

Sweeping Heterogeneity with Smart MoPs: Mixture of Prompts for LLM Task Adaptation

Chen Dun^{1*}, Mirian Del Carmen Hipolito Garcia^{2*}, Guoqing Zheng², Ahmed Hassan Awadallah²,
Robert Sim², Anastasios Kyrillidis^{1, 3}

¹Rice University, Computer Science Department

²Microsoft

³Ken Kennedy Institute at Rice University

Abstract

Prompt instruction tuning is a popular approach to better adjust pretrained LLMs for specific downstream tasks. How to extend this approach to simultaneously handle multiple tasks and data distributions is an interesting question. We propose *Mixture of Prompts* (MoPs) with smart gating functionality. Our proposed system identifies relevant skills embedded in different groups of prompts and dynamically weighs experts (i.e., collection of prompts) based on the target task. Experiments show that MoPs are resilient to model compression, data source, and task composition, making them highly versatile and applicable in various contexts. In practice, MoPs can simultaneously mitigate prompt training “interference” in multi-task, multi-source scenarios (e.g., task and data heterogeneity across sources) and possible implications from model approximations. Empirically, MoPs show particular effectiveness in compressed model scenarios, while maintaining favorable performance in uncompressed settings: MoPs can reduce final perplexity from 9% up to 70% in non-i.i.d. distributed cases and from 3% up to 30% in centralized cases, compared to baselines.

Introduction

Advances in large language models (LLMs) show that they are powerful general-purpose models (Bubeck et al. 2023), with the ability to solve different tasks out of the box. (Soft) prompt instruction tuning helps in this direction by finetuning deployed models to adjust to individual downstream tasks, without the need of full-model finetuning (Lester, Al-Rfou, and Constant 2021; Kenton et al. 2021; Tamkin et al. 2021).¹

Yet, vanilla finetuning procedures do not always behave well in some practical scenarios. For instance, consider the following: A company XYZ intends to develop a general-purpose office assistant application that solves different tasks at the same time. A desired solution would be to create a

family of prompts that can be combined on-the-fly and tackle a wide variety of incoming tasks, without heavy supervision. With that goal in mind, company XYZ uses both human “labelers” to generate demonstration data of related tasks (maybe, stored in a central server) and client data to utilize their own local demonstration.

Most existing approaches suggest the “centrally training and finetuning” solution, where all data (company-owned and private client) are gathered, and a pre-trained model is further finetuned. However, beyond the privacy concerns due to centrally collecting all data, such an approach, while it would work for the tasks and data included in the training dataset, it would not necessarily adapt to new incoming tasks. Further, company XYZ desires to decrease both training and inference costs of the final model by deploying highly compressed models and aiming for the generation of specialized “experts” that can be used on the fly and just-in-time for most incoming clients, without necessarily requiring further finetuning.²

Overview of our approach and contributions. Such scenarios suggest a *multi-source, multi-task prompt tuning* approach. We consider scenarios where the system does not just face data from a single data distribution but has to learn how to handle data and task heterogeneity. Following recent literature (Mittal, Bengio, and Lajoie 2022), our emphasis here is on training specialized prompts that operate in a modular way such that, when combined, they tackle tasks in a just-in-time manner.

We propose to use *Mixture of Prompts* (or MoPs) in *multi-source, multi-task prompt instruction tuning* to efficiently leverage all available data from both the central server and local clients while using highly compressed models.

The use of MoPs is guided by a gating functionality that can identify relevant skills embedded in different groups of prompts (“experts” in this work) based on the data domain of the current input. This approach dynamically selects a “soft” combination of relevant prompts, which we hypothesize is critical to avoid the appearance of implicit “interference” during training. Simply put, if prompts are not aware of multiple data sources and we allow all prompts to “follow” (=be trained toward) the streaming incoming tasks and data

*These authors contributed equally.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Similar attempts created the term *parameter-efficient finetuning* (PEFT) methods (Houlsby et al. 2019), including adapter tuning (Houlsby et al. 2019), prefix tuning (Li and Liang 2021), prompt tuning (Lester, Al-Rfou, and Constant 2021), low-rank adaptation (LoRA) (Hu et al. 2021), few-shot tuning (IA³) (Liu et al. 2022a) and compression aware prompts (Xu et al. 2023); in this work, without loss of generality and as a proof-of-concept, we focus on soft prompt instruction pruning.

²The definition of an “expert” here will be apparent later in the text; this should not necessarily be assumed as MLP experts in a sparse mixture of expert scenarios (Puigcerver et al. 2023).

distributions, then all prompts would learn the same things and lead to plain “average” solutions during training as data and tasks change. The latter is similar to “simple aggregation” rules using PEFT techniques, where one averages the updated prompts. In contrast, we propose:

- **Tackling task/data heterogeneity.** MoPs could utilize either solely centralized data collected by human “labelers”, heterogeneous local data (e.g., stored on edge devices), or a combination of those while being agnostic about the composition of instruction data.
- **Adaptable Computing.** The design of MoPs enables flexibility in the model architecture, as the prompts (experts) can be placed in any layer rather than just the first layer. This reduces the computational cost of training, as fewer layers need to be backpropagated.
- **Model compression resiliency.** We observe an emerging ability of MoPs: they often work *out of the box*, regardless of reasonable model compression ratios or techniques (i.e., pruning, quantization). We note that when the model is not as strong as the full precision model, it is even more difficult for prompts to be trained correctly and learn diverse things. MoPs improve existing baselines, demonstrating their effectiveness and robustness.
- **Empirical performance.** MoPs manage to decrease final perplexity from $\sim 9\%$ up to $\sim 70\%$, as compared to baselines in the federated scenario, and from $\sim 3\%$ up to $\sim 30\%$ in the centralized scenario. Our gains in the non-i.i.d. (federated) setup support our hypothesis that MoPs effectively improve upon data heterogeneity under skewed distributions, thus reducing the model drift problem.

Related Works

While multi-task learning and multi-source learning in LLMs have been considered in the past (Radford et al. 2021; Bubeck et al. 2023), to our knowledge, there is limited work on PEFT methods that satisfy the above desiderata.

From the federated learning perspective, (Babakniya et al. 2023; Chen et al. 2023) considers the federated version of LoRA (Hu et al. 2021); (Zhang et al. 2023c) considers the federated version of adapters; while (Jiang, Svoboda, and Lane 2023) suggests ongoing pretraining of the full model for better domain adaptation, based on the findings in (Gururangan et al. 2020). Yet, to our understanding, *these works focus primarily on the periodic aggregation and averaging of the PEFT-based parameters*, without targeting on specialized experts. Concurrent work on multiple prompts (Si et al. 2023; Asai et al. 2022) assumes a prior knowledge of skill-s/tasks and uses hand-designed “expert” prompts. These do not consider multi-source data heterogeneity.

Inspired by studies on the linear connectivity of trained models in a full finetuning setting, (Zhang et al. 2023b) considers composing PEFT modules via linear arithmetic operations, drawing connections with the word-embedding hypothesis (like the famous “queen = king - man + woman” equation (Mikolov et al. 2013)). Such efforts are interesting but orthogonal to this work and could be combined as a future direction; in comparison, our gating function weighs and combines prompts (“experts”) on the fly, where their weights can be added or subtracted in a learnable way,

potentially providing more flexibility in how modules are combined.

AdapterSoup (Chronopoulou et al. 2023) averages the weights of the adapters that are most related to the new domain to improve out-of-domain performance at test time. This resembles FedAvg in federated learning (McMahan et al. 2017), where the weights used to aggregate models (here, PEFT modules) are not necessarily uniform. (Chronopoulou et al. 2023) use text clustering and semantic similarity on language tasks to compute the aggregation weights during testing; in our scenario, we are learning the gating function weights to automatically decide which experts and how much they should be combined.

(Zhang et al. 2023a) considers the incremental parameter allocation in LoRA to accommodate finetuning in new tasks. To improve adapter capacity without increasing parameters or computational cost, AdaMix in (Wang et al. 2022) introduces multiple shared adapter components, with sparse random routing during training, that are later on merged via averaging to a single PEFT module; there, the authors do not consider the case of multiple-source/multiple-task scenarios, but focus on the efficiency component. AdaFusion in (Pfeiffer et al. 2021) learns to combine adapters specific to different tasks, as well as a shared pretrained model, by introducing an additional attention layer that fuses all the above during additional training. These parameters learn to combine the adapters as a dynamic function of the target task data. SMEAR in (Muqeeth, Liu, and Raffel 2023) uses a given router’s distribution to average the parameters of the corresponding experts and then routes the input through a single merged expert, which is better than discrete routing.

In (Mahabadi et al. 2021), the authors employ adapter modules within the layers of a pretrained model. Table 4 compares our method to a similar technique, such as LoRA, a widely adopted method that involves efficient finetuning by updating the model weights. Our results show that MoPs show merit compared to prior work.

In (Wu et al. 2023), the authors create a bank of prompts, which integrates cross-task features from diverse sources, serving as an essential component for initializing the prompts before finetuning specific tasks. In contrast, MoPs learn task similarities from scratch under the supervision of a gating function that learns to scale only the relevant experts according to the current input.

During the review process, we became aware of the work of (Wang et al. 2024)³. While there are similarities in using MoE-style routing, our work differs in: *i*) mid-network prompt injection enabling computational efficiency, *ii*) effectiveness in federated learning with heterogeneous data distribution, and *iii*) explicit focus on compressed models.

Background

LLMs and Decoder-only Transformers. The backbone of LLMs are decoder-only transformers (Vaswani et al. 2017; Liu et al. 2018). An LLM takes as input a question/context and performs next-word prediction to generate responses to

³Note that (Wang et al. 2024) paper has been an implementation that follows our approach chronologically on arXiv.

the question. The forward pass of the ℓ -th layer is shown below; for all head indices $h \in \{1, 2, \dots, H\}$:

$$\begin{aligned} \mathbf{Q}^{h,\ell} &= \mathbf{W}_q^{h,\ell} \mathbf{X}^\ell \in \mathbb{R}^{d_h \times n} \\ \mathbf{K}^{h,\ell} &= \mathbf{W}_k^{h,\ell} \mathbf{X}^\ell \in \mathbb{R}^{d_h \times n} \\ \mathbf{A}^{h,\ell} &= \text{Softmax}(\mathbf{M} \odot (\mathbf{Q}^{h,\ell \top} \mathbf{K}^{h,\ell})) \in \mathbb{R}^{n \times n}; \\ \mathbf{V}^{h,\ell} &= (\mathbf{W}_v^{h,\ell} \cdot \mathbf{X}^\ell) \cdot \mathbf{A}^{h,\ell} \in \mathbb{R}^{d_h \times n}; \\ \mathbf{O}^\ell &= \mathbf{W}_o^\ell \cdot \text{Concat}(\mathbf{V}^{1,\ell}, \dots, \mathbf{V}^{H,\ell}) \in \mathbb{R}^{d_o \times n}; \\ \mathbf{X}^{\ell+1} &= \mathbf{W}_{\text{ff2}}^\ell (\text{ReLU}(\mathbf{W}_{\text{ff1}}^\ell \mathbf{O}^\ell)) \in \mathbb{R}^{d_t \times n}. \end{aligned}$$

In particular, let d_h be the dimension of the attention head, d_t the dimension of the input token embedding, d the width of the feedforward layer, H the number of attention heads, and n the input sequence length. Here, $\text{Concat}(\mathbf{B}, \mathbf{C})$ —with \mathbf{B} and \mathbf{C} of appropriate dimensions—concatenates the two matrices columnwise. In the above expressions, $\mathbf{X}^\ell \in \mathbb{R}^{d_t \times n}$ is the input to the ℓ -th layer (i.e., when $\ell = 0$, this is the data input); $\mathbf{W}_q^{h,\ell}, \mathbf{W}_k^{h,\ell}, \mathbf{W}_v^{h,\ell} \in \mathbb{R}^{d_h \times d_t}$ are the weight matrices associated with the query, key and value embedding inputs, where we use for simplicity the same dimension d_h for all of them; $\mathbf{W}_o^\ell \in \mathbb{R}^{d_o \times (H \cdot d_h)}$ is the weight matrix associated with the output of the attention head before the FFN layer, $\mathbf{M} \in \mathbb{R}^{n \times n}$ is the decoder attention mask that prevents positions from attending to the future. Finally, $\mathbf{W}_{\text{ff1}}^\ell \in \mathbb{R}^{d_1 \times d_o}$ and $\mathbf{W}_{\text{ff2}}^\ell \in \mathbb{R}^{d_t \times d_1}$ are the weight matrices of the fully-connected layers.

LLMs with Trainable Prompts: Following (Ouyang et al. 2022; Kenton et al. 2021; Bender et al. 2021; Tamkin et al. 2021), we consider trainable (soft) prompts to perform efficient instruction tuning on LLMs. Using similar notation and additional K trainable prompts $\mathbf{P}^\ell \in \mathbb{R}^{d_t \times K}$, for some $\ell \in \{1, \dots, L\}$, the forward pass of the ℓ -th module where prompts apply can be formulated as below (blue text highlights the main differences with the above):

$$\begin{aligned} \mathbf{B}^\ell &= \text{Concat}(\mathbf{P}^\ell, \mathbf{X}^\ell) \in \mathbb{R}^{d_t \times (n+K)} \\ \mathbf{Q}^{h,\ell} &= \mathbf{W}_q^{h,\ell} \mathbf{B}^\ell \in \mathbb{R}^{d_h \times (n+K)} \\ \mathbf{K}^{h,\ell} &= \mathbf{W}_k^{h,\ell} \mathbf{B}^\ell \in \mathbb{R}^{d_h \times (n+K)} \\ \mathbf{A}^{h,\ell} &= \text{Softmax}(\mathbf{M}' \odot (\mathbf{Q}^{h,\ell \top} \mathbf{K}^{h,\ell})) \in \mathbb{R}^{(n+K) \times (n+K)}; \\ \mathbf{V}^{h,\ell} &= (\mathbf{W}_v^{h,\ell} \cdot \mathbf{B}^\ell) \cdot \mathbf{A}^{h,\ell} \in \mathbb{R}^{d_h \times (n+K)}; \\ \mathbf{O}^\ell &= \mathbf{W}_o^\ell \cdot \text{Concat}(\mathbf{V}^{1,\ell}, \dots, \mathbf{V}^{H,\ell}) \in \mathbb{R}^{d_o \times (n+K)}; \\ \text{Concat}(\mathbf{P}^{\ell+1}, \mathbf{X}^{\ell+1}) &= \mathbf{W}_{\text{ff2}}^\ell (\text{ReLU}(\mathbf{W}_{\text{ff1}}^\ell \mathbf{O}^\ell)) \in \mathbb{R}^{d_t \times (n+K)}, \end{aligned}$$

where $\mathbf{W}_q^{h,\ell}, \mathbf{W}_k^{h,\ell}, \mathbf{W}_v^{h,\ell}, \mathbf{W}_{\text{ff1}}^\ell, \mathbf{W}_{\text{ff2}}^\ell$ are all *frozen*. How \mathbf{P}^ℓ are treated (i.e., frozen or trainable) for different ℓ values will be described in the text. $\mathbf{M}' \in \mathbb{R}^{(n+K) \times (n+K)}$ is the modified decoder attention mask where all prompts are never masked out for all input tokens. We omit skip connections and layer normalization to simplify notations.

Injection of prompts. Inspired by (Li and Liang 2021; Liu et al. 2022b), we propose next to reduce communication costs by injecting trainable prompts in the intermediate layers. By enabling the insertion of these prompts on different layers, we can significantly reduce the backpropagation time, thus

improving the speed and efficiency of the training process. As illustrated in Figure 1, this approach allows for an adjustable design tailored to different setups.

Prompt-tuning in Federated Learning: Recent approaches adapt FedAvg (McMahan et al. 2017) to prompts tuning (Zhao et al. 2023; Babakniya et al. 2023): During synchronization, all updated copies of prompts are averaged on the server for the next round of training. *This is in contrast with this work: while the idea of mixing prompts is not new, we are focusing on learning relevant skills as expressed via selected subsets of prompts based on the data domain of the current input and dynamically selecting the combination of appropriate prompts to solve current and new tasks.*

MoPs With A Smart Gating Function

Our hypotheses in a nutshell. Training prompts to handle *universally* multi-source/multi-task scenarios might result in *prompt interference* across tasks and sources. As our hypothesis, a way *prompt interference* can be decomposed is:

- In centralized training, prompts might converge to poor-performing parameter configurations when heterogeneous tasks are considered due to *conflicting training signals from different tasks* (Žliobaitė 2010). This case is challenging when the tasks are distinctly diverse.
- The above resembles the scenario of *model drifting* in federated learning (Li et al. 2020; Karimireddy et al. 2020; Li, He, and Song 2021): when different clients “pull” the trainable model to work well on their local data, the aggregated model across clients results in a poor-performing final model. In such privacy-preserving scenarios, *heterogeneous data distributions* add more training interference across clients. The model can be biased towards the tasks with more data, losing its capability for generalization.
- For efficiency reasons, compressed LLMs (Frantar and Alistarh 2023; Frantar et al. 2022; Sun et al. 2023; Jaiswal et al. 2023; Ji, Cao, and Liu 2023; Ma, Fang, and Wang 2023; Liu et al. 2023; Kim et al. 2023, 2021; Dettmers et al. 2023b,a; Lin et al. 2023) are now widely used for both centralized and FL scenarios. Such model approximations could impose implicit prompt training interference (and thus leave room to improve performance), as trainable prompts are responsible for *both recovering model capacity loss and model adaptation for downstream tasks*.

Mixture of Experts (MoPs) Design

Half through the procedure: Basics & Pretraining. Consider $\mathbf{P}^1 \in \mathbb{R}^{d_t \times K}$ as some frozen pretrained prompts at the input layer, and \mathbf{X}^1 is the input data. I.e., we assume that either \mathbf{P}^1 is provided to us or we first train \mathbf{P}^1 for a few epochs using the basic neural network (i.e., without the use of mid-prompt injection or any gating function). E.g., the pretrained prompts can be obtained through a brief pretraining phase using the basic architecture without gating or mid-layer injection. This means that $\mathbf{B}^1 = \text{Concat}(\mathbf{P}^1, \mathbf{X}^1) \in \mathbb{R}^{d_t \times (n+K)}$ in the description in Section 2 represents the concatenation of the input tokens with these frozen pretrained prompts; see Figure 1(e). Note also that $\mathbf{W}_q^{h,\ell}, \mathbf{W}_k^{h,\ell}, \mathbf{W}_v^{h,\ell}$,

$\mathbf{W}_{\text{ff1}}^\ell, \mathbf{W}_{\text{ff2}}^\ell$ are all *pretrained and frozen*. As we show in the experiments, the pretrained prompts serve empirically two crucial purposes: *i*) They provide a stable initialization point for the gating mechanism, helping to avoid cold-start problems where the model must simultaneously learn both basic task adaptation and expert specialization. *ii*) They act as a form of regularization, preventing overfitting of the trainable experts to specific tasks or data distributions.

Each layer ℓ of our architecture follows the description in the **LLMs with Trainable Prompts** Section, leading to the output $\text{Concat}(\mathbf{P}^{\ell+1}, \mathbf{X}^{\ell+1})$, where the latter is split into “prompts” and regular output, based on their positions. We consider the representation $\mathbf{P}^\ell, \ell > 1$ as the “transformed prompt”: i.e., the result as we propagate prompts through the layers. *We use this split for ease of notation and consistency among layers*. Note that, while \mathbf{P}^1 is trainable during the pretraining phase, the rest $\mathbf{P}^\ell, \ell > 1$ represent the output of intermediate layers, and are not trainable. See also the first part of Algorithm 1.

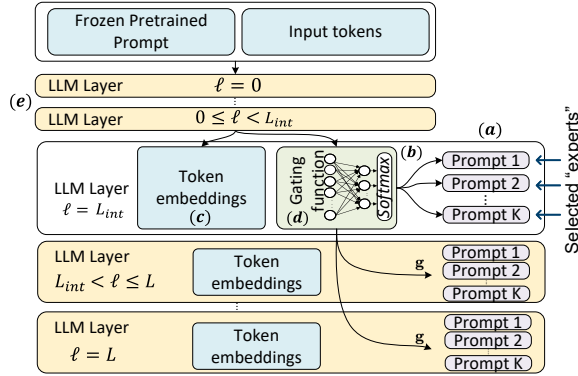


Figure 1: Mixture of prompts with a smart gating function on compressed LLMs overview.

During MoP training: Prompts as experts. Given this initial setting where we have $\mathbf{W}_q^{h,\ell}, \mathbf{W}_k^{h,\ell}, \mathbf{W}_v^{h,\ell}, \mathbf{W}_{\text{ff1}}^\ell, \mathbf{W}_{\text{ff2}}^\ell$ and \mathbf{P}^1 *pretrained and frozen*, we apply *prompt injection* (Li and Liang 2021; Liu et al. 2022b): at an intermediate layer level $L_{\text{int}} \in \{1, \dots, L\}$, we inject multiple (say, K) trainable prompts as experts to embed different skills across subtasks, each specializing in different skills; see Figure 1(a). These prompts are trainable in the following sense: At this training phase, they are weighted by a gating function (see Figure 1(b) and (d)), and as described below), depending on the current input. The input to the gating function is designed to be the embedding of the network’s input via the first $\ell < L_{\text{int}}$ layers; see Figure 1(c).⁴ Prompts are selected and updated

⁴Beyond a design choice, we inject the prompts in the intermediate layers as a means to reduce the training cost. An alternative –and fully valid– procedure would be to use two LLMs, where the first is used as an embedding mechanism, and the other uses the embeddings as inputs to the gating function. To perform fair comparisons in our experiments, we decided not to change the size of a given LLM but rather change its design in this way.

based on the output weights of the gating function; these weights mitigate the training interference between prompts. This allows us to use different combinations of skills for various input tasks.

The gating function. To dynamically weigh prompts based on the input, we have designed a gating function that embeds the present question. When $\ell = L_{\text{int}}$, the embedding $\text{Concat}(\mathbf{P}^{\ell+1}, \mathbf{X}^{\ell+1}) \in \mathbb{R}^{d_t \times (n+K)}$ gets averaged wrt its “sequence length” to obtain $\text{avg_emb} \in \mathbb{R}^{d_t}$; see Algorithm 1. I.e., to avoid incurring extra computation and memory costs, as is common in MoEs, our gating function directly utilizes the first part of the given model ($1 \leq \ell \leq L_{\text{int}}$) as the embedding network, without any additional cost.

The gating function comprises a shallow MLP network that takes the intermediate layer (averaged) embeddings as input, followed by a softmax layer that outputs expert scores:

$$\mathbf{g} = \text{Softmax}(\mathbf{W}_{\text{gff2}}(\text{Relu}(\mathbf{W}_{\text{gff1}}(\text{avg_emb})))) \in \mathbb{R}^K.$$

$\mathbf{W}_{\text{gff1}} \in \mathbb{R}^{d_g \times d_t}$ and $\mathbf{W}_{\text{gff2}} \in \mathbb{R}^{K \times d_g}$ are trainable weights.

After the intermediate layer. Scores in $\mathbf{g} \in \mathbb{R}^K$ are used to scale the network’s attention over the collection of experts in the subsequent layers ($L_{\text{int}} \leq \ell \leq L$), as shown in Figure 1(d). In particular, these scores scale accordingly the weights of the prompt embeddings $\mathbf{B}^\ell = \text{Concat}(\hat{\mathbf{P}}^\ell, \mathbf{X}^\ell)$, so that those related to the task at hand are more crucial than the others, shown mathematically as follows:

$$\hat{\mathbf{B}}^\ell = \mathbf{B}_{:,1:K}^\ell \odot \mathbf{g} \in \mathbb{R}^{d_t \times (n+K)}, \quad \forall L_{\text{int}} \leq \ell \leq L.$$

Here, \odot represents the Hadamard product; note that $\mathbf{B}_{:,1:K}^\ell$ represents the K first columns of \mathbf{B}^ℓ . We overload the Hadamard multiplication to indicate that \mathbf{g} applies to all the rows of this restricted matrix, recursively. This is highlighted in Algorithm 1. Using softmax output score, the gating function “forces” later layers to focus mostly on dominating selected prompts, which scales the updates for each prompt accordingly during backpropagation. Our gating function imposes a negligible computation overhead in total.

Pretraining the gating function. To improve the initial performance of our gating function, we assume we have unlabeled instruction data (instruction/question only) with domain/task labels on the server side. As such data are instruction only, we can collect such data in both centralized and federated learning cases beforehand.

To use this data, we manually assign a one-to-one relationship between each prompt group and each data domain/task. This provides a good initialization to the gating function, assuming that *i*) each subtask is drastically different and represents one distinct skill, and *ii*) each prompt embeds the corresponding skill. Such an assumption does not need to be accurate for the available dataset. As shown in the experiments, such initialization is good enough: the gating function, together with trainable prompts, can discover a more accurate relationship between subtasks.

Using compressed LLMs for efficient prompt tuning. Compressed LLMs (Frantar and Alistarh 2023; Frantar et al. 2022; Sun et al. 2023) are widely used for downstream instruction tuning due to training efficiency concerns. Overall,

Input: $\mathbf{W}_q^{h,\ell}, \mathbf{W}_k^{h,\ell}, \mathbf{W}_v^{h,\ell}, \mathbf{W}_{\text{ff1}}^\ell, \mathbf{W}_{\text{ff2}}^\ell, \ell \in [L]$ pretrained and frozen; \mathbf{P}^1 pretrained and frozen (see text); $L_{\text{int}} \in \{1, \dots, L\}$; input data \mathbf{X}^1 .

Details: \odot denotes row-wise element; we inject prompts $\hat{\mathbf{P}}^{L_{\text{int}}} \in \mathbb{R}^{d_t \times K}$ at the layer $\ell = L_{\text{int}}$.

♠ **Before intermediate layer & gating function** ♠

```

for  $1 \leq \ell < L_{\text{int}}$  do
   $\mathbf{B}^\ell = \text{Concat}(\mathbf{P}^\ell, \mathbf{X}^\ell)$ 
  for  $1 \leq h \leq H$  do
     $\mathbf{Q}^{h,\ell} = \mathbf{W}_q^{h,\ell} \mathbf{B}^\ell$ 
     $\mathbf{K}^{h,\ell} = \mathbf{W}_k^{h,\ell} \mathbf{B}^\ell$ 
     $\mathbf{A}^{h,\ell} = \text{Softmax}(\mathbf{M}^\ell \odot (\mathbf{Q}^{h,\ell \top} \mathbf{K}^{h,\ell}))$ ;
     $\mathbf{V}^{h,\ell} = (\mathbf{W}_v^{h,\ell} \cdot \mathbf{B}^\ell) \cdot \mathbf{A}^{h,\ell}$ 
  end for
   $\mathbf{O}^\ell = \mathbf{W}_o^\ell \cdot \text{Concat}(\mathbf{V}^{1,\ell}, \dots, \mathbf{V}^{H,\ell})$ 
   $\text{Concat}(\mathbf{P}^{\ell+1}, \mathbf{X}^{\ell+1}) = \mathbf{W}_{\text{ff2}}^\ell(\text{Relu}(\mathbf{W}_{\text{ff1}}^\ell \mathbf{O}^\ell))$ 
end for
    
```

♠ **Intermediate layer & gating function** ♠

if $\ell = L_{\text{int}}$ then
 Let $\text{avg_emb} \in \mathbb{R}^{d_t}$ such that:

$$\text{avg_emb} := \frac{1}{n+K} \sum_{i=1}^{n+K} \text{Concat}(\mathbf{P}^{L_{\text{int}}}, \mathbf{X}^{L_{\text{int}}})_{:,i}$$

Compute $\mathbf{g} \in \mathbb{R}^K$ through the gating function:

$$\mathbf{g} = \text{Softmax}(\mathbf{W}_{\text{gff2}}(\text{Relu}(\mathbf{W}_{\text{gff1}}(\text{avg_emb}))))$$

end if

♠ **After intermediate layer** ♠

```

for  $L_{\text{int}} \leq \ell < L$  do
   $\mathbf{B}^\ell = \text{Concat}(\hat{\mathbf{P}}^\ell, \mathbf{X}^\ell)$ 
   $\hat{\mathbf{B}}^\ell = \mathbf{B}_{:,1:K}^\ell \odot \mathbf{g}$ 
  for  $1 \leq h \leq H$  do
     $\hat{\mathbf{Q}}^{h,\ell} = \mathbf{W}_q^{h,\ell} \hat{\mathbf{B}}^\ell$ 
     $\hat{\mathbf{K}}^{h,\ell} = \mathbf{W}_k^{h,\ell} \hat{\mathbf{B}}^\ell$ 
     $\hat{\mathbf{A}}^{h,\ell} = \text{Softmax}(\mathbf{M}^\ell \odot (\hat{\mathbf{Q}}^{h,\ell \top} \hat{\mathbf{K}}^{h,\ell}))$ ;
     $\hat{\mathbf{V}}^{h,\ell} = (\mathbf{W}_v^{h,\ell} \cdot \hat{\mathbf{B}}^\ell) \cdot \hat{\mathbf{A}}^{h,\ell}$ 
  end for
   $\hat{\mathbf{O}}^\ell = \mathbf{W}_o^\ell \cdot \text{Concat}(\hat{\mathbf{V}}^{1,\ell}, \dots, \hat{\mathbf{V}}^{H,\ell})$ 
   $\text{Concat}(\hat{\mathbf{P}}^{\ell+1}, \mathbf{X}^{\ell+1}) = \mathbf{W}_{\text{ff2}}^\ell(\text{Relu}(\mathbf{W}_{\text{ff1}}^\ell \hat{\mathbf{O}}^\ell))$ 
end for
    
```

given the large scale of LLMs, using them out-of-the-box for deployment scenarios is often infeasible due to their resource-intensive requirements. Our system, depicted in Figure 1, follows this paradigm by utilizing *compressed LLMs*. We add prompts only to the middle layers, thus allowing flexibility to avoid backpropagation of the entire model during training.

The above are summarized in Algorithm 1. Briefly, given an input question, MoP first embeds the question using the first $< L_{\text{int}}$ layers of a given compressed LLM. We set $L_{\text{int}} = 10$ for a LLaMA-7B model with $L = 32$. Such choice of $L_{\text{int}} = 10$ is to balance two conflicting requirements: *i*) we want to inject prompts as late as possible to reduce back propagation cost during training; *ii*) prompts should be injected early to have more capacity in influencing the pretrained LLM network. In the Experiments Section, we provide a detailed analysis of the trade-offs of injecting prompts on different layers. At layer L_{int} , we inject K trainable prompts. The gating network uses the previous layer’s embedding to generate expert scores for each prompt based on the input question, which is used to re-scale attention weight from other tokens to those prompts. It follows the normal LLM forward propagation after layer L_{int} .

Computational overhead and scalability. The computational cost of our gating function is relatively modest, requiring only a shallow MLP network that processes the averaged embeddings from intermediate layers. Specifically, it requires $O(d_t \cdot d_g + d_g \cdot K)$ operations. For LLaMA-7B, this represents less than 0.1% overhead in total computation time. In federated settings, the gating computation occurs locally on each client, with only the expert weights being synchronized.

Dataset	Dolly-15K In-structions	Super-Natural Instructions
Subtasks	creative writing	quoref-question-generation
	closed QA	drop-question-generation
	open QA	essential-terms-identifying-essential-words
	summarization	add-integer-to-list
	information extraction	evaluation-semantic-relation-classification
	classification	ljsspeech-textmodification
	brainstorming	mmmlu-answer-generation-global-facts
Total	5000 samples	5000 samples

Table 1: Task categories used per dataset

Experiments

Datasets. We used two datasets: Databricks Dolly 15k (Conover et al. 2023) and Super-Natural Instructions (Mishra et al. 2022); see Table 1. These datasets challenge our method: MoPs must learn and select relevant skills from scratch without prior knowledge of the complex relationships between the subtasks. We split the original 5k samples from each dataset into 90% training and 10% testing sets. We used a batch size of 1 for both training and testing. In the federated scenario, we simulated an uneven distribution of data across 100 clients, resulting in different proportions and sizes of data. The batch size remained at 1. The distribution of data skew across clients is explained in Appendix A.

System. We used 4 NVIDIA RTX A6000 GPU with 46 GB Memory. The total training of the prompts took around 2.5 hours (when experts being injected on layer 10th). The federated training was performed in a distributed fashion using 1:1 relationship between expert and worker.

		Unstructured pruning (ratio)			Structured pruning (Type & Ratio)			
Dataset	Method	90%	85%	75%	7:8 (87.5%)	3:4 (75%)	2:4 (50%)	4:8 (50%)
Dolly-15K	Baseline	52.65	18.16	8.25	70.14	9.06	3.67	3.59
	MoPs	40.34	15.04	7.24	54.97	8.08	3.54	3.59
	Gain \pm	+12.31 (30%)	+3.12 (20%)	+1.01 (13%)	+15.17 (27%)	+0.98 (12%)	+0.13 (4%)	+0.17 (5%)
Super-Natural	Baseline	58.47	16.50	8.54	67.86	10.64	6.01	5.90
	MoPs	52.86	14.59	7.80	59.80	10.05	5.79	5.73
	Gain \pm	+5.61 (11%)	+1.91 (13%)	+0.74 (9%)	+8.06 (13%)	+0.59 (6%)	+0.22 (4%)	+0.17 (3%)

Table 2: Summary of final perplexities (\downarrow) for the centralized scenario on Dolly-15 and Super-Natural datasets.

		Unstructured pruning (Ratio)			Structured pruning (Type & Ratio)			
Dataset	Method	90%	85%	75%	7:8 (87.5%)	3:4 (75%)	2:4 (50%)	4:8 (50%)
Dolly-15K	FedPrompt	98.13	28.28	11.99	143.02	17.20	5.09	4.91
	MoPs	65.25	20.77	9.45	84.10	12.20	4.23	4.06
	Gain \pm	+32.88 (50%)	+7.51 (36%)	+2.54 (27%)	+58.92 (70%)	+5.00 (41%)	+0.86 (20%)	+0.85 (21%)
Super-Natural	FedPrompt	76.17	18.64	9.14	91.64	14.42	6.43	6.14
	MoPs	66.51	16.52	7.88	72.04	12.38	5.75	5.65
	Gain \pm	+9.66 (15%)	+2.12 (13%)	+1.26 (16%)	+19.6 (27%)	+2.04 (16%)	+0.68 (12%)	+0.49 (9%)

Table 3: Summary of final perplexities (\downarrow) for the federated scenario, using a pool of 100 available clients.

Setup. We use SparseGPT (Frantar and Alistarh 2023) to perform structured/unstructured pruning and LLM.int8() (Dettmers et al. 2022) to perform Int8 quantization of the LLaMA-7B model, creating different compression ratios of the LLM. Inspired by (Xu et al. 2023), we assign ten prompts to each single expert, totaling seven experts and 70 randomly initialized prompts, ensuring a 1:1 relationship between experts and tasks. Our experiments show that the 1:1 relationship between experts and tasks is not a strict constraint; the gating function adapts to task similarities by dynamically grouping tasks and allocating experts, often using fewer experts than initially assigned (see Appendices B and C for further discussion). We also study the impact of varying the number of prompts per expert.

To show our method can further recover/improve the performance of the pruned model, we add such pretrained prompts to both our baselines and our model. We trained these prompts from scratch in a preprocessing step over 20 training steps. These prompts are frozen during training.

In the centralized setting, we use 20000 steps with a learning rate 0.001. In the federated setting, we adapt FedAvg to average the updated prompts from all active clients during each synchronization round. We use 100 clients, with ten active clients per training round, and set each local training round to 250 training steps. Counting all clients, the total number of training steps is 50000 with a learning rate 0.001.

Baselines. A reasonable baseline is directly applying prompt tuning to centralized and federated training without any gating function. In centralized training, we use the method from (Xu et al. 2023). In federated training, we utilize FedPrompt from (Zhao et al. 2023), which adapts FedAvg to prompt training and periodically averaging the updated prompts from all clients. In both cases, to match computation and memory cost with our method during training, we add additional prompts in $L_{\text{int}} = 10$ and freeze the given pre-trained prompts in the first layer, thus eliminating the need to calculate gradients before L_{int} .

Centralized training results. See Table 2. $X\%$ denotes the pruned model with $X\%$ weight pruned out for unstructured pruning. For structured pruning, we follow (Frantar and Alistarh 2023) to use $N : M$ to denote pruning N elements out of consecutive M elements in the weight matrix. The uncompressed models are already trained extensively on massive data (and actually on data very similar to the scenarios we create for different tasks and data source distributions); thus, one expects these models to work well. Table 4 shows gains over other PEFT methods for scenarios where little space exists for improvements.

Our method reduces the final PPL for all cases, with an advantage for the highest pruning ratios. As the pruning ratio increases, more pressure is placed on prompts to recover skills lost due to the model loss. However, MoPs can help to reduce this burden by providing more capacity for task adaptation, thus alleviating the training interference.

The PPL reduction in the centralized case is more pronounced for unstructured pruning due to lower sparsity. Even if one could argue that the model improvements on higher levels might be less useful, in both Tables 2 and 3 we improve the quality of the models by 5 to 10 PPL units in scenarios that matter. Such an improvement is not trivial: as we do not investigate all possible design choices (best optimizer, best tuning, best transformer architecture, etc), we provide one component that can turn a 17PPL loss to 12PPL loss.

Influence of Injection Layer. MoP can select the layer where the prompts are injected. We explored this "hyperparameter" by injecting the experts on different layers. Figure 2 revealed that injecting the prompts in earlier layers improves the performance, albeit at the cost of "heavier" backpropagation. The plots also showed that if the prompts are injected too late, there are only a few layers after the gating function to adjust the prompts, resulting in poorer performance.

Influence of Number of Prompts Injected (Per expert). We analyzed how the performance is affected by increasing or decreasing the number of prompts per expert. As illustrated

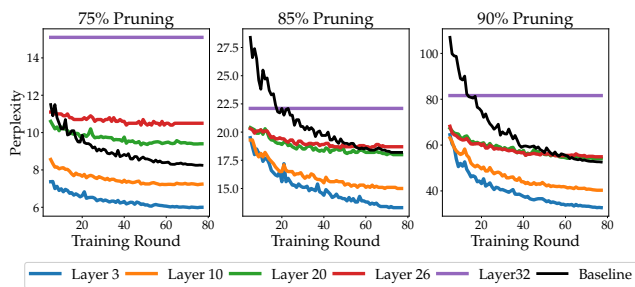


Figure 2: Layer Injection impact on Llama-7B for different unstructured pruning ratios (Dolly-15k) in the centralized setup. Injection on $L_{\text{int}} = 10$ outperforms the baseline.

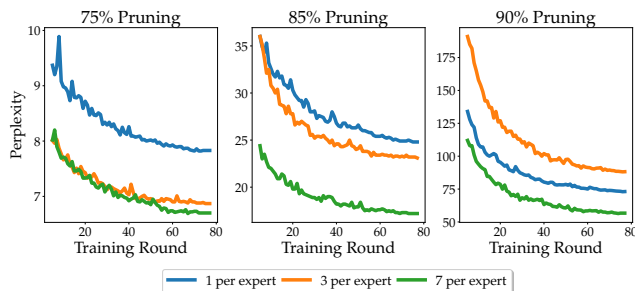


Figure 3: Prompt Injection impact on Llama-7B for different unstructured pruning ratios (Dolly-15k - centralized setup).

in Figure 3, the more prompts injected, the faster convergence is observed. Providing enough granularity through the prompts allows the gating function to learn more quickly and effectively. Thus, providing sufficient prompts to the experts to maximize the overall performance is essential.

Federated training results. Table 3 shows that MoPs is superior to the baseline (FedPrompt) for all pruning ratios. Comparing the relative gains (PPL decrease) with Table 2, our method yields superior gains in the federated setup. But, *why is MoP performing even better in FL settings?* Figure Table 3 demonstrates the gating function’s role in overcoming the data heterogeneity, showing improved performance compared to the baseline. Figure 4 shows how MoPs accentuate the convergence to a set of experts with relatively distinct skills. Starting from experts with relatively more mixed skills (after pretraining), it is interesting to note that, at the end of the training, Expert Groups 0 and 6 are specialized in creative writing, open QA, and brainstorming, which all are related to free text generation tasks; Expert Groups 3 and 4 are specialized on closed QA, summarization and information extraction, which all are related to more restricted contexts; finally, the Expert Group 5 is specialized on classification, which is a category by itself.

Evaluating MoPs against LoRA and AI³. Even though current PEFT approaches such as LoRA (Hu et al. 2021) and AI³ (Liu et al. 2022a) may not be completely compatible in terms of setup, it is still beneficial to compare against

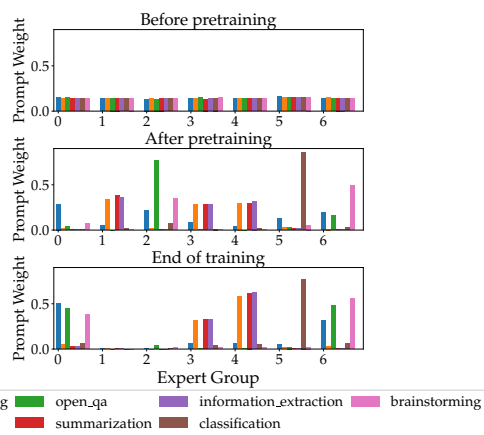


Figure 4: Averaged prompt weights for test dataset using 3:4 (75%) structured pruning Llama-7B. This verifies how MoPs learn experts with a specialized skill set.

them empirically. We conducted experiments to demonstrate how we can orchestrate the different elements of MoPs – injection layer, number of prompts, number of experts – to boost the expert skills and match similar performance as LoRA and AI³ under different pruning ratios. Table 4 reports final perplexities on Prompt-Tuning (Baseline: prompt-tuning without MoPs), MoPs, LoRA, and AI³ after 20k training steps in centralized training. Overall, each method has its strengths: existing methods excel at 90% compression, while MoPs show advantages at lower ratios.

Compression	Baseline	MoPs	LoRA	AI ³	MoPs Gain
Uncompressed	3.24	3.09	3.14	3.11	+0.15 (5%)
<i>Trainable Params</i>	0.01%				
75% Pruned	8.25	6.0	8.7	6.2	+2.25 (38%)
85% Pruned	18.16	12.4	12.4	12.4	+5.76 (46%)
90% Pruned	52.65	26.6	21.4	26.2	+26.05 (98%)

Table 4: Final perplexities; see Appendix E.

Conclusions

MoPs allow the identification of relevant skills for the current task and dynamically select and combine prompts accordingly. This overcomes prompt training interference from multi-tasks across centralized and federated learning scenarios. Results show how the gating functionality boosts soft-prompts to match similar performance as other state-of-the-art PEFT methods and suggest that the gating function helps to overcome model drift problems resulting from heterogeneous data distribution in multi-source learning scenarios.

Acknowledgments

This work is supported by NSF CMMI no. 2037545, NSF CAREER award no. 2145629, Welch Foundation Grant no.

A22-0307, a Microsoft Research Award, an Amazon Research Award, and a Ken Kennedy Research Cluster award (Rice K2I). AK was supported in part through seed funding from the Ken Kennedy Institute at Rice University.

References

- Asai, A.; Salehi, M.; Peters, M. E.; and Hajishirzi, H. 2022. Attentional mixtures of soft prompt tuning for parameter-efficient multi-task knowledge sharing. *arXiv preprint arXiv:2205.11961*, 3.
- Babakniya, S.; Elkordy, A. R.; Ezzeldin, Y. H.; Liu, Q.; Song, K.-B.; El-Khamy, M.; and Avestimehr, S. 2023. SLoRA: Federated Parameter Efficient Fine-Tuning of Language Models. *arXiv preprint arXiv:2308.06522*.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, 610–623. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Chen, C.; Feng, X.; Zhou, J.; Yin, J.; and Zheng, X. 2023. Federated large language model: A position paper. *arXiv preprint arXiv:2307.08925*.
- Chronopoulou, A.; Peters, M. E.; Fraser, A.; and Dodge, J. 2023. AdapterSoup: Weight Averaging to Improve Generalization of Pretrained Language Models. In *Findings of the Association for Computational Linguistics: EACL 2023*, 2009–2018.
- Conover, M.; Hayes, M.; Mathur, A.; Xie, J.; Wan, J.; Shah, S.; Ghodsi, A.; Wendell, P.; Zaharia, M.; and Xin, R. 2023. Free Dolly: Introducing the World’s First Truly Open Instruction-Tuned LLM.
- Dettmers, T.; Lewis, M.; Belkada, Y.; and Zettlemoyer, L. 2022. LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale. *arXiv preprint arXiv:2208.07339*.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023a. QLoRA: Efficient finetuning of quantized LLMs. *arXiv preprint arXiv:2305.14314*.
- Dettmers, T.; Svirschevski, R.; Egiazarian, V.; Kuznedelev, D.; Frantar, E.; Ashkboos, S.; Borzunov, A.; Hoefler, T.; and Alistarh, D. 2023b. SpQR: A Sparse-Quantized Representation for Near-Lossless LLM Weight Compression. *arXiv preprint arXiv:2306.03078*.
- Frantar, E.; and Alistarh, D. 2023. SparseGPT: Massive Language Models Can Be Accurately Pruned in One-Shot. *arXiv preprint arXiv:2301.00774*.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. GPTQ: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Gururangan, S.; Marasović, A.; Swayamdipta, S.; Lo, K.; Beltagy, I.; Downey, D.; and Smith, N. A. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*, 2790–2799. PMLR.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Jaiswal, A.; Liu, S.; Chen, T.; and Wang, Z. 2023. The Emergence of Essential Sparsity in Large Pre-trained Models: The Weights that Matter. *arXiv preprint arXiv:2306.03805*.
- Ji, Y.; Cao, Y.; and Liu, J. 2023. Pruning large language models via accuracy predictor. *arXiv preprint arXiv:2309.09507*.
- Jiang, L.; Svoboda, F.; and Lane, N. D. 2023. FDAPT: Federated Domain-adaptive Pre-training for Language Models. *arXiv preprint arXiv:2307.06933*.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. SCAFFOLD: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, 5132–5143. PMLR.
- Kenton, Z.; Everitt, T.; Weidinger, L.; Gabriel, I.; Mikulik, V.; and Irving, G. 2021. Alignment of Language Agents. *arXiv:2103.14659*.
- Kim, J.; Lee, J. H.; Kim, S.; Park, J.; Yoo, K. M.; Kwon, S. J.; and Lee, D. 2023. Memory-Efficient Fine-Tuning of Compressed Large Language Models via sub-4-bit Integer Quantization. *arXiv preprint arXiv:2305.14152*.
- Kim, S.; Gholami, A.; Yao, Z.; Mahoney, M. W.; and Keutzer, K. 2021. I-BERT: Integer-only BERT quantization. In *International conference on machine learning*, 5506–5518. PMLR.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045–3059.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4582–4597.
- Lin, J.; Tang, J.; Tang, H.; Yang, S.; Dang, X.; and Han, S. 2023. AWQ: Activation-aware Weight Quantization for LLM Compression and Acceleration. *arXiv preprint arXiv:2306.00978*.
- Liu, H.; Tam, D.; Muqeeth, M.; Mohta, J.; Huang, T.; Bansal, M.; and Raffel, C. 2022a. Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning. *arXiv:2205.05638*.

- Liu, P. J.; Saleh, M.; Pot, E.; Goodrich, B.; Sepassi, R.; Kaiser, L.; and Shazeer, N. 2018. Generating Wikipedia by Summarizing Long Sequences. In *International Conference on Learning Representations*.
- Liu, X.; Sun, T.; Huang, X.; and Qiu, X. 2022b. Late Prompt Tuning: A Late Prompt Could Be Better Than Many Prompts. *arXiv preprint arXiv:2210.11292*.
- Liu, Z.; Oguz, B.; Zhao, C.; Chang, E.; Stock, P.; Mehdad, Y.; Shi, Y.; Krishnamoorthi, R.; and Chandra, V. 2023. LLM-QAT: Data-Free Quantization Aware Training for Large Language Models. *arXiv preprint arXiv:2305.17888*.
- Ma, X.; Fang, G.; and Wang, X. 2023. LLM-Pruner: On the Structural Pruning of Large Language Models. *arXiv preprint arXiv:2305.11627*.
- Mahabadi, R. K.; Ruder, S.; Dehghani, M.; and Hender-son, J. 2021. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- Mishra, S.; Khashabi, D.; Baral, C.; and Hajishirzi, H. 2022. Cross-task generalization via natural language crowdsourcing instructions. In *ACL*.
- Mittal, S.; Bengio, Y.; and Lajoie, G. 2022. Is a modular architecture enough? *Advances in Neural Information Processing Systems*, 35: 28747–28760.
- Muqeeth, M.; Liu, H.; and Raffel, C. 2023. Soft Merging of Experts with Adaptive Routing. *arXiv preprint arXiv:2306.03745*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 27730–27744. Curran Associates, Inc.
- Pfeiffer, J.; Kamath, A.; Rücklé, A.; Cho, K.; and Gurevych, I. 2021. AdapterFusion: Non-Destructive Task Composition for Transfer Learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 487–503.
- Puigcerver, J.; Riquelme, C.; Mustafa, B.; and Houlsby, N. 2023. From Sparse to Soft Mixtures of Experts. *arXiv preprint arXiv:2308.00951*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Si, C.; Shi, W.; Zhao, C.; Zettlemoyer, L.; and Boyd-Graber, J. 2023. Mixture of Prompt Experts for Generalizable and Interpretable Question Answering. *arXiv preprint arXiv:2305.14628*.
- Sun, M.; Liu, Z.; Bair, A.; and Kolter, J. Z. 2023. A Simple and Effective Pruning Approach for Large Language Models. *arXiv preprint arXiv:2306.11695*.
- Tamkin, A.; Brundage, M.; Clark, J.; and Ganguli, D. 2021. Understanding the Capabilities, Limitations, and Societal Impact of Large Language Models. *arXiv:2102.02503*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, R.; An, S.; Cheng, M.; Zhou, T.; Hwang, S. J.; and Hsieh, C.-J. 2024. One Prompt is not Enough: Automated Construction of a Mixture-of-Expert Prompts. *arXiv preprint arXiv:2407.00256*.
- Wang, Y.; Mukherjee, S.; Liu, X.; Gao, J.; Awadallah, A. H.; and Gao, J. 2022. AdaMix: Mixture-of-adapter for parameter-efficient tuning of large language models. *arXiv preprint arXiv:2205.12410*, 1(2): 4.
- Wu, M.; Liu, W.; Xu, J.; Lv, C.; Ling, Z.; Li, T.; Huang, L.; Zheng, X.; and Huang, X.-J. 2023. Parameter efficient multi-task fine-tuning by learning to transfer token-wise prompts. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 8734–8746.
- Xu, Z.; Liu, Z.; Chen, B.; Tang, Y.; Wang, J.; Zhou, K.; Hu, X.; and Shrivastava, A. 2023. Compress, Then Prompt: Improving Accuracy-Efficiency Trade-off of LLM Inference with Transferable Prompt. *arXiv preprint arXiv:2305.11186*.
- Zhang, F.; Li, L.; Chen, J.; Jiang, Z.; Wang, B.; and Qian, Y. 2023a. IncreLoRA: Incremental parameter allocation method for parameter-efficient fine-tuning. *arXiv preprint arXiv:2308.12043*.
- Zhang, J.; Chen, S.; Liu, J.; and He, J. 2023b. Composing parameter-efficient modules with arithmetic operations. *arXiv preprint arXiv:2306.14870*.
- Zhang, X.; Li, M.; Chang, X.; Chen, J.; Roy-Chowdhury, A. K.; Suresh, A. T.; and Oymak, S. 2023c. FedYolo: Augmenting Federated Learning with Pretrained Transformers. *arXiv preprint arXiv:2307.04905*.
- Zhao, H.; Du, W.; Li, F.; Li, P.; and Liu, G. 2023. Fed-Prompt: Communication-Efficient and Privacy Preserving Prompt Tuning in Federated Learning. *arXiv:2208.12268*.
- Žliobaitė, I. 2010. Learning under concept drift: an overview. *arXiv preprint arXiv:1010.4784*.