
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Schultheis, Erik; Babbar, Rohit

Towards Memory-Efficient Training for Extremely Large Output Spaces : Learning with 670k Labels on a Single Commodity GPU

Published in:
Machine Learning and Knowledge Discovery in Databases

DOI:
[10.1007/978-3-031-43418-1_41](https://doi.org/10.1007/978-3-031-43418-1_41)

Published: 01/01/2023

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Schultheis, E., & Babbar, R. (2023). Towards Memory-Efficient Training for Extremely Large Output Spaces : Learning with 670k Labels on a Single Commodity GPU. In D. Koutra, C. Plant, M. Gomez Rodriguez, E. Baralis, & F. Bonchi (Eds.), *Machine Learning and Knowledge Discovery in Databases: Research Track - European Conference, ECML PKDD 2023, Proceedings* (pp. 689-704). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 14171 LNAI). Springer. https://doi.org/10.1007/978-3-031-43418-1_41

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Towards Memory-Efficient Training for Extremely Large Output Spaces – Learning with 670k Labels on a Single Commodity GPU

Erik Schultheis¹✉ and Rohit Babbar^{1,2}

¹ Aalto University, Espoo, Finland

`erik.schultheis@aalto.fi`, `rb2608@bath.ac.uk`

² University of Bath, Bath, UK

Abstract. In classification problems with large output spaces (up to millions of labels), the last layer can require an enormous amount of memory. Using sparse connectivity would drastically reduce the memory requirements, but as we show below, applied naïvely it can result in much diminished predictive performance. Fortunately, we found that this can be mitigated by introducing an intermediate layer of intermediate size. We further demonstrate that one can constrain the connectivity of the sparse layer to be of constant fan-in, in the sense that each output neuron will have the exact same number of incoming connections, which allows for more efficient implementations, especially on GPU hardware. The CUDA implementation of our approach is provided at <https://github.com/xmc-aalto/ecml23-sparse>.

1 Introduction

In this paper, we present findings towards employing sparse connectivity in order to reduce the memory consumption of the classification layer for problems with extremely large output spaces (XMC). Such problems arise in, e.g., tagging of text documents [8], next-word predictions [21], and different kinds of recommendation tasks [1, 5, 19, 24, 29]. In order to ensure computational tractability of these tasks, which can have up to several millions of labels, one typically builds a *hierarchical label tree* [14, 23, 30, 32], only exploring branches that are likely to contain relevant labels for the current instance. Even though this is very effective at reducing the computation (from linear to logarithmic in the number of labels), it does not help in addressing the memory consumption, which is still linear in the number of labels times the number of hidden units.

As an illustration consider the AMAZON-3M [18] dataset. If we were to map the inputs to a hidden representation of 1024 units, the fully connected last layer for this dataset would need about 2.9 billion parameters, corresponding to 10.7 GiB¹. Given that modern deep learning optimizers such as ADAM [16] need to

The original version of this chapter was previously published without open access. A correction to this chapter is available at

https://doi.org/10.1007/978-3-031-43418-1_42

¹ Assuming 32-bit floating point numbers.

keep track of the value, gradient, and first and second moment, this leads to an overall peak memory consumption of over 40 GiB, making it nigh impossible to train such models on commodity hardware.

Therefore, we want to investigate possibilities for memory efficient *sparse training* of this huge last layer. There are two pre-existing approaches that serve as an indication that this is an idea that could be successful: First, for DISMEC, a *linear model* applied to tf-idf representations of input text, it is known that the resulting layer can be sparsified *after training* to contain *less than 1%* non-zeros [2]. In a linear model, the different classifiers for each label can be trained *independently*. As a result, only the full weights of the label that is currently trained needs to be kept in memory, and can be pruned as soon as the training for that label has finished. For non-linear models, the MACH [19] algorithm can be interpreted as a special case of training with *static*, random sparsity. It works by hashing the labels into different buckets, and performing training and predictions only on the level of buckets. If enough independent hashes are used, this method allows to solve the original problem in the large output space. However, in practice, the results presented for MACH are not as good as for competing methods.

The contributions of this paper are as follows: We show that naïvely applying a dynamic sparse training algorithm to the last layer of an XMC problem results in strongly reduced predictive performance. Inspired by MACH, we then propose to alleviate this problem by inserting a penultimate layer that is larger than the hidden representation of the inputs, but still much smaller than the size of the label space. Such an increased layer size drastically improves the chances of dynamic sparse training finding a good subnetwork, and enables us to get results only slightly worse than training with a dense last layer. We demonstrate this on several large-scale datasets, for which we train a classification layer on a fixed set of pre-trained features. To ensure memory efficient and quick computations, we propose to restrict the sparsity structure to *constant fan-in*, such that each unit in the output layer receives exactly the same number of inputs. This has several important consequences: (i) it makes it impossible for the training to focus most non-zero weights on a few, prominent head labels, and instead ensures a more even distribution of the representational capacity, (ii) compared to [coordinate-format](#) this requires only half the memory to store the indices, and compared to [compressed row sparse](#) matrices the data layout is simpler, making it easier to implement the corresponding operations on a GPU, and (iii) it also means that changing the sparsity structure (redistribution of connections) can be implemented as a very cheap operation.

2 Setup and Background

We consider classification problems that map an input instance $x \in \mathcal{X}$ to a subset of a label set with m labels, represented as a binary vector $\mathbf{y} \in \{0, 1\}^m$. More precisely, we assume that $(x, \mathbf{y}) \sim \mathbb{P}$ are jointly distributed according to some probability measure. If almost surely $\|\mathbf{y}\|_1 = 1$, it is a multiclass setup,

otherwise a multilabel setup. We want to find a classifier $f: \mathcal{X} \rightarrow \{0, 1\}^m$ so that predicted labels $\hat{\mathbf{y}} = f(x)$ and actual labels are close. Usually, f can be decomposed into two operations: First, the inputs are *embedded* into a fixed-size vector space using a function $\psi: \mathcal{X} \rightarrow \mathbb{R}^e$ (e.g. a linear projection, multilayer perceptron, or transformer-based text model), and then a decoding $\mathbf{W} \in \mathbb{R}^{e \times m}$ is applied to extract scores for each label. The actual prediction is then generated by selecting the k highest scoring labels as positive, $\hat{\mathbf{y}} = \text{top}_k(\mathbf{W}^\top \psi(x))$. Consequently, performance is typically measured in terms of *precision-at- k* , defined as the fraction of correct predictions

$$\text{P@}k(\mathbf{y}, \hat{\mathbf{y}}) = k^{-1} \sum_{j=1}^m y_j \hat{y}_j \quad \text{for } \|\hat{\mathbf{y}}\|_1 = k. \quad (1)$$

In order to find the optimal \mathbf{W} that maximizes $\text{P@}k$, one often performs a *One-vs-All (OvA)* reduction [2, 3, 20]: A binary classification loss ℓ is applied to each label separately. As this involves evaluating the scores $\mathbf{W}^\top \psi(x)$ for each label, many methods select a subset $\mathcal{N} \subset [m]$ of *hard negatives* [7, 12, 14, 15, 26], to approximate the sum as

$$\begin{aligned} l(\mathbf{y}, x) &= \sum_{j=1}^m \ell(y_j, \mathbf{w}_j^\top \psi(x)) = \sum_{j:y_j=1} \ell(1, \mathbf{w}_j^\top \psi(x)) + \sum_{j:y_j=0} \ell(0, \mathbf{w}_j^\top \psi(x)) \\ &\approx \sum_{j:y_j=1} \ell(1, \mathbf{w}_j^\top \psi(x)) + \sum_{j \in \mathcal{N}} \ell(0, \mathbf{w}_j^\top \psi(x)). \end{aligned} \quad (2)$$

This is very effective at reducing the required computations, and could also be beneficial for accuracy because it effectively changes the distribution of labels seen by the classifier [25], but it does not decrease the enormous amount of memory required to store the weight matrix \mathbf{W} .

There are several established approaches to handle this problem: The most straightforward method is to place a *bottleneck* layer just before the final classification layer, so that the dimension of the embedding that \mathbf{W} operates on is comparatively low. For example, LightXML [14] project the 3280-dimensional representation used for determining hard negatives down to only 300 units for the extreme-level classification. This approach is limited in its effectiveness, as too small sizes start to severely affect the classification quality. A second strategy is to *prune* the matrix \mathbf{W} after training, turning it into a very sparse matrix. This can reduce the model size to only a tiny fraction of the dense equivalent, without negatively affecting its predictive power, but this does not solve the problem of memory consumption during the training itself. The only exception are linear models, where the weight vectors \mathbf{w}_j for different labels can be trained independently, and be sparsified immediately after training, so that the full matrix never has to materialize [2, 3]. Additionally, it is possible to exploit the relation between primal and dual of linear problems to achieve sparse training for max margin classifiers with appropriate loss functions [31]. Finally, MACH [19] has shown that it is possible to train an extreme classifier on the level of meta-labels,

obviating the need for the large weight matrix \mathbf{W} altogether. However, this corresponds, implicitly, to a multiplication by a sparse, fixed, binary matrix, which therefore limits the expressiveness of the model, and it requires keeping multiple copies of the embedding network ψ .

Thus, existing sparse training methods for XMC either use post-training sparsification, or a fixed sparsity structure. Here, we want to apply the *sparse evolutionary training* (SET) algorithm [22] to the classification layer, so that we have sparse training with dynamic sparsity structure. The SET algorithm follows a general prune-redistribute-regrowth cycle, which means that periodically, a subset of existing non-zero weights is selected to be removed (*pruned*), and new structural non-zeros will be inserted (*redistributed*). After that, the training of the sparse layer proceeds just as in any other gradient-based optimization, i.e., the structural non-zeros are updated according to their mini-batch gradient (*regrown*), and the structural zeros are left unchanged, until the next cycle.

This general algorithmic structure can be implemented in various ways, depending on how the pruned weights are selected, and how it is determined where they should be re-distributed. [11] The SET algorithm uses very simple heuristics: The set of least important connections is determined by sorting according to the absolute value of their weight, and removing the fraction α of connections with lowest weight. The same number of new connections is inserted after pruning, by choosing uniformly randomly from the structural zeros.

While there exist more elaborate schemes, they are generally more complex to implement and will require additional memory. For example, [4] chooses its pruning based on weights switching their sign, which means that it needs to store the previous signs of all structural non-zeros. To determine useful locations for inserting the redistributed connections, [9] uses a momentum term, which means that this requires the same amount of memory as the weights for the original dense layer, and thus is infeasible in our setting. This also excludes any strategy that requires, even if only intermittently, a full, dense gradient, such as [10].

A naïve application of SET to the last layer leads to unsatisfactory results, and an implementation using just the available tools in `tensorflow` turns out to be suboptimal in terms of speed and memory consumption. Thus, we present in the next section some modifications to the architecture and training algorithm, as well as insights into an efficient implementation, to alleviate these shortcomings.

3 Method

In principle, implementing a sparse layer in `tensorflow`² is straightforward: Replace the dense-dense matrix multiplication with a sparse-dense operation that is supplied by the framework, and the weight matrix with a `SparseTensor`.

There are four problems with this approach: First, it wastes memory due to `tensorflow`'s requirement that all indices be given as 64-bit integers. Second, completely unstructured sparsity makes efficient implementations challenging, especially on GPUs. Third, the `tensorflow` operations cannot exploit the sparsity in the gradient signal that arises naturally when training with hinge-like

² PyTorch still considers its sparse tensor support to be in beta.

losses. Finally, replacing the dense layer with a highly sparse layer results in underfitting. We will address these problems below.

3.1 Efficient 32-Bit Indexing

In `tensorflow`, sparse tensors are represented in *coordinate* (COO) format (Fig. 1a), which means that each structural nonzero in a sparse matrix is described by three numbers. Two 64-bit integers define the row and column of the structural nonzero, and a 32-bit floating point number its value. This means that a single sparse weight requires as much memory as five weights in the dense matrix.

Even for extreme-scale classification, however, 32-bit integers would be more than sufficient as column and row indices of \mathbf{W} . A maximum representable value of around 4 billion is still an order of magnitude larger than even very large scale proprietary problems [19] with 100s of millions of labels, and three orders of magnitude larger than publicly available benchmark datasets.

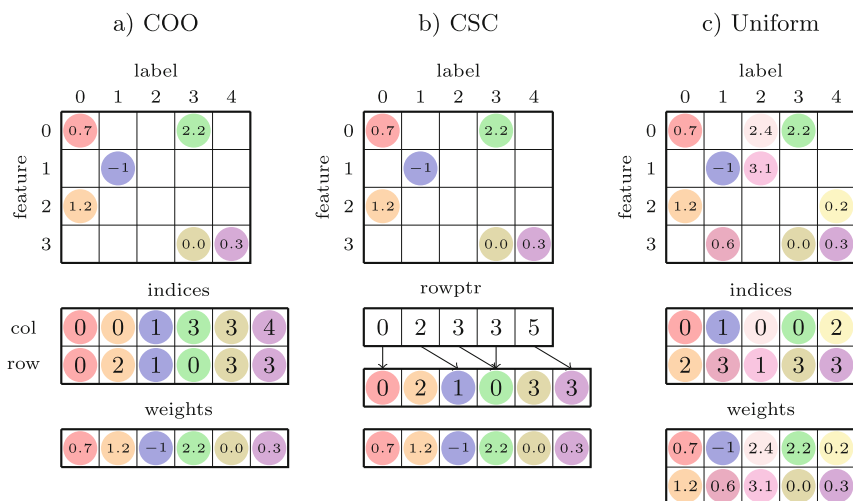


Fig. 1. Schematic depiction of different sparse matrix formats. Note that in COO format (a), the `indices` array in Algorithm 1 is of shape $2 \times \text{nnz}$. In uniform format (c) it is $\text{nnz per column} \times \text{labels}$, and hence only half as big, compared to the COO format, for the same number of nonzeros.

3.2 Compressed Indexing and Equitable Work Distribution Through Constant Fan-In

Even with 32-bit indices, a sparse weight still consumes three times as much memory as a dense weight, when represented in coordinate format. This could be made much more efficient by switching to *compressed sparse column* (CSC) format, where only row indices are saved directly, and for each column only

the offset of its first index is stored (Fig. 1b). While this drastically reduces the amount of memory needed to store the indices, it also increases the complexity of involved computations. For example, in COO format, one can assign each GPU thread to the same amount of structural non-zeros to handle during the matrix multiplication, as getting the corresponding row and column indices is a simple array lookup. In contrast, in CSC format, it is still trivial to assign one column to each thread (i.e., each thread will compute one output), but that can lead to a significant difference in the amount of work each thread has to do, and thus lead to inefficient use of GPU resources. Furthermore, redistribution becomes more involved, as inserting a new structural nonzero in an early column means that all the weights and indices that come after have to be shifted.

This can be simplified if we stipulate that each column should have the exact same amount of structural non-zeros, such that $\forall j : \|\mathbf{w}_j\|_0 = s$. Then, a single index array is sufficient, and the starting offset of each column can be calculated simply by multiplying the number of non-zeros per column with the column index, like in regular multidimensional array indexing (Fig. 1c). Distributing a multiplication with a *constant fan-in* sparse matrix across many threads is also easy, as we can simply assign one column (i.e., \mathbf{w}_j) to each thread, knowing that they correspond to the same amount of work. Finally, connection redistribution is cheaper, because the number of non-zeros stays constant for each column, and thus changes in one column never require moving around the data of other columns. As we will show in Sect. 4.2, the additional constraint on the number of connections per output does not negatively influence the models predictive performance in the overparametrized regime.

Broadly, the implementation works as follows: The sparse weights are represented by two matrices, `indices` $\in \mathbb{N}^{s \times m}$ and `weights` $\in \mathbb{R}^{s \times m}$. The input is given as a matrix `features` $\in \mathbb{R}^{b \times e}$, where b denotes the batch size, and the output is a matrix `output` $\in \mathbb{R}^{b \times m}$. CUDA threads are generated on a two dimensional grid, with one thread for each output. Thus, threads will be indexed by pairs, each of them consisting of `instance` $\in [b]$ and `label` $\in [m]$. Every thread performs the calculations given, schematically, in Algorithm 1.

Algorithm 1. Calculation of the score for a single label `label` and `instance` for uniform sparsity (see Fig. 1c) with `s` non-zeros for each label.

```

value = 0;
for weight_idx in range(s):
    source = indices[weight_idx, label]
    feature = features[instance, source]
    value += feature * weights[weight_idx, label]
output[instance, label] = value

```

3.3 Speeding up Backward Pass Through Implicit Negative Mining

Our experiments with a sparse last layer showed that the largest fraction of time was spent in the backward pass. This is not surprising, as the backward pass requires two sparse matrix multiplications: to calculate the gradient with respect to the inputs, and to calculate the gradient with respect to the weights.

Fortunately, certain margin-based losses can induce high amounts of sparsity in the gradient of XMC problems, which can be exploited to ensure considerable speed-up [27,31]. In the given enormous label space, each instance will have only a tiny subset of labels which are relevant to it, and many for which the decision that they are not relevant is “easy”. Thus, if the loss function gives zero penalty for these easy classifications (e.g., if the margin is large enough in hinge-like losses), then the error term to be back-propagated will be highly sparse. For the loss function that is mainly used in this paper, the squared-hinge loss $\ell(y, \hat{y}) = \max(0, 1 - y\hat{y})^2$, the gradient is $\partial\ell/\partial\hat{y} = -2y \max(0, 1 - y\hat{y})$, and thus exactly zero whenever $y\hat{y} \geq 1$.

Therefore, in the backward kernel, it becomes beneficial to explicitly check whether the backpropagated signal $\partial\ell/\partial\hat{y}$, denoted by `backward` $\in \mathbb{R}^{b \times m}$ in the algorithm, is already zero, and if so to skip the corresponding operations. In particular, this means not only that the multiplication with zero can be skipped, but also makes it unnecessary to load the second operand and to store the result. As sparse matrix operations are memory-bound, this can be highly beneficial.

In fact, if we distribute the threads in the same way as the forward pass for the calculation of the gradient with respect to the features (one thread assigned for each `label` and `instance`) then most threads can be skipped entirely.³ A schematic of the resulting implementation is given in Algorithm 2. Because multiple labels can contribute to the gradient of each input feature, in this case several threads need to update the same part of the gradient array. Therefore, we have to resort to using atomic addition operations here.

For calculating the gradient of the weight values, it is possible to arrange threads so that they can act independently, by using one thread for each gradient entry, i.e., for each `label` $\in [m]$ and `weight_idx` $\in [s]$. In this case, one cannot skip entire threads, but a zero in the backward signal still allows to skip the unpredictable, indirect memory lookup of `feature = features [instance, source]`, as shown in Algorithm 3.

3.4 Mitigating Underfitting by Adding an Intermediate Layer

Finally, we noticed that—even without constant fan-in—replacing the dense layer with a sparse layer results in diminished classification accuracy, which we

³ On a GPU, skipping a single thread might not be helpful, as threads are executed together in groups of 32 as a warp. However, with the very high level of sparsity in the backward signal, it becomes common that all threads within a warp can be skipped.

Algorithm 2. Contribution to the gradient for the input features caused by a given label and instance in the mini-batch.

```

out = backward[instance, label]
if out == 0:
    return

for weight_idx in range(s):
    source = indices[weight_idx, label]
    weight = weights[weight_idx, label]
    atomicAdd(gradient[instance, source], weight * out)

```

Algorithm 3. Calculation of the gradient for a given structural non-zero weight.

```

source = indices[weight_idx, label]
result = 0
for instance in range(batch_size):
    out = backward[instance, label]
    if out == 0: continue

    feature = features[instance, source]
    result += feature * out;
gradient[weight_idx, label] = result

```

attribute to underfitting. Thus, we propose to improve the expressiveness of the model by adding an intermediate layer between the embedding layer and the final classification layer. Because the last layer is sparse, its memory consumption is independent of the size of the preceding layer. Consequently, as long as this new intermediate layer is at least an order of magnitude smaller than the number of labels, this does not impede our goal of reducing memory requirements.

4 Experiments

In this section, we provide the experimental evidence showing that sparse last layers are a viable approach to extreme multilabel classification. We run experiments with several well-known benchmark datasets, measuring duration and peak GPU memory consumption, as well as P@k. After presenting results that justify the architectural choices we made, we provide additional data illustrating the trade-offs between memory consumption and classification accuracy by varying the sparsity and size of the intermediate layer. Then we present investigate the effect of implicit negative mining. The section concludes with a discussion of the results. Additional experiments are given in the supplementary at <https://github.com/xmc-aalto/ecml23-sparse>.

4.1 Experimental Setup

In this paper we focus on the setting of learning from fixed, low-dimensional representations of the instances. This enables us to do many more experiments than if we had to fine-tune an expensive transformer-based encoder for each run.

We use two different sources for the embeddings: 512-dimensional fast-text based representations as used for SLICE [12], and the final classification embeddings from a trained CASCADEXML [15] model with 768 dimensions. We present results on two datasets, [33], AMAZON-670K [17], and WIKIPEDIA-500K [6].

To update the network’s weights, we use the ADAM optimizer [16] with an initial learning rate of 1×10^{-3} that is decayed by 1/2 whenever validation P@3 stops improving, until reaching 1×10^{-4} . After that, training is stopped once P@3 stops increasing. For sparse layers, we initialize the connections uniformly randomly, potentially subject to the constraint that each label gets the same amount of connections. Every 1000 training steps, each consisting of 32 samples in a minibatch, the 10% lowest-magnitude weights are randomly redistributed. In order to mitigate overfitting, we apply dropout to the input features, dropping 10% for AMAZON-670K and WIKIPEDIA-500K-SLICE features, and 20% for WIKIPEDIA-500K-CASCADE.

The experiments are run on a NVIDIA V100. Even though we want to demonstrate the feasibility of XMC learning on a commodity GPU, in order to be able to make meaningful comparisons, we have to train on the same GPU for all settings, which means that the GPU needs to have enough memory to fit in a dense last layer. To quantify the memory benefits of sparse training, we record the peak memory consumption as reported by tensorflow (`tf.config.experimental.get_memory_info("GPU:0")['peak']`). Note, in particular, that all cases with our proposed architecture consume significantly less than 4 GiB of GPU memory, and thus will be feasible, albeit training more slowly, on cheap gaming GPUs.

4.2 Results with Varying Architecture

As a first step, we want to show that the architectural choices described in Sect. 3 are useful. To that end, we compare the training with a dense last layer to the following settings:

- A single, unstructured sparse layer,
- A single, constant fan-in sparse layer,
- An intermediate, dense layer, followed by an unstructured sparse layer,
- An intermediate, dense layer, followed by a constant fan-in sparse layer.

The number of structural non-zeros is chosen such that in the UNSTRUCTURED sparse layers, there are an average of 32 connections per label, and in the CONSTANT-FAN-IN sparse layers there are exactly 32 connections per label. As a baseline with comparable memory consumption, we also trained a BOTTLENECK architecture, that maps the input representation to a low-dimensional space of only 64 dimensions, before projecting into the label space.

Table 1. Comparison of different network architectures. **Con** denotes the (average) number of connections per label, **Int** the intermediate layer’s size, **Mem** the peak GPU memory consumption, **Eps** the number of training epochs, and **Time** the duration of a single epoch in seconds. Bold marks the best results in any sparse setting.

Setup	Test					Train			Mem.	Eps.	Time
Sparsity	Con.	Int.	P@1	P@3	P@5	P@1	P@3	P@5	GiB		sec
<u>WIKI500K-SLICE</u>											
DENSE	512	–	58.2	37.9	28.0	97.3	77.5	60.4	6.7	39.4	1 249
UNSTRUCTURED	32	–	45.5	27.3	19.9	58.3	37.7	28.1	4.8	78.0	3 612
CONSTANT-FAN-IN	32	–	37.5	23.2	17.6	42.6	28.0	21.9	0.7	54.8	659
UNSTRUCTURED	32	32k	59.0	38.5	28.9	83.7	61.4	47.7	4.8	40.0	3 977
CONSTANT-FAN-IN	32	32k	58.9	38.4	28.9	84.2	62.2	48.4	1.0	45.8	723
BOTTLENECK	64	64	56.5	36.5	27.5	71.8	50.0	38.5	1.0	41.8	639
<u>WIKI500K-CASCADE</u>											
DENSE	768	–	77.2	58.6	45.1	96.7	79.7	64.2	10.0	25.6	1 744
UNSTRUCTURED	32	–	65.2	43.7	31.4	78.3	54.7	39.8	4.8	100.0	3 870
CONSTANT-FAN-IN	32	–	58.7	42.0	32.2	69.1	51.8	40.4	0.7	59.4	715
UNSTRUCTURED	32	32k	73.7	54.7	42.0	92.4	73.5	58.1	4.9	58.0	4 423
CONSTANT-FAN-IN	32	32k	73.6	54.8	42.1	93.0	74.3	58.9	1.0	67.4	842
BOTTLENECK	64	64	71.9	50.7	37.9	86.4	64.8	49.5	1.0	47.6	678
<u>AMAZON670K-SLICE</u>											
DENSE	512	–	33.8	29.6	26.6	99.2	93.9	88.4	9.0	27.2	472
UNSTRUCTURED	32	–	14.5	11.5	9.5	64.8	49.4	39.0	6.4	73.0	1 357
CONSTANT-FAN-IN	32	–	7.1	6.3	5.6	16.2	13.9	12.4	1.0	24.8	223
UNSTRUCTURED	32	32k	32.7	28.7	25.8	98.8	93.4	87.5	6.4	45.0	1 619
CONSTANT-FAN-IN	32	32k	32.8	28.7	25.9	98.7	93.2	87.3	1.2	38.0	244
BOTTLENECK	64	64	30.7	27.3	24.6	96.4	88.9	80.1	1.1	33.6	219
<u>AMAZON670K-CASCADE</u>											
DENSE	768	–	47.5	42.3	38.3	99.8	94.5	89.0	13.4	28.4	624
UNSTRUCTURED	32	–	30.4	23.8	19.0	88.8	71.4	55.6	6.3	95.0	1 369
CONSTANT-FAN-IN	32	–	37.1	31.6	27.6	92.4	84.0	74.8	1.0	76.2	234
UNSTRUCTURED	32	32k	42.5	37.1	33.0	99.7	94.3	88.7	6.5	36.0	1 512
CONSTANT-FAN-IN	32	32k	42.6	37.1	33.1	99.7	94.3	88.7	1.4	36.4	271
BOTTLENECK	64	64	38.0	33.7	30.4	99.1	93.3	86.4	1.1	31.6	232

We repeated each experiment five times and report the average, except for the extremely slow settings with unstructured sparsity, which we ran only once.

The results of these experiments are presented in Table 1. Several facts are immediately obvious from the recorded data: First, the naive, tensorflow-based implementation for UNSTRUCTURED sparsity is very slow, to the degree that the sparse matrix multiplication ends up being 2-3× slower than dense multiplication on the large datasets. Second, with the intermediate layer, the classification performance of constant fan-in and unstructured sparsity is almost identical.

Third, without an intermediate layer, there is a significant drop in P@3, both in training and test performance, showing that naïve sparsification leads to severe overfitting. In cases where the unstructured sparse layer fails to perform well even on the training set, the additional constraint does lead to a further drop in performance. The BOTTLENECK baseline outperforms the direct sparse layer, but is significantly weaker than the combination of sparse and intermediate layer.

The measurements further show that for training based on SLICE features, the sparse implementation manages to attain and slightly surpass the classification performance of the equivalent dense layer, whereas for CASCADE features there still remains a noticeable gap between dense and sparse training. As a first possible explanation, one might argue that CASCADE features have been specifically trained so that they work well with a linear extreme classification layer, whereas SLICE are more general features. Therefore, it is not the sparse realizations that perform better, but instead the dense setting that performs disproportionately worse for SLICE features, as it does not have the benefit of the additional intermediate layer that allows non-linear classification boundaries. This argument does not hold up, though, as both features result in comparable model performance on the training set—it is the *generalization gap* that is much increased with SLICE features.

Looking at the memory consumption, we can see that sparsification of the last layer does lead to a noticeable reduction, but only becomes really effective if we use our implementation of constant fan-in sparsity. In this case, the memory consumption reduces to between one third and on tenth of the dense equivalent.

4.3 Results with Varying Network Size

In Table 2, we demonstrate the effect of varying the number of connections per label, and the size of the intermediate layer, for the uniformly sparse setup. Unsurprisingly, increasing the network size results in improved classification performance. For SLICE features, the sparse network can be considerably better than the dense counterpart. For CASCADE features, increasing the size of the sparse layer provides a way of shrinking the gap between sparse and dense performance, while still remaining much more memory efficient than the dense setup. In particular for WIKIPEDIA-500K, the change in memory consumption is only by a few percent, while the improvement in P@k is substantial. Except for AMAZON-670K with CASCADE features, increasing the model size results in reducing the number of training epochs.

The data also shows a clear qualitative difference between AMAZON-670K and WIKIPEDIA-500K: For AMAZON-670K, switching from dense to sparse does not lead to a noticeable decline in the ability of the classifier to fit the training set, whereas for WIKIPEDIA-500K the drop is dramatic, especially in the case of SLICE features. This suggests that for the smaller AMAZON-670K (490 449 instances), even the sparse architectures are overparametrized enough to interpolate the training set, whereas for WIKIPEDIA-500K (1 813 391 instances), this is no longer the case, especially for the smaller sparse models.

Table 2. Train and test P@k on AMAZON-670K with varying sparsity and intermediate-layer size, relative to dense performance. Results of a single run.

Setup			Test			Train			Mem.	Eps.	Time
Sparsity	Con.	Int.	P@1	P@3	P@5	P@1	P@3	P@5	GiB		sec
<u>SLICE FEATURES</u>											
DENSE	512	–	33.8	29.6	26.6	99.2	93.9	88.4	9.0	27.2	472
CONSTANT-FAN-IN	32	16k	–2.0	–1.8	–1.6	–1.0	–1.5	–2.4	1.1	42.0	259
CONSTANT-FAN-IN	32	32k	–1.0	–0.9	–0.7	–0.5	–0.7	–1.0	1.2	38.0	244
CONSTANT-FAN-IN	32	65k	–0.1	0.1	0.3	–0.2	–0.3	–0.4	1.3	36.0	309
CONSTANT-FAN-IN	32	100k	0.5	0.6	0.8	–0.1	–0.2	–0.2	1.8	35.0	302
CONSTANT-FAN-IN	64	16k	–0.6	–0.4	–0.2	–0.1	–0.1	–0.2	1.9	33.0	301
CONSTANT-FAN-IN	64	32k	0.2	0.3	0.5	–0.1	–0.1	–0.1	2.2	32.0	314
CONSTANT-FAN-IN	64	65k	0.8	0.9	1.1	–0.1	–0.1	–0.1	2.5	30.0	396
CONSTANT-FAN-IN	64	100k	1.3	1.4	1.5	–0.1	–0.2	–0.2	2.6	29.0	411
<u>CASCADE FEATURES</u>											
DENSE	768	–	47.5	42.3	38.3	99.8	94.5	89.0	13.4	28.4	624
CONSTANT-FAN-IN	32	16k	–6.2	–6.3	–6.4	–0.3	–0.5	–1.0	1.2	34.0	270
CONSTANT-FAN-IN	32	32k	–4.9	–5.1	–5.2	–0.1	–0.2	–0.4	1.4	36.4	271
CONSTANT-FAN-IN	32	65k	–3.8	–3.8	–3.9	–0.1	–0.1	–0.2	1.7	39.0	305
CONSTANT-FAN-IN	32	100k	–2.8	–3.0	–3.0	–0.1	–0.2	–0.3	2.4	34.0	334
CONSTANT-FAN-IN	64	16k	–4.2	–4.2	–4.1	–0.0	–0.1	–0.1	2.1	27.0	290
CONSTANT-FAN-IN	64	32k	–3.3	–3.3	–3.2	–0.0	–0.1	–0.1	2.4	31.0	306
CONSTANT-FAN-IN	64	65k	–2.3	–2.4	–2.4	–0.1	–0.1	–0.2	2.5	33.0	391
CONSTANT-FAN-IN	64	100k	–1.9	–1.9	–1.9	–0.1	–0.1	–0.2	2.9	31.0	435
CONSTANT-FAN-IN	72	65k	–2.3	–2.4	–2.3	–0.0	–0.1	–0.2	2.7	31.0	440

4.4 Quantifying the Effect of Implicit Negative Mining

Next, we show that the implicit negative mining effect discussed above can have a significant impact on the speed of training. To that end, we use the small model configuration with constant fan-in sparsity with 32 structural non-zeros per output and 16k intermediate units, and train it once using the squared hinge loss (SQH) and once using binary cross-entropy (BCE) loss function. As the BCE loss only goes to zero asymptotically, this means that there will not be many explicit zeros in the signal being back-propagated through the sparse layer, and thus all labels have to be processed.

As shown in Table 3, this has a strong effect on the training time per epoch: The implicit negative mining with SQH reduces the duration by about one third. Additionally, the squared hinge loss results in slightly better P@k, and fewer training epochs.

Table 3. Comparison of training with square hinge loss and binary cross-entropy.

Setup dataset	features	loss	Test			Train			Mem. GiB	Eps.	Time sec
			P@1	P@3	P@5	P@1	P@3	P@5			
WIKIPEDIA-500K	SLICE	SQH	58.0	37.7	28.4	80.3	58.2	45.1	0.9	59.0	946
WIKIPEDIA-500K	SLICE	BCE	57.1	37.1	28.0	77.0	53.7	41.1	1.0	52.0	1 121
WIKIPEDIA-500K	CASCADE	SQH	73.1	54.2	41.5	90.7	71.4	56.3	0.9	68.0	746
WIKIPEDIA-500K	CASCADE	BCE	71.6	52.4	40.1	89.2	68.6	53.4	1.0	79.0	1 247
AMAZON670K	SLICE	SQH	31.7	27.9	25.0	98.2	92.5	86.0	1.1	42.0	259
AMAZON670K	SLICE	BCE	30.9	27.1	24.4	96.2	89.6	82.0	1.2	54.0	400
AMAZON670K	CASCADE	SQH	41.3	35.9	31.9	99.5	94.0	88.0	1.2	34.0	270
AMAZON670K	CASCADE	BCE	38.4	33.2	29.4	98.2	92.3	85.3	1.3	63.0	423

4.5 Discussion

The results above show that sparsification of the extreme layer is possible without a strong decrease in classification performance, relative to a dense layer. However, it has to be noted that training the dense layer in the common experimental protocol employed here yields worse results than reported state-of-the-art for the same set of features. Thus, even in cases where the sparse architecture outperforms the dense layer, reported results from the literature are still better.

In Table 4, we present the results from SLICE [12] and CASCADE [15], compared against our largest setting with 64 nonzeros per label and 65k intermediate units. Compared to these methods, ours performs up to 4% worse, trading off a little classification accuracy versus a multifold reduction in memory consumption. For example, CASCADE runs for over a day on two NVIDIA A100 GPUs.

Table 4. Comparison of sparse results with state-of-the-art.

Dataset	Method	SLICE			CASCADE		
		P@1	P@3	P@5	P@1	P@3	P@5
WIKIPEDIA-500K	Literature	62.6	41.8	31.6	77.0	58.3	45.1
WIKIPEDIA-500K	Ours	60.5	39.8	29.8	74.5	56.0	43.2
AMAZON-670K	Literature	37.8	33.8	30.7	48.8	43.8	40.1
AMAZON-670K	Ours	34.6	30.5	27.7	45.3	39.8	35.9

5 Conclusion and Outlook

In this paper, we have shown that it is possible to replace an extreme-scale dense classification layer with a memory-efficient sequence of an intermediately-sized layer followed by a constant-fan-in sparsely connected layer, without a strong drop in classification performance, and in some cases even improved P@k.

The experiments performed so far investigate sparse layers in the context of a simple training procedure: Learning with the full label space, from fixed, pre-trained features. To achieve feature-parity with existing approaches, this needs to

be extended to allow for end-to-end training, where the featurizer ψ is learned jointly with the classifier. Secondly, even though the implicit negative mining effect allows to reduce the computation for the backward pass to be sub-linear in the overall number of labels, it still requires a full forward pass. In order to get to competitive training times, one thus has to integrate also explicit negative mining into the training pipeline. Finally, the datasets used in this work still do not exceed millions of labels.

We performed some initial experiments using AMAZON-3M [18], which indicate a decrease of memory consumption from 63 GiB to 12 GiB, at the cost of about 5% decrease in precision. While this is still too much memory consumption for cheap gaming GPUs, it is still well within the parameters of common workstation units. A more thorough investigation of this dataset is planned for future work.

We believe that this paper provides a good foundation, from which these goals can be achieved: First, by having the sparse multiplication implemented as a regular tensorflow layer, it can be readily included in a more general model, and automatic differentiation will ensure correct gradient calculations. Second, because we are constraining the sparsity to have constant fan-in, selecting a subset of labels for which scores shall be calculated becomes a trivial matrix slicing operation, similar to the fully-connected case. In follow-up works, we aim to incorporate our approach into existing end-to-end deep extreme classification frameworks while benefiting from explicit negative mining. Furthermore, from a statistical perspective, it is possible that constant-fan-in sparsity also leads to a better coverage of tail-labels, and improvements in the corresponding metrics [13,28], which should be investigated.

Acknowledgements. We acknowledge the support of computational resources provided by the Aalto Science-IT project, and CSC IT Center for Science, Finland. This work is funded in part by the Academy of Finland projects 347707 and 348215.

Ethical Statement. Our work does not provide new datasets or introduce conceptually new learning setups. Instead, we focus on enabling existing methods to run with much reduced hardware requirements. This will make large-scale classification more readily available also to people and organizations that cannot afford to buy latest-grade GPU hardware with enormous RAM, as they can reuse pre-existing and cheaper hardware.

References

1. Agrawal, R., Gupta, A., Prabhu, Y., Varma, M.: Multi-label learning with millions of labels: recommending advertiser bid phrases for web pages. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 13–24 (2013)
2. Babbar, R., Schölkopf, B.: Dismec: distributed sparse machines for extreme multi-label classification. In: Proceedings of the tenth ACM International Conference on Web Search and Data Mining, pp. 721–729 (2017)
3. Babbar, R., Schölkopf, B.: Data scarcity, robustness and extreme multi-label classification. *Mach. Learn.* **108**(8–9), 1329–1351 (2019)

4. Bellec, G., Kappel, D., Maass, W., Legenstein, R.: Deep rewiring: training very sparse deep networks (2017)
5. Beygelzimer, A., Langford, J., Lifshits, Y., Sorkin, G., Strehl, A.L.: Conditional probability tree estimation analysis and algorithms (2014)
6. Bhatia, K., et al.: The extreme classification repository: multi-label datasets and code (2016). <https://manikvarma.org/downloads/XC/XMLRepository.html>
7. Chang, W.C., Yu, H.F., Zhong, K., Yang, Y., Dhillon, I.S.: Taming pretrained transformers for extreme multi-label text classification. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3163–3171 (2020)
8. Dekel, O., Shamir, O.: Multiclass-multilabel classification with more classes than examples. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 137–144. JMLR Workshop and Conference Proceedings (2010)
9. Dettmers, T., Zettlemoyer, L.: Sparse networks from scratch: faster training without losing performance. arXiv preprint [arXiv:1907.04840](https://arxiv.org/abs/1907.04840) (2019)
10. Evci, U., Gale, T., Menick, J., Castro, P.S., Elsen, E.: Rigging the lottery: making all tickets winners. In: International Conference on Machine Learning, pp. 2943–2952. PMLR (2020)
11. Hoeffler, T., Alistarh, D., Ben-Nun, T., Dryden, N., Peste, A.: Sparsity in deep learning: pruning and growth for efficient inference and training in neural networks. *J. Mach. Learn. Res.* **22**(1), 10882–11005 (2021)
12. Jain, H., Balasubramanian, V., Chunduri, B., Varma, M.: Slice: scalable linear extreme classifiers trained on 100 million labels for related searches. In: WSDM, pp. 528–536 (2019)
13. Jain, H., Prabhu, Y., Varma, M.: Extreme multi-label loss functions for recommendation, tagging, ranking & other missing label applications. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 935–944 (2016)
14. Jiang, T., Wang, D., Sun, L., Yang, H., Zhao, Z., Zhuang, F.: Lightxml: transformer with dynamic negative sampling for high-performance extreme multi-label text classification, vol. 35, no. 9, pp. 7987–7994 (2021)
15. Kharbanda, S., Banerjee, A., Schultheis, E., Babbar, R.: Cascadexml: rethinking transformers for end-to-end multi-resolution training in extreme multi-label classification. In: Advances in Neural Information Processing Systems (2022)
16. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
17. McAuley, J., Leskovec, J.: Hidden factors and hidden topics: understanding rating dimensions with review text. In: Proceedings of the 7th ACM Conference on Recommender Systems, pp. 165–172 (2013)
18. McAuley, J., Pandey, R., Leskovec, J.: Inferring networks of substitutable and complementary products. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794 (2015)
19. Medini, T.K.R., Huang, Q., Wang, Y., Mohan, V., Shrivastava, A.: Extreme classification in log memory using count-min sketch: a case study of amazon search with 50m products, vol. 32 (2019)
20. Menon, A.K., Rawat, A.S., Reddi, S., Kumar, S.: Multilabel reductions: what is my loss optimising? In: Advances in Neural Information Processing Systems, vol. 32 (2019)
21. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)

22. Mocanu, D.C., Mocanu, E., Stone, P., Nguyen, P.H., Gibescu, M., Liotta, A.: Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nat. Commun.* **9**(1), 2383 (2018)
23. Prabhu, Y., Kag, A., Harsola, S., Agrawal, R., Varma, M.: Parabel: partitioned label trees for extreme classification with application to dynamic search advertising. In: *Proceedings of the 2018 World Wide Web Conference*, pp. 993–1002 (2018)
24. Prabhu, Y., Varma, M.: Fastxml: a fast, accurate and stable tree-classifier for extreme multi-label learning. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 263–272 (2014)
25. Rawat, A.S., et al.: Disentangling sampling and labeling bias for learning in large-output spaces. In: *International Conference on Machine Learning*, pp. 8890–8901. PMLR (2021)
26. Reddi, S.J., Kale, S., Yu, F., Holtmann-Rice, D., Chen, J., Kumar, S.: Stochastic negative mining for learning with large output spaces. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1940–1949. PMLR (2019)
27. Schultheis, E., Babbar, R.: Speeding-up one-versus-all training for extreme classification via mean-separating initialization. *Mach. Learn.* **111**, 1–24 (2022)
28. Schultheis, E., Wydmuch, M., Babbar, R., Dembczynski, K.: On missing labels, long-tails and propensities in extreme multi-label classification. In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1547–1557 (2022)
29. Weston, J., Makadia, A., Yee, H.: Label partitioning for sublinear ranking. In: *International Conference on Machine Learning*, pp. 181–189. PMLR (2013)
30. Wydmuch, M., Jasinska, K., Kuznetsov, M., Busa-Fekete, R., Dembczynski, K.: A no-regret generalization of hierarchical softmax to extreme multi-label classification. In: *Advances in Neural Information Processing Systems*, vol. 31 (2018)
31. Yen, I.E.H., Huang, X., Ravikumar, P., Zhong, K., Dhillon, I.: PD-sparse: a primal and dual sparse approach to extreme multiclass and multilabel classification. In: *International Conference on Machine Learning*, pp. 3069–3077. PMLR (2016)
32. You, R., Zhang, Z., Wang, Z., Dai, S., Mamitsuka, H., Zhu, S.: Attentionxml: label tree-based attention-aware deep model for high-performance extreme multi-label text classification, vol. 32 (2019)
33. Zubiaga, A.: Enhancing navigation on wikipedia with social tags (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

