which are indeed non-negative.

For any vector $\mathbf{v} \in \mathbb{R}_+^p$ with positive coordinates i.e. $\mathbf{v}_i \geq 0, i \in [p]$, we let $\text{rank}(\mathbf{v}) \in S_p$ denote the permutation ranking the p coordinates of \mathbf{v} in decreasing order i.e. if $\mathbf{r} := \text{rank}(\mathbf{v})$ then $\mathbf{v}_{\mathbf{r}_i} \geq \mathbf{v}_{\mathbf{r}_j}$ if $i > j$. For any positive vector $\mathbf{v} \in \mathbb{R}_+^p$ and any permutation $\mathbf{r} \in S_p$ is any permutation (not necessarily the one that ranks the coordinates of \mathbf{v}), then we define the *Discounted Cumulative Gain* (DCG) score of the permutation \mathbf{r} with respect to the vector \mathbf{v} as

$$\text{DCG}(\mathbf{r}, \mathbf{v}) := \sum_{j=1}^p \frac{\mathbf{v}_{\mathbf{r}_j}}{\log(1+j)}$$

We also define the maximum such score any ranking can achieve as the following

$$I(\mathbf{v}) := \left(\max_{\mathbf{r} \in S_p} \text{DCG}(\mathbf{r}, \mathbf{v}) \right)^{-1} = (\text{DCG}(\text{rank}(\mathbf{v}), \mathbf{v}))^{-1}$$

Given the above the *normalized* DCG score of any permutation $\mathbf{r} \in S_p$ with respect to the vector \mathbf{v} as

$$\text{nDCG}(\mathbf{r}, \mathbf{v}) := I(\mathbf{v}) \cdot \text{DCG}(\mathbf{r}, \mathbf{v})$$

Now, given a set of vectors $\mathbf{v}^1, \dots, \mathbf{v}^m \in \mathbb{R}^p$, a single “centroid” ranking that fits all of them the best can be found as $\max_{\mathbf{r} \in S_p} \sum_{i=1}^m \text{nDCG}(\mathbf{r}, \mathbf{v}^i)$ where

$$\sum_{i=1}^m \text{nDCG}(\mathbf{r}, \mathbf{v}^i) = \sum_{i=1}^m I(\mathbf{v}^i) \sum_{j=1}^p \frac{\mathbf{v}_{\mathbf{r}_j}^i}{\log(1+j)} = \sum_{j=1}^p \frac{1}{\log(1+j)} \sum_{i=1}^m I(\mathbf{v}^i) \cdot \mathbf{v}_{\mathbf{r}_j}^i$$

This implies that the best ranking is given by $\arg \max_{\mathbf{r} \in S_p} \sum_{i=1}^m \text{nDCG}(\mathbf{r}, \mathbf{v}^i) = \text{rank} \left(\sum_{i=1}^m I(\mathbf{v}^i) \cdot \mathbf{v}_{\mathbf{r}_j}^i \right)$. Algorithm 3 presents the DEFRAG-N clustering technique that uses the above rule to recompute cluster centroids.

A.2 Accelerated Clustering

Given the large size of the datasets used in our experiments, it was important to ensure that the feature clustering time of DEFRAG did not exceed the savings in training time it offered for various methods, so as to ensure that the total training time of DEFRAG (clustering + agglomeration + training) still remained smaller than that of the various extreme classification algorithms.

To do so we notice that, given the heavy tailed phenomenon exhibited by most large-scale datasets, clustering time can be reduced significantly by subsampling data points and labels. More specifically, we notice that in the execution of DEFRAG-X, performance is not affected even if we represent each feature using only its values in \tilde{n} most *voluminous data points* where the “volume” of a data point $i \in [n]$ is calculated as $\|\mathbf{x}^i\|_1$. Similarly, we noticed that the execution of DEFRAG-XY is not affected, i.e. clusters are not affected, even if we take into account only the \tilde{L} most popular labels. Such an effect (of performance not being affected by taking only “head” objects) has been observed before [Wei and Li, 2018] as well.

Although the above steps do not greatly affect the clustering performance, they do drastically reduce the clustering time of the DEFRAG variants. Thus, all our DEFRAG-X experiments are executed taking only the $\tilde{n} = 0.25n$ most voluminous documents. Thus, each feature is represented only as an \tilde{n} -dimensional vector. Similarly, all our DEFRAG-XY experiments are executed taking only the $\tilde{n} = 0.25n$ most voluminous documents and the $\tilde{L} = 0.05L$ most popular labels. Thus, each feature is represented only as an \tilde{L} -dimensional vector.

A.3 Ensemble Training

Several extreme classification methods, such as Parabel, SLEEC, PfastreXML construct ensemble classifiers by executing the algorithm independently a few times to obtain different classifiers and then using consensus voting techniques to aggregate the predictions of these different classifiers. Unless otherwise mentioned, we always executed DEFRAG variants afresh for each member of the ensemble as well.

For example, the Parabel algorithm trains an ensemble of 3 tree-based classifiers, each of which is independently capable of making classifications. We ran DEFRAG independently 3 times as well, once for each execution of Parabel. Since DEFRAG uses random initializations in its clustering routines, we found that the cluster partitions were not identical across the three executions. We did find this step to boost accuracy, presumably since it allowed different sets of features to be clustered together in different executions.

A.4 Cluster Averaging

We note that although DEFRAG simply sums up the feature values within a feature cluster, an alternative technique could be to average the feature values within the cluster. Although not a significant step in general, averaging does affect the performance of classifiers like DiSMEC or PPDsparse which practice *model trimming* i.e. setting model coordinates which have a value below a certain threshold, to zero in order to save model space. For such classifiers, simple agglomeration may produce features with inflated feature values which result in small model values which in turn get trimmed to zero. Cluster averaging may help in these settings.

For instance, on the EURLex dataset with DiSMEC, DEFRAG with cluster averaging can yield more than 1.3% boost in P1 accuracies. However, we note that this effect is not uniform and that on some datasets, averaging can actually hurt performance. For instance, on the Wiki10 dataset with DiSMEC, DEFRAG with cluster averaging actually causes minor dips of upto 0.4% in P1.

A.5 REFRAG Implementation Details

As mentioned in the main text, for a given test point $\mathbf{x}^i \in \mathbb{R}^d$, the pseudo co-occurrence model $C^{\mathcal{F}}$ was used to obtain a score $a_l = e^{-\frac{\gamma}{2} \cdot \|\mathbf{x}^i - \boldsymbol{\xi}^l\|_2^2}$ indicating the affinity of the test point \mathbf{x}^i to the label l . The base algorithm, say Parabel or PPDSParse, was used to obtain a separate score, say b_l indicating what did that algorithm think regarding the suitability of label l for the data point \mathbf{x}^i . These two scores were combined in the following manner.

$$c_l = \alpha \log b_l + (1 - \alpha) \log a_l,$$

and the labels with the highest score, as assigned by c_l were assigned to the data point \mathbf{x}^i . As Figure 3 indicates, this does not appreciably reduce the precision on popular labels as they keep getting predict as before. However, this does greatly increase the algorithm's ability to predict rare labels. We set $\alpha = 0.8$ for all experiments.

B Proofs from §5

We now provide complete proofs for the results mentioned in the main paper starting with the base lemma and the analysis for DEFrag-X followed by the analysis for DEFrag-XY.

Lemma 1. *Given any matrix $Z \in \mathbb{R}^{d \times p}$ and any K -partition $\mathcal{F} = \{F_1, \dots, F_K\}$ of $[d]$, suppose there exist vectors $\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^K \in \mathbb{R}^p$ such that $Z_k = \mathbf{1}_{d_k}(\boldsymbol{\mu}^k)^\top + \Delta_k$ where $\mathbf{1}_{d_k} := (1, \dots, 1)^\top \in \mathbb{R}^{d_k}$, then for every $\mathbf{w} \in \mathbb{R}^d$ and every $k \in [K]$, there must exist a real value $c_{\mathbf{w},k} \in \mathbb{R}$ such that*

$$(\mathbf{w}_{F_k} - c_{\mathbf{w},k} \cdot \mathbf{1}_{d_k})^\top Z_k Z_k^\top (\mathbf{w}_{F_k} - c_{\mathbf{w},k} \cdot \mathbf{1}_{d_k}) \leq \|\Delta_k^\top \mathbf{w}_{F_k}^\perp\|_2^2.$$

Proof. Consider a fixed k and for sake of notational simplicity, let us abbreviate $\mathbf{u} := \mathbf{w}_{F_k}$, $V := Z_k$, $c := c_{\mathbf{w},k}$ and $\mathbf{1} := \mathbf{1}_{d_k}$. We will show that if $V = \mathbf{1}(\boldsymbol{\mu}^k)^\top + \Delta_k$ as promised, then there must exist a $c \in \mathbb{R}$ such that $(\mathbf{u} - c \cdot \mathbf{1})^\top V V^\top (\mathbf{u} - c \cdot \mathbf{1}) \leq \|\Delta_k^\top \mathbf{u}^\perp\|_2^2$. To establish this result, first notice that the objective function $f(c) = (\mathbf{u} - c \cdot \mathbf{1})^\top V V^\top (\mathbf{u} - c \cdot \mathbf{1})$ is minimized at a value of $c_{\min} = \frac{\mathbf{u}^\top V V^\top \mathbf{1}}{\mathbf{1}^\top V V^\top \mathbf{1}}$. At this value, we have

$$f(c_{\min}) = (\mathbf{u}^\top V V^\top \mathbf{u}) - \frac{(\mathbf{u}^\top V V^\top \mathbf{1})^2}{(\mathbf{1}^\top V V^\top \mathbf{1})} = \frac{(\mathbf{u}^\top V V^\top \mathbf{u}) \cdot (\mathbf{1}^\top V V^\top \mathbf{1}) - (\mathbf{u}^\top V V^\top \mathbf{1})^2}{(\mathbf{1}^\top V V^\top \mathbf{1})}$$

Concentrating just on the numerator, upon using the orthogonal decomposition $\mathbf{u} = \mathbf{u}^\parallel + \mathbf{u}^\perp$ where $\mathbf{1}^\top \mathbf{u}^\perp = 0$ and letting $\mathbf{u}^\parallel := p \cdot \mathbf{1}$ for some $p \in \mathbb{R}$, we get

$$\begin{aligned} \mathbf{u}^\top V V^\top \mathbf{u} &= p^2 \cdot \mathbf{1}^\top V V^\top \mathbf{1} + 2p \cdot \mathbf{1}^\top V V^\top \mathbf{u}^\perp + (\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp \\ \mathbf{u}^\top V V^\top \mathbf{1} &= p \cdot \mathbf{1}^\top V V^\top \mathbf{1} + \mathbf{1}^\top V V^\top \mathbf{u}^\perp, \end{aligned}$$

which, upon inserting in the numerator expression, give us

$$\begin{aligned} & p^2 \cdot (\mathbf{1}^\top V V^\top \mathbf{1})^2 + 2p \cdot (\mathbf{1}^\top V V^\top \mathbf{u}^\perp)(\mathbf{1}^\top V V^\top \mathbf{1}) + ((\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp)(\mathbf{1}^\top V V^\top \mathbf{1}) \\ & - p^2 \cdot (\mathbf{1}^\top V V^\top \mathbf{1})^2 - 2p \cdot (\mathbf{1}^\top V V^\top \mathbf{u}^\perp)(\mathbf{1}^\top V V^\top \mathbf{1}) - (\mathbf{1}^\top V V^\top \mathbf{u}^\perp)^2 \\ & = ((\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp)(\mathbf{1}^\top V V^\top \mathbf{1}) - (\mathbf{1}^\top V V^\top \mathbf{u}^\perp)^2. \end{aligned}$$

Thus, we have

$$f(c_{\min}) = \frac{((\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp)(\mathbf{1}^\top V V^\top \mathbf{1}) - (\mathbf{1}^\top V V^\top \mathbf{u}^\perp)^2}{(\mathbf{1}^\top V V^\top \mathbf{1})} = (\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp - \frac{(\mathbf{1}^\top V V^\top \mathbf{u}^\perp)^2}{(\mathbf{1}^\top V V^\top \mathbf{1})}$$

Note that the second term in the above expression is always non-negative, albeit not one that is easy to lower-bound. Thus, we simply bound the function value as

$$f(c_{\min}) \leq (\mathbf{u}^\perp)^\top V V^\top \mathbf{u}^\perp = \|\mathbf{u}^\perp\|_2^2$$

Now, the preconditions of the lemma guarantee us that $V = \mathbf{1}(\boldsymbol{\mu}^k)^\top + \Delta_k$ and thus, we have

$$V^\top \mathbf{u}^\perp = \boldsymbol{\mu}^k \mathbf{1}^\top \mathbf{u}^\perp + \Delta_k^\top \mathbf{u}^\perp = \Delta_k^\top \mathbf{u}^\perp,$$

where in the last step we have used the fact that $\mathbf{1}^\top \mathbf{u}^\perp = 0$ by construction. This finishes the proof. \square

Theorem 2. *Upon executing DEFrag-X with a feature matrix $X = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ and label matrix $Y = [\mathbf{y}^1, \dots, \mathbf{y}^n]$, suppose we obtain a feature K -partition $\mathcal{F} = [F_1, \dots, F_K]$ with err_k denoting the Euclidean clustering error within the k^{th} cluster, then for any loss function $\ell(\cdot)$ that is L -Lipschitz, for every linear model $\mathbf{w} \in \mathbb{R}^d$, there must exist a model $\tilde{\mathbf{w}} \in \mathbb{R}^K$ such that for all subsets of data points $S \subseteq [n]$,*

$$\sqrt{\sum_{i \in S} (\ell(\mathbf{w}^\top \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2} \leq L \cdot \sum_{k=1}^K \|\mathbf{w}_{F_k}^\perp\|_2 \cdot err_k.$$

Proof. We will describe how to obtain the identity of $\tilde{\mathbf{w}}$ in a short while. For now, notice that since all terms in the summation on the left hand side are non-negative, we have for any $S \subseteq [n]$

$$\sum_{i \in S} (\ell(\mathbf{w}^\top \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2 \leq \sum_{i=1}^n (\ell(\mathbf{w}^\top \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2$$

Using the Lipschitz property of the loss function and using the partition structure gives us

$$\sum_{i=1}^n (\ell(\mathbf{w}^\top \mathbf{x}^i; \mathbf{y}^i) - \ell(\tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^i; \mathbf{y}^i))^2 \leq L^2 \cdot \sum_{i=1}^n (\mathbf{w}^\top \mathbf{x}^i - \tilde{\mathbf{w}}^\top \tilde{\mathbf{x}}^i)^2 = L^2 \cdot \sum_{i=1}^n \left(\sum_{k=1}^K (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \mathbf{x}_{F_k}^i \right)^2$$

Expanding the right hand side (and ignoring the L^2 term for now) gives us

$$\begin{aligned} & \sum_{i=1}^n \left(\sum_{k=1}^K ((\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \mathbf{x}_{F_k}^i)^2 + \sum_{k \neq l} ((\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \mathbf{x}_{F_k}^i) ((\mathbf{w}_{F_l} - \tilde{\mathbf{w}}_l \mathbf{1}_{d_l})^\top \mathbf{x}_{F_l}^i) \right) \\ &= \sum_{k=1}^K \sum_{i=1}^n ((\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \mathbf{x}_{F_k}^i)^2 + \sum_{k \neq l} \sum_{i=1}^n ((\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \mathbf{x}_{F_k}^i) ((\mathbf{w}_{F_l} - \tilde{\mathbf{w}}_l \mathbf{1}_{d_l})^\top \mathbf{x}_{F_l}^i) \\ &= \sum_{k=1}^K (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \left[\sum_{i=1}^n \mathbf{x}_{F_k}^i (\mathbf{x}_{F_k}^i)^\top \right] (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k}) + \sum_{k \neq l} (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top \left[\sum_{i=1}^n \mathbf{x}_{F_k}^i (\mathbf{x}_{F_l}^i)^\top \right] (\mathbf{w}_{F_l} - \tilde{\mathbf{w}}_l \mathbf{1}_{d_l}) \end{aligned}$$

Now, DEFrag-X represents each feature $j \in [d]$ as an n dimensional vector, $\mathbf{p}^j = [\mathbf{x}_j^1, \dots, \mathbf{x}_j^n]^\top \in \mathbb{R}^n$ and then performs clustering on these vectors to obtain a K -partition of the feature set $[d]$, say $\mathcal{F} = \{F_1, \dots, F_k\}$. Say the centroids of these K clusters are $\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^K \in \mathbb{R}^n$. Consider one of these clusters, say the k^{th} cluster F_k with, say d_k features in that cluster. If we denote $P_k = [\mathbf{p}^j]_{j \in F_k}^\top \in \mathbb{R}^{d_k \times n}$, then the following observations are immediate

1. For any $\mathbf{v} \in \mathbb{R}^{d_k}$, we have $\mathbf{v}^\top P_k = \mathbf{v}^\top \sum_{i=1}^n \mathbf{x}_{F_k}^i$
2. The Euclidean clustering error within the k^{th} cluster is given by

$$\text{err}_k^2 := \sum_{j \in F_k} \|\mathbf{p}^j - \boldsymbol{\mu}^k\|_2^2 = \|P - \mathbf{1}_{d_k} (\boldsymbol{\mu}^k)^\top\|_F^2$$

3. If we denote $P_k = \mathbf{1}_{d_k} (\boldsymbol{\mu}^k)^\top + \Delta_k$, then we must have $\|\Delta_k\|_F = \|P_k - \mathbf{1}_{d_k} (\boldsymbol{\mu}^k)^\top\|_F = \text{err}_k$

4. Lemma 1, when applied with $p = n$ show us that with the above notation, we have, for some real value $c_{\mathbf{w},k}$

$$\|P_k^\top (\mathbf{w}_{F_k} - c_{\mathbf{w},k} \cdot \mathbf{1}_{d_k})\| \leq \|\Delta_k^\top \mathbf{w}_{F_k}^\perp\|_2 \leq \|\Delta_k^\top\|_2 \|\mathbf{w}_{F_k}^\perp\|_2 = \|\Delta_k\|_2 \|\mathbf{w}_{F_k}^\perp\|_2 \leq \|\Delta_k\|_F \|\mathbf{w}_{F_k}^\perp\|_2 = \text{err}_k \cdot \|\mathbf{w}_{F_k}^\perp\|_2,$$

where $\|\Delta_k\|_2$ denotes the spectral norm (i.e. the largest singular value) of the matrix Δ_k and $\|\Delta_k\|_F$ denotes the Frobenius norm of the same matrix.

Note that although the inequality $\|\Delta_k\|_2 \leq \|\Delta_k\|_F$ may seem loose at first sight, notice that since $\text{rank}(\Delta_k) \leq d_k$, we also have $\|\Delta_k\|_F \leq \sqrt{d_k} \cdot \|\Delta_k\|_2$ and since d_k is typically a small number, this shows that this inequality is not too loose.

The above observations allow us to construct $\tilde{\mathbf{w}}$. All we need to do is, for every cluster $k \in [K]$, consult Theorem 2 to obtain a constant $c_{\mathbf{w},k}$ that offers the guarantees of the theorem. We then simply concatenate these constants to construct $\tilde{\mathbf{w}} = [c_{\mathbf{w},k}]_{k \in [K]} \in \mathbb{R}^K$. Using the first point in the list of observations, we can now also rewrite the last expression in our ongoing calculations as

$$\sum_{k=1}^K (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top P_k P_k^\top (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k}) + \sum_{k \neq l} (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})^\top P_k P_l^\top (\mathbf{w}_{F_l} - \tilde{\mathbf{w}}_l \mathbf{1}_{d_l})$$

Applying the Cauchy-Schwartz inequality and the rest of the observations allows us to upper-bound the above expression as

$$\begin{aligned} & \sum_{k=1}^K \|P_k^\top (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})\|_2^2 + \sum_{k \neq l} \|P_k^\top (\mathbf{w}_{F_k} - \tilde{\mathbf{w}}_k \mathbf{1}_{d_k})\|_2 \|P_l^\top (\mathbf{w}_{F_l} - \tilde{\mathbf{w}}_l \mathbf{1}_{d_l})\|_2 \\ & \leq \sum_{k=1}^K \text{err}_k^2 \cdot \|\mathbf{w}_{F_k}^\perp\|_2^2 + \sum_{k \neq l} \text{err}_k \cdot \|\mathbf{w}_{F_k}^\perp\|_2 \cdot \text{err}_l \cdot \|\mathbf{w}_{F_l}^\perp\|_2 = \left(\sum_{k=1}^K \text{err}_k \cdot \|\mathbf{w}_{F_k}^\perp\|_2 \right)^2 \end{aligned}$$

This finishes the proof upon putting back the L^2 factor we had omitted earlier and taking square roots on both sides. \square

B.1 Performance Guarantees for DEFRAG-XY

We now prove a similar result for DEFRAG-XY, specifically that DEFRAG-XY accurately preserves the performance of label clustering methods such as Parabel. Since Parabel performs ‘‘spherical k-means’’ which relies on scores of the form $(\mathbf{c}^+ - \mathbf{c}^-)^\top \mathbf{z}^l$ to decide cluster assignments, the result below assures that these scores remain preserved on the agglomerated data as well. We stress that the above result can be readily adapted to usual k-means which relies on scores of the form $\|\mathbf{c}^+ - \mathbf{z}^l\|_2^2 - \|\mathbf{c}^- - \mathbf{z}^l\|_2^2$.

Theorem 3. *Upon executing DEFRAG-XY with feature matrix X and label matrix Y , suppose we obtain a feature K -partition \mathcal{F} with Euclidean clustering errors $\{\text{err}_k\}_{k \in [K]}$. Suppose $\mathbf{z}^l = \sum_i \mathbf{y}_i^l \mathbf{x}^i$ and $\tilde{\mathbf{z}}^l = \sum_i \mathbf{y}_i^l \tilde{\mathbf{x}}^i$ are the original and agglomerated label features for $l \in [L]$. Then for every 2-means clustering of the original label features, with cluster centroids \mathbf{c}^+ and \mathbf{c}^- , there must exist centroids $\tilde{\mathbf{c}}^+$ and $\tilde{\mathbf{c}}^-$ that offer similar clustering error over the agglomerated features. Specifically, we have, for all subsets of labels $T \subseteq [L]$,*

$$\sqrt{\sum_{l \in [T]} ((\mathbf{c}^+ - \mathbf{c}^-)^\top \mathbf{z}^l - (\tilde{\mathbf{c}}^+ - \tilde{\mathbf{c}}^-)^\top \tilde{\mathbf{z}}^l)^2} \leq \sum_{k=1}^K \text{err}_k \cdot \|\mathbf{c}_{F_k}^+ - \mathbf{c}_{F_k}^-\|_2.$$

Proof. We will establish the identity of the modified centroids $\tilde{\mathbf{c}}^+, \tilde{\mathbf{c}}^-$ in a short while. For now, notice that as before, by positivity of all terms in the summation, we have for any $T \subseteq [L]$

$$\sum_{l \in [T]} ((\mathbf{c}^+ - \mathbf{c}^-)^\top \mathbf{z}^l - (\tilde{\mathbf{c}}^+ - \tilde{\mathbf{c}}^-)^\top \tilde{\mathbf{z}}^l)^2 \leq \sum_{l=1}^L ((\mathbf{c}^+ - \mathbf{c}^-)^\top \mathbf{z}^l - (\tilde{\mathbf{c}}^+ - \tilde{\mathbf{c}}^-)^\top \tilde{\mathbf{z}}^l)^2$$

For sake of notational simplicity, let us denote $\boldsymbol{\delta} = \mathbf{c}^+ - \mathbf{c}^-$ and $\tilde{\boldsymbol{\delta}} = \tilde{\mathbf{c}}^+ - \tilde{\mathbf{c}}^-$. This gives us

$$\sum_{l=1}^L (\boldsymbol{\delta}^\top \mathbf{z}^l - \tilde{\boldsymbol{\delta}}^\top \tilde{\mathbf{z}}^l)^2 = \sum_{l=1}^L \left(\sum_{k=1}^K (\boldsymbol{\delta}_{F_k} - \tilde{\boldsymbol{\delta}}_k \mathbf{1}_{d_k})^\top \mathbf{z}_{F_k}^l \right)^2$$

Expanding the right hand side and expanding similarly as before gives us

$$\sum_{k=1}^K (\boldsymbol{\delta}_{F_k} - \tilde{\boldsymbol{\delta}}_k \mathbf{1}_{d_k})^\top \left[\sum_{l=1}^L \mathbf{z}_{F_k}^l (\mathbf{z}_{F_k}^l)^\top \right] (\boldsymbol{\delta}_{F_k} - \tilde{\boldsymbol{\delta}}_k \mathbf{1}_{d_k}) + \sum_{k \neq j} (\boldsymbol{\delta}_{F_k} - \tilde{\boldsymbol{\delta}}_k \mathbf{1}_{d_k})^\top \left[\sum_{l=1}^L \mathbf{z}_{F_k}^l (\mathbf{z}_{F_j}^l)^\top \right] (\boldsymbol{\delta}_{F_j} - \tilde{\boldsymbol{\delta}}_j \mathbf{1}_{d_j})$$

Now, DEFRAG-XY represents each feature $j \in [d]$ as an L dimensional vector, $\mathbf{q}^j = \sum_{i=1}^n \mathbf{x}_j^i \mathbf{y}^i \in \mathbb{R}^L$ and then performs clustering on these vectors to obtain a K -partition of the feature set $[d]$, say $\mathcal{F} = \{F_1, \dots, F_k\}$. Say the centroids of these K clusters are $\boldsymbol{\nu}^1, \dots, \boldsymbol{\nu}^K \in \mathbb{R}^n$. Consider one of these clusters, say the k^{th} cluster F_k with, say d_k features in that cluster. If we denote $Q_k = [\mathbf{q}^j]_{j \in F_k}^\top \in \mathbb{R}^{d_k \times L}$, then the following observations are immediate

1. For any $\mathbf{v} \in \mathbb{R}^{d_k}$, we have $\mathbf{v}^\top Q_k = \mathbf{v}^\top \sum_{l=1}^L \mathbf{z}_{F_k}^l$
2. The Euclidean clustering error within the k^{th} cluster is given by

$$\text{err}_k^2 := \sum_{j \in F_k} \|\mathbf{q}^j - \boldsymbol{\nu}^k\|_2^2 = \|Q_k - \mathbf{1}_{d_k} (\boldsymbol{\nu}^k)^\top\|_F^2$$

3. If we denote $Q_k = \mathbf{1}_{d_k} (\boldsymbol{\nu}^k)^\top + \Delta_k$, then we must have $\|\Delta_k\|_F = \|Q_k - \mathbf{1}_{d_k} (\boldsymbol{\nu}^k)^\top\|_F = \text{err}_k$
4. Lemma 1, when applied with $p = L$ show us that with the above notation, we have for some real value $d_{\delta,k}$

$$\|Q_k^\top (\boldsymbol{\delta}_{F_k} - d_{\delta,k} \cdot \mathbf{1}_{d_k})\| \leq \|\Delta_k^\top \boldsymbol{\delta}_{F_k}^\perp\|_2 \leq \|\Delta_k^\top\|_2 \|\boldsymbol{\delta}_{F_k}^\perp\|_2 = \|\Delta_k\|_2 \|\boldsymbol{\delta}_{F_k}^\perp\|_2 \leq \|\Delta_k\|_F \|\boldsymbol{\delta}_{F_k}^\perp\|_2 = \text{err}_k \cdot \|\boldsymbol{\delta}_{F_k}^\perp\|_2,$$

where $\|\Delta_k\|_2$ denotes the spectral norm (i.e. the largest singular value) of the matrix Δ_k and $\|\Delta_k\|_F$ denotes the Frobenius norm of the same matrix.

Again, note that the inequality $\|\Delta_k\|_2 \leq \|\Delta_k\|_F$ is not extremely loose since, $\text{rank}(\Delta_k) \leq d_k$ gives us $\|\Delta_k\|_F \leq \sqrt{d_k} \cdot \|\Delta_k\|_2$ and since d_k is typically a small number, this shows that this inequality is not too loose.

The above observations allow us to construct $\tilde{\boldsymbol{\delta}}$. All we need to do is, for every cluster $k \in [K]$, consult Theorem 2 to obtain a constant $d_{\delta,k}$ that offers the guarantees of the theorem. We then simply concatenate these constants to construct

Data set	Train n	Features d	Labels L	Test	\hat{d}	\hat{L}
EURLex-4K	15539	5000	3993	3809	236.8	5.31
AmazonCat-13K	1186239	203882	13330	306782	71.2	5.04
Wiki10-31K	14146	101938	30938	6616	673.4	18.64
Delicious-200K	196606	782585	205443	100095	301.2	75.54
WikiLSHTC-325K	1778351	1617899	325056	587084	42.1	3.19
Wikipedia-500K	1813391	2381304	501070	783743	385.3	4.77
Amazon-670K	490449	135909	670091	153025	75.7	5.45
Amazon-3M	1717899	337067	2812281	742507	49.7	36.17

Table 4: Dataset Statistics

$\tilde{\delta} = [d_{\delta,k}]_{k \in [K]} \in \mathbb{R}^K$. Once we have $\tilde{\delta}$, we may construct $\tilde{\mathbf{c}}^+$ and $\tilde{\mathbf{c}}^-$ as any two vectors such that $\tilde{\mathbf{c}}^+ - \tilde{\mathbf{c}}^- = \tilde{\delta}$. Moreover, using the first point in the list of observations, we can now also rewrite the last expression in our ongoing calculations as

$$\sum_{k=1}^K (\delta_{F_k} - \tilde{\delta}_k \mathbf{1}_{d_k})^\top Q_k Q_k^\top (\delta_{F_k} - \tilde{\delta}_k \mathbf{1}_{d_k}) + \sum_{k \neq j} (\delta_{F_k} - \tilde{\delta}_k \mathbf{1}_{d_k})^\top Q_k Q_j^\top (\delta_{F_j} - \tilde{\delta}_j \mathbf{1}_{d_j})$$

Applying the Cauchy-Schwartz inequality and the rest of the observations allows us to upper-bound the above expression as

$$\begin{aligned} & \sum_{k=1}^K \left\| Q_k^\top (\delta_{F_k} - \tilde{\delta}_k \mathbf{1}_{d_k}) \right\|_2^2 + \sum_{k \neq j} \left\| Q_k^\top (\delta_{F_k} - \tilde{\delta}_k \mathbf{1}_{d_k}) \right\|_2 \left\| Q_j^\top (\delta_{F_j} - \tilde{\delta}_j \mathbf{1}_{d_j}) \right\|_2 \\ & \leq \sum_{k=1}^K \text{err}_k^2 \cdot \left\| \delta_{F_k}^\perp \right\|_2^2 + \sum_{k \neq j} \text{err}_k \cdot \left\| \delta_{F_k}^\perp \right\|_2 \cdot \text{err}_j \cdot \left\| \delta_{F_j}^\perp \right\|_2 = \left(\sum_{k=1}^K \text{err}_k \cdot \left\| \delta_{F_k}^\perp \right\|_2 \right)^2 \end{aligned}$$

This finishes the proof upon taking square roots on both sides. \square

C Experimental Details from §6

We provide additional details about experimental settings, as well as additional experimental results in this appendix.

C.1 Clustering Metrics

We used several clustering metrics to evaluate the various clustering algorithms compared in Table 1. Those metrics are defined formally below. Note that in that experiment, to be fair, all algorithms were asked to output the same number of clusters $d/8$ where d is the original dimensionality of the data features.

The notion of mutual information loss (LMI) is defined in [Dhillon *et al.*, 2003] for multi-class classification problems and adapted here for our setting of multilabel classification problems. LMI measures the loss in predictive capability due to clustering of features. The LMI score is a positive real value between 0 and 1. Thus, a smaller value of the LMI metric is better. It is notable that in Table 1, DEFRAG achieved the lowest LMI score of 0.37 followed by SCBC which achieved an LMI score of 0.39. All other algorithms achieved an LMI score of around 0.50 or more.

Definition 1 (Mutual Information Loss (LMI)). *Given a K -clustering \mathcal{F} of $[d]$ features, let $X = [\mathbf{x}^1, \dots, \mathbf{x}^n] \in \mathbb{R}^{d \times n}$ denote the original feature matrix, $\tilde{X} = [\tilde{\mathbf{x}}^1, \dots, \tilde{\mathbf{x}}^n] \in \mathbb{R}^{K \times n}$ denote the matrix of agglomerated features and $Y = [\mathbf{y}^1, \dots, \mathbf{y}^n] \in \{0, 1\}^{L \times n}$ denote the label matrix. Then the normalized loss in mutual information is defined as*

$$LMI(\mathcal{F}) = \frac{I(Y; X) - I(Y; \tilde{X})}{I(Y; X)},$$

where $I(Y; X)$ is the mutual information between the labels and the original features and $I(Y; \tilde{X})$ is the mutual information between the agglomerated features. These terms are defined below for the multi-label learning setting.

Definition 2 (Mutual Information (MI)). *Given a feature matrix $Z = [\mathbf{z}^1, \dots, \mathbf{z}^n] \in \mathbb{R}^{p \times n}$ with each data point having p features and $Y = [\mathbf{y}^1, \dots, \mathbf{y}^n] \in \{0, 1\}^{L \times n}$ denote the label matrix. Then the mutual information between the labels and the features is defined as*

$$I(Y; Z) = \text{trace}(\mathbf{\Pi}^\top (\ln(\mathbf{\Pi}) - \ln(\boldsymbol{\pi}_p) \mathbf{1}_L^\top - \mathbf{1}_p \ln(\boldsymbol{\pi}_L)^\top))$$




Figure 4: REFRAG can offer far superior performance on rare labels as indicated on the Delicious (left) and Wiki10 (right) datasets. The x-axis records the percentile range of labels. For example, the first bar considers labels in the top 0-1 percentile, the second bar considers labels in the 1-2 percentile. Subsequent bars consider labels in the 2-3, 3-5, 5-7, 7-10, 10-20, 20-30, 30-50, 50-70 and 70-100 percentiles respectively. The y-axis records the macro precision of labels within a percentile range. This is calculated by taking an average of label-wise precision scores. Note that this method gives equal weightage to each label, irrespective of its popularity.

At the top of each figure, P indicates the overall precision@1 of the methods, R indicates the overall recall@1 of the methods and C indicates the overall coverage@1 of the methods. Coverage@1 is calculated as the fraction of labels that are correctly predicted as the top-ranked label in at least one test document. Note that DEFRAG more than triples the coverage of the Parabel method on the Wiki10 dataset.

Note that on both datasets, REFRAG continues to offer meaningful predictions even as we consider very rare labels. The default Parabel algorithm stops offering predictions after the 7-10 percentile on Delicious and the 20-30 percentile on Wiki10. However, REFRAG continues to offer predictions upto the 20-30 percentile on Delicious and 70-100 percentile on Wiki10.

where $\mathbf{\Pi} = \frac{ZY^T}{\mathbf{1}_p^T ZY^T \mathbf{1}_L} \in \Delta_{(p \times L)-1}$ is the joint feature-label probability matrix, $\ln(\mathbf{\Pi}) \in \mathbb{R}^{p \times L}$ denotes the matrix obtained by taking entry wise logarithm of the $\mathbf{\Pi}$ matrix, $\boldsymbol{\pi}_p = \frac{Y\mathbf{1}_n}{\|Y\mathbf{1}_n\|_1} \in \Delta_{p-1}$ is the feature probability vector, $\boldsymbol{\pi}_L = \frac{Z\mathbf{1}_p}{\|Z\mathbf{1}_p\|_1} \in \Delta_{L-1}$ is the label probability vector, $\ln(\boldsymbol{\pi}_p)$ and $\ln(\boldsymbol{\pi}_L)$ are obtained by taking element-wise logarithm on the vectors $\boldsymbol{\pi}_p$ and $\boldsymbol{\pi}_L$ respectively, and $\mathbf{1}_k = (1, \dots, 1)^T \in \mathbb{R}^k$ for any $k > 0$. The above definition assumes that the feature matrix Z contains no negative values. However, this is true for bag-of-words representations used in extreme classification problems.

The notion of balance factor is a rather unforgiving metric measuring how balanced are the clusters output by a clustering technique. It is frequently seen that imbalanced feature clusters can lead to reduction in classification performance. The balance factor is positive real number greater than or equal to one. A smaller balance factor indicates a more balanced clustering with a balance factor of 1 denoting perfectly balanced clusters. It is notable that in Table 1, DEFRAG achieved the lowest balance factor of 1.11 and 1.08.

Definition 3 (Balance Factor). Given a K -clustering $\mathcal{F} = [F_1, \dots, F_K]$ of $[d]$ features where we define $d_k := |F_k|$, its balance factor is defined as

$$\text{Balance}(\mathcal{F}) = \frac{\max_k d_k}{\min_k d_k}$$

If we have $\min_k d_k = 0$, then the balance factor is defined to be infinity.

The normalized entropy is a more gentle metric of balance. It takes a value between 0 and 1 and a larger value indicates a more balanced clustering and a smaller value indicates the presence of a few very concentrated clusters. It is notable that in Table 1, DEFRAG achieved the highest entropy values of 0.99.

Definition 4 (Normalized Entropy). Given a K -clustering $\mathcal{F} = [F_1, \dots, F_K]$ of $[d]$ features where we define $d_k := |F_k|$, its normalized entropy is defined as

$$\text{NormEnt}(\mathcal{F}) = -\frac{1}{\ln d} \sum_{k=1}^K \frac{d_k}{d} \ln \frac{d_k}{d},$$

where we define $0 \ln 0 = 0$ for sake of avoiding singular points.

C.2 Additional Experimental Results

In this section we report additional results of the reductions DEFRAG variants offer on training and prediction time and model size along with the effect on prediction accuracy. We also report here results on using the nDCG-based splitting method DEFRAG-N, as well as additional results on the performance of the FIAT algorithm on settings with missing features.




Figure 5: The figure shows the performance of FIAT on two different datasets Wiki10 (first row) and Delicious-200K (second row) with respect to precision at 1, 3 and 5. A certain fraction (as indicated in the x-axis) of features that were present in each of the test data points was randomly removed. Thus, if a test data point had \hat{d} of the d features with non-zero values in its feature representation, then at 50% noise level, $\hat{d}/2$ of these features were removed and their value set to zero. FIAT is able to offer better resilience to removal of features from data points, especially when noise levels go higher. At low noise levels, FIAT is competitive with Parabel but as noise levels grow, the performance gap between the two methods expands.

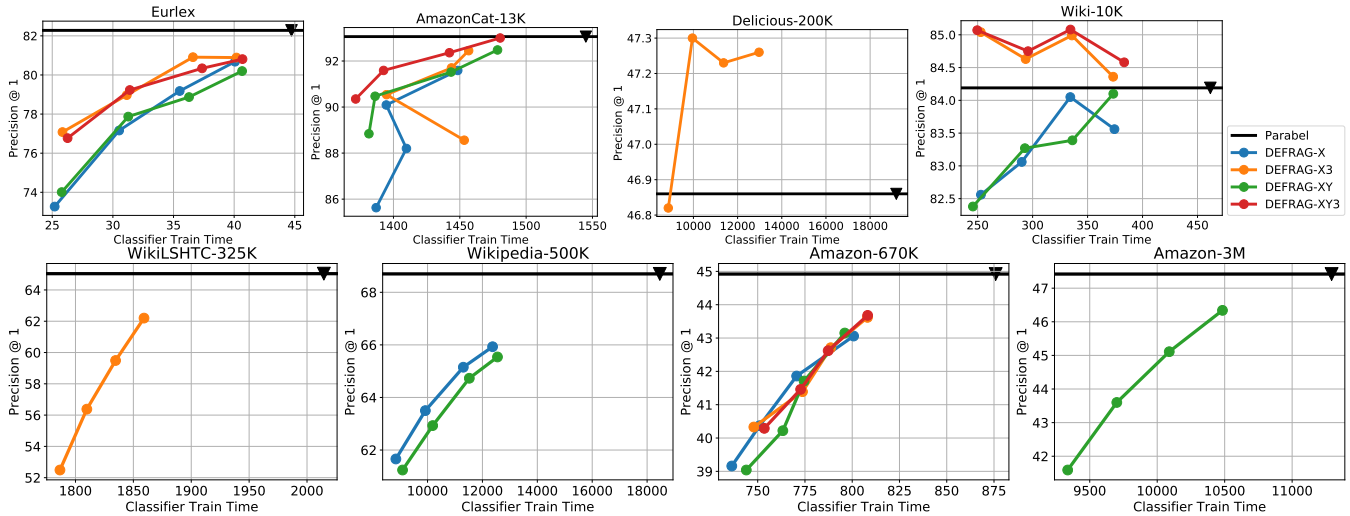


Figure 6: Effect of DEFRAG variants on reducing classifier training time. DEFRAG variants (DEFRAG-X, DEFRAG-XY, ensemble, no-ensemble) were executed with varying number of clusters by setting the minimum leaf size parameter (see Algorithm 1). The different markers in the figures correspond to $d/32$, $d/16$, $d/8$, $d/4$ balanced clusters. The black line shows Parabel's default performance with the black triangle marking its classifier training time. Aggressive clustering leads to much faster training times and smaller model sizes but also cause a drop in performance. DEFRAG-X3 and DEFRAG-XY3 refer to an ensemble of 3 independent realizations of DEFRAG. DEFRAG variants are able to bring about significant reductions in the training time of classifiers, for instance Delicious (more than 35% reduction), Wikipedia-500K (more than 33% reduction), Wiki10 (more than 20% reduction).

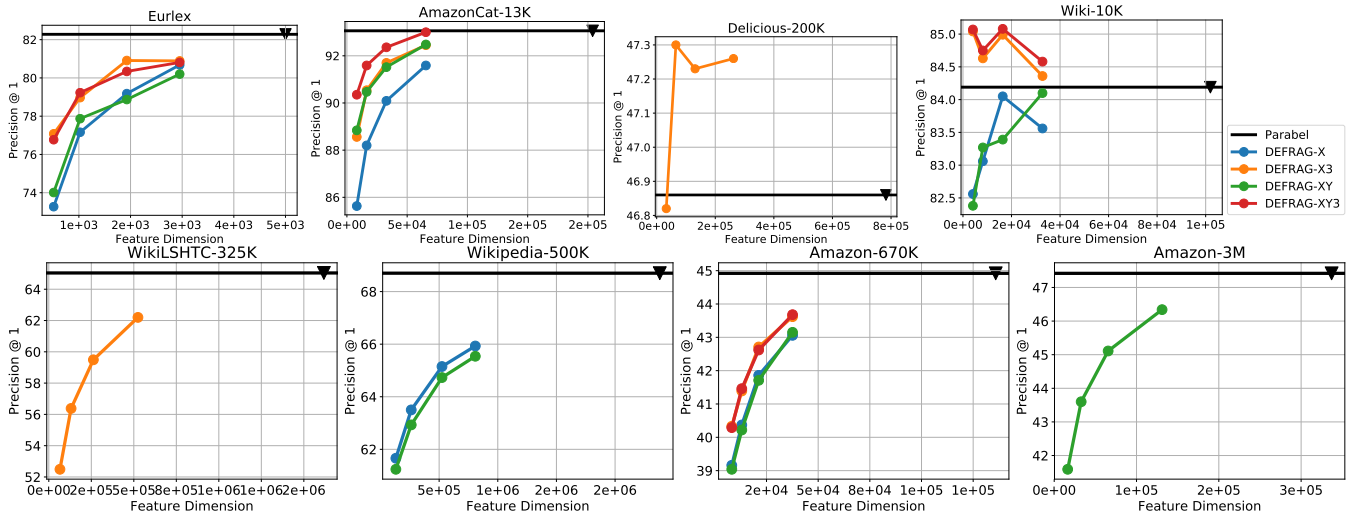


Figure 7: Effect of DEFRAG variants on reducing feature dimensionality. In all cases, DEFRAG variants are able to offer drastic reductions in feature dimensionalities of the datasets without significant reduction in prediction accuracies, as measured by precision@1, for example WikiLSHTC (4x reduction), AmazonCat (4x reduction with the same prediction accuracy), Delicious200K (16x reduction with even higher prediction accuracy), Amazon-3M (8x reduction).

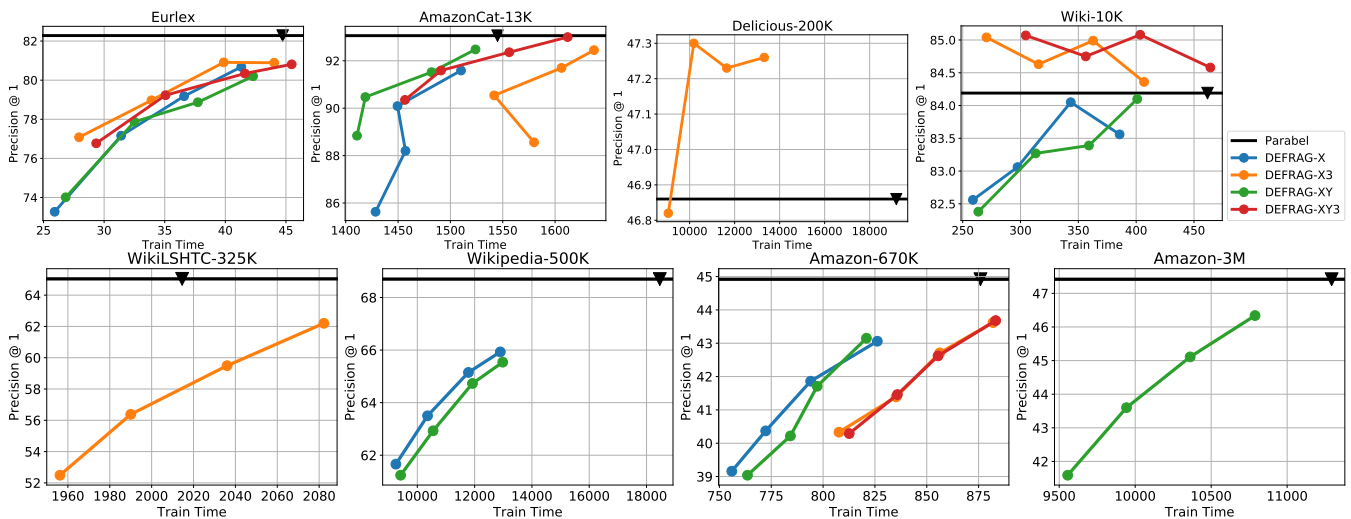


Figure 8: Effect of DEFRAG variants on reducing total training time. The total training time includes the classifier training time, as well as DEFRAG's clustering time. Note that the total times are very close to the classifier training time since the hierarchical clustering techniques adopted by DEFRAG operate very efficiently and put very little overhead on the total training time.




Figure 9: Effect of DEFrag variants on reducing prediction time. Note that the prediction time includes the feature agglomeration time as well as the actual prediction time. The performance of the classifier is indicated using the precision@5 metric. The feature agglomeration step being so inexpensive, allows DEFrag to offer drastic reductions in prediction times on most datasets, for example Delicious (70% reduction), Wiki10 (40% reduction), Wikipedia-500K (more than 40% reduction)

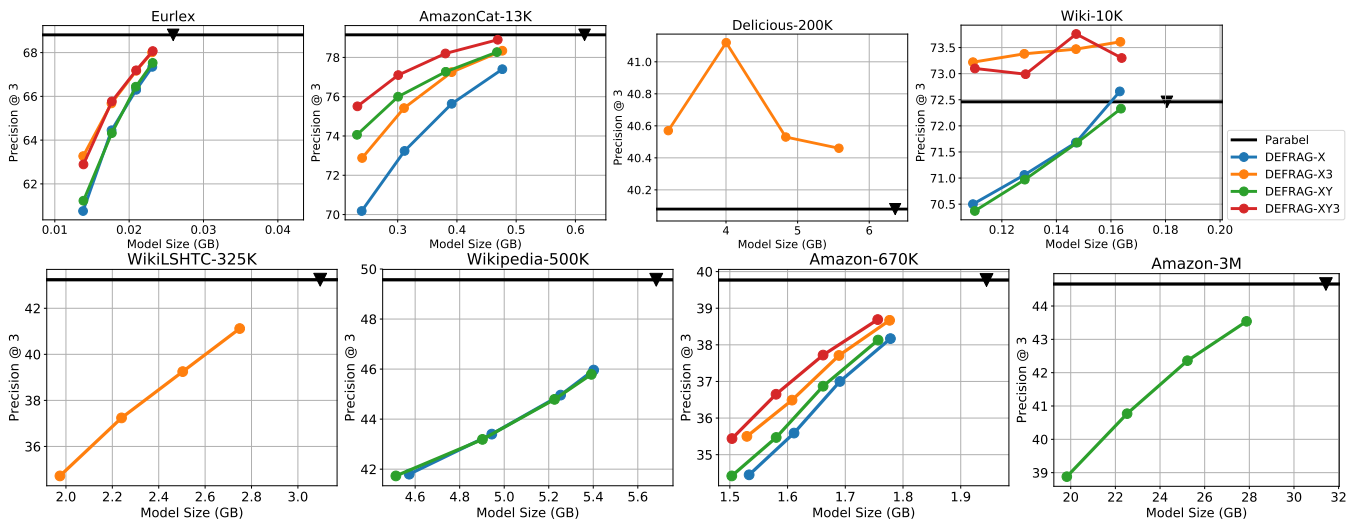


Figure 10: Effect of DEFrag variants on reducing model sizes. The performance of the classifier is indicated using the precision@3 metric. Since the model sizes of most algorithms depend on the dimensionality of the data, by performing feature agglomeration, DEFrag variants offer model size reductions as well.




Figure 11: A comparison between k-means based clustering, as followed by DEFRAG-X and nDCG based clustering, as followed by DEFRAG-N. Note that nDCG-based clustering is more expensive (see Algorithm 3). However, in several cases, e.g. Amazon-670K, AmazonCat, EURLex, nDCG clustering can offer slightly superior performance in terms of classification accuracy.