
PiKE: Adaptive Data Mixing for Large-Scale Multi-Task Learning Under Low Gradient Conflicts

Zeman Li^{1,2*} **Yuan Deng²** **Peilin Zhong²** **Meisam Razaviyayn^{1,2}** **Vahab Mirrokni²**

¹University of Southern California ²Google Research

{zemanli, razaviya}@usc.edu

{dengyuan, peilinz, mirrokni}@google.com

Abstract

Modern foundation models are trained on diverse datasets to enhance generalization across tasks and domains. A central challenge in this process is determining how to effectively mix and sample data from multiple sources. This naturally leads to a multi-task learning (MTL) perspective. While prior work in MTL has emphasized mitigating gradient conflicts, we observe that large-scale pretraining scenarios—such as multilingual or multi-domain training—often exhibit little to no gradient conflict. Motivated by this observation, we propose **PiKE** (**P**ositive gradient **i**nteraction-based **K**-task weights **E**stimator), an adaptive data mixing algorithm that dynamically adjusts sampling weights during training. PiKE exploits non-conflicting gradient interactions to minimize a near-tight upper bound on the average loss decrease at each step, while incurring negligible computational overhead. We provide theoretical convergence guarantees and show that PiKE outperforms static and non-adaptive mixing baselines. Furthermore, we extend PiKE to promote balanced learning across tasks. Extensive experiments on large-scale language model pretraining confirm that PiKE achieves faster convergence and improved downstream performance compared to existing approaches.

1 Introduction

Foundation models, such as large language models (LLMs), owe their strong generalization and multitask abilities to pretraining on diverse datasets spanning multiple domains [62, 36, 10]. The effectiveness of these models depend heavily on the composition of their training data [29, 18]. However, standardized practices for curating optimal pretraining data are lacking. Common approaches involve heuristic filtering, deduplication, and categorization into heterogeneous domains (e.g., The Pile [22] has 22 domains; GLaM [18] has 6). Even after such preprocessings, determining the optimal data mixing remains a key challenge—amplified by the scale of modern models and datasets.

A common strategy in pre-training is to use fixed data mixtures, typically chosen heuristically or via smaller proxy models. For example, mT5 [71] weights datasets by relative size, while GLaM [18] uses downstream performance from proxy models. DoReMi [69] also relies on proxy models, using group distributionally robust optimization (group DRO) to set dataset weights. However, these approaches have notable limitations. First, their optimality is unclear: heuristic methods lack theoretical backing, and policies from small models may not necessarily transfer to larger ones [74]. Second, proxy models introduce substantial computational overhead, often scaling linearly or worse with the number of domains. Third, static weights fixed at initialization may become suboptimal as training progresses [77, 34], a limitation discussed further in Section 2.

In this work, we frame adaptive data mixture selection problem as a multitask optimization problem, enabling a principled approach to dynamically adjusting data mixture. This view is natural: each data domain typically is related to a (set of) tasks and yields a distinct gradient. Prior multitask optimization methods [75, 66, 9, 16] focus on resolving gradient conflicts that impede convergence. However,

*Work done while interning at Google Research.

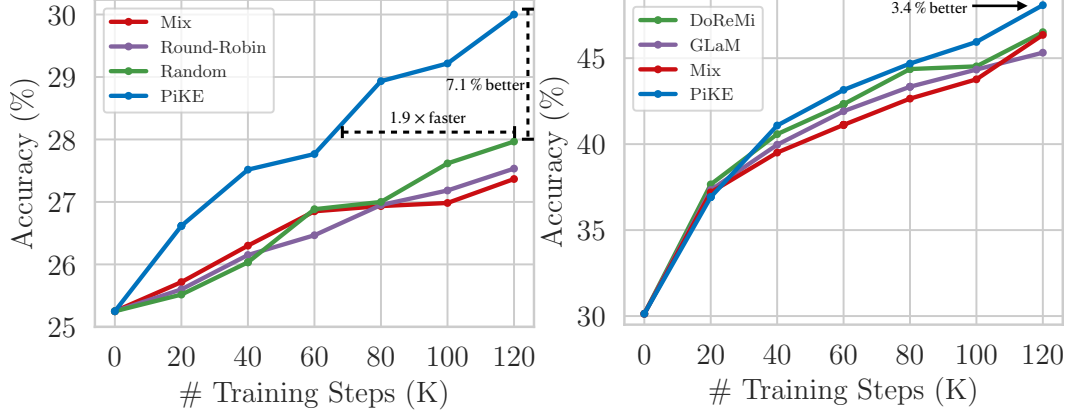


Figure 1: PiKE **adaptively optimizes task weights** in pre-training, outperforming baselines. **Left (1B models, multilingual C4 en/hi)**: PiKE boosts average downstream accuracy by 7.1% and reaches baseline accuracy $1.9\times$ faster. **Right (750M models, GLaM six domains)**: PiKE achieves a 3.4% higher accuracy over DoReMi [69]. PiKE scales efficiently with model size and number of tasks. Detailed results are in Tables 15 and 16.

most are impractical for current LLMs due to their $O(Kd)$ memory cost for storing K gradients or $O(K^2)$ computation from repeated backpropagation. A notable exception is FAMO [38], which scales more efficiently. Additionally, most MTL methods are tailored for vision tasks where gradient conflicts are prevalent, a condition less common in LLM pretraining: For example, GradVaccine [66] observed that in multilingual BERT (178M parameters), task gradients are mostly positively aligned or nearly orthogonal. We extend this finding to much larger autoregressive, decoder-only models (e.g., 1B parameters), which better reflect modern LLMs. Our results show minimal gradient conflict in multilingual and multi-domain pretraining (Section 2), suggesting that modern LLMs naturally exhibit cooperative (or sometimes nearly-orthogonal) gradient structures—potentially reducing the need for explicit conflict-mitigation. This motivates the central question of our work:

*Can we design an effective **adaptive data mixing strategy** based on **multitask optimization**, that **exploits non-conflicting nature of gradients** and **scales to large settings** with **minimal memory and computational overhead**?*

To address this, we propose *PiKE*, an adaptive data mixing that is **empirically effective** and **scalable** for pretraining LLMs. PiKE enjoys **theoretical guarantees** while remaining practical for large-scale settings—with many tasks, large models, and diverse data—at **negligible memory and compute overhead**. We summarize PiKE’s key features below; related work is discussed in Appendix 5.

Key Features of PiKE

1. Adaptively adjusts mixture weights using **per-task gradient magnitude** and **variance**.
2. Theoretically, achieves near-optimal per-iteration objective decrease and enjoys convergence guarantees (Section 3).
3. Incorporates tilted empirical risk minimization [35, 46], promoting balanced learning across tasks and preventing task under-representation (Section 3.3).
4. Scales efficiently (linearly) with model size and the number of tasks (Section 4).
5. Consistently outperforms existing methods across different scales (110M to 1B parameters) and scenarios (multilingual to domain mixing) (Figure 1 and Section 4).

Problem Definition and Notations We aim to train a *single* model with parameters $\theta \in \mathbb{R}^d$ to perform $K \geq 2$ tasks simultaneously. Each task k is associated with a smooth (possibly non-convex) loss $\ell_k(\theta, x) : \mathbb{R}^d \times \mathbb{R}^{d_x} \rightarrow \mathbb{R}$ where x is the data point. It is common to minimize the total loss:

$$\min_{\theta \in \mathbb{R}^d} \mathcal{L}(\theta) := \sum_{k=1}^K \mathbb{E}_{x \sim \mathcal{D}_k} [\ell_k(\theta; x)], \quad (1)$$

where \mathcal{D}_k represents the data distribution for task k . We define $\mathcal{L}_k(\theta) := \mathbb{E}_{x \sim \mathcal{D}_k} [\ell_k(\theta; x)]$. For notation, $\|\cdot\|$ represents the Euclidean norm, $\text{Tr}(\cdot)$ denotes the trace operator, and a function h is




Figure 2: **Left:** Average accuracy on four downstream tasks (ArcE, CSQA, HellaSwag, PIQA) for 750M GPT-2 style models pre-trained with *Mix*, *Round-Robin*, and *Random* strategies. See Appendix F.1 for more. **Right:** Task gradient cosine similarity for a 750M GPT-2 style model pre-trained on GLaM datasets. “*data1-data2*” indicates gradient similarity between tasks *data1* and *data2*. Further results in Appendix F.2.

L -Lipschitz if $\|h(\theta) - h(\theta')\| \leq L\|\theta - \theta'\|$ for any θ, θ' in the domain of $h(\cdot)$. A function $f(\cdot)$ is L -smooth if its gradient is L -Lipschitz continuous.

2 Main Building Blocks for PiKE

This section presents our main observations which form the main building blocks of PiKE.

Bulding Block #1: Mixing Domains per Batch Improves LLM Generalization

Optimizing equation (1) with stochastic methods, such as Adam or SGD, requires forming batches from one or more tasks at each step. The batch selection strategy strongly influences model performance [4, 24, 74, 69, 40]. Even with fixed data proportions, a key question remains: *how should one form a batch from K data domains at each step?*

Batch construction is a critical design choice in multitask training, as it directly impacts learning dynamics and final model performance [4, 24, 74, 69, 40]. We focus on three standard strategies: *Random*, *Round-Robin*, and *Mix*. Let us first define these methods assuming static, uniform sampling weights (i.e., $1/K$ per task). Let $\mathbf{b}_t = (b_{t,1}, \dots, b_{t,K})$ denote the number of samples from each task \mathcal{D}_k at iteration t , with total batch size $b = \sum_{k=1}^K b_{t,k}$. The strategies are defined as:

$$\mathbf{b}_t = \begin{cases} b \cdot \mathbf{e}_{k^*} & \text{Random, where } k^* \sim \text{Uniform}(\{1, \dots, K\}) \\ b \cdot \mathbf{e}_{(t \bmod K)+1} & \text{Round-Robin} \\ b \cdot (\frac{1}{K}, \dots, \frac{1}{K}) & \text{Mix} \end{cases}$$

where $\mathbf{e}_k \in \mathbb{R}^K$ is the k -th standard basis vector. That is, *Random* selects one domain per batch, *Round-Robin* cycles through domains, and *Mix* includes all tasks in each batch.

Historically, *Mix* has been widely used in computer vision multitask settings [14, 45, 9, 54, 75, 38], while earlier language modeling efforts favored *Random* or *Round-Robin* [41, 44, 42]. Recent studies on large-scale language models [17, 51, 7, 61] have revisited this question and found that *Mix* generally performs best—especially in diverse-data settings [18, 11, 69, 22, 65]. Our experiments reaffirm this trend: across a wide range of tasks and scales, *Mix* consistently outperforms the alternatives (Figure 2, Figure 4). This observation underpins the design of our proposed method, *PiKE*.

Bulding Block #2: Large-Scale Training: Low Gradient Conflict & MTL Scalability Challenges

The *Mix* batching strategy offers a natural lens for analyzing data mixing in multitask learning (MTL). When a batch of total size b is constructed with b_k samples from each task k , the overall gradient at

iteration t , denoted \mathbf{g}_t , is the average of the individual sample gradients:

$$\mathbf{g}_t = \frac{1}{b} \sum_{k=1}^K \sum_{i=1}^{b_k} \nabla \ell_k(\boldsymbol{\theta}_t; x_{k,i}) = \sum_{k=1}^K \frac{b_k}{b} \bar{\mathbf{g}}_{t,k}, \quad (2)$$

where $\bar{\mathbf{g}}_{t,k} = \frac{1}{b_k} \sum_{i=1}^{b_k} \nabla \ell_k(\boldsymbol{\theta}_t; x_{k,i})$ is the average gradient from task k , and $x_{k,i} \sim \mathcal{D}_k$. This formulation highlights how batch gradients blend contributions from different tasks and serves as the foundation for analyzing task interactions.

A central challenge in prior MTL literature is *gradient conflict* [75, 66, 47], where task gradients oppose the overall update direction. Formally, a conflict exists if: $\langle \mathbf{g}_t, \bar{\mathbf{g}}_{t,k} \rangle < 0$, indicating that the shared update could increase the loss for task k .

While such conflicts are well-documented in vision tasks, we observe they are rare in large-scale language model pretraining. In particular, as shown in Figures 2, 5, and 6, task gradients are generally non-conflicting in pretraining LLMs. This observation shifts our goal: instead of mitigating conflict, we can **leverage naturally non-conflicting gradients** to improve training efficiency.

This insight renders many conflict-aware MTL methods—such as PCGrad [75], AdaTask [73], MGDA [56], and NashMTL [47]—less effective for pre-training LLMs: PCGrad acts only when conflicts occur (rare in LLMs) and thus performs similarly to simple *Mix*. AdaTask treats different tasks as completely independent and thus does not utilize potential cooperative nature of tasks, leading to slower convergence. More fundamentally, these methods incur substantial memory and compute overhead, making them unsuitable for large-scale models (see Appendix A for discussions). A notable exception is the recent work FAMO [38], which is designed for scalability.

In summary, the lack of gradient conflict in LLM pretraining, combined with the limitations of existing MTL approaches, motivates the need for scalable methods that exploit non-conflicting gradient structures—rather than fixating on rare conflicts.




Figure 3: Adaptive mixing consistently outperforms static mixing in the Example 2.1.

Bulding Block #3: Adaptively Changing Mix Sampling Weights Is Necessary

Prior work using the Mix sampling strategy typically relies on fixed (static) sampling weights, keeping (b_1, \dots, b_K) constant throughout training. However, dynamically adjusting batch composition can significantly enhance efficiency. We illustrate this fact with a simple example:

Example 2.1. Consider training on $K = 2$ tasks with losses $\ell_1(\boldsymbol{\theta}; x_1) = \frac{1}{2}(\boldsymbol{\theta}^\top e_1)^2 + x_1^\top \boldsymbol{\theta}$ and $\ell_2(\boldsymbol{\theta}; x_2) = \frac{1}{2}(\boldsymbol{\theta}^\top e_2)^2 + x_2^\top \boldsymbol{\theta}$, where $e_1 = [1 \ 0]^\top$, $e_2 = [0 \ 1]^\top$, and $\boldsymbol{\theta} \in \mathbb{R}^2$. Data for task 1 follows $x_1 \sim \mathcal{N}(0, \sigma_1^2 I)$, while task 2 follows $x_2 \sim \mathcal{N}(0, \sigma_2^2 I)$. The loss for task k simplifies to $\mathcal{L}_k(\boldsymbol{\theta}) = \frac{1}{2}(\boldsymbol{\theta}^\top e_k)^2$. Using b_k samples from task k , the gradient at iteration t is given by

$$\mathbf{g}_t = \frac{1}{b_1 + b_2} (b_1 e_1 e_1^\top + b_2 e_2 e_2^\top) \boldsymbol{\theta}_t + \mathbf{z},$$

where $\mathbf{z} \sim \mathcal{N}(0, \frac{b_1 \sigma_1^2 + b_2 \sigma_2^2}{b^2} I)$ with $b = b_1 + b_2$. Updating $\boldsymbol{\theta}_t$ via SGD, $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \mathbf{g}_t$, we have

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1})] = \frac{1}{2} (1 - \eta \frac{b_1}{b})^2 \theta_{1,t}^2 + \frac{1}{2} (1 - \eta \frac{b_2}{b})^2 \theta_{2,t}^2 + \eta^2 \frac{b_1 \sigma_1^2 + b_2 \sigma_2^2}{b^2}, \quad (3)$$

where $\theta_{1,t}$ and $\theta_{2,t}$ denote the first and second component of the vector $\boldsymbol{\theta}_t$. The derivation details of equation (3) can be found in Appendix G.1. Letting $w_1 := \frac{b_1}{b}$, $w_2 := \frac{b_2}{b}$, and relaxing them to take real values, we can optimize the mixing weights w_1 and w_2 as

$$w_1^* = \Pi \left(\frac{b^{-1}(\sigma_2^2 - \sigma_1^2) + \eta^{-1}(\theta_{1,t}^2 - \theta_{2,t}^2) + \theta_{2,t}^2}{\theta_{1,t}^2 + \theta_{2,t}^2} \right) \quad (4)$$

and $w_2^* = 1 - w_1^*$ where $\Pi(\xi) = \min\{\max\{\xi, 0\}, 1\}$ is the projection onto $[0, 1]$. This result shows that **optimal batch composition should change over iterates** to maximize training efficiency.

Figure 3 illustrates the superiority of the adaptive mixing approach based on equation (4) over various static mixing strategies. Moreover, the adaptive mixing strategy in this example does not require any hyperparameter tuning, while finding the best static mixing requires tuning.

Despite its simplicity, this example mirrors key aspects of MTL in large models: 1) The optimal solution $\theta^* = 0$ minimizes all task losses simultaneously, reflecting the high expressive power of large models. 2) Task gradients are non-conflicting, resembling real-world gradient interactions observed in Building Block #2. Moreover, equation (4) further reveals that optimal data mixing depends on (1) the gradient norm squared per task $\|\nabla \mathcal{L}_1(\theta)\|^2 = \theta_1^2$, $\|\nabla \mathcal{L}_2(\theta)\|^2 = \theta_2^2$ and (2) gradient variance (σ_1^2, σ_2^2) . As we will see in the next section, these factors play a crucial role in defining optimal mixing strategies for more general settings.

3 Method

3.1 PiKE: Conceptual Version

To develop our method, we first start by quantifying gradient conflicts:

Definition 3.1. For a given point θ , we say gradients are \underline{c} -conflicted (with $\underline{c} \geq 0$) if, for all task pairs $j, k, j \neq k$,

$$-\underline{c}(\|\nabla \mathcal{L}_j(\theta)\|^2 + \|\nabla \mathcal{L}_k(\theta)\|^2) \leq \langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle.$$

The above definition is implied by a lower bound on the gradients cosine similarity. In particular, if $\frac{\langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle}{\|\nabla \mathcal{L}_j(\theta)\| \|\nabla \mathcal{L}_k(\theta)\|} \geq -\tilde{c}$, then the gradients are \underline{c} -conflicted for $\underline{c} = \tilde{c}/2$. Therefore, experiments in Section 2 and Figures 5 and 6 in Appendix F.2 show that \underline{c} is typically small for LLM training.

While task gradient conflict are rare in LLM training, we observed in Section 2 that task gradients are not fully aligned either. To quantify the level of alignment, we define the following concept:

Definition 3.2. For a given point θ , we say that the gradients are \bar{c} -aligned (with $\bar{c} \geq 0$) if, for all task pairs $j, k, j \neq k$,

$$\langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle \leq \bar{c} \|\nabla \mathcal{L}_j(\theta)\|_2 \|\nabla \mathcal{L}_k(\theta)\|_2.$$

Notice that Definition 3.1 and 3.2 always hold for $\bar{c} = 1$ and $\underline{c} = 1/2$. However, smaller values allow for more refined analysis and more desirable properties. Notably, when both \bar{c} and \underline{c} are small (as we observed in our experiments), the value of $\|\nabla \mathcal{L}(\theta)\|$ is small if and only if $\|\nabla \mathcal{L}_k(\theta)\|$ is small for all k (see Lemma G.1 in Appendix G). To proceed further, we make the following standard assumption:

Assumption 3.3. Per-task gradients are L -Lipschitz, unbiased, and have bounded variance, i.e., $\forall k$:

$$\|\nabla \mathcal{L}_k(\theta_1) - \nabla \mathcal{L}_k(\theta_2)\| \leq L \|\theta_1 - \theta_2\|, \quad \forall \theta_1, \theta_2 \quad (5)$$

$$\mathbb{E}_{x \sim \mathcal{D}_k}[\nabla \ell_k(\theta; x)] = \nabla \mathcal{L}_k(\theta), \quad \forall \theta \quad (6)$$

$$\mathbb{E}_{x \sim \mathcal{D}_k}[\|\nabla \ell_k(\theta; x) - \nabla \mathcal{L}_k(\theta)\|^2] \leq \sigma_k^2, \quad \forall \theta \quad (7)$$

Using a mix batch with b_k samples per task k , the estimated gradient follows equation (2). The next theorem characterizes the descent amount in one step of SGD under low conflict conditions:

Theorem 3.4. Assume Assumption 3.3 holds and the gradients at θ_t are \underline{c} -conflicted and \bar{c} -aligned, with $\underline{c} < \frac{1}{K-2+b/b_k}$ for all k . If gradients are computed using the Mix strategy equation (2), then:

$$\mathbb{E}[\mathcal{L}(\theta_{t+1})] \leq \mathcal{L}(\theta_t) + \sum_{k=1}^K b_k \left(-\frac{\eta}{b} \beta \|\nabla \mathcal{L}_k(\theta_t)\|^2 + \frac{L\eta^2}{2b^2} \sigma_k^2 \right) + \sum_{k=1}^K b_k^2 \frac{L\eta^2}{2b^2} \gamma \|\nabla \mathcal{L}_k(\theta_t)\|^2 \quad (8)$$

for SGD update $\theta_{t+1} = \theta_t - \eta \mathbf{g}_t$. Here, $\beta \triangleq \min_k (1 + \underline{c}(-K + 2 - \frac{b}{b_k}))$, $\gamma \triangleq 1 + \bar{c}(K - 1)$, and the expectation is over batch sampling randomness under the mix strategy (b_1, \dots, b_K) .

A formal proof is provided in Theorem G.3 (Appendix G). Theorem 3.4 provides an upper-bound on the decrease in the objective depending on the mix batch composition (b_1, \dots, b_K) . As we show in Theorem G.4 in the appendix, this upperbound is tight in the non-conflicting, non-aligned regime of $\bar{c} = \underline{c} = 0$ (which is a good approximation for pre-training LLMs according to our experiments).

According to Theorem 3.4, to maximize descent in Mix sampling, we need to minimize the RHS of equation (8). Thus, relaxing $w_k = b_k/b$ to continuous values, we need to solve:

$$\min_{w_1, \dots, w_K \geq 0} \sum_{k=1}^K w_k \lambda_k + \frac{1}{2} w_k^2 \kappa_k \quad \text{s.t.} \quad \sum_{k=1}^K w_k = 1 \quad (9)$$

where $\lambda_k \triangleq -\eta\beta\|\nabla\mathcal{L}_k(\boldsymbol{\theta})\|^2 + \frac{L\eta^2}{2b}\sigma_k^2$ and $\kappa_k \triangleq L\eta^2\gamma\|\nabla\mathcal{L}_k(\boldsymbol{\theta})\|^2$. Using KKT conditions, the optimal solution is given by $w_k^* = \max\left\{0, -\frac{\mu + \lambda_k}{\kappa_k}\right\}$, where μ is chosen such that $\sum_{k=1}^K w_k^* = 1$ (see Lemma G.2, Appendix G). This leads to the conceptual version of PiKE (Positive gradient Interactions-based K-task weight Estimator), summarized in Algorithm 2 in Appendix C.

The conceptual version of PiKE (Algorithm 2) adaptively adjusts sampling weights. This adaptive adjustment makes the stochastic gradients biased, i.e., $\mathbb{E}[\mathbf{g}_t] \neq \nabla\mathcal{L}(\boldsymbol{\theta}_t)$. Due to this introduced bias, the classical convergence results of SGD can no longer be applied. The following theorem establishes the convergence of conceptual PiKE:

Theorem 3.5. *Suppose the assumptions in Theorem 3.4 hold and the Conceptual PiKE Algorithm (Algorithm 2 in the Appendix) initialized at $\boldsymbol{\theta}_0$ with the SGD optimizer in Step 10 of the algorithm. Let $\Delta_L = \mathcal{L}(\boldsymbol{\theta}_0) - \min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta})$ and $\sigma_{\max} = \max_k \sigma_k$. Suppose $\delta > 0$ is a given constant and the stepsize $\eta \leq \frac{\beta\delta}{L\sigma_{\max}^2/b + L\gamma\delta}$. Then, after $T = \frac{2\Delta_L}{\eta\beta\delta}$ iterations, Algorithm 2 finds a point $\bar{\boldsymbol{\theta}}$ such that*

$$\mathbb{E}\|\nabla\mathcal{L}_k(\bar{\boldsymbol{\theta}})\|^2 \leq \delta, \quad \forall k = 1, \dots, K. \quad (10)$$

Particularly, if we choose $\eta = \frac{\beta\delta}{L\sigma_{\max}^2/b + L\gamma\delta}$, then the Conceptual PiKE algorithm requires at most $\bar{T} = \frac{2L\Delta_L(\sigma_{\max}^2/b + \gamma\delta)}{\delta^2\beta^2}$ iterations to find a point satisfying equation (10).

This theorem (restated as Theorem G.5) is proved in Appendix G.

Remark 1. Theorem 3.5 guarantees that, given enough iterations, all tasks are learned jointly, i.e., the gradients of all losses become small. Moreover, the convergence rate is $T = O(1/\delta^2)$, which matches the optimal iteration complexity for smooth, nonconvex stochastic optimization.

Remark 2. The Conceptual PiKE algorithm maximizes the expected decrease at each iteration by finding the optimal mix batching strategy. This property leads to improved iteration complexity bound than (static) uniform mix batching, as discussed in Appendix G.3, when \underline{c} and \bar{c} are small.

3.2 PiKE: Simplified Computationally Efficient Version

Solving equation (9) requires estimating $\{\|\nabla\mathcal{L}_k(\boldsymbol{\theta}_t)\|^2\}_{k=1}^K$ and $\{\sigma_k\}_{k=1}^K$; however, per-iteration estimation of these terms with sufficient accuracy often necessitates large batch computations, thereby impeding convergence. To speed up the algorithm, we can update these estimates every T_0 iterations. However, this can cause abrupt changes in (w_1, \dots, w_K) , leading to instability, especially with optimizers like Adam, where sudden shifts may disrupt momentum estimates. To mitigate this, we update (w_1, \dots, w_K) using a single mirror descent step on equation (9), ensuring gradual adjustments:

$$w_k \leftarrow w_k \exp \left(\alpha \eta (\beta - L\eta\gamma w_k) \frac{\text{gradient norm}}{\|\nabla\mathcal{L}_k(\boldsymbol{\theta})\|^2} - \frac{\alpha L\eta^2}{2b} \frac{\text{gradient variance}}{\sigma_k^2} \right)$$

followed by normalization: $\mathbf{w} \leftarrow \mathbf{w} / \|\mathbf{w}\|_1$, where α is the mirror descent step size.

Even after the above simplifications, fine-tuning L, γ, α , and β can be challenging in practice. We simplify this finetuning by noting two observations: 1) The coefficient in front of σ_k^2 is constant, independent of w_k . 2) For small η and $w_k < 1$, the coefficient in front of $\|\nabla\mathcal{L}_k(\boldsymbol{\theta})\|^2$ remains nearly constant: $\alpha\eta(\beta - L\eta\gamma w_k) \approx \alpha\eta\beta$. Therefore, in our practical implementation, PiKE employs tunable constant coefficients for terms related to task-specific gradient variance and gradient norms, which simplifies its application. The resulting algorithm is detailed in Algorithm 1.

Algorithm 1 PiKE: Positive gradient Interaction-based K-task weights Estimator

```
1: Input:  $\theta$ ,  $T_0$ , total batch size  $b$ , task  $k$  dataset  $\mathcal{D}_k$ , hyperparameters  $\zeta_1$  and  $\zeta_2$ , prior weights  $\mathbf{w}'$ 
2: Initialize:  $w_k \leftarrow 1/K$  or  $w_k \leftarrow w'_k$ 
3: for  $t = 0, 1, \dots$  do
4:   if  $t \bmod T_0 = 0$  then
5:     Estimate  $\|\nabla \mathcal{L}_k(\theta_t)\|^2$  and  $\sigma_k^2$  for every  $k$ 
6:      $w_k \leftarrow w_k \exp\left(\zeta_1 \|\nabla \mathcal{L}_k(\theta_t)\|^2 - \frac{\zeta_2}{2b} \sigma_k^2\right)$ 
7:      $\mathbf{w} \leftarrow \mathbf{w} / \|\mathbf{w}\|_1$ 
8:      $(b_1, \dots, b_K) \leftarrow \text{round}(b(w_1, \dots, w_K))$ 
9:   end if
10:  Sample  $b_k$  data points from each task  $k$ 
11:  Compute the gradient  $\mathbf{g}$  using the estimates samples
12:  Update:  $\theta_{t+1} \leftarrow \text{Optimizer}(\eta, \theta_t, \mathbf{g})$ 
13: end for
```

Remark 3. The inclusion of the variance-related term, σ_k^2 , is crucial in PiKE; its omission would cause PiKE to disproportionately prioritize tasks with currently large losses even if their gradients are very noisy (high variance). Our ablation study, presented in Appendix F.6, underscores the importance of this term for achieving robust and balanced multi-task performance.

Remark 4. PiKE updates sampling weights only every T_0 steps. In our experiments, this additional step incurs negligible memory and computational overhead (see Appendix F.4).

3.3 Balanced-PiKE: Balanced Learning Across Different Tasks

While Algorithm 1 optimizes the average loss across tasks equation (1), this objective can lead to unbalanced learning. In practice, some tasks may dominate due to larger datasets or easier loss reduction, causing others to be under-optimized. This imbalance is problematic when certain domains (e.g., code, math) are more important to user, or when balanced performance across all tasks is desired. To address this shortcoming, we consider a balance-promoting objective based on *tilted empirical risk minimization* [35], also known as the α -fairness utility [46]:

$$\min_{\theta} \tilde{\mathcal{L}}(\tau; \theta) := \frac{1}{\tau} \log \left(\sum_{k=1}^K e^{\tau \mathcal{L}_k(\theta)} \right). \quad (11)$$

This objective interpolates between average loss ($\tau \rightarrow 0$) and worst-case loss ($\tau \rightarrow \infty$), allowing users to control the trade-off between efficiency and fairness. Moderate values of τ ($\tau \approx 1 \sim 3$ in our experiments) encourage balanced learning without sacrificing overall performance.

To leverage our PiKE developments in solving equation (11), we need to connect it to our initial objective in equation (1). The following lemma achieves this through the use of Fenchel duality [53]:

Lemma 3.6. Assume $\mathcal{L}_k(\theta) > 0$, $\forall k$ and let $0 < \tau \in \mathbb{R}$. Then equation (11) is equivalent to solving

$$\min_{\theta} \max_{\substack{\mathbf{y} \in \mathbb{R}_+^K \\ \sum_{k=1}^K y_k = \tau}} \sum_{k=1}^K y_k \mathcal{L}_k(\theta) - \sum_{k=1}^K \frac{y_k}{\tau} \log \left(\frac{y_k}{\tau} \right). \quad (12)$$

Moreover, for any fixed θ , the inner maximization problem is maximized at $y_k^* = \frac{\tau e^{\tau \mathcal{L}_k(\theta)}}{\sum_{j=1}^K e^{\tau \mathcal{L}_j(\theta)}}$, $\forall k$.

This lemma, which is proved in Appendix G.4, provides a natural alternative minimization algorithm for solving equation (11): Fixing \mathbf{y} , equation (12) reduces to a weighted minimization over tasks, where *regular PiKE sampling* with proper weights y_k in front of each loss can be applied to determine the optimal mixing strategy. On the other hand, fixing θ , the optimal solution y_k^* can be computed according to Lemma 3.6. This leads to *Balanced-PiKE* algorithm, described in Appendix D, which balances overall loss minimization and balanced/fair learning of all tasks. Although prior MTL literature explored fair learning [3, 47], such methods often grapple with scalability limitations, a critical aspect where PiKE is designed to excel through efficient large-scale performance.

Table 1: We report the perplexities (lower the better) on the validation split of multilingual C4 datasets. We also report the accuracies (% , higher the better) of different models on HellaSwag and its corresponding translated version. **Bolding** indicates the best model in the task; Metrics means the average across different tasks. Additional results can be found in Table 15.

	C4 (en)		C4 (hi)	C4 (de)		HellaSwag (en)	HellaSwag (hi)	HellaSwag (de)
	Perplexity ↓	Perplexity ↓	Perplexity ↓	Perplexity ↓	Accuracy(%) ↑	0-shot ↑	0-shot ↑	0-shot ↑
C4 (en), C4 (hi), and C4 (de) datasets, GPT-2 large style, 1B params, 36 Layers default, 120K training steps								
Mix	8.29	11.13	4.45	9.29	27.5	28.1	27.1	27.6
Round-Robin	8.41	11.31	4.97	9.46	26.5	27.6	26.7	26.3
Random	8.48	11.38	4.54	9.55	26.6	27.0	26.9	26.1
FAMO [38]	8.25	11.04	4.48	9.23	27.2	27.3	26.9	27.3
ADO [30]	8.30	11.12	4.45	9.31	27.5	27.7	27.5	27.2
PiKE	9.56	9.49	5.32	13.87	28.7	33.0	27.2	26.2
Balanced-PiKE ($\tau = 1$)	8.29	11.12	4.46	9.31	27.9	28.3	27.4	28.0
Balanced-PiKE ($\tau = 3$)	8.18	10.14	4.93	9.49	28.9	31.3	27.3	28.1
Balanced-PiKE ($\tau = 5$)	8.42	10.02	6.30	8.94	28.9	31.2	26.9	28.6

Table 2: We report perplexity (lower is better) on the validation split of the GLaM datasets, averaging perplexities across six domains. We also compare the accuracies (% , higher the better) of different models on four different Q/A tasks. PiKE (Uniform) means PiKE using initial uniform sampling weights and PiKE (GLaM) means PiKE using GLaM tuned weights as initial weights. **Bolding** indicates the best model in the task, underlining indicates PiKE beating Mix, Round-Robin, Random methods. Additional detailed (per-domain) results with different model sizes can be found in Table 16.

	GLaM		ArcE	CSQA	HellaSwag	PIQA
	Perplexity ↓	Accuracy(%) ↑	7-shot ↑	7-shot ↑	7-shot ↑	7-shot ↑
Six domains of GLaM dataset, GPT-2 large style, 750M params, 36 layers default						
Mix	12.77	46.4	47.2	39.6	37.9	60.9
Round-Robin	12.98	44.3	43.5	36.7	36.8	60.3
Random	12.99	42.7	41.7	34.2	36.6	58.2
FAMO [38]	13.25	45.0	43.7	40.0	36.4	59.8
ADO [30]	12.77	45.9	45.5	38.7	38.1	61.1
GLaM [18]	13.20	45.3	46.9	39.8	38.0	56.4
DoReMi [69]	13.25	46.5	48.6	40.1	37.5	59.6
PiKE (Uniform)	13.22	<u>47.6</u>	<u>49.6</u>	<u>43.2</u>	37.2	60.4
PiKE (GLaM)	13.35	<u>48.1</u>	49.8	43.5	<u>38.0</u>	61.2
Balanced-PiKE ($\tau = 1$)	13.21	<u>47.5</u>	<u>48.8</u>	<u>42.5</u>	37.6	<u>61.2</u>
Balanced-PiKE ($\tau = 3$)	13.26	<u>47.2</u>	<u>48.8</u>	<u>41.5</u>	37.2	<u>61.3</u>
Balanced-PiKE ($\tau = 5$)	13.19	48.2	<u>49.3</u>	<u>42.6</u>	38.5	62.4

4 Experiments

We evaluate PiKE in two multitask pretraining scenarios: 1) *Pretraining language models on multilingual mC4 dataset* [71], a dataset covering diverse languages from Common Crawl corpus. 2) *Pretraining language models on the GLaM dataset* [18], an English dataset spanning six domains. As we will show, across multiple model sizes (110M, 270M, 750M, and 1B parameters), PiKE consistently outperforms all existing methods, including heuristic or static data mixing, multitask optimization, and adaptive data mixture approaches. We first start by explaining our setup.

4.1 Experiment Setup

Baselines: We evaluate a range of sampling strategies: (1) *(Uniform) Mix*, (2) *Round-Robin*, (3) *Random*, (4) *FAMO* [38], (5) *ADO* [30], (6) *GLaM* [18], (7) *DoReMi* [69], (8) *PiKE*, and (9) *Balanced-PiKE*. Among these, DoReMi, GLaM, and ADO are specifically designed for LLM pretraining. We also include FAMO, a recent multitask learning (MTL) method, because it is scalable to large model sizes and has demonstrated competitive performance—serving as a useful comparison point for our work. DoReMi estimates task weights by training a small proxy model, while GLaM assigns static weights based on downstream performance from smaller models. However, since both methods report weights only for the GLaM dataset and do not provide configurations for multilingual C4, we exclude them from our multilingual experiments. In contrast to these static methods, *PiKE* dynamically updates task sampling weights during training using gradient information, enabling adaptive optimization throughout training. PiKE is scalable to both large models and many tasks, with minimal overhead, as detailed in Appendix F.4.

Datasets: For multilingual experiments, we use mC4 [71], focusing on English (en), Hindi (hi), and German (de). An overview of these datasets is provided in Table 3. For GLaM-based experiments, we use the six-domain GLaM dataset [18]. Additional details are presented in Table 4.

Evaluation: Perplexity is measured on held-out validation data. Downstream evaluation follows the OLMES suite [26]. For multilingual downstream tasks, we use multilingual HellaSwag [13], covering 26 languages. For models trained on GLaM, we evaluate on downstream tasks ARC-Easy [12], CommonsenseQA [59], PIQA [5], and HellaSwag [76]. HellaSwag and ArcE tasks have 4 choices, CSQA has 5 choices, and PIQA has 2 choices.

Further details on our experimental setup and evaluation are in Appendix E.

4.2 Main Observations on Multilingual Pretraining Experiments

Table 1 presents results for pretraining a 1B multilingual GPT-2 model [49] on English, Hindi, and German. Mix batching outperforms Round-Robin and Random strategies, supporting our analysis in Section 2 and justifying our focus on Mix as the foundation for PiKE. Additional results with different model sizes (270M and 1B) and language settings are reported in Table 15. Our main observation is that *PiKE and its Balanced variant achieve the highest average downstream accuracy* across all language settings and model sizes, demonstrating their effectiveness for multilingual pretraining.

We also observe that *Balanced-PiKE promotes more fair learning across tasks*. In particular, we pre-trained 1B models using Balanced-PiKE with different values of parameter $\tau \in \{1, 3, 5\}$. As τ increases, task losses become more uniform, reflecting improved balanced learning. At $\tau = 5$, perplexity becomes more balanced across languages, while $\tau = 3$ offers the best trade-off—achieving both the lowest perplexity and highest downstream accuracy. These results highlight the importance of incorporating fairness/balanced learning into data mixing strategies during pretraining.

4.3 Main Observations on Pretraining Experiments with GLaM Datasets

Table 2 shows results for pretraining a 750M GPT-2 model on the GLaM dataset. Additional results with both 110M and 750M models across six domains are provided in Table 16. Across both 110M and 750M model sizes, *PiKE consistently outperforms DoReMi, GLaM, and Mix* in downstream accuracy. For the 750M model, PiKE improves average accuracy by **3.4%** over DoReMi, **6.2%** over GLaM, **7.1%** over FAMO, and **4.8%** over ADO. In the 110M setting, PiKE achieves **37.8%** accuracy, exceeding DoReMi (**36.0%**), GLaM (**35.3%**), and FAMO (**35.9%**). Unlike DoReMi, which requires a separate proxy model, or GLaM, which depends on tuning with smaller models, PiKE delivers these gains with negligible additional overhead.

PiKE benefits from apriori downstream-tuned weights. We evaluate PiKE with two initializations: (1) uniform weights $b_k = b/K$ and (2) GLaM-tuned weights. In both small and large GPT-2 configurations, PiKE benefits from utilizing already fine tuned weights as initialization, achieving **48.1%** accuracy with GLaM-tuned weights vs. **47.6%** with uniform initialization. This shows that PiKE can effectively leverage pre-existing fine-tuned weights while still outperforming other methods with uniform initialization.

Mixing datasets improves language model generalization. We compare models trained on individual domains to those trained on mixed-domain datasets. Table 16 shows that single-domain training underperforms compared to mixed-domain training, even with simple Mix sampling. This reinforces the importance of diverse data for pretraining and aligns with prior work [40, 29].

Perplexity versus downstream performance. Table 16 reveals that validation perplexity does not always align with downstream performance. For instance, while Mix sampling yields lower perplexity in 750M models, PiKE achieves better downstream accuracy. This aligns with prior findings [23, 60, 39, 67], suggesting that perplexity alone is not a reliable performance metric.

5 Related Work

Data Curation and Selection. The effectiveness of language models heavily depends on the quality of the pre-training corpus. Consequently, significant efforts have been made to enhance pre-training data. These efforts include heuristic-based filtering [51, 50, 32, 48, 57] and deduplication [1, 33, 10, 19].

Recently, [64] proposed an automated method for constructing large, diverse, and balanced datasets for self-supervised learning by applying hierarchical k-means clustering. [55] introduced techniques that leverage instruction-tuned models to assess and select high-quality training examples, along with density sampling to ensure diverse data coverage by modeling the data distribution. Additionally, [27] simulated training runs to model the non-additive effects of individual training examples, enabling the analysis of their influence on a model’s predictions.

Multitask Learning Optimization Multitask learning (MTL) optimization is closely related to our framework. Among different MTL approaches, one idea is to modify gradient updates to mitigate gradient conflicts—situations where task gradients point in opposing directions, slowing down optimization [63, 75]. The Multiple Gradient Descent Algorithm (MGDA) [16, 56] updates the model by optimizing the worst improvement across all tasks, aiming for equal descent in task losses. Projected Conflicting Gradient Descent (PCGrad) [75] modifies task gradients by iteratively removing conflicting components in a randomized order, ensuring that updates do not interfere destructively across tasks. Conflict-Averse Gradient Descent (CAGRAD) [37] optimizes for the worst task improvement while ensuring a decrease in the average loss. NASHMTL [47] determines gradient directions by solving a bargaining game that maximizes the sum of log utility functions. While these methods improve performance, they introduce computational and memory overhead, making them impractical for large-scale models with numerous tasks [70]. Similar challenges exist in AdaTask [73], which improves multitask learning by balancing parameter updates using task-wise adaptive learning rates, mitigating task dominance, and enhancing overall performance. Unlike previous approaches that requires $O(K)$ storage for task gradients (e.g. PCGrad) or optimizer states (e.g. AdaTask), FAMO [38] balances task loss reductions efficiently using $O(1)$ space and time. However, these methods fail to exploit the *non-conflicting* interactions among tasks, focusing instead on resolving conflicts that seldom arise. This motivates our work that actively leverages lack of gradient conflicts to enhance training efficiency.

Data Mixture Problem Another closely related line of work to our paper is data mixture problem. Early works relied on manual heuristics for forming data compositions to construct datasets [18, 71, 7, 22] for training LLMs. To create more principled and less labor-intensive methods, subsequent research has focused on optimizing these mixtures. One line of work uses offline optimization with proxy models, where a smaller, less expensive model is trained to find optimal data weights that are then used for the final large model training. Methods in this area include DoReMi [69], which optimizes domain weights to accelerate training convergence, and DOGE [20], which re-weights data by relying on the inner product between task gradients. A third category involves adaptive online data mixing, where the data distribution is adjusted dynamically during the training process. This includes methods that efficiently adjust sampling rates based on training signals ADO [30] and skill-based curriculum learning like Skill-it! [8], which dynamically samples data corresponding to specific skills to improve model proficiency. Our work PiKE is also in this category. In particular, PiKE dynamically leverages the near-orthogonality of task gradients, making it far more computational and memory-efficient than offline approaches like DoGE by only needing two scalars per task: per-task gradient norm and variance. PiKE also uniquely incorporates gradient variance or noise, a fundamental factor impacting the performance of the training, which was largely overlooked in recent data mixing literature such as DoGE. Overall, PiKE offers a reliable way to enhance final model performance.

6 Conclusion

In this work, we introduced *PiKE*, an adaptive data mixing algorithm for multitask learning that dynamically adjusts task sampling based on observed gradient interactions. Unlike prior methods that aim to resolve gradient conflicts, PiKE exploits the predominantly non-conflicting gradients seen in large-scale language model pretraining. We provided theoretical analysis and showed, through extensive experiments, that PiKE improves both convergence speed and downstream performance across multilingual and multi-domain settings. To promote balanced task learning, we extended PiKE with a fairness-aware objective, resulting in *Balanced-PiKE*, which reduces task-level performance gaps without sacrificing overall accuracy.

One limitation of PiKE is its lack of sensitivity to dataset size when assigning sampling weights. Future work could incorporate data abundance or downstream performance feedback into the sampling strategy. Additionally, extending PiKE to other domains beyond language modeling remains a promising direction for further research.

References

- [1] A. Abbas, K. Tirumala, D. Simig, S. Ganguli, and A. S. Morcos. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *arXiv preprint arXiv:2303.09540*, 2023.
- [2] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] H. Ban and K. Ji. Fair resource allocation in multi-task learning. *arXiv preprint arXiv:2402.15638*, 2024.
- [4] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.
- [5] Y. Bisk, R. Zellers, R. Le Bras, J. Gao, and Y. Choi. Reasoning about physical commonsense in natural language, 2019.
- [6] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [7] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [8] M. F. Chen, N. Roberts, K. Bhatia, J. WANG, C. Zhang, F. Sala, and C. Re. Skill-it! a data-driven skills framework for understanding and training language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=Ioizw01NLf>.
- [9] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, pages 794–803. PMLR, 2018.
- [10] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [11] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [12] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- [13] V. Dac Lai, C. Van Nguyen, N. T. Ngo, T. Nguyen, F. Dernoncourt, R. A. Rossi, and T. H. Nguyen. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. *arXiv e-prints*, pages arXiv–2307, 2023.
- [14] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3150–3158, 2016.
- [15] M. Dehghani, J. Djolonga, B. Mustafa, P. Padlewski, J. Heek, J. Gilmer, A. Steiner, M. Caron, R. Geirhos, I. Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. *arXiv preprint arXiv:2302.05442*, 2023.
- [16] J.-A. Désidéri. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathématique*, 350(5-6):313–318, 2012.
- [17] J. Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

- [18] N. Du, Y. Huang, A. M. Dai, S. Tong, D. Lepikhin, Y. Xu, M. Krikun, Y. Zhou, A. W. Yu, O. Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR, 2022.
- [19] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [20] S. Fan, M. Pagliardini, and M. Jaggi. Doge: Domain reweighting with generalization estimation. *arXiv preprint arXiv:2310.15393*, 2023.
- [21] C. Gaffney, D. Li, R. Sang, A. Jain, and H. Hu. Orbax, 2023. URL <http://github.com/google/orbax>.
- [22] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [23] T. Gao, A. Wettig, L. He, Y. Dong, S. Malladi, and D. Chen. Metadata conditioning accelerates language model pre-training. *arXiv preprint arXiv:2501.01956*, 2025.
- [24] C. Ge, Z. Ma, D. Chen, Y. Li, and B. Ding. Data mixing made efficient: A bivariate scaling law for language model pretraining. *arXiv preprint arXiv:2405.14908*, 2024.
- [25] Google. Grain - feeding jax models, 2023. URL <http://github.com/google/grain>.
- [26] Y. Gu, O. Tafjord, B. Kuehl, D. Haddad, J. Dodge, and H. Hajishirzi. Olmes: A standard for language model evaluations. *arXiv preprint arXiv:2406.08446*, 2024.
- [27] K. Guu, A. Webson, E. Pavlick, L. Dixon, I. Tenney, and T. Bolukbasi. Simfluence: Modeling the influence of individual training examples by simulating training runs. *arXiv preprint arXiv:2303.08114*, 2023.
- [28] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee. Flax: A neural network library and ecosystem for JAX, 2023. URL <http://github.com/google/flax>.
- [29] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark, et al. An empirical analysis of compute-optimal large language model training. *Advances in Neural Information Processing Systems*, 35:30016–30030, 2022.
- [30] Y. Jiang, A. Zhou, Z. Feng, S. Malladi, and J. Z. Kolter. Adaptive data optimization: Dynamic sample selection with scaling laws. *arXiv preprint arXiv:2410.11820*, 2024.
- [31] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*, pages 1–12, 2017.
- [32] H. Laurençon, L. Saulnier, T. Wang, C. Akiki, A. Villanova del Moral, T. Le Scao, L. Von Werra, C. Mou, E. González Ponferrada, H. Nguyen, et al. The bigscience roots corpus: A 1.6 tb composite multilingual dataset. *Advances in Neural Information Processing Systems*, 35: 31809–31826, 2022.
- [33] K. Lee, D. Ippolito, A. Nystrom, C. Zhang, D. Eck, C. Callison-Burch, and N. Carlini. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*, 2021.
- [34] C. Li, H. Farkhor, R. Liu, and J. Yosinski. Measuring the intrinsic dimension of objective landscapes. In *International Conference on Learning Representations (ICLR)*, 2018. <https://arxiv.org/abs/1804.08838>.
- [35] T. Li, A. Beirami, M. Sanjabi, and V. Smith. Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162*, 2020.

- [36] A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
- [37] B. Liu, X. Liu, X. Jin, P. Stone, and Q. Liu. Conflict-averse gradient descent for multi-task learning. *Advances in Neural Information Processing Systems*, 34:18878–18890, 2021.
- [38] B. Liu, Y. Feng, P. Stone, and Q. Liu. Famo: Fast adaptive multitask optimization. *Advances in Neural Information Processing Systems*, 36, 2024.
- [39] H. Liu, S. M. Xie, Z. Li, and T. Ma. Same pre-training loss, better downstream: Implicit bias matters for language models. In *International Conference on Machine Learning*, pages 22188–22214. PMLR, 2023.
- [40] Q. Liu, X. Zheng, N. Muennighoff, G. Zeng, L. Dou, T. Pang, J. Jiang, and M. Lin. Regmix: Data mixture as regression for language model pre-training. *arXiv preprint arXiv:2407.01492*, 2024.
- [41] X. Liu, J. Gao, X. He, L. Deng, K. Duh, and Y.-Y. Wang. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In *Association for Computational Linguistics*, 2015.
- [42] X. Liu, P. He, W. Chen, and J. Gao. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*, 2019.
- [43] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [44] M.-T. Luong, Q. V. Le, I. Sutskever, O. Vinyals, and L. Kaiser. Multi-task sequence to sequence learning. *arXiv preprint arXiv:1511.06114*, 2015.
- [45] I. Misra, A. Shrivastava, A. Gupta, and M. Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3994–4003, 2016.
- [46] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on networking*, 8(5):556–567, 2000.
- [47] A. Navon, A. Shamsian, I. Achituve, H. Maron, K. Kawaguchi, G. Chechik, and E. Fetaya. Multi-task learning as a bargaining game. *arXiv preprint arXiv:2202.01017*, 2022.
- [48] G. Penedo, Q. Malartic, D. Hesslow, R. Cojocaru, H. Alobeidli, A. Cappelli, B. Pannier, E. Almazrouei, and J. Launay. The refinedweb dataset for falcon llm: Outperforming curated corpora with web data only. *Advances in Neural Information Processing Systems*, 36:79155–79172, 2023.
- [49] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language Models are Unsupervised Multitask Learners, 2019. <https://openai.com/blog/better-language-models/>.
- [50] J. W. Rae, S. Borgeaud, T. Cai, K. Millican, J. Hoffmann, F. Song, J. Aslanides, S. Henderson, R. Ring, S. Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.
- [51] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [52] J. Ren, S. Rajbhandari, R. Y. Aminabadi, O. Ruwase, S. Yang, M. Zhang, D. Li, and Y. He. {ZeRO-Offload}: Democratizing {Billion-Scale} model training. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 551–564, 2021.
- [53] R. T. Rockafellar. Convex analysis:(pms-28). *Princeton university press*, 2015.

- [54] S. Ruder, J. Bingel, I. Augenstein, and A. Søgaard. Latent multi-task architecture learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4822–4829, 2019.
- [55] N. Sachdeva, B. Coleman, W.-C. Kang, J. Ni, L. Hong, E. H. Chi, J. Caverlee, J. McAuley, and D. Z. Cheng. How to train data-efficient llms. *arXiv preprint arXiv:2402.09668*, 2024.
- [56] O. Sener and V. Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018.
- [57] L. Soldaini, R. Kinney, A. Bhagia, D. Schwenk, D. Atkinson, R. Authur, B. Bogin, K. Chandu, J. Dumas, Y. Elazar, et al. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv preprint arXiv:2402.00159*, 2024.
- [58] J. Su, Y. Lu, S. Pan, A. Murtadha, B. Wen, and Y. Liu. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021.
- [59] A. Talmor, J. Herzig, N. Lourie, and J. Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.
- [60] Y. Tay, M. Dehghani, J. Rao, W. Fedus, S. Abnar, H. W. Chung, S. Narang, D. Yogatama, A. Vaswani, and D. Metzler. Scale efficiently: Insights from pre-training and fine-tuning transformers. *arXiv preprint arXiv:2109.10686*, 2021.
- [61] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [62] G. Team, P. Georgiev, V. I. Lei, R. Burnell, L. Bai, A. Gulati, G. Tanzer, D. Vincent, Z. Pan, S. Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [63] S. Vandenhende, S. Georgoulis, W. Van Gansbeke, M. Proesmans, D. Dai, and L. Van Gool. Multi-task learning for dense prediction tasks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3614–3633, 2021.
- [64] H. V. Vo, V. Khalidov, T. Darcet, T. Moutakanni, N. Smetanin, M. Szafraniec, H. Touvron, C. Couprie, M. Oquab, A. Joulin, et al. Automatic data curation for self-supervised learning: A clustering-based approach. *arXiv preprint arXiv:2405.15613*, 2024.
- [65] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.
- [66] Z. Wang, Y. Tsvetkov, O. Firat, and Y. Cao. Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models. *arXiv preprint arXiv:2010.05874*, 2020.
- [67] A. Wettig, A. Gupta, S. Malik, and D. Chen. Qurating: Selecting high-quality data for training language models. *arXiv preprint arXiv:2402.09739*, 2024.
- [68] M. Wortsman, P. J. Liu, L. Xiao, K. Everett, A. Alemi, B. Adlam, J. D. Co-Reyes, I. Gur, A. Kumar, R. Novak, et al. Small-scale proxies for large-scale transformer training instabilities. *arXiv preprint arXiv:2309.14322*, 2023.
- [69] S. M. Xie, H. Pham, X. Dong, N. Du, H. Liu, Y. Lu, P. S. Liang, Q. V. Le, T. Ma, and A. W. Yu. Doremi: Optimizing data mixtures speeds up language model pretraining. *Advances in Neural Information Processing Systems*, 36, 2024.
- [70] D. Xin, B. Ghorbani, J. Gilmer, A. Garg, and O. Firat. Do current multi-task optimization methods in deep learning even help? *Advances in neural information processing systems*, 35: 13597–13609, 2022.
- [71] L. Xue. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*, 2020.

- [72] L. Xue, A. Barua, N. Constant, R. Al-Rfou, S. Narang, M. Kale, A. Roberts, and C. Raffel. Byt5: Towards a token-free future with pre-trained byte-to-byte models. corr, abs/2105.13626. *arXiv preprint arXiv:2105.13626*, 2021.
- [73] E. Yang, J. Pan, X. Wang, H. Yu, L. Shen, X. Chen, L. Xiao, J. Jiang, and G. Guo. Adatask: A task-aware adaptive learning rate approach to multi-task learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 10745–10753, 2023.
- [74] J. Ye, P. Liu, T. Sun, Y. Zhou, J. Zhan, and X. Qiu. Data mixing laws: Optimizing data mixtures by predicting language modeling performance. *arXiv preprint arXiv:2403.16952*, 2024.
- [75] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020.
- [76] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- [77] Y. Zhang, C. Chen, T. Ding, Z. Li, R. Sun, and Z.-Q. Luo. Why transformers need adam: A hessian perspective. *arXiv preprint arXiv:2402.16788*, 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Yes, the claims made in the abstract and introduction accurately reflect the paper's contributions and scope. Each claim is clearly supported in the main text, with appropriate pointers provided for readers to easily verify and locate the corresponding results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Yes we discuss the limitations of the work in the conclusion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: Yes we provide the full set of assumptions and a complete and correct proof.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Yes we disclose all the information needed to reproduce the main results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We do not include the full code for pre-training language models because it is tailored to the specific framework and system we use. However, our method is general, and a professional programmer should be able to reproduce it within one to two days using the detailed implementation description provided in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Yes, we have specify all the training and test details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Due to the high computational cost of each experiment, which takes approximately 1–2 days to run, we did not perform multiple runs and therefore do not report error bars or other statistical significance measures.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Yes we report the information on the computer resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Yes we strictly follow the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Yes our work should only have positive societal impacts and no negative societal impacts. There are many potential positive societal consequences of our work, none of which we feel must be specifically highlighted here.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A The Scalability Problem of Existing MTL Methods

Many existing multitask learning (MTL) methods [16, 56, 66, 75, 37, 47, 73, 3] require computing and storing all K task-specific gradients at each training iteration, where K is the number of tasks. While exceptions like FAMO [38] are designed for efficient scaling, the aforementioned general approach typically leads to $O(Kd)$ space complexity for storing gradients and $O(Kd)$ time complexity per iteration for their computation, where d denotes the number of model parameters. This cost is further exacerbated when these methods also involve solving auxiliary optimization problems to combine or re-weight gradients [70]. In stark contrast, when the overall loss is a simple average of task-specific losses (e.g., $\mathcal{L} = \frac{1}{K} \sum_{k=1}^K \mathcal{L}_k$), its gradient $\nabla \mathcal{L}$ can be computed via a single backward pass. This results in $O(d)$ space and time complexity for the gradient computation, with the task-aggregation step being effectively independent of K .

While early MTL research often focused on relatively small-scale vision models or language models with fewer than 100 million parameters, modern language models frequently exceed 100 billion parameters (100B). Storing the gradients or weights for such a model (assuming 32-bit floating-point precision) requires approximately 400 GB of memory. This capacity already surpasses that of a single high-performance GPU like the NVIDIA H100 (80 GB), necessitating at least five such GPUs for storing just one set of gradients or weights.

Applying typical prior MTL methods in this large-scale context becomes impractical. Storing K full gradients for a 100B parameter model would demand approximately $400K$ GB of memory, equivalent to $5K$ H100 GPUs. This scaling of memory and computational requirements—at least linear with the number of tasks K for gradient operations and often compounded by super-linear costs from auxiliary optimization steps—renders naive applications of many existing MTL methods infeasible for current large language models. Indeed, even if storing per-task gradients were feasible, Xin et al. [70] reported a significant reduction in training throughput (steps per second) for such MTL techniques even on small-scale vision models, let alone for modern LLMs (e.g., those with 1B+ parameters). Collectively, these challenges highlight the pressing need for novel MTL methodologies engineered for large-scale pre-training, demanding both exceptional scalability and minimal additional computational and memory overhead.

B Discussion of Gradient Conflicts

GradVaccine [66] made a similar observation regarding low gradient conflicts among task gradients in multilingual BERT (178M parameters). We extend this finding to substantially larger autoregressive, decoder-only models (up to 1B parameters), which are more characteristic of current large-scale language modeling paradigms.

Our experiments demonstrate that task gradients in large-scale models indeed exhibit minimal conflicts. To illustrate, we conduct two pre-training experiments: (i) a 1B parameter GPT-2-style model [49] on the multilingual mC4 dataset [71] (covering six languages: English, Hindi, German, Chinese, French, and Arabic), and (ii) a 750M parameter model on The GLaM dataset [18] (English text from six diverse domains). Experimental details are provided in Appendix E. Figures 2 and 5 depict cosine similarity trends for task gradients, revealing several key observations: 1) Gradient similarity is initially high but generally decreases as training progresses. 2) In multilingual settings, gradient similarity correlates with linguistic proximity (e.g., English-German gradients align more closely), whereas gradients from GLaM’s diverse domains exhibit more uniform positive alignment. 3) Task gradients rarely conflict—multilingual cosine similarities seldom drop below -0.1, and GLaM domain gradients remain predominantly positive.

C PiKE: Conceptual Version

Here, we present the conceptual (basic) version of PiKE. As discussed in the main text, this approach lacks computational efficiency due to the frequent estimation of the norm and the variance of the per-task gradient.

Algorithm 2 Conceptual version of PiKE: Positive gradient Interaction-based K-task weights Estimator

```

1: Input:  $\theta$ , total batch size  $b$ , stepsize  $\eta$ , task  $k$  dataset  $\mathcal{D}_k$ , constants  $\beta, L, \gamma$ , and prior weights  $\mathbf{w}'$ 
2: Initialize:  $w_k \leftarrow 1/K$  or  $w_k \leftarrow w'_k, \forall k$ 
3: for  $t = 0, 1, \dots$  do
4:   Estimate  $\|\nabla \mathcal{L}_k(\theta_t)\|^2$  and  $\sigma_k^2$  for every  $k$ 
5:   Compute  $\lambda_k \triangleq -\eta\beta\|\nabla \mathcal{L}_k(\theta_t)\|^2 + \frac{L\eta^2}{2b}\sigma_k^2$  and  $\kappa_k \triangleq L\eta^2\gamma\|\nabla \mathcal{L}_k(\theta_t)\|^2$ 
6:   set  $w_k^* = \max\{0, -\frac{\mu+\lambda_k}{\kappa_k}\}$  where  $\mu$  is found (by bisection) such that  $\sum_{k=1}^K w_k^* = 1$ 
7:   Set  $(b_1, \dots, b_K) \leftarrow \text{round}(b(w_1^*, \dots, w_K^*))$ 
8:   Sample  $b_k$  data points from each task  $k$ 
9:   Compute the gradient  $\mathbf{g}$  using the estimates samples
10:  Update:  $\theta_{t+1} \leftarrow \text{Optimizer}(\eta, \theta_t, \mathbf{g})$ 
11: end for

```

As discussed in section 3.2, this algorithm is computationally inefficient as it requires estimating $\nabla \mathcal{L}_k(\theta_t)$ and σ_k at each iteration. To improve efficiency, we introduced modifications that led to the development of the PiKE algorithm (Algorithm 1 in the main body).

D Balanced-PiKE: Fairness Considerations Across Tasks

Here, we present the *Balanced-PiKE* algorithm in more detail. As discussed in the main body, the main difference with PiKE is that the fair version requires the computation of the coefficients

$$y_k^* = \frac{\tau e^{\tau \mathcal{L}_k(\theta)} - 1}{\sum_{k=1}^K e^{\tau \mathcal{L}_k(\theta)} - 1}, \forall k$$

Then updating the sampling weights by

$$w_k \leftarrow w_k \exp \left((y_k^*)^2 \zeta_1 \|\nabla \mathcal{L}_k(\mathbf{w})\|^2 - (y_k^*)^2 \frac{\zeta_2}{2b} \sigma_k^2 \right), \quad \forall k$$

The overall algorithm is summarized in Algorithm 3. For our experiments, we evaluate three different values of τ : 1, 3, and 5. A larger τ results in a stronger balancing effect between different tasks.

Algorithm 3 *Balanced-PiKE*: Balanced considerations across tasks

```

1: Input:  $\theta, T_0$ , total batch size  $b$ , task  $k$  dataset  $\mathcal{D}_k$ , hyperparameters  $\zeta_1, \zeta_2, \tau$ , prior weights  $\mathbf{w}'$ 
2: Initialize:  $w_k \leftarrow 1/K$  or  $w_k \leftarrow w'_k$ 
3: for  $t = 0, 1, \dots$  do
4:   if  $t \bmod T_0 = 0$  then
5:     Estimate  $\|\nabla \mathcal{L}_k(\theta_t)\|^2, \sigma_k^2$ , and  $\mathcal{L}_k(\theta_t)$  for every  $k$ 
6:      $y_k^* = \frac{\tau e^{\tau \mathcal{L}_k(\theta)}}{\sum_{k=1}^K e^{\tau \mathcal{L}_k(\theta)}}$ 
7:      $w_k \leftarrow w_k \exp \left( (y_k^*)^2 \zeta_1 \|\nabla \mathcal{L}_k(\mathbf{w})\|^2 - (y_k^*)^2 \frac{\zeta_2}{2b} \sigma_k^2 \right)$ 
8:      $\mathbf{w} \leftarrow \mathbf{w} / \|\mathbf{w}\|_1$ 
9:      $(b_1, \dots, b_K) \leftarrow \text{round}(b(w_1, \dots, w_K))$ 
10:   end if
11:   Sample  $b_k$  data points from each task  $k$ 
12:   Compute the gradient  $\mathbf{g}$  using the estimates samples
13:   Update:  $\theta_{t+1} \leftarrow \text{Optimizer}(\eta, \theta_t, \mathbf{g})$ 
14: end for

```

E Experiments Setup

E.1 Dataset Details

Our experiments construct two primary scenarios for multitask learning: multilingual tasks and diverse task mixtures spanning multiple domains. We consider two widely-used datasets for our study: mC4 [71] and GLaM [18].

mC4 Dataset The mC4 dataset [71] is a multilingual text corpus derived from the Common Crawl web archive, covering a diverse range of languages. It has been widely used for pretraining multilingual models, such as mT5 [71] and ByT5 [72]. The dataset is curated by applying language-specific filtering to extract high-quality text, ensuring a balanced representation across languages. Mixture weights for training models on mC4 are often chosen based on token counts. In our cases, we mainly focus on English (en), Hindi (hi), and German (de). We report their details in Table 3.

Table 3: Partial statistics of the mC4 corpus, totaling 6.3T tokens.

ISO code	Language	Tokens (B)
en	English	2,733
hi	Hindi	24
de	German	347

GLaM Dataset The GLaM dataset [18] comprises English text from six distinct sources and has been used to train the GLaM series models and PaLM [11]. Mixture weights for GLaM training were determined based on small model performance [18], while [69] employed group distributionally robust optimization (Group DRO) to compute domain-specific weights. Table 4 summarizes the six domains in the GLaM dataset and the mixture weights selected by GLaM and DoReMi. We use these weights as oracle baselines for comparison with PiKE, which dynamically adjusts task weights over time using gradient information, unlike the fixed weights employed by GLaM and DoReMi.

Table 4: GLaM dataset [18] and fixed mixture weights used in GLaM [18] and DoReMi [69].

Dataset	Tokens (B)	Weight chosen by GLaM [18]	Weight chosen by DoReMi [69]
Filtered Webpages	143	0.42	0.51
Wikipedia	3	0.06	0.05
Conversations	174	0.28	0.22
Forums	247	0.02	0.04
Books	390	0.20	0.20
News	650	0.02	0.02

E.2 Training Details

Our experiments explore two distinct scenarios for multitask learning: multilingual training and diverse task mixtures spanning multiple domains. To achieve optimal results, we customize the training setups for each scenario and present them separately in this section. All training is performed from scratch.

Multilingual Training To address the complexities of tokenizing multilingual data, we utilize the mT5 tokenizer [71], which features a vocabulary size of 250K. Both GPT-2 small and GPT-2 large models are trained with a context length of 1024 and a batch size of 256. The AdamW optimizer [43] is employed with consistent hyperparameters and a learning rate scheduler. Additional details on hyperparameter configurations are provided in Appendix E.5.

GLaM Training For GLaM training, we use the T5 tokenizer [51], implemented as a SentencePiece tokenizer trained on the C4 dataset with a vocabulary size of 32,000. Both GPT-2 small and GPT-2 large models are trained with a context length of 1024 and a batch size of 256. The AdamW optimizer [43] is used, and additional details on hyperparameters is in Appendix E.5.

E.3 Model Architecture

The detailed architecture is summarized in Table 5. Our implementation utilizes pre-normalization [49] Transformers with qk-layernorm [15]. Consistent with [10], we omit biases, and the layernorm [2] value remains set to the Flax [28] default of 1e-6. Additionally, we incorporate rotary positional embeddings [58].

Table 5: Architecture hyperparameters for different model scales used in the paper. All models are GPT-2-like decoder-only architectures. The multilingual models employ a vocabulary size of 250K, whereas GLaM training uses a vocabulary size of 32K. Differences in the total number of parameters arise due to the variation in vocabulary sizes.

Size	# Params	Layers	Attention heads	Attention head dim	Hidden dim
GPT-2 small	110M/270M	12	12	64	768
GPT-2 large	750M/1B	36	20	64	1280

Table 6: Hyperparameter settings for our experiments.

Hyperparameters	Values
Optimizer	AdamW ($\beta_1 = 0.95, \beta_2 = 0.98$)
Initial and final learning rate	$7e - 6$
Peak learning rate	$7e - 4$
Weight decay	0.1
Batch size	256
Context length	1024
Gradient clipping norm	1.0
Training step	120,000
Warm-up step	10,000
Schedule	Linear decay to final learning rate

Table 7: Hyperparameter settings for running PiKE (Algorithm 1).

Hyperparameters	Values
PiKE hyperparameter (ζ_1, ζ_2)	$\{(1e - 2, 1e - 3), (1.5e - 2, 1e - 3), (5e - 2, 5e - 3), (7.5e - 2, 5e - 3), (1e - 1, 1e - 2), (1.5e - 1, 1e - 2)\}$
Check interval T_0	1,000
Batch size for estimation	256

E.4 Experimental Resource

All experiments are conducted on 8 Google TPuv5. The training time for GPT-2 small and GPT-2 large models for 120K steps are approximately 1 day and 2 days per run, respectively.

E.5 Hyper-parameters

Table 6 shows the detailed hyperparameters that we used in all our experiments. We also report our hyperparameters grid for tuning PiKE in Table 7.

E.6 Implementation Details

Our implementation builds upon the Nanodo training infrastructure [68], incorporating enhancements for efficiency. This framework relies on Flax [28], JAX [6], and TPUs [31].

To enable training of larger models, we shard both model and optimizer states, following the methodology of FSDP [52], and define these shardings during JIT compilation. Checkpointing is handled using Orbx [21], while deterministic data loading is facilitated by Grain [25].

For data loading, sequences are packed to avoid padding. When a sequence contains fewer tokens than the context length hyperparameter, an end-of-sequence token is appended. This differs from Nanodo [68], where both begin-of-sequence and end-of-sequence tokens are added.

E.7 Evaluation


Our evaluation adheres to the OLMES suite [26]. For multilingual downstream performance, we utilize the multilingual version of HellaSwag [13], which supports evaluations across 26 languages. English downstream tasks are assessed using ARC-Easy [12], CommonsenseQA [59], PIQA [5],

and HellaSwag [76]. Unless specified otherwise, multilingual evaluations are performed in a 0-shot setting, while GLaM pretraining evaluations employ 7-shot in-context learning, with demonstration candidates separated by two line breaks. For HellaSwag and its translated variants, we evaluate the first 3,000 examples. For all other downstream tasks, evaluations are conducted on their respective validation sets. In the case of multiple-choice tasks, different candidates are included in the prompt, and the average log-likelihood for each candidate is computed. The candidate with the highest score is then selected as the predicted answer.

F Additional Experiment Results

F.1 Comparison of Performance Using Mix, Random, and Round-Robin Sampling Strategies

Figure 4 presents the average downstream accuracies of language models pre-trained using Mix, Random, and Round-Robin sampling strategies. In both multilingual pre-training and GLaM pre-training, the Mix sampling strategy consistently outperforms the other two. This motivates its use in pre-training large language models.



(a): 1B models on multilingual C4 (en), C4 (hi), (b): 750M models on GLaM datasets with six domains and C4 (de) datasets

Figure 4: Average downstream task accuracy of pretraining language models using Mix, Round-Robin, and Random sampling strategies. Mix and Random use equal batch size for each task ($b_k = b/K, \forall k \in K$).

F.2 Cosine Similarity and \underline{c} -Conflicted Gradients

Figures 5 and 6 show the cosine similarity, defined as $\frac{\langle \mathcal{L}_j(\theta), \mathcal{L}_k(\theta) \rangle}{\|\mathcal{L}_j(\theta)\| \|\mathcal{L}_k(\theta)\|}$ and the “ratio,” defined as $\frac{\langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle}{\|\nabla \mathcal{L}_j(\theta)\| \|\nabla \mathcal{L}_k(\theta)\|}$. In particular, if $\frac{\langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle}{\|\nabla \mathcal{L}_j(\theta)\| \|\nabla \mathcal{L}_k(\theta)\|} \geq -\tilde{c}$, then the gradients are \underline{c} -conflicted for $\underline{c} = \tilde{c}/2$, which aligns with the observations in Figures 5 and 6.




Figure 5: 1B models trained on multilingual mC4 datasets. **Left:** Cosine similarity between task gradients during language model pre-training over time. **Right:** The “ratio,” which defined as $\frac{\langle \nabla \mathcal{L}_j(\theta), \nabla \mathcal{L}_k(\theta) \rangle}{\|\nabla \mathcal{L}_j(\theta)\| \|\nabla \mathcal{L}_k(\theta)\|}$, between task gradients during language model pre-training over time. “data1-data2” denotes the cosine similarity or ratio between the gradient of data1 and the gradient of data2.

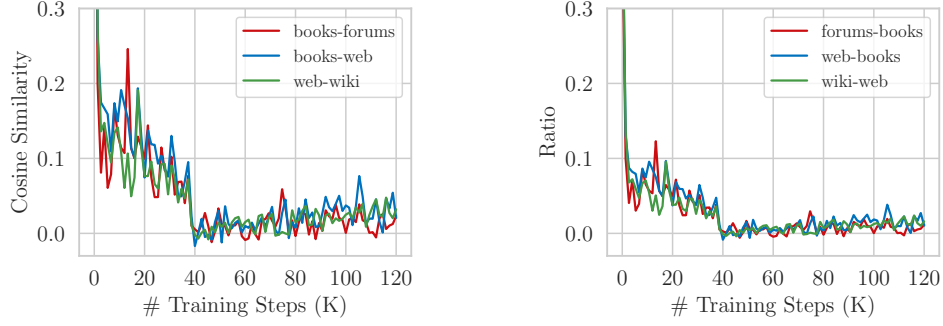


Figure 6: 750M models on GLaM datasets with six domains. **Left:** Cosine similarity between task gradients during language model pre-training over time. **Right:** The “ratio,” which defined as $\frac{\langle \mathcal{L}_j(\theta), \mathcal{L}_k(\theta) \rangle}{\|\mathcal{L}_j(\theta)\|^2 + \|\mathcal{L}_k(\theta)\|^2}$, between task gradients during language model pre-training over time. “*data1-data2*” denotes the cosine similarity or ratio between the gradient of *data1* and the gradient of *data2*.

F.3 Comparison of Performance Using PCGrad, AdaTask, and Mix

Figure 7 presents the average downstream task performance on HellaSwag (en) and HellaSwag (hi) for 270M multilingual language models pre-trained using PCGrad, AdaTask, and Mix. As shown in Figure 7: 1) PCGrad performs similarly to Mix, as it only adjusts gradients when conflicts occur—which is rare. 2) AdaTask converges more slowly due to noisy gradients and suboptimal optimizer state updates. Additionally, both methods are memory-intensive, requiring $O(K)$ storage for task gradients (PCGrad) or optimizer states (AdaTask), making them impractical for large-scale models such as the 540B PaLM [10].




Figure 7: Eval perplexity of pretraining 270M GPT-2 style multilingual language models on mC4 datasets (English and Hindi) using Mix, PCGrad, and AdaTask.

F.4 PiKE: Minimal Training and Memory Overhead

Table 8 reports the total training time and the computational overhead incurred by PiKE. In our experiments, per-task gradient squared norms and variances are estimated only once every $T_0 = 1,000$ training steps. Profiling results indicate that this infrequent estimation introduces minimal overhead during the pre-training of LLMs. Specifically, for a 1B parameter multilingual model pre-trained for approximately 71 hours, this estimation process accounted for only about 2.3% of the total training time (roughly 1.6 hours). This overhead percentage can be further reduced by using a larger estimation interval, T_0 (e.g., 5,000 steps), or by adopting a more efficient parallelization implementation for handling multiple tasks.

Regarding memory requirements, PiKE only needs to store the periodically estimated scalar values for each task’s gradient variance and squared norm. This amounts to an additional memory footprint of $O(K)$, which is negligible compared to the model parameters ($O(d)$) and especially to the $O(Kd)$ memory demanded by prior MTL methods that store K full per-task gradients. This efficiency starkly

contrasts with those earlier approaches, which also often involve solving auxiliary optimization problems that further increase computational and memory burdens [70]. Unlike approaches such as DoReMi or GLaM, PiKE simplifies deployment due to its automated learning of task sampling weights during the training procedure. This dynamic weight adaptation facilitates straightforward implementation and seamless integration into existing workflows, eliminating the need for manual intervention or training disruptions.

Table 8: We report the total training time (hrs) and how much overhead time that running PiKE (hrs). Compared with the total training, the overhead for running PiKE is minimal only taking from 1.2% to 2.4% in training large models.

	Total Training Time (hrs)	Overhead time for running PiKE (hrs)	Overhead time for running PiKE (%)
GPT-2 small (110M, GLaM)	10	0.12	1.2
GPT-2 small (270M, Multilingual)	12	0.24	2
GPT-2 large (750M, GLaM)	41	1.0	2.4
GPT-2 large (1B, Multilingual)	71	1.6	2.3

F.5 Additional Zero-shot and Seven-shot Results

In the main text, our evaluation focused on 0-shot performance for multilingual pre-training and 7-shot performance for GLaM pre-training. This appendix provides a more complete picture by presenting both 0-shot and 7-shot evaluation results for the checkpoints from these respective experimental setups. The results are in the Table 9 and 10.

F.6 Details on Estimating Gradient Variance and Magnitude

We estimate the per-task gradient variance and magnitude (measured using the L_2 norm) every $T_0 = 1000$ training steps. At each estimation step, the number of samples used for each task corresponds to its assigned portion (b_k) of the standard mixed training batch. Our JAX-based implementation enables efficient computation of these statistics with negligible overhead relative to total training time.

To evaluate the impact of estimation accuracy on PiKE’s performance, we conduct a sensitivity analysis while keeping all other experimental settings and hyperparameters fixed. Specifically, we compare two strategies. The first, used in our main experiments, computes statistics for each task once using samples from a single training batch. The second estimates statistics multiple times using different subsets of samples (e.g., distinct micro-batches) and averages them to produce the final values used by PiKE. Results are shown in Table 11. We find that PiKE is robust to the estimation

Table 9: We compare the accuracies (% , higher the better) of different models on HellaSwag and its corresponding translated version using 0- and 7-Shot settings. **Bolding** indicates the best model in the task; Metrics means the average across different tasks.

	HellaSwag (en)		HellaSwag (hi)		HellaSwag (de)		HellaSwag (en)		HellaSwag (hi)		HellaSwag (de)	
	Accuracy \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	Accuracy(%) \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow
C4 (en), C4 (hi), and C4 (de) datasets, GPT-2 large style, 1B params, 36 Layers default, 120K training steps												
Mix	27.5	28.1	27.1	27.6	27.6	27.6	32.6	33.9	31.4	31.4	32.5	32.5
Round-Robin	26.5	27.6	26.7	26.3	26.3	26.3	32.0	34.0	29.7	29.7	32.4	32.4
Random	26.6	27.0	26.9	26.1	26.1	26.1	31.9	34.0	30.0	30.0	31.7	31.7
PiKE	28.7	33.0	27.2	26.2	26.2	26.2	33.2	39.0	29.6	29.6	30.9	30.9

Table 10: We compare the accuracies (% , higher the better) of different models on four different Q/A tasks using 0- and 7-Shot settings. **Bolding** indicates the best model in the task, Metrics means the average across different tasks.

	ArcE		CSQA		HellaSwag		PIQA		ArcE		CSQA		HellaSwag		PIQA	
	Accuracy \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	0-shot \uparrow	Accuracy(%) \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow	7-shot \uparrow
Six domains of GLaM dataset, GPT-2 large style, 750M params, 36 layers default																
Mix	33.6	30.3	20.8	29.5	53.8	53.8	46.4	47.2	39.6	37.9	60.9	60.9	60.9	60.9	60.9	60.9
Round-Robin	32.7	30.4	20.3	26.1	53.9	53.9	44.3	43.5	36.7	36.8	60.3	60.3	60.3	60.3	60.3	60.3
Random	32.0	28.9	20.5	26.2	52.3	52.3	42.7	41.7	34.2	36.6	58.2	58.2	58.2	58.2	58.2	58.2
GLaM	31.7	28.8	19.9	26.3	51.9	51.9	45.3	46.9	39.8	38.0	56.4	56.4	56.4	56.4	56.4	56.4
DoReMi	35.6	33.0	23.8	30.0	55.7	55.7	46.5	48.6	40.1	37.5	59.6	59.6	59.6	59.6	59.6	59.6
PiKE (Uniform)	37.9	37.4	24.2	33.9	56.1	56.1	47.6	49.6	43.2	37.2	60.4	60.4	60.4	60.4	60.4	60.4
PiKE (GLaM)	35.5	33.5	20.4	31.2	56.8	56.8	48.1	49.8	43.5	38.0	61.2	61.2	61.2	61.2	61.2	61.2

Table 11: Mean accuracies (% , higher is better) of GLaM 740M models pretrained with PiKE on four question answering (Q/A) tasks, evaluated under different PiKE hyperparameter settings. The table compares the effect of estimating gradient statistics for each task either once (using a single batch) or multiple times using different subsets of samples.

Number of Estimation	1	2	3	4
Accuracy	46.8	47.4	47.1	46.8

Table 12: Mean Accuracies (% , higher values indicate better performance) of GLaM 740M models pre-trained with PiKE on four Question/Answering (Q/A) tasks, under different PiKE hyperparameter settings. PiKE’s ζ_1 hyperparameter was fixed constant at 0.1, while ζ_2 was varied. A value of $\zeta_2 = 0$ means an ablation where the per-task gradient variance term in PiKE is omitted.

ζ_2	0	0.001	0.005	0.01	0.05	0.1
Accuracy	45.3	46.2	47.0	46.4	46.9	45.3

Table 13: Pre-training 110M GLaM for 20K steps with different ζ_1 and fix $\zeta_2 = 0.01$

ζ_1	0.08	0.09	0.1	0.11	0.12
Accuracy	37.18	37.75	37.92	37.81	37.46

noise: computing gradient statistics from a single batch is sufficient to achieve strong performance, making the method both effective and efficient.

F.7 Impact of K , ζ_1 , and ζ_2 on Performance

K (Number of Domains) Our experiments show that PiKE achieves robust performance across a varying number of domains ($K = 2$ to 6), consistently outperforming baselines. A key advantage of PiKE is its linear computational complexity with respect to K , in contrast to some existing methods with quadratic complexity. This makes PiKE highly scalable for large K . However, PiKE’s formulation is based on the observation that per-task gradients are nearly orthogonal. This assumption typically holds in high-dimensional settings where the number of parameters d is much larger than the number of tasks K (i.e., $d \gg K$). Therefore, we anticipate that PiKE’s advantages may diminish when training smaller models on a very large number of tasks.

ζ_1, ζ_2 (Hyperparameters of the PiKE algorithm) As derived from our theoretical analysis, PiKE’s data sampling rate is designed to prioritize tasks with a high gradient norm (indicating difficulty) and low gradient variance (indicating training stability). The hyperparameters ζ_1 and ζ_2 control the model’s sensitivity to these two terms, respectively.

To underscore the importance of the gradient variance component, we conducted an ablation study by fixing $\zeta_1 = 0.1$ and varying ζ_2 (Table 12). The case $\zeta_2 = 0$ represents the complete omission of this term. Optimal performance (47.0% mean accuracy) was achieved at $\zeta_2 = 0.005$. In contrast, removing the variance component ($\zeta_2 = 0$) caused performance to drop to 45.3%. This degradation highlights the critical role of the per-task gradient variance in regulating PiKE’s sampling weights and achieving optimal results.

We also conducted a sensitivity analysis for ζ_1 , which weights tasks by their gradient norm to prioritize difficulty. Due to computational constraints, these experiments were run for a limited number of training steps. The results (Table 13) demonstrate that PiKE is robust across a range of ζ_1 values.

Based on our ablation studies, we recommend setting ζ_1 in the range 0.05–0.15 and ζ_2 in the range 0.005–0.015 for practical use.

F.8 Advantages of Adaptive over Static Data Mixtures

To demonstrate the advantages of an adaptive data mixture, we conducted an exhaustive search over various static (fixed-weight) data mixtures of English and Hindi. We then compared the performance

of these static baselines against our adaptive method, PiKE. The results, presented in Table 14, show that PiKE consistently outperforms all static configurations. This finding empirically highlights the limitations of fixed mixing strategies and validates the necessity of an adaptive approach to achieve optimal pre-training performance.

Table 14: Pre-training 270M GPT-2 style LMs with context length 2048 for 20K steps using C4 (English and Hindi) datasets.

Method	W_{en}	W_{hi}	Mean Accuracy
Mix	0.2	0.8	27.43
Mix	0.4	0.6	27.38
Mix	0.5	0.5	27.65
Mix	0.6	0.6	27.90
Mix	0.8	0.2	28.67
PiKE	Adaptive	Adaptive	29.27

F.9 Full Pre-training Results

Tables 15 and 16 present the complete results of pre-training language models across various scales (110M, 270M, 750M, and 1B) and scenarios (Multilingual and GLaM datasets). PiKE consistently outperforms all baselines across all scales and scenarios.

Table 15: We report the perplexities (lower the better) on the validation split of multilingual C4 datasets. We also compare the accuracies (% , higher the better) of different models on HellaSwag and its corresponding translated version. HellaSwag and its translated versions have 4 choices. **Bolding** indicates the best model in the task, Metrics means the average across different tasks.

	C4 (en)		C4 (hi)	C4 (de)		HellaSwag (en)		HellaSwag (hi)	HellaSwag (de)
	Perplexity ↓	Perplexity ↓	Perplexity ↓	Perplexity ↓	Accuracy(%) ↑	0-shot ↑	0-shot ↑	0-shot ↑	0-shot ↑
Single dataset, GPT-2 small style, 270M params, 12 layers default, 120K training steps									
C4 (en)	13.25	13.25	*	*	26.5	26.5	*	*	*
C4 (hi)	4.97	*	4.97	*	26.4	*	26.4	*	*
C4 (de)	11.27	*	*	11.27	26.1	*	*	*	26.1
C4 (en) and C4 (hi) datasets, GPT-2 small style, 270M params, 12 layers default, 120K training steps									
Mix	10.50	15.46	5.55	*	25.5	24.4	26.5	*	*
Round-Robin	10.57	15.57	5.57	*	25.6	25.2	26.0	*	*
Random	10.57	15.57	5.57	*	25.3	24.3	26.3	*	*
FAMO [38]	10.38	15.18	5.57	*	25.7	24.8	26.5	*	*
ADO [30]	10.45	15.39	5.52	*	25.1	24.3	25.8	*	*
PiKE	10.15	14.31	5.99	*	26.5	26.0	27.0	*	*
C4 (en), C4 (hi), and C4 (de) datasets, GPT-2 small style, 270M params, 12 layers default, 120K training steps									
Mix	12.00	16.30	5.88	13.83	25.3	24.4	26.0	25.5	
Round-Robin	12.10	16.44	5.91	13.95	25.1	24.3	26.0	24.9	
Random	12.16	16.49	5.95	14.03	25.1	24.7	26.6	23.9	
FAMO [38]	11.92	16.19	6.00	13.57	24.8	24.5	25.2	24.8	
ADO [30]	12.01	16.31	5.88	13.84	24.9	24.4	25.4	24.8	
PiKE	12.01	15.48	5.92	14.64	25.6	25.4	26.4	24.8	
Single dataset, GPT-2 large style, 1B params, 36 Layers default, 120K training steps									
C4 (en)	9.30	9.30	*	*	33.6	33.6	*	*	*
C4 (hi)	3.87	*	3.87	*	27.5	*	27.5	*	*
C4 (de)	7.72	*	*	7.72	28.1	*	*	*	28.1
C4 (en) and C4 (hi) datasets, GPT-2 large style, 1B params, 36 Layers default, 120K training steps									
Mix	7.41	10.60	4.22	*	27.3	28.2	26.5	*	*
Round-Robin	7.49	10.72	4.25	*	27.5	28.0	27.0	*	*
Random	7.52	10.76	4.28	*	28.0	28.9	27.0	*	*
FAMO [38]	7.33	10.44	4.22	*	26.8	27.1	26.5	*	*
ADO [30]	7.41	10.59	4.23	*	26.5	26.0	26.9	*	*
PiKE	7.21	9.63	4.80	*	30.0	32.7	27.3	*	*
C4 (en), C4 (hi), and C4 (de) datasets, GPT-2 large style, 1B params, 36 Layers default, 120K training steps									
Mix	8.29	11.13	4.45	9.29	27.5	28.1	27.1	27.6	
Round-Robin	8.41	11.31	4.97	9.46	26.5	27.6	26.7	26.3	
Random	8.48	11.38	4.54	9.55	26.6	27.0	26.9	26.1	
FAMO [38]	8.25	11.04	4.48	9.23	27.2	27.3	26.9	27.3	
ADO [30]	8.30	11.12	4.45	9.31	27.5	27.7	27.5	27.2	
PiKE	9.56	9.49	5.32	13.87	28.7	33.0	27.2	26.2	
Balanced-PiKE ($\tau = 1$)	8.29	11.12	4.46	9.31	27.9	28.3	27.4	28.0	
Balanced-PiKE ($\tau = 3$)	8.18	10.14	4.93	9.49	28.9	31.3	27.3	28.1	
Balanced-PiKE ($\tau = 5$)	8.42	10.02	6.30	8.94	28.9	31.2	26.9	28.6	

Table 16: We report perplexity (lower is better) on the validation split of the GLaM datasets, averaging perplexities across six domains when applicable or reporting a single perplexity when only training with a single domain. We also compare the accuracies (% , higher the better) of different models on four different Q/A tasks. HellaSwag and ArcE tasks have 4 choices, CSQA has 5 choices, and PIQA has 2 choices. PiKE (Uniform) means PiKE using initial sampling weights of 1/6 for each task and PiKE (GLaM) means PiKE using GLaM tuned weights as initial task weights. **Bolding** indicates the best model in the task, Metrics means the average across different tasks, underlining indicates PiKE beating Mix, Round-Robin, Random methods

	GLaM		ArcE	CSQA	HellaSwag	PIQA
	Perplexity ↓	Accuracy(%) ↑	7-shot ↑	7-shot ↑	7-shot ↑	7-shot ↑
Single domain of GLaM dataset, GPT-2 small style, 110M params, 12 layers default						
Wikipedia	9.96	33.5	32.5	20.9	27.3	53.3
Filtered Webpage	16.05	37.2	38.4	26.8	27.6	55.8
News	9.33	33.8	31.1	22.7	27.0	54.5
Forums	22.87	35.5	32.1	23.4	28.7	57.6
Books	16.81	34.7	34.3	22.1	27.8	54.7
Conversations	18.27	36.1	32.6	25.6	28.6	57.6
Six domains of GLaM dataset, GPT-2 small style, 110M params, 12 layers default						
Mix	18.27	36.2	35.6	24.1	28.5	56.7
Round-Robin	18.45	35.9	35.8	24.2	27.5	56.0
Random	18.48	35.5	34.3	22.4	28.4	56.8
FAMO [38]	18.19	35.9	35.3	24.2	27.8	56.4
ADO [30]	18.27	36.2	35.7	24.8	27.7	56.4
GLaM [18]	18.91	35.8	35.3	24.1	28.5	55.1
DoReMi [69]	18.98	37.0	36.0	28.3	28.2	55.3
PiKE (Uniform)	18.44	<u>37.4</u>	<u>36.8</u>	<u>27.5</u>	28.5	57.0
PiKE (GLaM)	19.34	37.8	39.0	<u>27.0</u>	28.0	57.0
Single domain of GLaM dataset, GPT-2 large style, 750M params, 36 layers default						
Wikipedia	7.24	35.9	35.1	24.0	30.5	53.9
Filtered Webpage	11.12	40.9	36.7	33.2	34.2	56.5
News	6.62	37.4	33.6	24.7	34.1	57.3
Forums	16.29	43.6	38.0	35.8	39.7	60.7
Books	11.83	41.3	40.0	33.0	34.5	57.8
Conversations	13.50	42.2	36.9	33.2	39.2	59.6
Six domains of GLaM dataset, GPT-2 large style, 750M params, 36 layers default						
Mix	12.77	46.4	47.2	39.6	37.9	60.9
Round-Robin	12.98	44.3	43.5	36.7	36.8	60.3
Random	12.99	42.7	41.7	34.2	36.6	58.2
FAMO [38]	13.25	45.0	43.7	40.0	36.4	59.8
ADO [30]	12.77	45.9	45.5	38.7	38.1	61.1
GLaM [18]	13.20	45.3	46.9	39.8	38.0	56.4
DoReMi [69]	13.25	46.5	48.6	40.1	37.5	59.6
PiKE (Uniform)	13.22	<u>47.6</u>	<u>49.6</u>	<u>43.2</u>	37.2	60.4
PiKE (GLaM)	13.35	<u>48.1</u>	49.8	43.5	<u>38.0</u>	<u>61.2</u>
Balanced-PiKE ($\tau = 1$)	13.21	<u>47.5</u>	<u>48.8</u>	<u>42.5</u>	37.6	<u>61.2</u>
Balanced-PiKE ($\tau = 3$)	13.26	<u>47.2</u>	<u>48.8</u>	<u>41.5</u>	37.2	<u>61.3</u>
Balanced-PiKE ($\tau = 5$)	13.19	48.2	<u>49.3</u>	<u>42.6</u>	38.5	62.4

F.10 Adaptive Sampling Weights of PiKE During Pre-training

Figure 8 illustrates how the adaptive sampling weights of PiKE evolve during language model pre-training. Compared to the Mix sampling strategy, which assigns equal sampling weights to each task, PiKE adaptively adjusts the sampling weights among English, German, and Hindi by leveraging the positive interaction of task gradients. This adaptive data selection allows PiKE to achieve superior performance compared to fixed or heuristic-based baselines.




Figure 8: The sampling weights for each dataset during the pre-training of 1B GPT-2-style multilingual language models on mC4 (English), mC4 (Hindi), and mC4 (German). Here, w_{en} represents the sampling weight for the English dataset, w_{hi} for the Hindi dataset, and w_{de} for the German dataset.

G Derivations and Proofs

G.1 Detailed Derivation of equation (3)

Recall that

$$\mathbf{g}_t = \frac{1}{b_1 + b_2} (b_1 \mathbf{e}_1 \mathbf{e}_1^\top + b_2 \mathbf{e}_2 \mathbf{e}_2^\top) \boldsymbol{\theta}_t + \mathbf{z},$$

Then

$$\begin{aligned} \boldsymbol{\theta}_{t+1} &= \boldsymbol{\theta}_t - \eta \frac{1}{b_1 + b_2} (b_1 \mathbf{e}_1 \mathbf{e}_1^\top + b_2 \mathbf{e}_2 \mathbf{e}_2^\top) \boldsymbol{\theta}_t - \eta \mathbf{z} \\ &= \boldsymbol{\theta}_t - \frac{\eta}{b} \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} \boldsymbol{\theta}_t - \eta \mathbf{z} \end{aligned}$$

Now consider the loss functions for task 1, $\mathcal{L}_1(\boldsymbol{\theta}_{t+1})$, and task 2, $\mathcal{L}_2(\boldsymbol{\theta}_{t+1})$, separately, taking the expectation over the randomness of \mathbf{z}

$$\begin{aligned} \mathbb{E}[\mathcal{L}_1(\boldsymbol{\theta}_{t+1})] &= \mathbb{E} \left[\frac{1}{2} (\mathbf{e}_1^\top \boldsymbol{\theta}_{t+1})^2 \right] \\ &= \mathbb{E} \left[\frac{1}{2} \left(\mathbf{e}_1^\top \left[\begin{array}{cc} 1 - \frac{\eta b_1}{b} & 0 \\ 0 & 1 - \frac{\eta b_2}{b} \end{array} \right] \boldsymbol{\theta}_t - \mathbf{e}_1^\top \eta \mathbf{z} \right)^2 \right] \\ &= \frac{1}{2} \left(\left[1 - \frac{\eta b_1}{b} \quad 0 \right] \boldsymbol{\theta}_t \right)^2 + \frac{1}{2} \eta^2 \mathbf{e}_1^\top \mathbf{Q} \mathbf{e}_1 \\ &= \frac{1}{2} \left(\left(1 - \frac{\eta b_1}{b} \right) \theta_{1,t} \right)^2 + \frac{1}{2} \eta^2 \mathbf{e}_1^\top \mathbf{Q} \mathbf{e}_1 \end{aligned}$$

Similarly, for task 2, we have

$$\mathbb{E}[\mathcal{L}_2(\boldsymbol{\theta}_{t+1})] = \frac{1}{2} \left(\left(1 - \frac{\eta b_2}{b} \right) \theta_{2,t} \right)^2 + \frac{1}{2} \eta^2 \mathbf{e}_2^\top \mathbf{Q} \mathbf{e}_2$$

where $\theta_{1,t}$ and $\theta_{2,t}$ denote the first and second component of the vector $\boldsymbol{\theta}_t$. Combining the losses for both tasks, the total expected loss becomes

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1})] &= \mathbb{E}[\mathcal{L}_1(\boldsymbol{\theta}_{t+1})] + \mathbb{E}[\mathcal{L}_2(\boldsymbol{\theta}_{t+1})] \\ &= \frac{1}{2} \left(\left(1 - \frac{\eta b_1}{b} \right) \theta_{1,t} \right)^2 + \frac{1}{2} \left(\left(1 - \frac{\eta b_2}{b} \right) \theta_{2,t} \right)^2 + \eta^2 \frac{b_1 \sigma_1^2 + b_2 \sigma_2^2}{b^2} \\ &= \frac{1}{2} \left(1 - \frac{\eta b_1}{b} \right)^2 \theta_{1,t}^2 + \frac{1}{2} \left(1 - \frac{\eta b_2}{b} \right)^2 \theta_{2,t}^2 + \eta^2 \frac{b_1 \sigma_1^2 + b_2 \sigma_2^2}{b^2}, \end{aligned}$$

which completes the derivations.

G.2 PiKE: Main Theoretical Results

Lemma G.1. Assume $\frac{1}{2(K-1)} > \underline{c}$. If $\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 \leq \epsilon$, we have

$$\sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 \leq \frac{\epsilon}{1 - 2\underline{c}(K-1)}.$$

Conversely, if $\|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 \leq \delta_k$, $\forall k$, then

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 \leq (1 - \bar{c}) \sum_{k=1}^K \delta_k + \bar{c} \left(\sum_{k=1}^K \sqrt{\delta_k} \right)^2$$

Proof: We first prove the first direction. Notice that

$$\begin{aligned} \|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 &= \left\| \sum_{k=1}^K \nabla \mathcal{L}_k(\boldsymbol{\theta}) \right\|^2 \\ &= \sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 + \sum_{k=1}^K \sum_{j \neq k} \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}), \nabla \mathcal{L}_k(\boldsymbol{\theta}) \rangle \leq \epsilon \end{aligned}$$

where we use the definition of $\nabla \mathcal{L}(\boldsymbol{\theta})$ and expand the term. Then we have

$$\begin{aligned} \sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 + \sum_{k=1}^K \sum_{j \neq k} \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}), \nabla \mathcal{L}_k(\boldsymbol{\theta}) \rangle &\stackrel{(a)}{\geq} \sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 - \underline{c} \sum_{k=1}^K \sum_{j \neq k} (\|\nabla \mathcal{L}_j(\boldsymbol{\theta})\|^2 + \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2) \\ &\stackrel{(b)}{\geq} \sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 (1 - 2\underline{c}(K-1)) \end{aligned}$$

where (a) uses the Definition 3.1, (b) uses symmetric identity. Thus we get

$$\sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 \leq \frac{\epsilon}{1 - 2\underline{c}(K-1)}$$

This completes the proof of the first inequality. We now prove the second inequality. Notice that

$$\begin{aligned} \|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 &= \left\| \sum_{k=1}^K \nabla \mathcal{L}_k(\boldsymbol{\theta}) \right\|^2 = \sum_{k=1}^K \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 + \sum_{k=1}^K \sum_{j \neq k} \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}), \nabla \mathcal{L}_k(\boldsymbol{\theta}) \rangle \\ &\stackrel{(a)}{\leq} \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 + \bar{c} \sum_{k=1}^K \sum_{j \neq k} \|\nabla \mathcal{L}_j(\boldsymbol{\theta})\|^2 \|\mathcal{L}_k(\boldsymbol{\theta})\|^2 \\ &= (1 - \bar{c}) \|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 + \bar{c} \|\nabla \mathcal{L}_k\|^2 + \bar{c} \sum_{k=1}^K \sum_{j \neq k} \|\nabla \mathcal{L}_j(\boldsymbol{\theta})\|^2 \|\mathcal{L}_k(\boldsymbol{\theta})\|^2 \\ &\stackrel{(b)}{\leq} (1 - \bar{c}) \sum_{k=1}^K \delta_k + \bar{c} \left(\sum_{k=1}^K \sqrt{\delta_k} \right)^2 \end{aligned}$$

where (a) use the Definition 3.2 and (b) combines the second and third terms and use the condition that $\|\nabla \mathcal{L}_k(\boldsymbol{\theta})\|^2 \leq \delta_k$. This completes the proof of the second inequality.

Lemma G.2. For the optimization problem

$$\begin{aligned} \min_{w_1, \dots, w_K} \quad & \sum_{k=1}^K w_k \lambda_k + \frac{1}{2} w_k^2 \kappa_k \\ \text{s.t.} \quad & \sum_{k=1}^K w_k = 1, \quad w_k \geq 0, \quad \forall k \end{aligned} \tag{13}$$

the optimal solution is

$$w_k^* = \max \left\{ 0, -\frac{\mu + \lambda_k}{\kappa_k} \right\} \quad (14)$$

where μ is chosen such that $\sum_{k=1}^K w_k^* = 1$

Proof: Consider the Lagrangian function

$$\mathcal{L}(w_1, \dots, w_K, \mu, \alpha_1, \dots, \alpha_K) = \sum_{k=1}^K w_k \lambda_k + \frac{1}{2} w_k^2 \kappa_k + \mu \left(\sum_{k=1}^K w_k - 1 \right) - \sum_{k=1}^K \alpha_k w_k$$

where μ is Lagrange multiplier for the equality constraint for the constraint $\sum_{k=1}^K w_k = 1$ and $\alpha_k \geq 0$ are Lagrange multipliers for the nonnegativity constraints w_k . Take the partial derivative of \mathcal{L} with respect to w_k and set it to 0:

$$\frac{\partial \mathcal{L}}{\partial w_k} = \lambda_k + w_k \kappa_k + \mu - \alpha_k = 0$$

From the Karush-Kuhn-Tucker (KKT) conditions, we also have $w_k^* \geq 0$, $\alpha_k \geq 0$, and $\alpha_k w_k^* = 0$. If $w_k^* > 0$, then $\alpha_k = 0$, which implies

$$0 = \lambda_k + w_k^* \kappa_k + \mu \implies w_k^* = -\frac{\mu + \lambda_k}{\kappa_k}$$

If $-(\mu + \lambda_k) / \kappa_k$ is negative, then $w_k^* = 0$ must hold. Combining these, we get

$$w_k^* = \max \left\{ 0, -\frac{\mu + \lambda_k}{\kappa_k} \right\}$$

Finally, the Lagrange multiplier μ is determined by enforcing the equality constraint:

$$\sum_{k=1}^K w_k^* = 1$$

with μ chosen so that the w_k^* sum to 1. This completes the proof.

Theorem G.3. (Formal Statement of Theorem 3.4) Suppose Assumption 3.3 is satisfied. Assume that at the given point θ_t the gradients are \underline{c} -conflicted and \bar{c} -aligned with $\underline{c} < \frac{1}{K-2+b/b_k}, \forall k$. Moreover, assume batching is performed based on the mix strategy equation (2), i.e.,

$$\mathbf{g}_t = \frac{1}{b} \sum_{k=1}^K \sum_{i=1}^{b_k} \nabla \ell_k(\theta_t; x_{k,i}) = \sum_{k=1}^K \frac{b_k}{b} \bar{\mathbf{g}}_{t,k},$$

Then, we have

$$\mathbb{E}[\mathcal{L}(\theta_t - \eta \mathbf{g}_t) | \mathcal{F}_t] \leq \mathcal{L}(\theta_t) + \sum_{k=1}^K b_k \left(-\frac{\eta}{b} \beta \|\nabla \mathcal{L}_k(\theta_t)\|^2 + \frac{L\eta^2}{2b^2} \sigma_k^2 \right) + \sum_{k=1}^K b_k^2 \frac{L\eta^2}{2b^2} \gamma \|\nabla \mathcal{L}_k(\theta_t)\|^2$$

where $0 \leq \beta \triangleq \min_k (1 + \underline{c}(-K + 2 - \frac{b}{b_k}))$ and $\gamma \triangleq 1 + \bar{c}(K - 1)$, the expectation is taken over batch sampling randomness under the mix strategy (b_1, \dots, b_K) , and $|\mathcal{F}_t$ denotes the natural filtration of our process.

Proof: We begin by revisiting the multi-task optimization problem under consideration. The objective is defined as:

$$\min_{\theta \in \mathbb{R}^d} \mathcal{L}(\theta) := \sum_{k=1}^K \mathbb{E}_{x \sim \mathcal{D}_k} [\ell_k(\theta; x)], \quad (15)$$

where $\mathcal{L}(\theta)$ is the aggregate loss over all tasks. Assume we mix the gradients with taking b_k i.i.d. samples from task k for $k = 1, \dots, K$. Then under, Assumption 3.3 and based on equation (2), the estimated gradient direction is given by

$$\mathbf{g}_t = \frac{1}{\sum_{k=1}^K b_k} \left(\sum_{k=1}^K \sum_{\substack{i=1 \\ x_i \sim \mathcal{D}_k}}^{b_k} \nabla \ell_k(\boldsymbol{\theta}_t; x_i) \right) \quad (16)$$

Let $\boldsymbol{\theta}_{t+1}$ be the updated point after gradient descent with $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \mathbf{g}_t$. By the descent lemma, the following inequality holds for the updated parameter $\boldsymbol{\theta}^+$:

$$\mathcal{L}(\boldsymbol{\theta}_{t+1}) \leq \mathcal{L}(\boldsymbol{\theta}_t) - \eta \mathbf{g}_t^\top \nabla \mathcal{L}(\boldsymbol{\theta}_t) + \frac{L\eta^2}{2} \|\mathbf{g}_t\|^2, \quad (17)$$

Conditioned on the natural filtration process \mathcal{F}_t , we take the expectation over the randomness of the samples draw and obtain:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] &\leq \mathcal{L}(\boldsymbol{\theta}_t) - \eta \mathbb{E}[\mathbf{g}_t | \mathcal{F}_t]^\top \nabla \mathcal{L}(\boldsymbol{\theta}_t) + \frac{L\eta^2}{2} \mathbb{E}(\|\mathbf{g}_t\|^2 | \mathcal{F}_t) \\ &\stackrel{(a)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \eta \left(\frac{1}{b} \sum_{k=1}^K b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right)^\top \left(\sum_{k=1}^K \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\left\| \sum_{k=1}^K b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right\|^2 + \sum_{k=1}^K (b_k \sigma_k^2) \right) \\ &\stackrel{(b)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K \sum_{j \neq k} b_k \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}_t), \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \rangle \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\left\| \sum_{k=1}^K b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right\|^2 + \sum_{k=1}^K (b_k \sigma_k^2) \right), \end{aligned}$$

where (a) substitutes the definition of \mathbf{g}_t and uses the Assumption 3.3, and (b) expands the terms. The notation $| \mathcal{F}_t$ denotes the natural filtration of the random process (loosely speaking, conditioned on the current time t and the past history). Continuing our simplifications, we have

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] &\stackrel{(a)}{\leq} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 - \sum_{k=1}^K \sum_{j \neq k} b_k \mathbb{C}(\|\nabla \mathcal{L}_j(\boldsymbol{\theta}_t)\|^2 + \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2) \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K \sum_{j \neq k} b_j b_k \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}_t), \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \rangle + \sum_{k=1}^K b_k \sigma_k^2 \right), \\ &\stackrel{(b)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K \left(b_k - \mathbb{C} b_k (K-1) - \mathbb{C} \sum_{j \neq k} b_j \right) \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K \sum_{j \neq k} b_j b_k \langle \nabla \mathcal{L}_j(\boldsymbol{\theta}_t), \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \rangle + \sum_{k=1}^K b_k \sigma_k^2 \right), \\ &\stackrel{(c)}{\leq} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K (b_k - \mathbb{C}(K-1)b_k - \mathbb{C}(b-b_k)) \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K \sum_{j \neq k} \bar{c} b_j b_k \|\nabla \mathcal{L}_j(\boldsymbol{\theta}_t)\|^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K b_k \sigma_k^2 \right) \\ &\stackrel{(d)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k (1 + \mathbb{C}(-K+2-b/b_k)) \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \end{aligned}$$

$$\begin{aligned}
& + \frac{L\eta^2}{2b^2} \left(\bar{c} \left(\sum_{k=1}^K b_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\| \right)^2 + (1 - \bar{c}) \sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K b_k \sigma_k^2 \right) \\
& \stackrel{(e)}{\leq} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k (1 + \underline{c}(-K + 2 - b/b_k)) \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \\
& + \frac{L\eta^2}{2b^2} \left(\bar{c} K \sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + (1 - \bar{c}) \sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K b_k \sigma_k^2 \right) \\
& = \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k (1 + \underline{c}(-K + 2 - b/b_k)) \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \\
& + \frac{L\eta^2}{2b^2} \left((1 - \bar{c} + \bar{c}K) \sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K b_k \sigma_k^2 \right) \tag{18}
\end{aligned}$$

where (a) applies Definition 3.1 to the second term and expands the third term, (b) expands the summation in the second term, (c) uses the identity $\sum_{k=1}^K \sum_{j \neq k} b_j = \sum_{k=1}^K (b - b_k)$ in the second term and applies Definition 3.2 to the third term, (d) combines terms in the third term, and (e) uses the inequality $\|\sum_{i=1}^N u_i\|^2 \leq N \sum_{i=1}^N u_i^2$, where \mathbf{u} is a column vector. We define β and γ such that

$$\begin{aligned}
\beta &= \min_k (1 + \underline{c}(-K + 2 - \frac{b}{b_k})) \\
\gamma &= 1 + \bar{c}(K - 1)
\end{aligned}$$

Then using the definition of β and γ , substituting back we have

$$\begin{aligned}
\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] &\leq \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta\beta}{b} \left(\sum_{k=1}^K b_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) + \frac{L\eta^2}{2b^2} \left(\gamma \sum_{k=1}^K b_k^2 \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K b_k \sigma_k^2 \right) \\
&= \mathcal{L}(\boldsymbol{\theta}_t) + \sum_{k=1}^K b_k \left(-\frac{\eta\beta}{b} \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2}{2b^2} \sigma_k^2 \right) + \sum_{k=1}^K b_k^2 \frac{L\eta^2}{2b^2} \gamma \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2,
\end{aligned}$$

which completes the proof.

Theorem G.4. *There exist loss functions $\{\ell_k(\cdot, \cdot)\}_{k=1}^K$ satisfying Assumption 3.3, whose gradients are \underline{c} -conflicted and \bar{c} -aligned with $\underline{c} = \bar{c} = 0$. For these losses, the upper bound in equation (8) is tight when gradients are computed using the mix strategy equation (2); that is, the inequality in equation (8) holds with equality.*

Proof. The proof is by simply generalizing Example 2.1 and showing the steps in the proof of Theorem G.3 are tight for this example. Consider the multi-task learning problem where the loss of task k is given by

$$\ell_k(\boldsymbol{\theta}, x_k) = \frac{L}{2} (\mathbf{e}_k^\top \boldsymbol{\theta})^2 + \mathbf{x}_k^\top \boldsymbol{\theta},$$

with $\mathbf{x}_k \sim \mathcal{N}(\mathbf{0}, \frac{\sigma_k^2}{d} \mathbf{I})$ is the data for task k . It is easy to verify that this loss is smooth with smoothness parameter L . Moreover, we can show that Assumption 3.3 is satisfied. Notice that

$$\mathcal{L}_k(\boldsymbol{\theta}) = \mathbb{E}[\ell_k(\boldsymbol{\theta}, x_k)] = \frac{L}{2} (\mathbf{e}_k^\top \boldsymbol{\theta})^2$$

and

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{k=1}^K \mathcal{L}_k(\boldsymbol{\theta}) = \frac{L}{2} \|\boldsymbol{\theta}\|^2.$$

First notice that it is easy to check that $\underline{c} = \bar{c} = 0$. Let \mathbf{g}_t be the direction obtained by equation (2) and assume $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \mathbf{g}_t$. Using the form of $\mathcal{L}(\boldsymbol{\theta})$, one can easily show that (by expanding the ℓ_2 -loss):

$$\mathcal{L}(\boldsymbol{\theta}_{t+1}) = \mathcal{L}(\boldsymbol{\theta}_t) - \eta \mathbf{g}_t^\top \nabla \mathcal{L}(\boldsymbol{\theta}_t) + \frac{L\eta^2}{2} \|\mathbf{g}_t\|^2. \tag{19}$$

Conditioned on the natural filtration process \mathcal{F}_t , we obtain:

$$\begin{aligned}
\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] &= \mathcal{L}(\boldsymbol{\theta}_t) - \eta \mathbb{E}[\mathbf{g}_t | \mathcal{F}_t]^\top \nabla \mathcal{L}(\boldsymbol{\theta}_t) + \frac{L\eta^2}{2} \mathbb{E}(\|\mathbf{g}_t\|^2 | \mathcal{F}_t) \\
&\stackrel{(a)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \eta \left(\frac{1}{b} \sum_{k=1}^K b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right)^\top \left(\sum_{k=1}^K \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right) \\
&\quad + \frac{L\eta^2}{2b^2} \left(\left\| \sum_{k=1}^K b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \right\|^2 + \sum_{k=1}^K (b_k \sigma_k^2) \right) \\
&\stackrel{(b)}{=} \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\eta}{b} \left(\sum_{k=1}^K b_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \right) \\
&\quad + \frac{L\eta^2}{2b^2} \left(\sum_{k=1}^K \|b_k \nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \sum_{k=1}^K (b_k \sigma_k^2) \right),
\end{aligned}$$

where (a) substitutes the definition of \mathbf{g}_t and uses the Assumption 3.3, and (b) is because $\langle \nabla \mathcal{L}_j(\boldsymbol{\theta}_t), \nabla \mathcal{L}_k(\boldsymbol{\theta}_t) \rangle = 0, \forall j \neq k$. Since $\beta = \gamma = 1$ in this example, we can write

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] = \mathcal{L}(\boldsymbol{\theta}_t) + \sum_{k=1}^K b_k \left(-\frac{\eta\beta}{b} \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2}{2b^2} \sigma_k^2 \right) + \sum_{k=1}^K b_k^2 \frac{L\eta^2}{2b^2} \gamma \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2,$$

which shows our upper-bound is tight in this example and completes the proof. \square

Theorem G.5. (Theorem 3.5 in the main body) Suppose the assumptions in Theorem G.3 is satisfied and we run the Conceptual PiKE Algorithm (Algorithm 2) initialized at $\boldsymbol{\theta}_0$ with the SGD optimizer in Step 10 of the algorithm. Let $\Delta_L = \mathcal{L}(\boldsymbol{\theta}_0) - \min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta})$ and $\sigma_{\max} = \max_k \sigma_k$. Suppose $\delta > 0$ is a given constant and the stepsize $\eta \leq \frac{\beta\delta}{L\sigma_{\max}^2/b + L\eta\delta}$. Then, after $T = \frac{2\beta\Delta_L}{\eta\delta}$ iterations, Algorithm 2 finds a point $\bar{\boldsymbol{\theta}}$ such that

$$\mathbb{E}\|\nabla \mathcal{L}_k(\bar{\boldsymbol{\theta}})\|^2 \leq \delta, \quad \forall k = 1, \dots, K. \tag{20}$$

Moreover, if we choose $\eta = \frac{\beta\delta}{L\sigma_{\max}^2/b + L\eta\delta}$, then the Conceptual PiKE algorithm requires at most

$$\bar{T} = \frac{2L\Delta_L(\sigma_{\max}^2/b + \gamma\delta)}{\delta^2\beta^2}$$

iterations to find a point satisfying equation (20).

Proof: We prove this by contradiction. Assume that $\max_k \mathbb{E}\|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 > \delta$ for $t = 0, \dots, T-1$. First notice that Theorem G.3 implies that for all t, we have

$$\begin{aligned}
\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) | \mathcal{F}_t] &\leq \mathcal{L}(\boldsymbol{\theta}_t) + \sum_{k=1}^K w_k^* \left(-\eta\beta \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2\sigma_{\max}^2}{2b} \right) \\
&\quad + \sum_{k=1}^K \frac{w_k^*}{2} (L\eta^2\gamma \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2)
\end{aligned} \tag{21}$$

where $\{w_k^*\}_{k=1}^K$ is the minimizer of the RHS of the equation (21) on the constrained set $\{(w_1, \dots, w_K) | \sum_{k=1}^K w_k = 1, w_k \geq 0 \forall k \in K\}$. Since w_k^* is the minimizer of the RHS of equation (21), we have

$$\begin{aligned}
w_k^* \left(-\eta\beta \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2}{2b} \sigma_{\max}^2 \right) + \frac{w_k^*}{2} L\eta^2\gamma \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2 \\
\leq \left(-\eta\beta \|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2}{2b} \sigma_{\max}^2 \right) + \frac{L\eta^2}{2} \gamma \|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2
\end{aligned} \tag{22}$$

where $k_t^* \in \arg \max_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2$. Moreover, we have

$$\eta \leq \frac{\beta\delta}{L\frac{\sigma_{\max}^2}{b} + L\gamma\delta} \leq \frac{\beta\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta})\|^2}{L\frac{\sigma_{\max}^2}{b} + L\gamma\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2},$$

Therefore,

$$\left(-\eta\beta\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2 + \frac{L\eta^2}{2b}\sigma_{\max}^2\right) + \frac{L\eta^2}{2}\gamma\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2 \leq -\frac{\beta\eta}{2}\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2 \quad (23)$$

Combining equation (21), (22), and (23), we obtain

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) \mid \mathcal{F}_t] \leq \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\beta\eta}{2}\|\nabla \mathcal{L}_{k_t^*}(\boldsymbol{\theta}_t)\|^2$$

Or equivalently,

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_{t+1}) \mid \mathcal{F}_t] \leq \mathcal{L}(\boldsymbol{\theta}_t) - \frac{\beta\eta}{2} \max_k \|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2$$

Taking the expectation over all sources of randomness in the algorithm, summing the above inequality from $t = 0$ to $t = T - 1$, and simplifying the resulting telescoping sum, we obtain:

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_T)] \leq \mathcal{L}(\boldsymbol{\theta}_0) - \frac{\beta\eta}{2} \sum_{t=1}^{T-1} \max_k \mathbb{E}[\|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2]$$

Recalling the contradiction assumption that $\delta < \mathbb{E}[\|\nabla \mathcal{L}_k(\boldsymbol{\theta}_t)\|^2]$, we get

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}_T)] \leq \mathcal{L}(\boldsymbol{\theta}_0) - \frac{\beta\eta}{2}T\delta$$

Using the definition $\Delta_L \triangleq \mathcal{L}(\boldsymbol{\theta}_0) - \min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta})$, we get

$$T \leq \frac{2\Delta_L}{\beta\eta\delta}$$

Finally notice that by setting $\eta = \frac{\beta\delta}{L\frac{\sigma_{\max}^2}{b} + L\gamma\delta}$, we get

$$T \leq \bar{T} = \frac{2\Delta_L}{\beta\eta} = \frac{2L\Delta_L}{\beta\delta^2} \left(\frac{\sigma_{\max}^2}{b} + \gamma\delta \right),$$

which completes the proof.

G.3 Conceptual PIKE versus Static Uniform Mix Batching

We now perform a standard analysis of Stochastic Gradient Descent (SGD), adapted to our specific setup where uniform (static) mini-batching is used (i.e., when $b_k = b/K$, $\forall k = 1, \dots, K$). As is common in the analysis of gradient-based algorithms in smooth, nonconvex settings, we begin by quantifying the expected decrease in the objective function at each iteration. This is known as the descent per iteration. Once this is established, we use a telescoping sum argument to derive the iteration complexity of the algorithm, which tells us how many steps are needed to achieve a desired level of accuracy, given a properly chosen learning rate.

Let's begin by analyzing the descent that occurs in a single iteration. Recall that the objective is defined as:

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathcal{L}(\boldsymbol{\theta}) := \sum_{k=1}^K \mathbb{E}_{x \sim \mathcal{D}_k} [\ell_k(\boldsymbol{\theta}; x)]. \quad (24)$$

Under Assumption 3.3, and using uniform static mini-batch sampling—where each source k contributes equally to the batch with size $b_k = b/K$ —the estimated gradient direction at iteration t is given by:

$$\mathbf{g}_t = \frac{1}{b} \left(\sum_{k=1}^K \sum_{\substack{i=1 \\ x_i \sim \mathcal{D}_k}}^{b/K} \nabla \ell_k(\boldsymbol{\theta}_t; x_i) \right) \quad (25)$$

This expression aggregates gradients from all sources using equally sized sub-batches, providing a mini-batch estimate of the full gradient. Let θ_{t+1} denote the updated parameter after performing a gradient descent step, given by $\theta_{t+1} = \theta_t - \eta \mathbf{g}_t$. By the descent lemma, the following inequality holds for the updated parameter: θ_{t+1} :

$$\mathcal{L}(\theta_{t+1}) \leq \mathcal{L}(\theta_t) - \eta \mathbf{g}_t^\top \nabla \mathcal{L}(\theta_t) + \frac{L\eta^2}{2} \|\mathbf{g}_t\|^2, \quad (26)$$

Conditioned on the natural filtration \mathcal{F}_t , we take the expectation with respect to the randomness in the sampled data and obtain:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta_{t+1}) | \mathcal{F}_t] &\leq \mathcal{L}(\theta_t) - \eta \mathbb{E}[\mathbf{g}_t | \mathcal{F}_t]^\top \nabla \mathcal{L}(\theta_t) + \frac{L\eta^2}{2} \mathbb{E}(\|\mathbf{g}_t\|^2 | \mathcal{F}_t) \\ &= \mathcal{L}(\theta_t) - \eta \left(\frac{1}{K} \sum_{k=1}^K \nabla \mathcal{L}_k(\theta_t) \right)^\top \left(\sum_{k=1}^K \nabla \mathcal{L}_k(\theta_t) \right) \\ &\quad + \frac{L\eta^2}{2b^2} \left(\left\| \sum_{k=1}^K \frac{b}{K} \nabla \mathcal{L}_k(\theta_t) \right\|^2 + \frac{b}{K} \sum_{k=1}^K \sigma_k^2 \right) \end{aligned}$$

Observing the fact that $\mathcal{L}(\theta) = \sum_{k=1}^K \mathcal{L}_k(\theta)$, we can write

$$\mathbb{E}[\mathcal{L}(\theta_{t+1}) | \mathcal{F}_t] \leq \mathcal{L}(\theta_t) - \frac{\eta}{K} \|\nabla \mathcal{L}(\theta_t)\|^2 + \frac{L\eta^2}{2b^2} \left(\frac{b^2}{K^2} \|\nabla \mathcal{L}(\theta_t)\|^2 + \frac{b}{K} \sum_{k=1}^K \sigma_k^2 \right) \quad (27)$$

Define $\sigma^2 := \sum_{k=1}^K \sigma_k^2$. By summing and Rearranging the terms, we obtain:

$$\mathbb{E}[\mathcal{L}(\theta_{t+1}) | \mathcal{F}_t] \leq \mathcal{L}(\theta_t) + \left(-\frac{\eta}{K} + \frac{L\eta^2}{2K^2} \right) \|\nabla \mathcal{L}(\theta_t)\|^2 + \frac{L\eta^2}{2bK} \sigma^2 \quad (28)$$

Taking the expectation over all sources of randomness in the algorithm, summing the above inequality from $t = 0$ to $t = T - 1$, and simplifying the resulting telescoping sum, we obtain:

$$\mathbb{E}[\mathcal{L}(\theta_T)] \leq \mathcal{L}(\theta_0) + \left(-\frac{\eta}{K} + \frac{L\eta^2}{2K^2} \right) \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\theta_t)\|^2] + \frac{LT\eta^2}{2bK} \sigma^2. \quad (29)$$

Let $\Delta_L := \mathcal{L}(\theta_0) - \min_{\theta} \mathcal{L}(\theta)$ denote the initial suboptimality gap. As is standard in iteration complexity analysis, we are interested in determining the number of iterations required to reach a point θ_{t_0} where $\mathbb{E}[\|\nabla \mathcal{L}(\theta_{t_0})\|^2] \leq \delta$. Suppose, for the moment, that this condition has not yet been satisfied. Then, we can rewrite the above inequality as

$$0 \leq \Delta_L + \left(-\frac{\eta}{K} + \frac{L\eta^2}{2K^2} \right) T\delta + \frac{LT\eta^2}{2bK} \sigma^2. \quad (30)$$

Equivalently, under a properly chosen step size (to be defined later), we obtain the following upper bound on the number of iterations

$$T \leq \frac{\Delta_L}{\left(\frac{\eta}{K} - \frac{L\eta^2}{2K^2} \right) \delta - \frac{L\eta^2}{2bK} \sigma^2}.$$

By optimizing the choice of η to minimize the right-hand side, we derive the following bound on the iteration complexity

$$T \leq \frac{2L\Delta_L \left(\delta + \frac{\sigma^2}{b} K \right)}{\delta^2}.$$

In summary, for the uniform mix batching strategy, the standard analysis yields the following upper bound on the iteration complexity:

$$\bar{T}_{\text{Uniform}} = \frac{2L\Delta_L \left(\delta + \frac{\sigma^2}{b} K \right)}{\delta^2}. \quad (31)$$

In contrast, when the task gradients are nearly orthogonal—i.e., when the alignment/conflict parameters satisfy $\bar{c}, \underline{c} \approx 0$ —Theorem 3.5 gives the following iteration complexity bound for the Conceptual PiKE algorithm:

$$\bar{T}_{\text{PiKE}} = \frac{2L\Delta_L (\delta + \sigma_{\max}^2/b)}{\delta^2}, \quad (32)$$

which clearly illustrates the advantage of PiKE for this regime in terms of iteration complexity upper-bound.

G.4 Balanced-PiKE: Theoretical Developments and Derivations

Here we prove Lemma 3.6 which is the idea behind the Balanced-PiKE algorithm. Before proving it, let us, for the sake of completeness, re-state the lemma

Lemma G.6 (Restatement of Lemma 3.6). *Let $0 < \tau \in \mathbb{R}$. Then, the tilted empirical risk minimization problem*

$$\min_{\boldsymbol{\theta}} \frac{1}{\tau} \log \left(\sum_{k=1}^K e^{\tau \mathcal{L}_k(\boldsymbol{\theta})} \right)$$

is equivalent to

$$\min_{\boldsymbol{\theta}} \max_{\substack{\mathbf{y} \in \mathbb{R}_+^K \\ \sum_{k=1}^K y_k = \tau}} \sum_{k=1}^K y_k \mathcal{L}_k(\boldsymbol{\theta}) - \sum_{k=1}^K \frac{y_k}{\tau} \log \left(\frac{y_k}{\tau} \right). \quad (33)$$

Moreover, for any fixed $\boldsymbol{\theta}$, the inner maximization problem is maximized at $y_k^ = \frac{\tau e^{\tau \mathcal{L}_k(\boldsymbol{\theta})}}{\sum_{j=1}^K e^{\tau \mathcal{L}_j(\boldsymbol{\theta})}}$, $\forall k$.*

The above lemma is the direct consequence of applying Lemma G.8 and Lemma G.7.

Lemma G.7. *For the problem*

$$\max_{\substack{\mathbf{y} \in \mathbb{R}_+^K \\ \sum_{k=1}^K y_k = \tau}} \left(\sum_{k=1}^K y_k x_k - \sum_{k=1}^K \frac{y_k}{\tau} \log \left(\frac{y_k}{\tau} \right) \right),$$

the optimal \mathbf{y} is given by

$$y_k^* = \frac{\tau e^{\tau x_k - 1}}{\sum_{k=1}^K e^{\tau x_k - 1}}$$

Proof: We start by forming and maximizing the Lagrangian function

$$\max_{\mathbf{y} \in \mathbb{R}_+^K} \left(\sum_{k=1}^K y_k x_k - \sum_{k=1}^K \frac{y_k}{\tau} \log \left(\frac{y_k}{\tau} \right) + \mu \left(\sum_{k=1}^K y_k - \tau \right) \right)$$

where μ is a free variable. Taking the partial derivative of the objective with respect to y_k and setting it to zero gives

$$y_k^* = \alpha e^{\tau x_k},$$

where the coefficient α is independent of the index k and should be chosen such that $\sum_k y_k^* = 1$, implying

$$y_k^* = \frac{\tau e^{\tau x_k}}{\sum_{j=1}^K e^{\tau x_j}}.$$

Lemma G.8. *Let $\mathbf{x} \in \mathbb{R}^K$ and $\tau > 0$. Then,*

$$\log \left(\sum_{k=1}^K e^{\tau x_k} \right) = \max_{\substack{\mathbf{y} \in \mathbb{R}_+^K \\ \sum_{k=1}^K y_k = \tau}} \left(\sum_{k=1}^K y_k x_k - \sum_{k=1}^K \frac{y_k}{\tau} \log \left(\frac{y_k}{\tau} \right) \right)$$

Proof. The proof is by simply plugging in the optimal value y^* obtained by Lemma G.7 in the objective function. \square