

## MIT Open Access Articles

*COMET: Learning Cardinality Constrained Mixture of Experts with  
Trees and Local Search*

The MIT Faculty has made this article openly available. **Please share**  
how this access benefits you. Your story matters.

**Citation:** Ibrahim, Shibal, Chen, Wenyu, Hazimeh, Hussein, Ponomareva, Natalia, Zhao, Zhe et al. 2023. "COMET: Learning Cardinality Constrained Mixture of Experts with Trees and Local Search."

**Published Version:** <https://doi.org/10.1145/3580305.3599278>

**Publisher:** ACM|Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining

**Permanent Link:** <https://hdl.handle.net/1721.1/152039>

**Version:** Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

**Terms of use:** <https://creativecommons.org/licenses/by/4.0/>



# COMET: Learning Cardinality Constrained Mixture of Experts with Trees and Local Search

Shibal Ibrahim  
Massachusetts Institute of Technology  
Cambridge, MA, USA  
shibal@mit.edu

Wenyu Chen  
Massachusetts Institute of Technology  
Cambridge, MA, USA  
wenyu@mit.edu

Hussein Hazimeh  
Google Research  
New York, NY, USA  
hazimeh@google.com

Natalia Ponomareva  
Google Research  
New York, NY, USA  
nponomareva@google.com

Zhe Zhao  
Google DeepMind  
Mountain View, CA, USA  
zhezha@google.com

Rahul Mazumder  
Massachusetts Institute of Technology  
Cambridge, MA, USA  
rahulmaz@mit.edu

## ABSTRACT

The sparse Mixture-of-Experts (Sparse-MoE) framework efficiently scales up model capacity in various domains, such as natural language processing and vision. Sparse-MoEs select a subset of the “experts” (thus, only a portion of the overall network) for each input sample using a sparse, trainable gate. Existing sparse gates are prone to convergence and performance issues when training with first-order optimization methods. In this paper, we introduce two improvements to current MoE approaches. First, we propose a new sparse gate: *COMET*, which relies on a novel tree-based mechanism. *COMET* is differentiable, can exploit sparsity to speed up computation, and outperforms state-of-the-art gates. Second, due to the challenging combinatorial nature of sparse expert selection, first-order methods are typically prone to low-quality solutions. To deal with this challenge, we propose a novel, permutation-based local search method that can complement first-order methods in training *any* sparse gate, e.g., Hash routing, Top-k, DSelect-k, and *COMET*. We show that local search can help networks escape bad initializations or solutions. We performed large-scale experiments on various domains, including recommender systems, vision, and natural language processing. On standard vision and recommender systems benchmarks, *COMET+* (*COMET* with local search) achieves up to 13% improvement in ROC AUC over popular gates, e.g., Hash routing and Top-k, and up to 9% over prior differentiable gates e.g., DSelect-k. When Top-k and Hash gates are combined with local search, we see up to 100× reduction in the budget needed for hyperparameter tuning. Moreover, for language modeling, our approach improves over the state-of-the-art MoEBERT model for distilling BERT on 5/7 GLUE benchmarks as well as SQuAD dataset.

## CCS CONCEPTS

• **Computing methodologies** → **Neural networks; Ensemble methods; Natural language processing.**



This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '23, August 6–10, 2023, Long Beach, CA, USA  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0103-0/23/08.  
<https://doi.org/10.1145/3580305.3599278>

## KEYWORDS

Sparse Mixture of Experts; Conditional Computation; Trees; Local Search

### ACM Reference Format:

Shibal Ibrahim, Wenyu Chen, Hussein Hazimeh, Natalia Ponomareva, Zhe Zhao, and Rahul Mazumder. 2023. COMET: Learning Cardinality Constrained Mixture of Experts with Trees and Local Search. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, Long Beach, CA, USA, 13 pages. <https://doi.org/10.1145/3580305.3599278>

## 1 INTRODUCTION

The Sparse Mixture of Experts (Sparse-MoE) framework has led to state-of-the-art performance in various applications such as natural language processing (NLP) [2, 10, 11, 42, 49], vision [40, 46], time-series analysis [22], multi-task learning [17, 28, 32], and multimodal learning [35]. Sparse-MoE consists of a set of trainable experts (neural networks) and a trainable sparse gate. The sparse gate in Sparse-MoE selects an appropriate subset of experts on a per-sample basis, which allows for faster computation [42] and enhances interpretability [11, 22].

The literature on Sparse-MoE has traditionally focused on Top-k gating, which selects  $k$  out of  $n$  experts using a Top-k operation [11, 42, 49]. Top-k gating is simple and efficient because it allows sparse training. However, as highlighted by prior literature [11, 17, 49], the non-continuous nature of Top-k makes it susceptible to stability and convergence issues. Alternative gating strategies exist in the literature, based on reinforcement learning [4] or post-processing via linear assignment [7, 30]. However, these strategies also face challenges in terms of efficiency and interpretability; see related work in Section 2 for more details. Random routing strategies [39, 50] alternatively bypass learning of the gating function altogether. Although computationally efficient, these strategies lead to performance degradation [7]. Recent work [17] demonstrates that differentiable gating in Sparse-MoE can improve stability and performance compared to popular non-differentiable gates. However, it suffers from expert collapse in some cases as we observed in our experiments.

In this paper, we propose two new approaches for improving routing in Sparse-MoE. First, we introduce a novel differentiable sparse gate *COMET*<sup>1</sup> that improves over existing state-of-the-art

<sup>1</sup> *COMET*: This stands for Cardinality constrained Mixture of Experts with Trees.

sparse gates [11, 17, 39, 42, 49]. Second, we argue that the combinatorial nature of expert selection in Sparse-MoE presents a serious challenge for first-order methods. In particular, the performance of these methods is highly dependent on initialization, and they can get stuck in low-quality routing solutions. Thus, we propose a new permutation-based local search method for Sparse-MoEs, which can help first-order methods escape low-quality initializations or solutions. Our local search approach is general and can be applied to any sparse gate, including Top-k [42], Hash routing [39], DSelect-k [17], and our proposed gate *COMET*.

*COMET*. Our proposed *COMET* gate is the first decision-tree-based selection mechanism for sparse expert selection — decision trees naturally perform per-sample routing (i.e., each sample follows a root-to-leaf path). Our gate has several advantages: (i) it is differentiable and can be optimized using first-order optimization methods e.g., stochastic gradient descent; (ii) it allows (partially) conditional training, i.e., dense-to-sparse training; (iii) it enforces a cardinality constraint, i.e., selects (at most)  $k$  out of the  $n$  experts; (iv) it has superior predictive performance over state-of-the-art gates such as Hash routing, Top-k, and DSelect-k.

*Local Search*. The learning problem underlying Sparse-MoEs is of combinatorial nature, which poses additional challenges compared to non-MoE machine learning models. Popularly used optimization methods, such as SGD, may lead to low-quality solutions in Sparse-MoE, as we demonstrate in our numerical experiments in Section 5. To this end, we propose a permutation-based local search method, which can help first-order methods escape bad initializations and lead to better sample routing for any sparse gate e.g., Top-k, Hash routing, DSelect-k and even *COMET*. To the best of our knowledge, we are the first to explore local search methods in the context of Sparse-MoE. We provide empirical evidence through ablation studies and large-scale experiments to demonstrate permutation-based local search (i) pushes learning towards better gate/expert initializations in early optimization stages (see Section 4.4); (ii) effectively reduces the budget needed for hyperparameter tuning by up to 100× for some popular gates e.g., Hash Routing and Top-k (see Section 5.1.3); (iii) leads to SOTA performance in terms of prediction and expert selection when combined with *COMET*, across various applications (see Section 5).

**Contributions.** As discussed earlier, it is well-known in the literature that popular sparse gates are challenging to train and may suffer from stability and performance issues. In this context, our contributions can be summarized as follows:

- We propose *COMET*, a novel tree-based sparse gate that simultaneously has the following desirable properties: (a) differentiable, (b) allows (partially) conditional training i.e., dense-to-sparse training, and sparse inference, (c) satisfies per-sample cardinality constraint (selects at most  $k$  out of the  $n$  experts per-sample, where  $k$  is a user-specified parameter).
- Popular first-order methods used to optimize Sparse-MoEs are heavily influenced by expert and gate initializations, and may get stuck in low-quality solutions. Hence, we introduce a novel permutation-based local search method that can complement first-order methods by helping them escape bad initializations or

solutions. Our local search method is general and can be applied to any gate, e.g., Hash routing, Top-k, and *COMET*.

- We perform extensive experiments on recommender systems, vision and natural language processing tasks to highlight that *COMET* and permutation-based local search can give boosts in predictive performance and reduce hyperparameter tuning.

## 2 RELATED WORK

*Sparse-Mixture-of-Experts*. The MoE framework was introduced by [24], and since then has been extensively studied — see e.g., [23, 25, 26]. More recently, [42] proposed a *Sparse-MoE* framework, based on the Top-k gate, and showed good performance on natural language processing tasks. It was further improved upon by [11, 47, 49]. However, Top-k gate does not optimize the core expert selection problem as pointed out by [7]. Additionally, as highlighted by prior literature [11, 17, 49], the non-continuous nature of Top-k makes it vulnerable to training stability and convergence issues.

With *BASE* Layers, [7, 30] formulate Sparse-MoE as an assignment problem where they post-process the gate output for balanced expert selection. [4] formulates the expert selection as a reinforcement learning problem. Others [39, 50] proposed random routing strategies that do not learn the gating function during training. These methods are also promising as they have been shown to outperform models that learn routing through Top-k, e.g., in Switch Transformers [11]. Lastly, [17] introduced DSelect-k, a differentiable gate based on binary encodings, which improves over Top-k in terms of stability and statistical performance.

*Conditional Computation*. In addition to the Sparse-MoE framework, there are other related works that also study conditional computation, i.e., the setup where only some parts of neural network are activated based on the input — see e.g., [4, 5, 20, 44]. These works rely on heuristics where the training and inference models are different. More recently, [16] introduced conditional computation in differentiable (a.k.a. soft) trees [12, 16, 18, 19, 21, 24]. Their proposal allows routing samples through small parts of the tree; thus allowing for conditional computation with customized algorithms. Our work builds upon this approach to solve the cardinality-constrained expert selection problem in Sparse-MoE. Note that [16] does not address sparse expert selection in Sparse-MoE.

*Local Search and Permutation Learning*. There is an extensive optimization literature on local search, e.g., [3, 15]. However, such methods have not been used in Sparse-MoE. Here, we survey permutation learning methods that are most relevant to our proposal. This work uses differentiable relaxations of permutation via Sinkhorn operators [1, 33]. These earlier works use these relaxations in other contexts e.g., ranking in [1] and sorting in [33]. We use permutation learning as a local search to complement first-order optimization methods to improve sample routing in Sparse-MoE.

## 3 LEARNING SPARSE MIXTURE OF EXPERTS WITH DECISION TREES

*Problem Setup of Sparse-MoE*. We first review the Sparse-MoE objective. We assume that the task has an input space  $\mathcal{X} \subseteq \mathbb{R}^p$  and an output space  $\mathcal{Y} \subseteq \mathbb{R}^u$ . Denote the  $n$ -dimensional simplex by  $\Delta_n = \{w \in \mathbb{R}^n : \sum_{i \in [n]} w_i = 1, w \geq 0\}$ . In the MoE framework, the prediction function has two components: (i) a set of  $n$  experts

(parametrized by neural networks)  $f_i : \mathcal{X} \rightarrow \mathbb{R}^u$  for any  $i \in [n] := \{1, 2, \dots, n\}$ , and (ii) a gate  $g : \mathcal{X} \rightarrow \Delta_n$  that outputs weights in the probability simplex. Given a sample  $x \in \mathcal{X}$ , the corresponding output of the MoE is a convex combination of the experts with weights  $g(x) : \sum_{i \in [n]} f_i(x)g(x)_i$ .

The goal of Sparse-MoE paradigm is to develop a gate that selects a convex combination of at most  $k$  out of the  $n$  experts. The output of the gate can be thought of as a probability vector  $g$  with at most  $k$  nonzero entries, where  $g(\cdot)_i$  is the weight assigned to the expert  $f_i$ . The underlying optimization problem (also in [17]) is:

$$\min_{f_1, \dots, f_n, g} \frac{1}{N} \sum_{(x, y) \in \mathcal{D}} \ell \left( y, \sum_{i \in [n]} f_i(x)g(x)_i \right), \quad (1a)$$

$$\text{s.t. } \|g(x)\|_0 \leq k, g(x) \in \Delta_n, \forall x \in \mathcal{X}. \quad (1b)$$

$\|g(\cdot)\|_0$  denotes the number of nonzero entries in the vector  $g(\cdot)$ ,  $\ell(\cdot, \cdot)$  is the associated loss function such that  $\ell : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}$ , and  $N$  denotes the size of training samples  $\mathcal{D} = \{(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}\}_{i=1}^N$ .


The cardinality constraint in (1b) ensures that the gate selects at most  $k$  experts. Some popular gates e.g., Top- $k$  impose exact cardinality constraint instead of an inequality constraint in (1b). However, the inequality constraint can allow for sparser expert selection as observed in prior work [17] and in our experiments (Section 5). Problem (1) is a combinatorial optimization problem that is not amenable to stochastic gradient descent due to the cardinality constraint in (1b). In the next sections, we discuss our formulation that ensures the cardinality constraint and the simplex constraints are satisfied despite optimization with gradient-based methods.

The rest of this section is organized as follows. In Section 3.1, we discuss a high-level overview of our novel tree-based framework, which equivalently sets up the cardinality-constrained objective in problem (1) as a weighted sum of decision trees. Next in Section 3.2, we provide background on a single decision tree that selects a single expert per-sample while (i) allowing for smooth optimization, and (ii) conditional computation support – routing samples to a single leaf. Later in Section 3.3, we dive deeper into our novel tree-based framework that combines such trees to satisfy the cardinality constraint for  $k \geq 1$  without violating the simplex constraint. We additionally highlight important aspects regarding leaf parameterization and regularization. Next, in Section 3.4, we discuss how our method handles settings where experts are non-powers of 2. We then discuss in Section 3.5 an implementation of *COMET* for numerically stable training.

### 3.1 Sparse-MoE with $k$ decision trees

The cardinality constrained MoE objective (1) can be formulated equivalently using a set of decision trees. Classical decision trees are naturally suited to route each sample to a single leaf with a chain of hierarchical decisions. In the case of  $k = 1$ , we propose a single decision tree to route samples, where each leaf node is associated with an expert. In cases where  $k > 1$ , we instantiate  $k$  different decision trees and combine their output in a way that enforces the cardinality and simplex constraints in (1b).

Given that classical decision trees are not amenable to differentiable training with first-order methods, we use a variant [16] of differentiable (a.k.a. soft) decision trees [16, 18, 19, 24, 27]. We build upon this work to solve the cardinality-constrained problem (1).



**Figure 1: *COMET* for 8 experts. Note  $z_q(w_q \cdot x)$  denotes the binary state  $\{0, 1\}$  for  $h(w_q \cdot x)$  achieved due to smooth-step activation function and entropic regularization.**

We first provide a summary of a single soft tree (with conditional computation support) in Section 3.2. This serves as a building block for selecting a single expert per-sample.

### 3.2 Preliminaries: Differential Decision Tree with Conditional Computation

In this section, we provide a brief summary of a variant [16] of a differentiable (a.k.a. soft) tree [18, 21, 24, 26, 27], which we use to enable single-expert selection in Sparse-MoEs. We extend it in the next section to solve the cardinality constrained problem for a general case  $k \geq 1$ .

Differentiable decision trees are similar to classical decision trees with hyperplane splits [34]. However, they route each sample to left and right with different proportions, i.e., each sample reaches all leaves. Traditionally, differentiable decision trees have been unamenable to conditional computation as they cannot route a sample exclusively to the left or to the right. Recent work [16] introduced a variant of the differentiable tree model that supports conditional computation. Here, we discuss a brief summary of this variant.

We denote a single tree by  $v : \mathcal{X} \rightarrow \Delta_n$ , which maps an input sample  $x \in \mathcal{X}$  to a probability vector  $v$  over  $\Delta_n$ . Here,  $n$  corresponds to the number of root-to-leaf paths (also equal to number of experts in the MoE paradigm). Let  $v$  be a binary tree with depth  $d$  – note our framework can naturally support cases where number of experts is non-powers of 2, see Section 3.4 for more details. Let  $\mathcal{I}$  and  $\mathcal{L}$  denote sets of the internal (split) nodes and the leaves of the tree, respectively. For any node  $q \in \mathcal{I} \cup \mathcal{L}$ , we define  $T(q)$  as its set of ancestors. Let  $\{x \rightarrow q\}$  denote that a sample  $x \in \mathbb{R}^p$  reaches  $q$ .

*Sample Routing.* Following prior work [16, 18, 27], we will discuss sample routing using a probabilistic model. While sample routing is discussed using probability, differentiable trees are deterministic. Differentiable trees are based on hyperplane splits [34], where a *linear* combination of the features is used in making routing decisions. In particular, we assign a trainable weight vector  $w_q \in \mathbb{R}^p$  with each internal node, which parameterizes the node’s hyperplane split. Let  $h : \mathbb{R} \rightarrow [0, 1]$  be an activation function. Given a sample  $x \in \mathbb{R}^p$ , the probability that internal node  $q$  routes  $x$  to the left is defined by  $h(w_q \cdot x)$ .

Now we summarize how to model the probability that  $x$  reaches a certain leaf  $l$  [16]. Let  $[l \prec q]$  (resp.  $[q \succ l]$ ) denote the event that leaf  $l$  belongs to the left (resp. right) subtree of node  $q \in \mathcal{I}$ . The probability that  $x$  reaches  $l$  is given by:  $\Pr(\{x \rightarrow l\}) = \prod_{q \in \mathcal{T}(l)} r_{q,l}(x)$ , where  $r_{q,l}(x)$  is the probability of node  $q$  routing  $x$  towards the subtree containing leaf  $l$ , i.e.,  $r_{q,l}(x) := h(w_q \cdot x)^{1[l \prec q]}(1 - h(w_q \cdot x))^{1[q \succ l]}$ . Note that the vector  $v(x)$  given by

$$v(x) = [\Pr(\{x \rightarrow l_1\}), \dots, \Pr(\{x \rightarrow l_n\})] \in \Delta_n, \quad (2)$$

defines a per-sample distribution over the  $n$  leaves (or experts).

Next, we discuss how the split probabilities  $\{h(w_q \cdot x), 1 - h(w_q \cdot x)\}$  can achieve binary state with a particular choice of activation function  $h$  — this is crucial for achieving sparse expert selection (and conditional computation) in the Sparse-MoE paradigm.

*Smooth-Step Activation Function.* The common choice for activation function  $h$  in soft tree literature is a logistic function [12, 18, 26, 27]. However, it can not perform hard routing i.e., output exact zeros. This implies that any sample  $x$  will reach every node in the tree with a positive probability, leading to a dense  $v$ . [16] proposed a smooth-step activation function for a variant of soft trees — see Appendix A for details. Despite being continuously differentiable, smooth-step activation function can produce a sparse  $v$  (after an initial warm-up period of soft routing) for hard routing. This is crucial for a sparse expert selection in Sparse-MoE paradigm. Additionally, this choice of activation function also allows for (partially) conditional training with customized sparse backpropagation algorithms in soft trees (as shown in [16]), which is an important consideration for training large-scale Sparse-MoE models.

For cardinality-constrained Sparse-MoE learning with trees (not studied in [16]), the goal for each tree is to perform hard routing for all samples. Therefore, we add additional regularization on  $\{h(w_q \cdot x), 1 - h(w_q \cdot x)\}$  to encourage convergence of  $v$  to a one-hot state (discussed in more detail in Section 3.3).

### 3.3 Cardinality constraint with $k$ trees

Next, we discuss how to achieve the cardinality constraint ( $k \geq 1$ ) in Sparse-MoE with decision trees in the presence of simplex constraint. This key ideas are given as follows:

- We consider  $k$  decision trees, where each tree  $j$  selects a single expert via  $v^{(j)}(\cdot)$  as defined in (2).
- With the experts selected as above, we need to decide the relative weights assigned to each expert. We use auxiliary functions  $\alpha^{(j)}(\cdot)$ , where  $\alpha_i^{(j)}$  is a linear function  $\beta_i^{(j)} \cdot x$  of the input.  $\alpha_i^{(j)}$  reflects a linear weighting function (in the log space) for  $i$ -th expert (or leaf) in  $j$ -th tree.

See Figure 1 as an example. Next, we define the prediction function for Sparse-MoE with  $k$  decision trees to form *COMET*.

*COMET Prediction with  $k$  Out of  $n$  Experts.* The prediction function for Sparse-MoE with  $k \geq 1$  is a weighted sum of the predictions of  $i$ -th expert (or leaf) across  $k$  trees. To this end, we define the weight for  $i$ -th expert as follows

$$g(x; \alpha, v)_i = \frac{\sum_{j \in [k]} \exp(\alpha_i^{(j)}(x)) v_i^{(j)}(x)}{\sum_{j \in [k]} \sum_{i \in [n]} \exp(\alpha_i^{(j)}(x)) v_i^{(j)}(x)}, \quad (3)$$

where  $v_i^{(j)}(x)$  is the probability that a sample  $x$  will reach expert  $f_i$  in the  $j$ -th tree. Using (3), the prediction function for Sparse-MoE with  $k \geq 1$  is given by  $\hat{y} = \sum_{i \in [n]} f_i(x) g(x; \alpha, v)_i$ .

We present the following proposition (proof in Appendix B):

**PROPOSITION 3.1.** *For any  $\alpha$ , if  $v^{(j)}$  outputs a binary vector for every  $j$ , the function  $g(x; \alpha, v)$  satisfies the cardinality and simplex constraints in (1b).*

*Accelerating Convergence of  $v^{(j)}$  to One-Hot Encoding with Entropic Regularization.* In the Sparse-MoE setup, the goal is to achieve a one-hot vector state for  $v^{(j)}$  quickly — this ensures the cardinality constraint (i.e., to select at most  $k$  experts) is respected by the  $k$  trees. To encourage faster convergence towards a one-hot vector, we add a per-tree entropy regularizer,  $\lambda \Omega(v^{(j)}(x))$  to the loss objective, where  $\Omega(v^{(j)}(x)) = -\sum_{i \in [n]} v_i^{(j)}(x) \log(v_i^{(j)}(x))$ . Entropy regularizers are used in [17, 33] to get binary representations.

*Dense-to-Sparse Learning.* *COMET* supports conditional training only partially. At the start of training, it uses all the available experts as  $v^{(j)}$  is completely dense, so conditional training is not possible. As training proceeds,  $v^{(j)}$  becomes sparser due to smooth-step function and entropic regularization, eventually achieving binary state. From this stage onwards, the gate satisfies the cardinality constraint per-sample, i.e., each sample gets routed to at most  $k$  experts. Hence, sparse training can proceed to refine the solution quality. Empirically, we observe that a small number of epochs are sufficient for the optimizer to reach the sparse training phase.

### 3.4 Non-powers of 2

Typically, in Sparse-MoE, each expert is assigned to a separate machine for efficiency [11, 49]. This may mean that the number of experts could be defined by the number of machines — machines may not necessarily be available in powers of 2. Our gate naturally handles cases where the number of experts are not chosen to be powers of 2. We propose merging child nodes at the leaf level. In such instances, we have imperfect binary decision trees (Fig. 4 in Appendix) with  $n$  nodes, with  $2^d - n$  nodes in the  $(d - 1)$ -th level, and  $2n - 2^d$  nodes in the  $d$ -th level. Additional details are in Appendix C. In contrast to other differentiable gates (e.g., DSelect- $k$  [17]), our proposed gate *COMET* does not require any additional regularization to encourage the simplex constraint in (1b).

### 3.5 Stable numerical implementation

Next, we discuss a stable numerical implementation of *COMET* gate. *COMET* introduces additional exponential functions in the expert weights (or leaf nodes of the decision trees) — see (3). More exponential functions are known to cause instabilities in Sparse-MoE models. For example, [49] introduced router z-loss in Switch Transformers to encourage smaller logits. However, this may have a performance tradeoff. In our implementation of *COMET*, we can mitigate instability issues arising from additional exponential functions using the following approach: (i) convert root-to-leaf probabilities to the log-space,  $\log v_i^{(j)}(x)$ , (ii) compute  $\alpha_i^{(j)} + \log v_i^{(j)}(x)$ , (iii) subtract the maximum, i.e.,  $\max_{i,j} (\alpha_i^{(j)} + \log v_i^{(j)}(x))$  from each element, (iv) apply a two-way softmax operation to get  $g(x)$ .

## 4 LOCAL SEARCH

Expert selection is a challenging combinatorial problem that is known to be NP-hard. Although first-order heuristics can usually provide fast solutions, they rely heavily on initialization and are sometimes prone to arriving at low-quality solutions. To this end, we propose a permutation-based local search method that complements first-order methods in optimizing Sparse-MoEs. In both large-scale experiments and ablation studies, we see that the incorporation of local search can improve the performance of *any* gating method and can significantly reduce the number of tuning trials.

Our approach derives inspiration from the local search methods commonly used along with the first-order methods to help escape local minima in sparse linear models [3, 15]. We note that this is the first attempt in the literature to incorporate local search methods in the context of Sparse-MoE. Moreover, unlike common local search methods in literature, our proposed search method is differentiable. We want to highlight that our local search method is useful for any existing sparse gate, e.g., Hash routing, Top-k, and our proposed *COMET*. We hypothesize that our permutation-based approach can help navigate the optimization loss surface for various gates.

The rest of the section is organized as follows. In section 4.1, we formulate a refined cardinality-constrained Sparse-MoE objective with additional binary variables to add support for permutation-based local search. Then, in section 4.2, we provide background on permutation and its differentiable relaxation. Next in section 4.3, we outline our differentiable optimization approach for the refined Sparse-MoE objective and some additional practical considerations for computational efficiency. Later, in Section 4.4, we provide an ablation study to support our hypothesis that the local search can help escape bad initializations.

### 4.1 Permutation-based Local Search

In this section, we formulate a refined objective for the cardinality-constrained Sparse-MoE objective that adds support for permutation-based local search.

Let us denote by  $\mathcal{S}_n$  the set of all permutations of the set  $[n]$ . Given any permutation  $\sigma \in \mathcal{S}_n$ , we permute the  $n$  experts accordingly and assign  $i$ -th weight  $g(x)_i$  to  $\sigma(i)$ -th expert instead of  $i$ -th expert. With this permutation, the prediction for Sparse-MoE could be written as:  $\hat{y} = \sum_{i \in [n]} f_{\sigma(i)}(x)g(x)_i$ . We note that due to symmetry between experts and weights, permuting the experts is essentially same as permuting the weights. To see this, we can write  $\sum_{i \in [n]} f_{\sigma(i)}(x)g(x)_i = \sum_{j \in [n]} f_j(x)g(x)_{\sigma^{-1}(j)}$ , where  $\sigma^{-1}$  is the inverse map of  $\sigma$ , which is also a permutation.

For a permutation  $\sigma$ , we can define a corresponding permutation matrix  $P^\sigma$ , by setting  $P^\sigma[i, j] = \mathbf{1}\{\sigma(j) = i\}$ , where  $\mathbf{1}\{\cdot\}$  is an indicator function. Then it is easy to see that  $\sum_{j \in [n]} f_j(x)g(x)_{\sigma^{-1}(j)} = \sum_{j \in [n]} f_j(x)(P^\sigma g(x))_j$ . The refined Sparse-MoE problem is

$$\min_{f_1, \dots, f_n, g, P} \frac{1}{N} \sum_{(x, y) \in \mathcal{D}} \ell \left( y, \sum_{i \in [n]} f_i(x)(Pg(x))_i \right), \quad (4a)$$

$$\text{s.t. } \|g(x)\|_0 \leq k, \quad g(x) \in \Delta_n, \quad \forall x \in \mathcal{X}, \quad (4b)$$

$$P \in \mathcal{P}_n^{\text{local}}, \quad (4c)$$

where  $\mathcal{P}_n^{\text{local}}$  is a localized set of permutations in the full set of permutations, which we denote by  $\mathcal{P}_n$ . For example, one may only

allow for  $\mathcal{P}_n^{\text{local}} = \mathcal{P}_2$ , which only allows interchanging (swapping) two columns similar to “swap” operations shown to be useful in the sparse regression literature [15]. Besides optimizing the gates and experts, formulation (4) performs local search by optimizing over the permutation matrix. Specifically, the goal of local search here is to find a permutation  $P$  that leads to a better solution, i.e., one with a lower objective. Intuitively, if SGD is stuck at a low-quality solution, the permutation may be able to escape the solution by a better reordering of the experts. Standard local search, e.g., brute-force search may be computationally expensive. Therefore, we resort to a differentiable method that can be optimized efficiently.

### 4.2 Preliminaries: Permutation and a differentiable relaxation

In this section, we briefly summarize how the permutation learning problem is parameterized and later optimized. To parametrize the permutation matrix in the problem, a natural consideration is through the linear assignment problem [29]. To illustrate this, consider  $n$  people are to complete  $n$  tasks and a matrix  $U \in \mathbb{R}_{\geq 0}^{n \times n}$ , the goal is to assign each task to one person so as to maximize the utility given that the utility of assigning task  $j$  to person  $i$  is  $U_{ij}$ . This leads to the following optimization problem

$$M(U) = \arg \max_{P \in \mathcal{P}_n} \langle P, U \rangle_F := \sum_{i \in [n]} \sum_{j \in [n]} P_{ij} U_{ij}. \quad (5)$$

The operator  $M$  here is called the Matching operator, which maps a nonnegative matrix  $U$  to a permutation matrix  $P$ .

Problem (5) is a combinatorial optimization problem, which admits the following linear relaxation [6]:

$$\max_{B \in \mathcal{B}_n} \langle B, U \rangle_F := \sum_{i \in [n]} \sum_{j \in [n]} B_{ij} U_{ij}, \quad (6)$$

where  $\mathcal{B}_n$  denotes the set of double stochastic matrices  $\mathcal{B}_n = \{B \in \mathbb{R}^{n \times n} : \sum_{i \in [n]} B_{ij} = 1, \sum_{j \in [n]} B_{ij} = 1, B_{ij} \in [0, 1]\}$ , which is a convex hull of the set of permutation matrices  $\mathcal{P}_n$ .

However, this is still not a differentiable parametrization as problem (6) might end up with multiple solutions. To this end, Mena et al. [33] proposes a smooth version<sup>2</sup> of the permutation learning objective in (6):

$$S(U/\tau) = \arg \max_{B \in \mathcal{B}_n} \langle B, U \rangle_F - \tau \sum_{i, j \in [n]} B_{ij} \log B_{ij}, \quad (7)$$

and solves it using Sinkhorn operator  $S(\cdot)$  [1], defined by the following recursion:

$$S^0(U) = \exp(U), \quad (8a)$$

$$S^r(U) = \mathcal{T}_{col}(\mathcal{T}_{row}(S^{r-1}(U))), \quad (8b)$$

$$S(U) = \lim_{r \rightarrow \infty} S^r(U), \quad (8c)$$

where  $\mathcal{T}_{row}(U) = U \oslash (U \mathbf{1}_n \mathbf{1}_n^T)$ , and  $\mathcal{T}_{col}(U) = U \oslash (\mathbf{1}_n \mathbf{1}_n^T U)$  are the row and column-wise normalization operators of a matrix, with  $\oslash$  denoting the element-wise division and  $\mathbf{1}_n$  a column vector of ones. The sinkhorn procedure in (8) allows differentiable training with first-order methods, making it appealing as a local search method for Sparse-MoE.

<sup>2</sup>Note that (7) is an entropy-regularized version of (6). Since the entropy term is strictly concave, problem (7) has a unique minimizer and thus the parametrization is differentiable.

As shown in [33],  $M(\mathbf{U})$  can be obtained as  $\lim_{\tau \rightarrow 0^+} S(\mathbf{U}/\tau)$ , and thus  $\lim_{\tau \rightarrow 0^+, R \rightarrow \infty} S^R(\mathbf{U}/\tau)$ . In practice, we set a max number of iterations  $R$  for normalization in (8b) as well as a small positive number  $\tau > 0$ , and use  $S^R(\mathbf{U}/\tau)$  to approximate the limit (8c). In this way, we are able to parametrize the permutation matrix  $P$  in (4) as a differentiable function  $S^R(\mathbf{U}/\tau)$  of learnable matrix  $\mathbf{U}$ . However, additional considerations are needed to ensure that a hard permutation matrix can be achieved quickly in a few epochs — this is important in Sparse-MoE paradigm for computational reasons and a well-defined measure of sparsity. We discuss these in the next section.

### 4.3 Practical considerations for optimization

Next, we discuss some empirical considerations for the end-to-end learning approach that are important for Sparse-MoE.

*Need for a hard permutation matrix.* We would like to have a hard permutation matrix at inference time and ideally during the course of training, for exact sparsity and computational efficiency considerations. First, the gate does not perform sparse inference if the learnt permutation matrix is not a hard matrix. For example, even if  $g(\cdot)$  is sparse, the refined weights  $P \cdot g(\cdot)$  are not a sparse vector if  $P$  is not a binary matrix. This would result in a dense mixture of experts. Second, some sparse gates perform dense-to sparse-training (partially conditional training), e.g., DSelect-k, COMET, or variants of Top-k [37]. If the learnt permutation matrix is not hard, then sparse training cannot proceed in the later stages of optimization. To this end, we employ a two-stage optimization approach: (i) in the first stage, we simultaneously train the network (experts and gates) and the permutation for a small number of epochs. (ii) In the second phase, the permutation matrix is fixed and only the remaining network (experts and gate) is trained. Therefore, local search is only used in the early stages of training. Empirically, we observe that a small number of epochs (1 – 10) is sufficient to learn a good permutation in the first stage and improve solution quality. Since local search is restricted to the first stage, the computational efficiency of gates that perform dense-to-sparse training is not affected by much — please refer to Appendix D.3 for additional discussion.

In the two-stage approach outlined above, there is a transition from a soft to a hard matrix between the two stages. As we mentioned earlier, we use  $S^R(\mathbf{U}/\tau)$  to approximate  $M(\mathbf{U})$  as a limit of  $R \rightarrow \infty, \tau \rightarrow 0^+$ . In practice, the transition could be not continuous, as this approximation does not always reach a hard permutation matrix given that  $R$  is finite and  $\tau$  is nonzero. Therefore, at the transition point, we propose to convert the “soft” permutation matrix  $S^R(\mathbf{U}/\tau)$  to a hard one via the linear assignment problem given in (5), by invoking  $\mathbf{U}$  as  $S^R(\mathbf{U}/\tau)$ . In addition, empirically, small  $R$  can lead to numerical instabilities for small  $\tau$  [33]. Therefore, to decrease deviance of  $S^R(\mathbf{U}/\tau)$  from the closest hard permutation matrix, we introduce two schedulers on  $R$  and  $\tau$  that increase  $R$  for decreased  $\tau$ : (i) Ramp up (linearly)  $R$  from 20 to 150, (ii) Ramp down (linearly in log-scale)  $\tau$  from  $10^{-3}$  to  $10^{-7}$ .

Although the above schedulers decrease the deviance between soft and its closest hard permutation matrix at the transition point, the method still appears to suffer from pseudo-convergence. In particular, we observed, some row-columns can converge to fractional entries i.e., a  $2 \times 2$  sub-block having all entries with 0.5. Therefore, we introduce small separate row-wise and column-wise entropic

regularizations to mitigate such degenerate cases:  $\zeta \sum_{i \in [n]} (\Omega(S^R(\mathbf{U}/\tau)_i) + \Omega(\mathcal{T}_{row}(S^R(\mathbf{U}/\tau))_i))$ , where  $\zeta \geq 0$ .

*Implicit localization.* In the spirit of common local search approaches, a potential optimization approach could alternate between optimization of network (experts and gates) and permutation matrix. However, this is unnecessary because the differentiable relaxation of permutation is also amenable to first-order methods. Therefore, our approach jointly optimizes both the network and the permutation matrix. We noted earlier that the search space for permutation is “localized” out of the full set of permutation matrices  $\mathcal{P}_n$ . This localization is implicitly imposed through the smooth optimization of the permutation matrix via Sinkhorn. The permutation matrix learning relies on the initialization for  $\mathbf{U}$  and at each gradient step the  $\mathbf{U}^{(t)}$  is naturally expected to not deviate drastically from  $\mathbf{U}^{(t-1)}$ . Since the permutation matrix is updated for a limited number of steps in first stage, intuitively it cannot deviate significantly from the initial permutation matrix. This also defines an implicit neighborhood.

### 4.4 Ablation study for local search

In this section, we provide an ablation study to provide evidence that the permutation-based local search can complement first-order optimization methods for routing in Sparse-MoE. The study highlights that local search can improve solution quality through escape out of bad initializations in the first stages of optimization for different types of routing strategies: (a) fixed gates, (b) trainable gates. We perform this study on a subsampled (200k) MovieLens dataset and use the same MoE architecture with 16 experts as the one described in Supplement. We trained models for only 10 epochs without/with local search, where in the latter case we fixed the number of epochs for permutation learning to 5 epochs and  $\zeta = 10^{-5}$ . We used a batch size of 512 and learning rate of  $2.5 \times 10^{-5}$ . We repeat the training with 100 different random initializations and compute averages along with their standard errors.

*Fixed Gates.* In fixed gating strategies e.g., random hash routing (Hash-r), the samples are pre-assigned to experts. For example, in natural language processing tasks, tokens or words in vocabulary are clustered randomly [39] before training begins into groups and each group of words are assigned to a random expert in the set of experts. In our experiments on recommender systems, we randomly pre-assigned samples to experts based on user index for Hash-r (and Hash-r+). It is possible that the same group of users could be better aligned with another expert based on expert and user embedding initializations. Permutation-based local search can potentially find better assignment of each group to a more suited expert. We provide empirical evidence to demonstrate that local search indeed can find better loss. We report the average out-of-sample loss achieved by both Hash-r and Hash-r+ in Table 1. Learning permutation appears to help map each pre-assigned cluster of users to a more suitable expert based on expert initialization for second stage of optimization.

*Trainable Gates.* For trainable gates, we also study the effect of local search on non-differentiable (Top-k) and differentiable gates (COMET). We fixed  $k = 2$  for both types of gates and followed the same training protocol for 10 epochs. For COMET (and COMET+),

**Table 1: Test loss ( $\times 10^{-2}$ ) achieved for different gates without and with (marked with +) local search in early stages of optimization. Asterisk(\*) indicates statistical significance (p-value<0.05) over the corresponding gate without permutation with a one-sided unpaired t-test.**

Strategy	Smoothness	Gate	Test Loss ↓
Pre-assigned	-	Hash-r	57.434 ± 0.025
		<i>Hash-r+</i>	*57.000 ± 0.037
Trainable	Non-differentiable	Top-k	53.345 ± 0.033
		<i>Top-k+</i>	*53.140 ± 0.031
	Differentiable	COMET	52.034 ± 0.007
		<i>COMET+</i>	*52.017 ± 0.005

we fixed  $\gamma = 0.01$  (for smooth-step) and  $\lambda = 1$  (for entropic regularization). For *Top-k* and *COMET+*, we fixed the number of epochs for permutation learning as 5. We repeated this exercise for 100 different random initializations of the experts and gates. We report the average out-of-sample objective achieved by both types of gates in Table 1. We can observe that local search appears to complement first-order optimization methods by learning better initializations in the first stage of Sparse-MoE optimization for later learning.

The practical significance of local search achieving a better test objective across many initializations for various gates can be seen in terms of reducing hyperparameter tuning overhead as discussed in Section 5.1.3.

## 5 EXPERIMENTS

We study the performance of *COMET* and *COMET+* in recommender systems and image datasets in Section 5.1 and *COMET-BERT* in natural language processing tasks in 5.2. We also study the effect of local search for various gates. We denote our methods in *italics*. Supplement containing more details is available at <https://github.com/mazumder-lab/COMET/blob/main/COMET-Supplement.pdf>.

### 5.1 Experiments on Recommender Systems and Image Datasets

We study the performance of *COMET* and *COMET+* in recommender systems and image datasets. We compare with state-of-the-art gates and baselines including Softmax, Top-k, DSelect-k and Hash routing (Hash-r) on recommender systems (MovieLens [14], Jester [13], Books [48]) and image datasets (Digits [8, 36], MultiMNIST [41], MultiFashionMNIST [17], CelebA [31]). We also include an ablation study in Section 5.1.2 that shows that *COMET* achieves good performance with much less trials than existing popular gates e.g., Hash routing and Top-k. Additionally, in Section 5.1.3, we show that *Hash-r+*, *Top-k+*, and *COMET+* with local search can potentially achieve good performance with much less trials than Hash-r, Top-k and *COMET* respectively.

*Implementation.* We provide an open-source implementation of *COMET* and *COMET+*: <https://github.com/mazumder-lab/COMET>.

*Experimental Setup.* Although our exposition in Section 3 was for a single-task setting, the same gate can also be used in multi-task learning — multi-task requires multi-gate MoE architecture [32], where each task has a separate trainable gate, but tasks have to

select from a common set of experts. We briefly summarize the key aspects for each dataset. For MovieLens/Books/Jester we have two tasks: classification task predicts whether user watches/read/rates a particular movie/book/joke, regression problem predicts user’s rating. Loss is the convex combination of the two binary cross-entropy (for classification) and mean squared error (for regression) with task weights:  $\{\alpha, 1 - \alpha\}$ . We separately present results for two different  $\alpha$ ’s:  $\alpha \in \{0.1, 0.9\}$ . For MultiMNIST/MultiFashionMNIST, there are two multi-class classification tasks, which are equally weighted. For CelebA, there are 10 binary classification problems, which are equally weighted. Lastly, for Digits dataset, we have a

**Table 2: Tess Loss ( $\times 10^{-2}$ ) and actual no. of experts per sample ( $k_a$ ) at inference for *COMET*, *COMET+* and benchmark gates across various recommender system datasets. Asterisk(\*) indicates statistical significance (p-value<0.05) over the best existing gate, using a one-sided unpaired t-test.**

Dataset	$n$	Model	Test Loss ↓	$k_a$ ↓
Books ( $\alpha = 0.1$ )	9	Softmax	244.47 ± 0.14	9.00 ± 0.00
		Hash-r	247.43 ± 0.14	1.00 ± 0.00
		<i>Hash-r+</i>	247.33 ± 0.23	1.00 ± 0.00
		Top-k	247.87 ± 0.17	4.00 ± 0.00
		<i>Top-k+</i>	247.88 ± 0.14	4.00 ± 0.00
		DSelect-k	246.43 ± 0.36	1.09 ± 0.00
		COMET	*240.79 ± 0.14	2.81 ± 0.11
		<i>COMET+</i>	240.82 ± 0.19	3.03 ± 0.08
Books ( $\alpha = 0.9$ )	9	Softmax	73.88 ± 0.02	9.00 ± 0.00
		Hash-r	75.02 ± 0.03	1.00 ± 0.00
		<i>Hash-r+</i>	75.06 ± 0.03	1.00 ± 0.00
		Top-k	74.78 ± 0.03	4.00 ± 0.00
		<i>Top-k+</i>	74.86 ± 0.03	4.00 ± 0.00
		DSelect-k	75.98 ± 0.13	1.07 ± 0.00
		COMET	73.62 ± 0.03	2.94 ± 0.10
		<i>COMET+</i>	*73.55 ± 0.02	3.15 ± 0.08
MovieLens ( $\alpha = 0.9$ )	16	Softmax	42.26 ± 0.01	16.00 ± 0.00
		Hash-r	46.91 ± 0.02	1.00 ± 0.00
		<i>Hash-r+</i>	46.84 ± 0.03	1.00 ± 0.00
		Top-k	41.83 ± 0.02	2.00 ± 0.00
		<i>Top-k+</i>	41.74 ± 0.02	2.00 ± 0.00
		DSelect-k	40.82 ± 0.02	1.94 ± 0.06
		COMET	40.76 ± 0.02	1.76 ± 0.06
		<i>COMET+</i>	*40.69 ± 0.02	1.66 ± 0.06
MovieLens ( $\alpha = 0.1$ )	16	Softmax	75.52 ± 0.02	16.00 ± 0.00
		Hash-r	79.41 ± 0.02	1.00 ± 0.00
		<i>Hash-r+</i>	78.92 ± 0.05	1.00 ± 0.00
		Top-k	76.52 ± 0.04	2.00 ± 0.00
		<i>Top-k+</i>	75.12 ± 0.04	2.00 ± 0.00
		DSelect-k	73.91 ± 0.05	1.94 ± 0.03
		COMET	73.91 ± 0.04	1.94 ± 0.03
		<i>COMET+</i>	*73.67 ± 0.04	1.98 ± 0.03
Jester ( $\alpha = 0.1$ )	16	Softmax	68.17 ± 0.03	16.00 ± 0.00
		Hash-r	67.47 ± 0.01	1.00 ± 0.00
		Top-k	68.38 ± 0.05	2.00 ± 0.00
		<i>Top-k+</i>	68.00 ± 0.07	2.00 ± 0.00
		DSelect-k	67.06 ± 0.03	1.96 ± 0.02
		COMET	67.12 ± 0.04	1.98 ± 0.02
		<i>COMET+</i>	*66.91 ± 0.03	2.00 ± 0.00
		Jester ( $\alpha = 0.9$ )	16	Softmax
Hash-r	22.083 ± 0.004			1.00 ± 0.00
Top-k	21.958 ± 0.007			2.00 ± 0.00
<i>Top-k+</i>	21.961 ± 0.006			2.00 ± 0.00
DSelect-k	21.930 ± 0.005			2.00 ± 0.00
COMET	21.946 ± 0.005			2.00 ± 0.00
<i>COMET+</i>	*21.906 ± 0.005			2.00 ± 0.00

multi-class single-task classification cross-entropy objective. Full details about datasets and MoE architectures are in Supplement.

We used Adam for optimization, and we tuned the key hyperparameters using random grid search. Note that for *Hash-r+*, *COMET+* and *Top-k+*, we only allocate a very small portion of the epochs (1-10) for permutation learning. Full details about the hyperparameter tuning are given in Supplement.

**5.1.1 Performance of *COMET* and *COMET+*.** In Tables 2 and 3, we report the (average) test loss and the average number of selected experts per sample across multiple recommender and vision datasets. The results indicate that *COMET* and *COMET+* lead on many datasets, outperforming popular state-of-the-art gating methods e.g., Hash-r, Top-k and DSelect-k in test loss. Our proposed gate *COMET* can outperform standard routing techniques (without local search). Even without local search, *COMET* is getting relatively good solutions. We hypothesize that the good performance of *COMET* is due to a combination of factors including differentiability, and  $k$ -decision trees formulation. With local search, *COMET+* can sometimes further enhance solution quality. We also provide task-specific metrics (AUC/Accuracy/MSE) in Tables 7 in Appendix E. We observe *COMET+* can improve AUC by up to 13% over Hash routing and Top-k, and 9% over DSelect-k. We observe that Top-k gate does not uniformly outperform the Softmax across multiple datasets. However, *Top-k+* significantly improves the performance of Top-k across multiple datasets. In fact with the permutation module, *Top-k+* outperforms Softmax in all cases, so sparsity in gating seems to be beneficial on all these datasets.

**Inference Sparsity.** We see that *COMET* and *COMET+* can sometimes lead to a smaller number of experts selected than that for

**Table 3: Test Loss ( $\times 10^{-2}$ ) and actual no. of experts per sample ( $k_a$ ) at inference for *COMET*, *COMET+* and benchmark gates across various image datasets. Asterisk(\*) indicates statistical significance ( $p$ -value $<0.05$ ) over the best existing gate, using a one-sided unpaired t-test.**

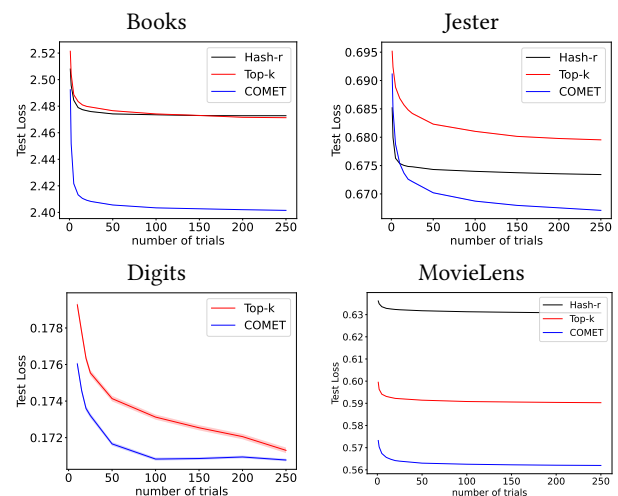
Dataset	$n$	Model	Test Loss $\downarrow$	$k_a \downarrow$
MultiFashionMNIST	5	Softmax	34.21 $\pm$ 0.09	5.00 $\pm$ 0.00
		Top-k	33.82 $\pm$ 0.09	2.00 $\pm$ 0.00
		<i>Top-k+</i>	*33.62 $\pm$ 0.08	2.00 $\pm$ 0.00
		DSelect-k	35.49 $\pm$ 0.10	1.00 $\pm$ 0.00
		<i>COMET</i>	33.70 $\pm$ 0.09	1.49 $\pm$ 0.07
		<i>COMET+</i>	*33.67 $\pm$ 0.09	1.54 $\pm$ 0.07
CelebA	6	Softmax	35.10 $\pm$ 0.32	6.00 $\pm$ 0.00
		Top-k	34.48 $\pm$ 0.24	2.00 $\pm$ 0.00
		<i>Top-k+</i>	34.54 $\pm$ 0.23	2.00 $\pm$ 0.00
		DSelect-k	35.39 $\pm$ 0.12	1.00 $\pm$ 0.00
		<i>COMET</i>	33.96 $\pm$ 0.15	1.00 $\pm$ 0.08
		<i>COMET+</i>	*33.93 $\pm$ 0.16	1.00 $\pm$ 0.08
Digits	8	Softmax	17.48 $\pm$ 0.07	8.00 $\pm$ 0.00
		Top-k	17.46 $\pm$ 0.09	2.00 $\pm$ 0.00
		<i>Top-k+</i>	17.29 $\pm$ 0.08	2.00 $\pm$ 0.00
		DSelect-k	17.18 $\pm$ 0.06	1.15 $\pm$ 0.06
		<i>COMET</i>	17.19 $\pm$ 0.06	1.07 $\pm$ 0.04
		<i>COMET+</i>	17.08 $\pm$ 0.09	1.06 $\pm$ 0.04
MultiMNIST	16	Softmax	6.88 $\pm$ 0.06	16.00 $\pm$ 0.00
		Top-k	6.84 $\pm$ 0.05	4.00 $\pm$ 0.00
		<i>Top-k+</i>	6.70 $\pm$ 0.08	4.00 $\pm$ 0.00
		DSelect-k	6.64 $\pm$ 0.07	3.40 $\pm$ 0.09
		<i>COMET</i>	6.48 $\pm$ 0.07	3.49 $\pm$ 0.09
		<i>COMET+</i>	6.49 $\pm$ 0.06	3.53 $\pm$ 0.08

Top-k. This leads to smaller number of FLOPs at inference time (see Appendix D.2). For some settings, DSelect-k appears to arrive at a sparser selection than *COMET+*; however, in these cases, DSelect-k loses significantly in terms of performance. We observed expert collapsing in DSelect-k in such cases.

**Timing Discussion.** For cost complexity of *COMET*, please see Appendix D.1. Additionally, we discuss the computational aspects of the local search in Appendix D.3.

**5.1.2 Reducing Hyperparameter Search with *COMET*.** Here, we study how our differentiable *COMET* gate (that performs dense-to-sparse training) can be beneficial in terms of hyperparameter tuning over popular gates such as Hash routing and Top-k. We perform a large set of tuning trials and perform a bootstrapping procedure (discussed in Appendix F) to see whether *COMET* helps in reducing the hyperparameter tuning overload. *COMET* can achieve the same level of performance as popular gates with much lesser number of hyperparameter trials. This indicates that *COMET* is not too heavily dependent on a very restricted set of hyperparameter values. We visualize this for various datasets in Fig. 2. We see tuning reduction by a factor of  $5\times$ – $100\times$  for *COMET* over popular gates.

**5.1.3 Effect of Local Search on Hyperparameter Tuning.** Here, we study how local search can be beneficial in terms of hyperparameter tuning. We study this effect for Hash-r, Top-k and *COMET*. We visualize this in Fig. 3 for MovieLens for both *Hash-r+*, *Top-k+* and *COMET+*. We observe that we can achieve comparable performance with much smaller number of trials. We see tuning reduction by a factor of  $3\times$ – $100\times$  for *Hash-r+*,  $20\times$ – $100\times$  for *Top-k+* and  $2\times$ – $5\times$  for *COMET+*. This suggests that permutation-based local search helps escape out of bad initializations. Such favorable properties of




**Figure 2: Sensitivity of *COMET* to hyperparameter tuning.** *COMET* can achieve the same level of performance as popular gates (e.g., Hash-r and Top-k) with significantly lesser number of hyperparameter trials. We see tuning reduction by  $5\times$ – $100\times$  for *COMET* over Top-k and Hash routing.

local search in terms of reducing the hyperparameter tuning load for existing gates can be beneficial for Large Language Models.

## 5.2 Experiments on NLP Tasks

In this section, we consider a setting where a pretrained large model (non-MoE based) needs to be distilled for a more efficient inference while preserving or improving the best performance. Following [51], we study a distillation setting, where BERT [9] is distilled into its Sparse-MoE based variant. Specifically, the FFN layers are replaced with MoE layers — this can result in a  $\sim 2\times$  smaller number of (effective) parameters with per-sample sparse routing (for  $k = 1$ ), thus allowing for more efficient inference.

Following [51], we use an importance-weight guided distillation strategy: (i) Finetune BERT on a downstream task. (ii) Compute importance weights in FFN layers to construct an MoE-based variant of BERT. (iii) Distill BERT into MoE-based variant on the downstream task with a layer-wise discrepancy loss. [51] used Hash routing in their MoEBERT model. We propose *COMET-BERT* (MoE based BERT model with *COMET/COMET+* gating) and evaluate the performance on the GLUE benchmarks [43] and SQuAD benchmark [38]. More details about the benchmarks are given in Supplement.



**Figure 3: Effect of local search on hyperparameter tuning. Comparison of Hash-r+ vs Hash-r, Top-k+ vs Top-k and COMET+ vs COMET on MovieLens with two different task weight settings. Local search appears to achieve the same level of performance with much lesser number of hyperparameter trials. We see tuning reduction by a factor of  $3\times-100\times$  for Hash-r+,  $20\times-100\times$  for Top-k+ and  $2\times-5\times$  for COMET+.**

*Implementation.* We implemented *COMET-BERT* in HuggingFace [45] and adapted the codebase of [51]. Unlike Hash routing, our gates can also cater to  $k \geq 1$ . However, for consistent comparison in terms of inference, we set  $k = 1$ . Tuning details are outlined in Supplement. Code for *COMET-BERT* is available at <https://github.com/mazumder-lab/COMET-BERT>.

*Results.* We report the performance metrics in Table 4 for 7 GLUE datasets and SQuAD dataset. *COMET-BERT* outperforms MoEBERT in 5/7 benchmarks on GLUE datasets. *COMET-BERT* also outperform MoEBERT significantly on SQuADv2.0. Notably, in 5 of these datasets (CoLA, MRPC, QNLI and MNLI, SQuAD v2.0), *COMET-BERT* achieves SOTA performance when distilling BERT, (when compared with all distillation methods in literature with same number of effective parameters for inference).

**Table 4: Performance metrics on the GLUE and SQuAD development sets. Models are trained without data augmentation. Both models have 66M (effective) parameters for inference.**

	GLUE							SQuAD
	RTE Acc	CoLA Mcc	MRPC F1	SST-2 Acc	QNLI Acc	QQP F1/Acc	MNLI m/mm	v2.0 F1/EM
MOEBERT <sup>3</sup>	70.8	55.4	91.0	93.2	90.9	88.5/91.4	84.7	76.8/73.6
COMET-BERT	71.1	57.0	91.3	93.0	91.2	88.4/91.3	85.5	78.4/75.3

## 6 CONCLUSION

In summary, we propose two new approaches for improving routing in Sparse-MoE. First, we introduce a new differentiable gate *COMET*, which relies on a novel tree-based sparse expert selection mechanism. *COMET* allows optimization with first-order methods, offers explicit control over the number of experts to select, allows (partially) conditional training and sparse inference. Second, in this work, we argue that combinatorial nature of expert selection in Sparse-MoE makes sparse routing optimization challenging with first-order methods. Thus, we propose a new local search method that can help any gate including ours (*COMET*) escape “bad” initializations. Our large-scale experiments on recommender systems, vision and natural language processing tasks show *COMET* and *COMET+*: (i) achieve statistically significant improvements in prediction (up to 13% improvement in AUC) and expert selection over popular sparse gates. (ii) reduce tuning up to a factor of  $100\times$  to achieve the same level of performance as popular gates e.g., Top-k and Hash routing. (iii) help *COMET-BERT* achieve state-of-the-art results for distilling BERT on GLUE and SQuAD benchmarks.

## ACKNOWLEDGMENTS

This research is supported in part by grants from Google and the Office of Naval Research (ONR N000142112841). We thank Mathieu Sibue for his initial help in setting up the baseline experiments on MovieLens. We thank Paul Theron for his help with the tuning experiments for *COMET-BERT* on NLP tasks. The authors acknowledge both MIT SuperCloud and Google Cloud Platform for HPC resources for experiments.

<sup>3</sup>Numbers are based on re-run of the official codebase (<https://github.com/SimiaoZuo/MoEBERT>) with hash routing with the optimal hyperparameters reported in [47].

## REFERENCES

- [1] Ryan Prescott Adams and Richard S. Zemel. 2011. Ranking via Sinkhorn Propagation. <https://doi.org/10.48550/ARXIV.1106.1925>
- [2] Mikel Artetxe, Shruti Bhosale, Naman Goyal, et al. 2022. Efficient Large Scale Language Modeling with Mixtures of Experts. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- [3] Amir Beck and Yonina C. Eldar. 2013. Sparsity Constrained Nonlinear Optimization: Optimality Conditions and Algorithms. *SIAM Journal on Optimization* 23, 3 (2013), 1480–1509. <https://doi.org/10.1137/120869778> arXiv:<https://doi.org/10.1137/120869778>
- [4] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. 2016. Conditional Computation in Neural Networks for faster models. In *International Conference on Learning Representations Workshop Tract*. <https://openreview.net/forum?id=B1ckMDqgl>
- [5] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation. *CoRR abs/1308.3432* (2013). arXiv:1308.3432 <http://arxiv.org/abs/1308.3432>
- [6] Dimitris Bertsimas and John N Tsitsiklis. 1997. *Introduction to linear optimization*. Vol. 6. Athena Scientific Belmont, MA.
- [7] Aidan Clark, Diego De Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, George Bm Van Den Driessche, Eliza Rutherford, Tom Hennigan, Matthew J Johnson, Albin Cassirer, Chris Jones, Elena Buchatskaya, David Budden, Laurent Sifre, Simon Osindero, Oriol Vinyals, Marc Aurelio Ran-zato, Jack Rae, Erich Elsen, Koray Kavukcuoglu, and Karen Simonyan. 2022. Unified Scaling Laws for Routed Language Models. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 4057–4086. <https://proceedings.mlr.press/v162/clark22a.html>
- [8] Li Deng. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 6 (2012), 141–142.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. <http://arxiv.org/abs/1810.04805> cite arxiv:1810.04805Comment: 13 pages.
- [10] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten P Bosma, Zongwei Zhou, Tao Wang, Emma Wang, Kellie Webster, Marie Pellat, Kevin Robinson, Kathleen Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc Le, Yonghui Wu, Zhifeng Chen, and Claire Cui. 2022. GLaM: Efficient Scaling of Language Models with Mixture-of-Experts. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 5547–5569. <https://proceedings.mlr.press/v162/du22c.html>
- [11] William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research* 23, 120 (2022), 1–39. <http://jmlr.org/papers/v23/21-0998.html>
- [12] Nicholas Frosst and Geoffrey Hinton. 2017. Distilling a Neural Network Into a Soft Decision Tree. arXiv:1711.09784
- [13] Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. 2001. Eigentaste: A Constant Time Collaborative Filtering Algorithm. *Information Retrieval* 4, 2 (2001), 133–151.
- [14] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (dec 2015), 19 pages. <https://doi.org/10.1145/2827872>
- [15] Hussein Hazimeh and Rahul Mazumder. 2020. Fast Best Subset Selection: Coordinate Descent and Local Combinatorial Optimization Algorithms. *Oper. Res.* 68, 5 (Sept. 2020), 1517–1537. <https://doi.org/10.1287/opre.2019.1919>
- [16] Hussein Hazimeh, Natalia Ponomareva, Petros Mol, Zhenyu Tan, and Rahul Mazumder. 2020. The tree ensemble layer: Differentiability meets conditional computation. In *International Conference on Machine Learning*. PMLR, 4138–4148.
- [17] Hussein Hazimeh, Zhe Zhao, Aakanksha Chowdhery, Maheswaran Sathiamoorthy, Yihua Chen, Rahul Mazumder, Lichan Hong, and Ed Chi. 2021. DSelect-k: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=tKYQJLYN8v>
- [18] Thomas M. Hehn, Julian F. P. Kooij, and Fred A. Hamprecht. 2019. End-to-End Learning of Decision Trees and Forests. *International Journal of Computer Vision* 128 (2019), 997–1011.
- [19] Shibal Ibrahim, Hussein Hazimeh, and Rahul Mazumder. 2022. Flexible Modeling and Multitask Learning Using Differentiable Tree Ensembles. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Washington DC, USA) (KDD '22)*. Association for Computing Machinery, New York, NY, USA, 666–675. <https://doi.org/10.1145/3534678.3539412>
- [20] Yani Ioannou, Duncan P. Robertson, Darko Zikic, Peter Kotschieder, Jamie Shotton, Matthew Brown, and Antonio Criminisi. 2016. Decision Forests, Computational Networks and the Models in-Between. *CoRR abs/1603.01250* (2016). arXiv:1603.01250 <http://arxiv.org/abs/1603.01250>
- [21] Ozan Irsoy, O. T. Yildiz, and Ethem Alpaydin. 2012. Soft decision trees. *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)* (2012), 1819–1822.
- [22] Aya Abdelsalam Ismail, Sercan O Arik, Jinsung Yoon, Ankur Taly, Soheil Feizi, and Tomas Pfister. 2023. Interpretable Mixture of Experts. *Transactions on Machine Learning Research* (2023). <https://openreview.net/forum?id=DdZoPUPm0a>
- [23] Robert A. Jacobs. 1997. Bias/Variance Analyses of Mixtures-of-Experts Architectures. *Neural Computation* 9, 2 (1997), 369–383. <https://doi.org/10.1162/neco.1997.9.2.369>
- [24] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. Adaptive Mixtures of Local Experts. *Neural Computation* 3, 1 (1991), 79–87. <https://doi.org/10.1162/neco.1991.3.1.79>
- [25] W. Jiang and M.A. Tanner. 1999. On the identifiability of mixtures-of-experts. *Neural Networks* 12, 9 (1999), 1253–1258. [https://doi.org/10.1016/S0893-6080\(99\)00066-0](https://doi.org/10.1016/S0893-6080(99)00066-0)
- [26] M.I. Jordan and R.A. Jacobs. 1993. Hierarchical mixtures of experts and the EM algorithm. In *Proceedings of 1993 International Conference on Neural Networks (IJCNN-93-Nagoya, Japan)*, Vol. 2. 1339–1344 vol.2. <https://doi.org/10.1109/IJCNN.1993.716791>
- [27] Peter Kotschieder, Madalina Fiterau, Antonio Criminisi, et al. 2015. Deep Neural Decision Forests. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 1467–1475.
- [28] Sneha Kudugunta, Yanping Huang, Ankur Bapna, Maxim Krikun, Dmitry Lepikhin, Minh-Thang Luong, and Orhan Firat. 2021. Beyond Distillation: Task-level Mixture-of-Experts for Efficient Inference. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, Punta Cana, Dominican Republic, 3577–3599. <https://doi.org/10.18653/v1/2021.findings-emnlp.304>
- [29] H. W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97. <https://doi.org/10.1002/nav.3800020109>
- [30] Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. 2021. BASE Layers: Simplifying Training of Large, Sparse Models. In *International Conference on Machine Learning*.
- [31] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- [32] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1930–1939.
- [33] Gonzalo Mena, David Belanger, Scott Linderman, and Jasper Snoek. 2018. Learning Latent Permutations with Gumbel-Sinkhorn Networks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Byt3oJ0W>
- [34] Sreerama K. Murthy, Simon Kasif, and Steven Salzberg. 1994. A System for Induction of Oblique Decision Trees. *J. Artif. Int. Res.* 2, 1 (aug 1994), 1–32.
- [35] Basil Mustafa, Carlos Riquelme Ruiz, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. 2022. Multimodal Contrastive Learning with LIMoE: the Language-Image Mixture of Experts. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). <https://openreview.net/forum?id=Qy1D9JyMBg0>
- [36] Yuval Netzer, Tao Wang, Adam Coates, A. Bissacco, Bo Wu, and A. Ng. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning.
- [37] Xiaonan Nie, Xupeng Miao, Shijie Cao, Lingxiao Ma, Qibin Liu, Jilong Xue, Youshan Miao, Yi Liu, Zhi Yang, and Bin Cui. 2021. EvoMoE: An Evolutional Mixture-of-Experts Training Framework via Dense-To-Sparse Gate. <https://doi.org/10.48550/ARXIV.2112.14397>
- [38] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Melbourne, Australia, 784–789. <https://doi.org/10.18653/v1/P18-2124>
- [39] Stephen Roller, Sainbayar Sukhbaatar, Arthur Szlam, and Jason E Weston. 2021. Hash Layers For Large Sparse Models. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=IMgDDWb1ULW>
- [40] Carlos Riquelme Ruiz, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. 2021. Scaling Vision with Sparse Mixture of Experts. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). [https://openreview.net/forum?id=NGPmH3vbAA\\_](https://openreview.net/forum?id=NGPmH3vbAA_)

- [41] Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. 2017. Dynamic Routing between Capsules. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 3859a–3869.
- [42] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1ckMDqlg>
- [43] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rJ4km2R5t7>
- [44] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E. Gonzalez. 2018. SkipNet: Learning Dynamic Routing in Convolutional Networks. In *The European Conference on Computer Vision (ECCV)*.
- [45] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- [46] Lemeng Wu, Mengchen Liu, Yinpeng Chen, Dongdong Chen, Xiyang Dai, and Lu Yuan. 2022. Residual Mixture of Experts. *ArXiv abs/2204.09636* (2022).
- [47] Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Y Zhao, Andrew M. Dai, Zhifeng Chen, Quoc V Le, and James Laudon. 2022. Mixture-of-Experts with Expert Choice Routing. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). <https://openreview.net/forum?id=jDJo1HIVmI>
- [48] Cai-Nicolas Ziegler, Sean M. McNeel, Joseph A. Konstan, and Georg Lausen. 2005. Improving Recommendation Lists through Topic Diversification. In *Proceedings of the 14th International Conference on World Wide Web (Chiba, Japan) (WWW '05)*. Association for Computing Machinery, New York, NY, USA, 22a–32. <https://doi.org/10.1145/1060745.1060754>
- [49] Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam M. Shazeer, and William Fedus. 2022. ST-MoE: Designing Stable and Transferable Sparse Expert Models.
- [50] Simiao Zuo, Xiaodong Liu, Jian Jiao, Young Jin Kim, Hany Hassan, Ruofei Zhang, Jianfeng Gao, and Tuo Zhao. 2022. Taming Sparsely Activated Transformer with Stochastic Experts. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B72HXs80q4>
- [51] Simiao Zuo, Qingru Zhang, Chen Liang, Pengcheng He, Tuo Zhao, and Weizhu Chen. 2022. MoEBERT: from BERT to Mixture-of-Experts via Importance-Guided Adaptation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 1610–1623. <https://doi.org/10.18653/v1/2022.naacl-main.116>

## APPENDIX

### A SMOOTH-STEP ACTIVATION FUNCTION

In smooth routing, the internal nodes of a soft tree use an activation function  $h$  in order to compute the routing probabilities. The common choice for  $h$  — logistic (a.k.a. sigmoid) function in soft trees literature [12, 18, 26, 27] — can not output exact zeros. This implies that any sample  $x$  will reach every node in the tree with a positive probability. Thus, computing the output of mixture of experts will require computation over every expert. The smooth-step activation function proposed in [16] can output exact zeros and ones, thus allowing for conditional computation. Let  $\gamma$  be a non-negative scalar parameter. The smooth-step function is:

$$h(t) = \begin{cases} 0 & \text{if } t \leq -\gamma/2 \\ -\frac{2}{\gamma^3}t^3 + \frac{3}{2\gamma}t + \frac{1}{2} & \text{if } \gamma/2 \leq t \leq \gamma/2 \\ 1 & \text{if } t \geq \gamma/2 \end{cases} \quad (9)$$

The smooth-step function is continuously differentiable, similar to the logistic function. Additionally, it performs hard routing, i.e., outside  $[-\gamma/2, \gamma/2]$ , the function produces exact zeros and ones. For

cardinality-obeying Sparse-MoE learning with trees (not studied in [16]), the goal for each tree is to perform hard routing for all samples. Therefore, we propose additional regularization on  $\{h(w_q \cdot x), 1 - h(w_q \cdot x)\}$  to encourage convergence of  $v$  to a one-hot state (discussed in more detail in Section 3.3).

### B PROOF FOR PROPOSITION 3.1

PROOF. First off, it is straightforward to see that  $g(x; \alpha, v)$  satisfies the simplex constraint in (1b):

$$\sum_{i=1}^n g(x; \alpha, v)_i = \sum_{i=1}^n \frac{\sum_{j=1}^k \exp(\alpha_i^{(j)}(x)) \cdot v_i^{(j)}(x)}{\sum_{j=1}^k \sum_{i=1}^n \exp(\alpha_i^{(j)}(x)) v_i^{(j)}(x)} = 1. \quad (10)$$

It remains to show that  $\|g(x; \alpha, v)\|_0 \leq k$  under the given conditions. Recall the hierarchical binary encoding  $v^{(j)}$  outputs a one-hot vector for each sample  $x$  as the routing decision. Let us denote by  $\hat{i}_j$  the expert number selected by  $j$ -th tree. For now, let us assume that  $\hat{i}_j$  are different for any  $j$ . Then, we have

$$g(x; \alpha, v)_i = \exp(\alpha_i^{(j)}(x)) \Big/ \sum_{j \in [k]} \exp(\alpha_i^{(j)}(x)), \quad (11)$$

i.e., the weights are restricted on the  $k$  experts selected by the  $k$  trees, and the weights form a softmax activation of logits  $\alpha_i^{(j)}$ 's. Given that the  $j$ -th tree selects  $\hat{i}_j$ -th expert, we know that if  $i \notin \{\hat{i}_1, \hat{i}_2, \dots, \hat{i}_k\}$ ,  $v_i^{(j)} = 0$  for any  $j$ , and thus  $g(x; \alpha, v)_i = 0$ , and this means that the support of  $g(x; \alpha, v)$  is contained in the set  $\{\hat{i}_1, \hat{i}_2, \dots, \hat{i}_k\}$ , which has at most  $k$  elements. Therefore, the cardinality constraint in (1b) holds.  $\square$

### C COMET FOR NON-POWERS OF 2

COMET has a natural way to cater to settings where number of experts are non-powers of 2. Recall we have  $n$  experts. Let  $d = \lceil \log_2 n \rceil$ , then  $2^{d-1} < n \leq 2^d$ . Let  $\mathcal{T}$  be a full binary tree with depth  $d$  and  $2^d$  leaf nodes. We collapse  $2(2^d - n)$  leaf nodes to their parents in the  $(d-1)$ -th level. Each time we collapse two leaf nodes, we get a new node in the  $(d-1)$ -th level, and the total number of nodes reduce by one. Therefore, we get a tree with  $n$  nodes, with  $2^d - n$  nodes in the  $(d-1)$ -th level, and  $2n - 2^d$  nodes in the  $d$ -th level. This is visualized in Figure 4 for number of experts equal to 5.




Figure 4: COMET for 5 (non-powers of 2) experts.

## D TIMING DISCUSSION

### D.1 Cost Complexity

We compare the cost of forward pass *COMET* vs Top-k. We consider a single-task,  $n$  experts ( $i$ th expert is  $f_i$ ),  $p$  features, desired sparsity  $k \ll n$ , and some shared layers.

**Table 5: Time complexity of Top-k and *COMET*.**

Cost breakdown	Top-k	<i>COMET</i> (training)	<i>COMET</i> (inference)
Experts	$kO(f_i)$	$k_a O(f_i)$	$k_a O(f_i)$
Gate	$O(pn)+O(n+k \log n)$	$O(kpn)$	$O(kp \log n)$
Shared layers	$O(\text{Shared Layers})$	$O(\text{Shared Layers})$	$O(\text{Shared Layers})$

*Cost for Top-k.* Here the desired sparsity ( $k$ ) is always equal to the actual sparsity ( $k_a$ ) during training and inference, i.e.  $k_a \equiv k$ . The cost of computing  $g(\cdot)$  is given by  $O(pn) + O(k \log n + n)$ . If the cost of computing expert  $f_i$  is  $O(f_i)$ , then the cost of the full model is  $kO(f_i) + O(pn) + O(k \log n + n) + O(\text{shared layers})$ . Here the main cost is from how many  $f_i$ 's are evaluated per sample.

*Cost for *COMET*.* For *COMET*, the cost of the gate is  $O(kpn)$ . The cost of the full model is  $k_a O(f_i) + O(kpn) + O(\text{shared layers})$ . After a very few epochs (e.g., 4-5), when *COMET* has reached the desired sparsity, we have  $k_a \leq k$ .

Note that during inference, *COMET* has smaller cost: All root-to-leaf paths in *COMET* do not require evaluation, hence the gate cost reduces from  $O(kpn)$  to  $O(kp \log(n))$ . Hence, due to fewer expert evaluations and smaller gate cost, *COMET* is more efficient at inference time than Top-k. Also see Table 6 for FLOP counts at inference time for different gating methods.

### D.2 FLOPs Comparison

We compare the FLOP counts of different gating methods (across different models/datasets) to compare inference speed—see Table 6 below. Inference FLOPs per sample - number of floating point operations that a model performs per sample - is a standard measure to evaluate the inference speed for Sparse-MoE e.g., in [11] etc.

**Table 6: #FLOPs per-sample at inference time for *COMET* against benchmarks (Softmax, Top-k, DSelect-k) across various datasets.**

Dataset	Model	FLOPs
Books	Softmax	1195K
	Top-k	402K
	DSelect-k	<b>214K</b>
	<i>COMET</i>	326K
MovieLens	Softmax	2255K
	Top-k	413K
	DSelect-k	399K
	<i>COMET</i>	<b>362K</b>
MultiFashionMNIST	Softmax	7.49M
	Top-k	3.03M
	DSelect-k	<b>1.58M</b>
	<i>COMET</i>	2.35M
CelebA	Softmax	22.02M
	Top-k	8.82M
	DSelect-k	5.64M
	<i>COMET</i>	<b>5.47M</b>

**Table 7: Test AUC/Accuracy/MSE for *COMET+* and benchmark gates on recommender systems and image datasets.**

Recommender Systems			
	Gate	Task 1 (Test AUC)	Task 2 (Test MSE)
Books (alpha=0.1)	Softmax	56.70±0.16	2.6470±0.0016
	Hash-r	54.55±0.07	2.6791±0.0017
	Top-k	55.28±0.07	2.6783±0.0026
	DSelect-k	59.19±0.36	2.6667±0.0038
	<i>COMET+</i>	<b>68.18±0.24</b>	<b>2.6063±0.0018</b>
Books (alpha=0.9)	Softmax	77.85±0.01	2.6195±0.0019
	Hash-r	77.32±0.05	2.7152±0.0050
	Top-k	77.46±0.03	2.6942±0.0016
	DSelect-k	77.07±0.09	2.7581±0.0092
	<i>COMET+</i>	<b>77.95±0.02</b>	<b>2.6158±0.0021</b>
MovieLens (alpha=0.9)	Softmax	90.92±0.01	0.7585±0.0005
	Hash-r	88.95±0.02	0.8065±0.0004
	Top-k	91.25±0.01	0.7635±0.0008
	DSelect-k	91.65±0.02	0.7455±0.0006
	<i>COMET+</i>	<b>91.70±0.01</b>	<b>0.7437±0.0006</b>
MovieLens (alpha=0.1)	Softmax	85.50±0.03	0.7867±0.0029
	Hash-r	84.27±0.07	0.8279±0.0003
	Top-k	87.12±0.04	0.8005±0.0004
	DSelect-k	<b>88.16±0.07</b>	0.7734±0.0005
	<i>COMET+</i>	88.02±0.01	<b>0.7707±0.0005</b>
Jester (alpha=0.9)	Softmax	97.350±0.004	0.7460±0.0003
	Hash-r	97.323±0.003	0.7530±0.0003
	Top-k	97.346±0.004	0.7456±0.0006
	DSelect-k	97.361±0.004	0.7464±0.0005
	<i>COMET+</i>	<b>97.362±0.004</b>	<b>0.7439±0.0006</b>
Jester (alpha=0.1)	Softmax	97.22±0.01	0.7380±0.0003
	Hash-r	97.01±0.01	0.7301±0.0002
	Top-k	97.38±0.01	0.7412±0.0005
	DSelect-k	<b>97.45±0.00</b>	0.7273±0.0003
	<i>COMET+</i>	<b>97.45±0.00</b>	<b>0.7257±0.0004</b>
Image Datasets			
	Gate	Test Accuracy (Averaged across tasks)	
MultiFashionMNIST	Softmax	87.99±0.04	
	Top-k	88.03±0.03	
	DSelect-k	87.42±0.04	
	<i>COMET+</i>	<b>88.12±0.04</b>	
CelebA	Softmax	83.84±0.15	
	Top-k	83.95±0.17	
	DSelect-k	83.53±0.06	
	<i>COMET+</i>	<b>84.27±0.08</b>	
Digits	Softmax	93.53±0.12	
	Top-k	95.34±0.04	
	DSelect-k	95.41±0.03	
	<i>COMET+</i>	<b>95.45±0.04</b>	
MultiMNIST	Softmax	98.01±0.03	
	Top-k	98.01±0.02	
	DSelect-k	98.03±0.02	
	<i>COMET+</i>	<b>98.07±0.02</b>	

We show that gates that learn sparse routing decisions per sample, e.g., Top-k, DSelect-k, *COMET*, significantly reduce the number of FLOPs ( $3\times-6\times$ ) at inference time in comparison to dense gates e.g., Softmax. Additionally, we see that in all 4 cases, *COMET* has smaller number of FLOPs ( $1.1\times-1.6\times$ ) than the highly popular Topk gate. We also outperform DSelect-k in some cases in number of FLOPs. While in some cases, we have larger number of FLOPs than DSelect-k, our AUC is higher (up to 9%) in these cases.

### D.3 Effect of local search on computation

*Inference.* Note that the permutation matrix is global and not sample specific. At inference time, multiplying permutation matrix  $P$  with  $g(x)$  amounts to a reordering of the expert indices – hence, additional cost for this permutation is negligible compared to evaluation of  $f(x)$  and  $g(x)$ .

*Training.* In the first stage of *COMET* training (a few epochs  $\sim 5$ ), the training is dense (requiring all experts per sample). For *COMET+*, we also learn the permutation matrix during this stage. There is a small additional computational cost: (a) permutation matrix of size  $n \times n$ , where  $n$  is the number of experts, e.g., 16; (b) cost of Sinkhorn operator which constitutes row/column sum normalizations. This cost is marginal compared to the cost of evaluating the experts  $f'_i$ 's, each of which is an MLP/CNN. In the second stage of training, where the samples are being routed to a small  $k(= 2)$  subset of experts per-sample, there is no additional cost for *COMET* vs *COMET+*. To show an example, for MovieLens 200k, where we learn permutation matrix in first 5 epochs, the total time for 50 epochs (on 4 GPUs) is given by: 494s for *COMET* and 496s for *COMET+*. Note 50 epochs were sufficient to achieve convergence for both gates.

## E TASK-SPECIFIC METRICS CORRESPONDING TO TABLES 2 AND 3

We provide task-specific metrics for all recommender systems and image datasets in Table 7. We observe *COMET+* can give superior AUC performance by up to 13% over Hash routing and Top-k, and 9% over DSelect-k.

## F BOOTSTRAPPING PROCEDURE FOR STUDYING HYPERPARAMETER TUNING

We performed 500 tuning trials and performed a bootstrapping procedure as outlined below:

- Randomly sample  $s$  ( $s \in \{1, 2, 5, 10, 15, \dots, 250\}$ ) trials from the bag of a larger set of 500 trials.
- Find the trial with the best validation loss.
- Compute the test loss for that trial.
- Repeat this exercise for 1000 times.
- Compute the average test loss across the best selected trials.