

Rethinking Data Mixture for Large Language Models: A Comprehensive Survey and New Perspectives

Yajiao Liu¹, Congliang Chen², Junchi Yang^{1,3}, Ruoyu Sun^{*1,3}

¹The Chinese University of Hong Kong, Shenzhen

²Shenzhen Loop Area Institute

³Shenzhen Research Institute of Big Data

yajiaoliu@link.cuhk.edu.cn, chencongliang@slai.edu.cn

{yangjunchi, sunruoyu}@cuhk.edu.cn

Abstract

Training large language models with data collected from various domains can improve their performance on downstream tasks. However, given a fixed training budget, the sampling proportions of these different domains significantly impact the model’s performance. How can we determine the domain weights across different data domains to train the best-performing model? In this paper, we provide a comprehensive overview of existing data mixture methods. First, we propose a fine-grained categorization of existing methods, extending beyond the previous offline and online classification. Offline methods are further grouped into heuristic-based, algorithm-based, and function fitting-based methods. For online methods, we categorize them into three groups—online min-max optimization, online mixing law, and other approaches—by drawing connections with the optimization frameworks underlying offline methods. Second, we summarize the problem formulations, representative algorithms for each subtype of offline and online methods, and clarify the relationships and distinctions among them. Finally, we discuss the advantages and disadvantages of each method and highlight key challenges in the field of data mixture.

1 Introduction

Recently, many works have focused on improving the generalization capability of large language models (LLMs) on downstream tasks by increasing the diversity of training data. This is typically achieved by collecting data from various domains or different sources (e.g., CommonCrawl, Wikipedia, and Github), either during the pre-training phase (Radford et al., 2019; Gao et al., 2020), or instruction fine-tuning phase (Wei et al., 2022; Chung et al., 2022). However, given a fixed training budget, the sampling proportions of these different domains or data sources significantly impact the model’s

performance (Xie et al., 2023; Albalak et al., 2023; Shen et al., 2024). This raises an important question: how can we determine the **domain weights** (the sampling probabilities for each domain) across different data domains to train the best-performing model within constrained computational resources? **Data mixture**, also known as domain reweighting, is an emerging research area dedicated to addressing this challenge. It focuses on optimizing the domain weights to sample different domain data, thereby improving the effectiveness and efficiency of language model training.

Many early works rely on heuristically chosen or manually assigned domain weights for training language models (Radford et al., 2019; Gao et al., 2020). However, subsequent studies (Xie et al., 2023; Fan et al., 2023; Albalak et al., 2023) argue that such approaches are likely suboptimal and propose more principled methods to determine the domain weights automatically. Albalak et al. (2024) provide an overview of this rapidly developing field, categorizing existing data mixture methods into two main types: **offline methods**, which determine domain weights separately from the training of the target model, and **online methods**, which dynamically adjust domain weights during the target model’s training process. As the cost of training LLMs continues to rise with the increasing scale and diverse sources of data, data mixture has become increasingly important. Meanwhile, many new studies—beyond the scope of Albalak et al. (2024)—have gradually emerged. Accordingly, we believe it is timely to provide a more comprehensive and updated review for the community. A detailed comparison between our survey and Albalak et al. (2024) is provided in Appendix D.

In this paper, we propose a more fine-grained categorization of existing data mixture methods, building on the framework proposed by Albalak et al. (2024). Offline methods assume that, for a given corpus composed of multiple domains, there

*Corresponding Author.

Algorithm 1 Framework of Offline methods

- 1: **Input:** Proxy model ϕ , target model θ , training data $\mathcal{D}_{\text{train}}$, validation data \mathcal{D}_{val} (which can be optional);
 - 2: **Initialization:** w^0 ;
 - 3: $w^* = \text{Get_Weights}(w^0, \phi, \mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}})$;
 - 4: $\theta^* = \arg \min_{\theta} \sum_{i=1}^k w_i^* \ell_{\text{train},i}(\theta)$;
 - 5: **Output:** w^* and θ^* .
-

exist optimal static domain weights that lead to the best-performing model. Based on this assumption, they typically estimate these weights using proxy models, which are typically smaller than the target model to reduce the additional training costs, and then apply the estimated weights to train the target model, as described in Algorithm 1. We divide offline methods into three subcategories: heuristic-based methods, algorithm-based methods, and function fitting-based methods. **Heuristic-based methods**, such as uniform sampling and proportional sampling, determine domain weights based on simple heuristics. Although they do not rely on proxy models, we treat them as degenerate cases of offline methods for completeness in our categorization. **Algorithm-based methods** apply specific algorithms to determine the optimal domain weights with proxy models. Specifically, we realize that data mixture should be regarded as a subfield of Multi-Domain Learning (MDL), which is often closely associated with Multi-Task Learning (MTL) ¹. Some methods proposed for MTL can also address MDL, and vice versa (Yang and Hospedales, 2015). Therefore, we systematically analyze the connections between existing algorithm-based methods and two prominent optimization frameworks widely used in MTL: **min-max optimization** and **bi-level optimization**, aiming to offer a deeper understanding and valuable insights for optimization strategies in data mixture. **Function fitting-based methods** empirically learn a mapping from domain weights to model performance by conducting multiple training runs on proxy models using different domain weights. Once this function—also referred to as a **mixing law**—is learned, it can be used to infer the optimal domain weights for training the target model.

Online methods jointly optimize the target model and the domain weights during training, as illus-

¹One major difference between Data Mixture and the previous methods like MDL and MTL is that, in Data Mixture, all domains share the same set of neural network parameters.

Algorithm 2 Framework of Online methods

- 1: **Input:** Target model θ , training data $\mathcal{D}_{\text{train}}$, validation data \mathcal{D}_{val} (optional), the total number of training steps T , the update interval t_{update} ;
 - 2: **Initialization:** The target model θ^0 , and w^0 ;
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: $\theta^t = \text{Update_}\theta(w^{t-1}, \theta^{t-1}, \mathcal{D}_{\text{train}})$;
 - 5: **if** $t \bmod t_{\text{update}} = 0$ **then**
 - 6: $w^t = \text{Update_}w(w^{t-1}, \theta^t, \mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}})$;
 - 7: **end if**
 - 8: **end for**
 - 9: **Output:** θ^T .
-

trated in Algorithm 2. During the target model training, the domain weights are dynamically optimized to effectively and efficiently steer the target model toward optimal performance. Most online methods can be regarded as online variants of existing offline approaches. For example, certain offline methods can be transformed into their online counterparts by replacing the proxy models used for optimizing domain weights with the target model itself, thereby enabling joint optimization of both domain weights and the target model during training. From this perspective, we categorize existing online methods into three groups: **online min-max optimization** methods, **online mixing law** methods, and other approaches.

Our contributions can be summarized as follows:

1. We propose a fine-grained categorization of existing data mixture methods, extending beyond the previous offline and online classification. Specifically, we divide offline methods into heuristic-based, algorithm-based, and function fitting-based methods. For online methods, we categorize them into three groups—online min-max optimization, online mixing law, and other approaches—by drawing connections with the optimization frameworks underlying offline methods.
2. For each subtype of offline and online methods, we summarize their problem formulations, representative algorithms, and clarify the relationships and distinctions among them, aiming to offer a deeper understanding and valuable insights for optimization strategies in data mixture to the community.
3. Finally, we discuss the advantages and disadvantages of each method and highlight key challenges and future directions in this field.

2 Notations

Given k data domains, each domain $i \in [k] = \{1, 2, \dots, k\}$ is associated with a training dataset $D_{\text{train},i}$ and a validation dataset $D_{\text{val},i}$. We represent a composition of the training data $\mathcal{D}_{\text{train}} = \{D_{\text{train},1}, \dots, D_{\text{train},k}\}$ using domain weights $\mathbf{w} = [w_1, w_2, \dots, w_k]^T$, where w_i denotes the weight assigned to domain i and $\sum_{i=1}^k w_i = 1, w_i \geq 0$ (i.e., $\mathbf{w} \in \mathbb{R}_{\geq 0}^k$ lies within a probabilistic simplex Δ^k). Our ultimate goal is to obtain the best-performing model θ^* , where θ is used to parameterize the target model, by optimizing the domain weights \mathbf{w} across the training data $\mathcal{D}_{\text{train}}$. We use ϕ to represent the parameters of the proxy model.

The best-performing model refers to the one that achieves the highest overall performance across various downstream tasks. However, during the training of a model θ , we typically optimize a loss function $\ell(\theta) : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ as the surrogate objective for its downstream task performance, which is often the negative log-likelihood in the context of LLMs. For example, in the pretraining setting, the loss can be written as $\ell(x; \theta) = -\log p_{\theta}(x) = -\frac{1}{|x|} \sum_{z=1}^{|x|} \log p_{\theta}(x_z | x_{1:z-1})$, where $|x|$ denotes the number of tokens in a given example x . There are typically two settings for computing this loss: on the training data or the validation data. We denote the average loss of model θ on each training domain $D_{\text{train},i}$, $i \in [k]$ as $\ell_{\text{train},i}(\theta)$ (i.e., $\ell_{\text{train},i}(\theta) = \ell(D_{\text{train},i}; \theta) = \frac{1}{|D_{\text{train},i}|} \sum_{x \in D_{\text{train},i}} \ell(x; \theta)$). For validation data, we consider two cases: **in-domain** validation data $\mathcal{D}_{\text{val}}^{\text{in}} = \{D_{\text{val},1}, \dots, D_{\text{val},k}\}$, which correspond to the same domains as the training data $\mathcal{D}_{\text{train}}$, and **out-of-domain** validation data $\mathcal{D}_{\text{val}}^{\text{ood}} = \{\hat{D}_{\text{val},1}, \dots, \hat{D}_{\text{val},m}\}$, consisting of m domains completely unrelated to $\mathcal{D}_{\text{train}}$ ². We denote the average loss of θ on $D_{\text{val},j}$, $j \in [k]$ as $\ell_{\text{val},j}(\theta) = \ell(D_{\text{val},j}; \theta)$, on $\hat{D}_{\text{val},j}$, $j \in [m]$ as $\hat{\ell}_{\text{val},j}(\theta) = \ell(\hat{D}_{\text{val},j}; \theta)$.

We provide a detailed explanation of **static weights** and **dynamic weights** in Appendix A to clarify how different methods leverage them to approximate the best-performing model θ^* .

²In fact, the **out-of-domain** validation data $\mathcal{D}_{\text{val}}^{\text{ood}}$ could contain several domains related to the training data $\mathcal{D}_{\text{train}}$. For simplicity, we assume that all domains in $\mathcal{D}_{\text{val}}^{\text{ood}}$ are completely unrelated to $\mathcal{D}_{\text{train}}$.

3 Offline Methods

Offline methods assume the existence of optimal static domain weights \mathbf{w}^* for a given corpus composed of multiple domains. Based on the strategies used to find \mathbf{w}^* (i.e., line 3 of Algorithm 1), we categorize offline methods into three groups: heuristic methods (Section 3.1), algorithm-based methods (Section 3.2), and function fitting-based methods (Section 3.3).

3.1 Heuristic-based Methods

Heuristic-based methods determine \mathbf{w}^* using predefined heuristics. For example, uniform sampling assigns equal weights to all k domains, i.e., $\mathbf{w}^* = [\frac{1}{k}, \dots, \frac{1}{k}]$. Since these methods can be implemented by omitting line 3 in Framework 1, we consider them as degenerate cases of offline methods, even though they do not rely on proxy models. Further details are provided in Appendix B.1.

3.2 Algorithm-based Methods

Next, we systematically survey existing algorithm-based methods for data mixture under two optimization frameworks: **min-max optimization** and **bi-level optimization**. The definition of the best-performing model θ^* can vary depending on the specific setting—when the model’s loss is computed solely using training data or incorporating validation data. This distinction leads to different optimization algorithms being employed to find \mathbf{w}^* .

3.2.1 Min-Max Optimization

Given all k domains, we could determine the \mathbf{w}^* by focusing on whether the proxy model ϕ is well-optimized on each domain of the training data $\mathcal{D}_{\text{train}}$ without considering any validation data. One surrogate objective to achieve this is to ensure that the worst-case performance across domains remains satisfactory, which can be formulated as the following min-max problem:

$$\min_{\phi} \max_{\mathbf{w} \in \Delta^k} \sum_{i=1}^k w_i f_i(\phi), \quad (1)$$

where $f_i(\phi)$ is often defined based on $\ell_{\text{train},i}(\phi)$, the average training loss of ϕ on the i -th domain $D_{\text{train},i}$, for $i \in [k]$.

Min-max optimization. Equation 1 represents a nonconvex-concave min-max optimization problem (Lin et al., 2020; Razaviyayn et al., 2020). The well-known gradient descent ascent (GDA) algorithm (Daskalakis and Panageas, 2018; Lin et al., 2020), a simple extension of gradient descent (GD)

	Category	Method	Training phase(s)
	Heuristic	Uniform / Proportional sampling	PT / CPT / SFT
Offline	Algorithm-based (Min-max)	DoReMi (Xie et al., 2023)	PT
	Algorithm-based (Bi-level)	DoGE (Fan et al., 2023) ScaleBiO (Pan et al., 2024b)	PT SFT
	Function fitting-based	DML (Ye et al., 2024)	PT
		BiMix (Ge et al., 2024)	PT
D-CPT (Que et al., 2024)		CPT	
AutoScale (Kang et al., 2025)		PT	
RegMix (Liu et al., 2024)		PT	
MDE (Belenki et al., 2025)	PT		
Online	Online min-max	MFTCoder (Liu et al., 2023a)	SFT
		ShearedLLaMa (Xia et al., 2024)	CPT
		Velocitune (Luo et al., 2024)	CPT
	Online mixing laws	Skill-It (Chen et al., 2023) Aoili (Chen et al., 2024)	CPT / SFT PT
Others	ODM (Albalak et al., 2023)	PT	
	ADO (Jiang et al., 2024)	PT	

Table 1: Overview of existing data mixture methods discussed in this paper. In the last column, we report the specific training setting for each method, as described in the corresponding original paper. SFT denotes Supervised Fine-Tuning for instruction tuning; PT refers to Pretraining, and CPT stands for Continued Pretraining.

to the min-max setting, is one of the most popular algorithms for solving this problem, capable of finding a reasonably good solution for the proxy model ϕ and domain weights w .

Xie et al. (2023) propose Domain Reweighting with Minimax Optimization (DoReMi) to achieve the same goal in the LLMs pretraining setting with $f_i(\phi) = \ell_{\text{train},i}(\phi) - \ell_{\text{train},i}(\phi_{\text{ref}})$, where ϕ_{ref} represents a pre-trained reference model, and its loss $\ell_{\text{train},i}(\phi_{\text{ref}})$ serves as the baseline for the loss $\ell_{\text{train},i}(\phi)$. The objective is to minimize the worst-case excess loss across domains, which is the gap in loss between ϕ and ϕ_{ref} . They adopt the Group DRO (Sagawa et al., 2020) to solve this problem, which can be viewed as a variant of GDA. Since w lies within the probabilistic simplex Δ^k , at each step t , w^t is updated using exponential gradient ascent (i.e., mirror ascent). Then, the average weights are returned as the optimized domain weights (i.e., $w^* = \frac{1}{T} \sum_{i=1}^T w^t$). Finally, they use w^* to re-sample the training data $\mathcal{D}_{\text{train}}$ for training a larger target model θ . There are other variants of GDA, such as Optimistic GDA (Daskalakis et al., 2018) and Smoothed GDA (Zhang et al., 2022), which can also be used to solve Equation 1.

3.2.2 Bi-Level Optimization

Ensuring that the model is well-optimized on each domain in the training data $\mathcal{D}_{\text{train}}$ implicitly assumes that good optimization leads to good gen-

eralization. However, we can explicitly enhance the model’s generalization ability by finding w^* that maximize its performance on the validation data \mathcal{D}_{val} . In this case, whether evaluated using in-domain or out-of-domain validation loss, this objective can be formulated as the following standard bi-level optimization problem (Franceschi et al., 2017; Lorraine et al., 2019):

$$\begin{aligned}
 w^* \in \arg \min_{w \in \Delta^k} \ell_{\text{val}}(\phi^*(w)) \\
 \text{s.t. } \phi^*(w) \in \arg \min_{\phi} \sum_{i=1}^k w_i \ell_{\text{train},i}(\phi),
 \end{aligned} \tag{2}$$

where $\ell_{\text{val}}(\phi^*(w))$ can be $\frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(\phi^*(w))$ or $\frac{1}{m} \sum_{j=1}^m \hat{\ell}_{\text{val},j}(\phi^*(w))$ for considering the in-domain validation data $\mathcal{D}_{\text{val}}^{\text{in}}$ or the out-of-domain validation data $\mathcal{D}_{\text{val}}^{\text{ood}}$.

Bi-level optimization. A bi-level optimization problem involves a two-level hierarchical structure (i.e., outer and inner levels). The inner-level problem is auxiliary since its solution supports the outer-level problem in finding a better solution. In equation 2, the inner-level problem involves obtaining a good model $\phi^*(w)$ trained on a combination of multiple domain data $\mathcal{D}_{\text{train}}$, weighted by domain weights w . The outer-level problem aims to find the optimal domain weights w^* by evaluating $\phi^*(w)$ on the validation data \mathcal{D}_{val} . There are three primary frameworks for solving bi-level optimiza-

Method	Mixing law $f_j, j \in [k]$	The learning parameters	Objective function (including aggregation function g)
DML	$\ell_{\text{val},j}(\mathbf{w}) = c_j + b_j \exp\left(\sum_{i=1}^k A_{ij} w_i\right)$	$\mathbf{c}, \mathbf{b} \in \mathbb{R}^k, \mathbf{A} \in \mathbb{R}^{k \times k}$	$\mathbf{w}^* = \min_{\mathbf{w}} \sum_{j=1}^k h_j \ell_{\text{val},j}(\mathbf{w})$
BiMix	$\ell_{\text{val},j}(w_j, s) = \frac{A_j}{w_j^{\alpha_j}} \left(\frac{B_j}{s^{\beta_j}} + C_j \right)$	$A, B, C, \alpha, \beta \in \mathbb{R}^k$	$\mathbf{w}^* = \min_{\mathbf{w}} \frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(w_j, s)$ s.t. $\sum w_j = 1$
AutoScale	$\ell_{\text{val},j}(w_j, D) = (D_0^j + w_j D)^{-\gamma_j} + c_j$	$D_0, \gamma, c \in \mathbb{R}^k$	$\mathbf{w}^* = \min_{\mathbf{w}} \frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(w_j, D)$ s.t. $\sum w_j = 1$
RegMix	$\ell_{\text{val},j}(\mathbf{w}) = \mathbf{M}_j(\mathbf{w})$	Regression model $\{\mathbf{M}_j\}_{j=1}^k$	$\mathbf{w}^* = \min_{\mathbf{w}} \ell_{\text{val},j}(\mathbf{w})$
MDE	$\ell_{\text{val},j}(\mathbf{w}; \text{MDE features}) = \mathbf{M}_j(\mathbf{w}; L_{\text{MDE}}^1, \dots, L_{\text{MDE}}^k)$, where $L_{\text{MDE}}^j = -\frac{1}{ \mathcal{D}_{\text{val},j} } \sum_{x \in \mathcal{D}_{\text{val},j}} \sum_{i=1}^k w_i \log p_{\phi_i^*}(x)$	Regression model $\{\mathbf{M}_j\}_{j=1}^k$	$\mathbf{w}^* = \min_{\mathbf{w}} \frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(\mathbf{w}, \text{MDE features})$

Table 2: The comparison of function fitting-based methods in data mixture.

tion (BLO) problems (Zhang et al., 2023): (1) the Implicit Function (IF)-based approach, (2) the Gradient Unrolling (GU)-based approach, and (3) the Value Function (VF)-based approach. More details are provided in Appendix B.2.1.

Fan et al. (2023) propose DOrain reweighting with Generalization Estimation (DoGE) to find the optimal domain weights for LLMs pretraining, based on validation data, including $\mathcal{D}_{\text{val}}^{\text{in}}$ and $\mathcal{D}_{\text{val}}^{\text{out}}$. To solve Equation 2 cheaply, they provide a single-loop Gradient Unrolling (GU)-based approach, which only updates ϕ with a single stochastic step in the inner problem. For the outer-level problem, considering the in-domain validation data $\mathcal{D}_{\text{val}}^{\text{in}} = \{D_{\text{val},1}, \dots, D_{\text{val},k}\}$, they update \mathbf{w} at each step t using a simple first-order update rule: $\mathbf{w}^t = \mathbf{w}^{t-1} \odot \exp(\eta^t \mathcal{W}^t)$, where $\mathcal{W}^t = [W_1^t, \dots, W_k^t]$ is a k -dimension vector in which each component $W_i^t = \langle \nabla \ell_{\text{train},i}(\phi^t), \sum_{j \in [k]} \nabla \ell_{\text{val},j}(\phi^t) \rangle$, $\ell_{\text{val},j}(\phi^t)$ is the validation loss of ϕ^t on $D_{\text{val},j}$. Similar to DoReMi, the final domain weights are averaged over all middle domain weights (i.e., $\mathbf{w}^* = \frac{1}{T} \sum_{i=1}^T \mathbf{w}^t$) and then used to resample the training data $\mathcal{D}_{\text{train}}$ for training the target model θ .

Pan et al. (2024b) propose ScaleBiO to find the optimal domain weights with proxy model ϕ during the supervised instruction fine-tuning (SFT) phase, which is also based on $\mathcal{D}_{\text{val}}^{\text{in}}$ or $\mathcal{D}_{\text{val}}^{\text{out}}$. They adopt the Fully First-order Stochastic Approximation (F²SA) method introduced by Kwon et al. (2023) to solve Equation 2, which is a value function (VF)-based bi-level optimization approach. By integrating F²SA with the memory-efficient training technique LISA (Pan et al., 2024a), they successfully scale the bi-level optimization technique to 34-billion-parameter LLMs on eight A40 GPUs.

3.3 Function Fitting-based methods

Inspired by scaling law (Kaplan et al., 2020; Hoffmann et al., 2022), several studies attempt to empirically learn a mapping from domain weights to

model performance by conducting multiple training runs on proxy models using different domain weights. The resulting function enables the prediction of model performance trained with arbitrary domain weights before actual training, thereby facilitating the identification of optimal domain weights \mathbf{w}^* for a given corpus composed of multiple domains. We refer to these approaches as function fitting-based methods and the learned functions as **mixing laws**.

In general, function fitting-based methods evaluate model performance using validation loss. For simplicity, we illustrate these methods using in-domain validation loss (the out-of-domain case is analogous), computed over the in-domain validation data $\mathcal{D}_{\text{val}}^{\text{in}} = \{D_{\text{val},1}, \dots, D_{\text{val},k}\}$. These methods typically consist of three key components for identifying \mathbf{w}^* : **(1)** fitting a mixing law f_j for each validation domain $j, j \in [k]$:

$$f_j \in \arg \min_{f_j} \sum_{l=1}^R \text{erf}(f_j(\mathbf{w}_l), \ell_{\text{val},j}(\phi^*(\mathbf{w}_l))), \quad (3)$$

$$\text{where } \phi^*(\mathbf{w}_l) \in \arg \min_{\phi} \sum_{i=1}^k w_{l,i} \ell_{\text{train},i}(\phi),$$

where $\text{erf}(\cdot, \cdot)$ denotes an error function (e.g., mean squared error), and each f_j is fitted using R input-output pairs $\{(\mathbf{w}_l, \ell_{\text{val},j}(\phi^*(\mathbf{w}_l)))\}_{l=1}^R$. Each pair is obtained by training a proxy model ϕ on $\mathcal{D}_{\text{train}}$ using domain weights \mathbf{w}_l , and then evaluating the resulting model $\phi^*(\mathbf{w}_l)$ on $D_{\text{val},j}$ to compute its average loss $\ell_{\text{val},j}(\phi^*(\mathbf{w}_l))$; **(2)** defining an aggregation function g to combine all mixing laws into a single objective: $g(f_1(\mathbf{w}), \dots, f_k(\mathbf{w}))$ to predict model performance for any given \mathbf{w} ; **(3)** optimizing the objective using appropriate algorithms to identify the optimal domain weights $\mathbf{w}^* = \min_{\mathbf{w} \in \Delta^k} g(f_1(\mathbf{w}), \dots, f_k(\mathbf{w}))$.

Function fitting-based methods reduce the bi-level optimization problem to a single-level one by fitting mixing laws, which can be solved using gradient-based methods. We summarize these

methods in Table 2. For notational convenience, we use f_j and $\ell_{\text{val},j}$ interchangeably in the following discussion. Ye et al. (2024) propose the Data Mixing Law (DML), modeling each mixing law as an exponential function of the domain weights \mathbf{w} , as shown in Table 2. They adopt a weighted sum as the aggregation function g over all mixing laws f_j , $j \in [k]$: $g(f_1(\mathbf{w}), \dots, f_k(\mathbf{w})) = \sum_{j=1}^k h_j \ell_{\text{val},j}(\mathbf{w})$, where h_j denotes the proportion of validation domain j in the overall validation dataset. Finally, they perform a grid search over a set of domain weights to identify the one that minimizes the aggregation function g as the optimal domain weights \mathbf{w}^* . BIMIX (Ge et al., 2024) and AutoScale (Kang et al., 2025) are bivariate data mixing laws that jointly consider domain weights \mathbf{w} and training data size. BIMIX represents data size using training steps s , whereas AutoScale uses the number of tokens to quantify training data size D . Both methods adopt the average sum as the aggregation function and formulate a constrained optimization problem: $\mathbf{w}^* = \min_{\mathbf{w}} \frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(w_j, s \text{ or } D)$ s.t. $\sum w_j = 1$. BIMIX solves this problem using Lagrange multipliers and numerical methods (Virtanen et al., 2020), while AutoScale solves it using projected gradient descent. We introduce D-CPT law (Que et al., 2024) in Appendix B.3.1.

Two recent works learn each mixing law f_j , $j \in [k]$ through regression models. RegMix (Liu et al., 2024) adopts a linear regression model or LightGBM—a gradient boosting algorithm that builds an ensemble of decision trees. Instead of minimizing a weighted or an average sum of all mixing laws, their objective is to minimize the mixing law fitted on a single validation domain, PileCC. Belenki et al. (2025) also employ regression models to learn the mixing laws. In addition to domain weights, they propose a Mixture of Data Experts (MDE) to approximate the cross-entropy losses associated with domain weights and use these approximations (i.e., the MDE features) as additional input features. To obtain the MDE features, they first train a set of data experts ϕ_i^* , where $i \in [k]$, each on a corresponding training domain $\mathcal{D}_{\text{train},i}$. Then, for each validation domain $j \in [k]$, the MDE feature is computed as: $L_{\text{MDE}}^j = \frac{1}{|\mathcal{D}_{\text{val},j}|} \sum_{x \in \mathcal{D}_{\text{val},j}} P_{\text{MDE}}(x, \mathbf{w})$, where $P_{\text{MDE}}(x, \mathbf{w}) := -\sum_{i=1}^k w_i \log p_{\phi_i^*}(x)$. They adopt the average sum as the aggregation function and solve the optimization problem: $\mathbf{w}^* =$

$\min_{\mathbf{w}} \frac{1}{k} \sum_{j=1}^k \ell_{\text{val},j}(\mathbf{w}, \text{MDE features})$ using the Vizier framework (Song et al., 2024).

3.4 Advantages and Disadvantages

Offline methods explicitly identify the optimal domain weights \mathbf{w}^* for a corpus composed of multiple domains, which can be beneficial in certain scenarios. For instance, publicly releasing \mathbf{w}^* would allow the research community or practitioners to reuse them to train their own models on the same corpus (Mehta et al., 2024).

However, all offline methods face (1) **model transfer issue**: Albalak et al. (2023) demonstrates that the optimal domain weights \mathbf{w}^* may not transfer well across models when adopting new model architectures or different tokenizers. Algorithm-based methods also face (2) **model scale issue**: the optimal weights derived from a small proxy model fail to generalize to larger target models, even if they share the same architecture and tokenizer (Jiang et al., 2024; Kang et al., 2025); (3) **data scale issue**: the optimal weights can shift as the scale of the training data increases (Kang et al., 2025). Function fitting-based methods can potentially address issues (2) and (3) by leveraging scaling laws to extrapolate the optimal weights obtained from small-scale training to larger model sizes and larger training data sizes, though only DML and AutoScale explicitly incorporate this idea into their designs. Nonetheless, function fitting-based methods also face a scalability issue: as the number of domains increases, the cost of fitting mixing laws grow linearly or even exponentially, making small-scale experiments expensive.

4 Online Methods

During target model θ training, online methods employ dynamic domain weights \mathbf{w}^t to effectively and efficiently steer the target model θ toward optimal performance. Following the framework outlined in Algorithm 2, most online methods update θ at each training step t (line 4 in Algorithm 2) in a similar manner: $\theta^t = \theta^{t-1} - \eta \sum_{i=1}^k w_i^{t-1} \nabla \ell_{\text{train},i}(\theta^{t-1})$, where η is the learning rate. The main difference among these methods lies in how they update the dynamic weights \mathbf{w}^t (line 6 in Algorithm 2). Based on this difference, we categorize existing online methods into three groups: online min-max optimization, online mixing law, and other methods.

Method	The update rules of \mathbf{w}^t	The value of t_{update}	The learning progress on each domain i
DoReMi _{online}	$w_i^t = \frac{w_i^{t-t_{\text{update}}} \exp(\zeta V_i^t)}{\sum_{i=1}^k w_i^{t-t_{\text{update}}} \exp(\xi V_i^t)}$	$t_{\text{update}} = 1$	$V_i^t = \max\{\ell_{\text{train},i}(\boldsymbol{\theta}^t) - \ell_{\text{train},i}(\boldsymbol{\theta}_{\text{ref}}), 0\}$
ShearedLLaMa	$w_i^t = \frac{w_i^{t-t_{\text{update}}} \exp(V_i^t)}{\sum_{i=1}^k w_i^{t-t_{\text{update}}} \exp(V_i^t)}$	$t_{\text{update}} > 1$	$V_i^t = \max\{\ell_{\text{val},i}(\boldsymbol{\theta}^t) - \ell_{\text{ref},i}, 0\}$
Velocitune	$w_i^t = \frac{w_i^{t-t_{\text{update}}} \exp(V_i^t)}{\sum_{i=1}^k w_i^{t-t_{\text{update}}} \exp(V_i^t)}$	$t_{\text{update}} > 1$	$V_i^t = \text{clamp}\{0, \frac{\ell_{\text{val},i}(\boldsymbol{\theta}^t) - \ell_{\text{ref},i}}{\ell_{\text{init},i} - \ell_{\text{ref},i}}, 1\}$

Table 3: Overview of three online min-max optimization methods: DoReMi_{online}, ShearedLLaMa, and Velocitune. These methods share a common update rule for the target model $\boldsymbol{\theta}$: $\boldsymbol{\theta}^t = \boldsymbol{\theta}^{t-1} - \eta \sum_{i=1}^k w_i^{t-1} \nabla \ell_{\text{train},i}(\boldsymbol{\theta}^{t-1})$, where η is the learning rate. The dynamic domain weights \mathbf{w}^t are updated every t_{update} steps based on V_i^t , with ζ denoting the learning rate for updating \mathbf{w}^t . Velocitune (Luo et al., 2024) refers to V_i^t as the learning velocity of the target model on domain i , which we interpret as a measure of learning progress.

Method	The update rules of \mathbf{w}^t	The learning parameter A_{ij}^t
Skill-It	$w_i^t = \frac{w_i^{t-t_{\text{update}}} \exp(\zeta \sum_{j=1}^k A_{ij}^{t-t_{\text{update}}} \ell_{\text{val},j}(\boldsymbol{\theta}^t))}{\sum_{i=1}^k w_i^{t-1} \exp(\zeta \sum_{j=1}^k A_{ij}^{t-t_{\text{update}}} \ell_{\text{val},j}(\boldsymbol{\theta}^t))}$	$A_{ij}^t = \ell_{\text{val},i}^t(\mathbf{w}) \left(\ell_{\text{val},i}^{T+1}(\mathbf{1}_j) - \ell_{\text{val},i}^1(\mathbf{1}_j) \right) / \ell_{\text{val},i}^1(\mathbf{1}_j)$
Aioli	$w_i^t = \frac{w_i^{t-t_{\text{update}}} \exp(\zeta \sum_{j=1}^k A_{ij}^{t-t_{\text{update}}})}{\sum_{i=1}^k w_i^{t-1} \exp(\zeta \sum_{j=1}^k A_{ij}^{t-t_{\text{update}}})}$	A_{ij} is fitted from $\ell_{\text{val},j}^{t+1}(\mathbf{w}) = \ell_{\text{val},j}^t(\mathbf{w}) - \sum_{i=1}^k A_{ij} w_i^t$

Table 4: Overview of two online mixing law methods: Skill-It and Aioli. We provide a simplified version of each method to highlight its main components. Both approaches update the dynamic domain weights \mathbf{w}^t with $t_{\text{update}} > 1$. ζ denotes the learning rate for updating \mathbf{w}^t . T denotes the total number of updates to the domain weights \mathbf{w}^t .

4.1 Online Min-Max Optimization Methods

Gradient manipulation methods in MTL, such as MGDA (Désidéri, 2012) and FAMO (Liu et al., 2023b), solve the min-max problem in Equation 1 in an online fashion; further details are provided in Appendix C.1.1. Therefore, these methods can be viewed as online approaches applicable to the data mixture problem. Building on this insight, MFTCoder (Liu et al., 2023a) applies FAMO to multi-task instruction tuning for code LLMs, replacing the training loss with the validation loss for each domain when updating the domain weights \mathbf{w}^t across different domains.

Similarly, DoReMi (Xie et al., 2023) can also be operated in an online fashion by replacing the proxy models used for optimizing domain weights with the target model itself, enabling joint optimization of both domain weights and the target model throughout training. Consequently, some recent works (Chen et al., 2024; Belenki et al., 2025) classify it as an online method. ShearedLLaMa (Xia et al., 2024) improves upon this online version of DoReMi (referred to as DoReMi_{online}) by using the loss predicted by a fitted scaling law on each domain as reference loss and also replacing training loss with validation loss for each domain when updating \mathbf{w}^t during continued pretraining of LLMs.

Luo et al. (2024) proposed Velocitune, a method that improves upon ShearedLLaMa. They introduce the concept of learning velocity, defined as: $V_i^t = \frac{\ell_{\text{val},i}(\boldsymbol{\theta}^t) - \ell_{\text{ref},i}}{\ell_{\text{init},i} - \ell_{\text{ref},i}}$, $i \in [k]$, to dynamically quantify how much the model has yet learned from each domain and then will be used to adjust the domain weights \mathbf{w}^t accordingly. $\ell_{\text{init},i}$ denotes the initial validation loss of the model $\boldsymbol{\theta}$ on each domain i , while $\ell_{\text{ref},i}$ is the reference loss, representing the final loss that the model is expected to achieve on each domain. As in ShearedLLaMa, the value of $\ell_{\text{ref},i}$ is predicted using a scaling law on each domain, providing a cost-efficient estimation of the final loss. We summarize different rules for updating \mathbf{w}^t across these methods in Table 3 to facilitate comparison; see Appendix C.1.2 for more details.

4.2 Online Mixing Law Methods

Chen et al. (2024) provide a unified perspective on several data mixture methods—including DML (Ye et al., 2024), DoReMi_{online}, and Skill-It (Chen et al., 2023)—by interpreting them as instances of optimizing a class of online mixing laws. These online mixing laws take the form of $\ell_{\text{val},j}^{t+1}(\mathbf{w}) = c_j^t + b_j^t \sigma \left(\sum_{i=1}^k -A_{ij}^t w_j^t \right)$, $j \in [k]$, where the matrix $A^t \in \mathbb{R}^{k \times k}$ and the vectors $\mathbf{b}^t, \mathbf{c}^t \in \mathbb{R}^k$ are mixing law parameters specified by the respec-

tive methods. The function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is either the identity function or the exponential function. By empirically examining their mixing law parameters with this unified formulation, [Chen et al. \(2024\)](#) observe that these approaches specify the mixing law parameters inaccurately and result in poor performance. To address this problem, they propose a simple online data mixing method AIOLI, directly fitting linear mixing laws (i.e., σ is the identity function) with the history of validation losses and domain weights during training: $\ell_{\text{val},j}^{t+1}(\mathbf{w}) = \ell_{\text{val},j}^t(\mathbf{w}) - \sum_{i=1}^k A_{ij}^t w_i^t$, $\forall j \in [k]$ and uses the fitted laws to dynamically adjust domain weights \mathbf{w}^t . We provide a comparison between Skill-It and Aioli in Table 4.

4.3 Others

ODM ([Albalak et al., 2023](#)) and ADO ([Jiang et al., 2024](#)) adopt alternative strategies for dynamically adjusting domain weights during the training of the target model, differing from the two categories discussed earlier. Due to space constraints, we describe these methods in detail in Appendix C.2.

4.4 Advantages and Disadvantages

Compared to offline methods, online methods do not suffer from scale-up issues. When updating \mathbf{w}^t , however, if the update rule involves computing model gradients, the computational overhead can be high. If it relies only on the model’s loss, the overhead is generally lower, though it can still be significant depending on the specific implementation. This overhead can be mitigated by updating the weights intermittently—every t_{update} steps, where $t_{\text{update}} > 1$ —instead of every step. Another limitation is that when out-of-domain loss is used to evaluate model performance, only online mixing law methods are applicable, while most other online methods (e.g., online min-max optimization methods) cannot be used in this setting.

5 Challenges and Future Directions

5.1 Difficulty in finding the optimal domain weights

Offline methods assume that there exist optimal static domain weights \mathbf{w}^* for a given corpus $\mathcal{D}_{\text{train}}$ composed of different domains that lead to the best-performing model. [Jiang et al. \(2024\)](#) argue that the optimal domain weights \mathbf{w}^* on $\mathcal{D}_{\text{train}}$ form a fixed, stationary curriculum throughout the training process, where models are progressively exposed

to each domain in a curated order, which is closely related to curriculum learning ([Bengio et al., 2009](#)) and empirically demonstrate the existence of effective training curricula. However, identifying the globally optimal curriculum is computationally challenging, making it highly unlikely to be found. This raises the question of how to identify a good curriculum in a computationally efficient manner. In response, [Jiang et al. \(2024\)](#) suggest adopting an online strategy for optimizing domain weights and training large language models, which could effectively identify good curricula while minimizing computational costs.

5.2 The natural domain

Most existing works optimize domain weights to the natural domains of the data, which could bring two issues: (1) The natural domains are typically at a coarse-grained level, which should be divided into fine-grained domains to allow for gains from applying various algorithms. For example, the GLaM dataset ([Du et al., 2022](#)) has 8 domains, while the Pile ([Gao et al., 2020](#)) is constructed from 22 diverse domains. The DoReMi algorithm ([Xie et al., 2023](#)) performs better on the Pile than on the GLaM dataset. (2) How can we find fine-grained domains in the training data? One direct way is through clustering ([Oren et al., 2019](#); [Gururangan et al., 2023](#); [Chen et al., 2023](#)). In [Gururangan et al. \(2023\)](#), the authors split a corpus into segments and use the embeddings of these segments as sample features to run k -means clustering algorithms for unsupervised domain discovery. [Chen et al. \(2023\)](#) explore several sample features for clustering the data, including the embedding of the last token, the average token, the last token of the training sample, and the loss trajectory of the training sample over multiple runs. They find that using the loss trajectory as the sample feature significantly outperforms the other metrics.

5.3 Loss metrics vs. Task performance

To determine the optimal domain weights, loss metrics, such as training loss or validation loss, are often used as surrogate objectives for the downstream task performance to evaluate the model during training. However, does the model trained with the optimal domain weights found by evaluating the loss really perform well on the downstream tasks? [Gadre et al. \(2024\)](#) explore this issue within the framework of scaling laws. They find that the perplexity of a language model can be related to

its downstream task performance via a power law, demonstrating that it is reasonable to use loss metrics as surrogate objectives for downstream task performance when determining the best training variables, such as model size or dataset size. However, the optimal domain weights found by most methods in data mixture are only slightly better than uniform sampling (Jiang et al., 2024), or even worse than uniform sampling (Chen et al., 2024). We think it is worth exploring the relationship between loss metrics and task performance under the data mixture setting.

6 Conclusion

We provide an overview of existing data mixture methods from an algorithm perspective in this paper. We first propose a fine-grained categorization of existing methods, extending beyond the previous offline/online classification. We divide offline methods into heuristic-based, algorithm-based, and function fitting-based methods. We categorize online methods into three groups—online min-max optimization, online mixing law, and other approaches—by drawing connections with the optimization frameworks underlying offline methods. We then summarize the problem formulations, representative algorithms for each subtype of offline and online methods, and clarify the relationships and distinctions among them. Finally, we discuss the advantages and disadvantages of each method and highlight key challenges in data mixture.

Limitation

We only provide insights into data mixture problems from theoretical perspectives and do not conduct any empirical studies. Most existing methods emphasize their own strengths, but their comparisons are not aligned in terms of experimental settings, such as model type, model scale, and training data scale. Even though we discuss the advantages and disadvantages of each method, we have not sufficiently compared them from an experimental perspective.

Ethics Statement

This paper is a theoretical and analytical study, primarily providing a systematic review and analysis of existing data mixture methods. All related resources referenced are publicly available. Consequently, this work does not raise any specific ethical concerns.

Acknowledgment

This research is supported by Longgang District Special Funds for Science and Technology Innovation under Grant LGKCSPT2023002. We would also like to thank the anonymous reviewers for their valuable suggestions.

References

- Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang, Niklas Muennighoff, Bairu Hou, Liangming Pan, Hae-won Jeong, Colin Raffel, Shiyu Chang, Tatsunori Hashimoto, and William Yang Wang. 2024. [A survey on data selection for language models](#). *Preprint*, arXiv:2402.16827.
- Alon Albalak, Liangming Pan, Colin Raffel, and William Yang Wang. 2023. [Efficient online data mixing for language model pre-training](#). *Preprint*, arXiv:2312.02406.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 2002. The nonstochastic multi-armed bandit problem. *SIAM journal on computing*, 32(1):48–77.
- Lior Belenki, Alekh Agarwal, Tianze Shi, and Kristina Toutanova. 2025. [Optimizing pre-training data mixtures with mixtures of data expert models](#). *Preprint*, arXiv:2502.15950.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Mayee F. Chen, Michael Y. Hu, Nicholas Lourie, Kyunghyun Cho, and Christopher Ré. 2024. [Aioli: A unified optimization framework for language model data mixing](#). *Preprint*, arXiv:2411.05735.
- Mayee F. Chen, Nicholas Roberts, Kush Bhatia, Jue Wang, Ce Zhang, Frederic Sala, and Christopher Ré. 2023. [Skill-it! a data-driven skills framework for understanding and training language models](#). *Preprint*, arXiv:2307.14430.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. 2018. [Training gans with optimism](#). *Preprint*, arXiv:1711.00141.

- Constantinos Daskalakis and Ioannis Panageas. 2018. [The limit points of \(optimistic\) gradient descent in min-max optimization](#). *Preprint*, arXiv:1807.03907.
- Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten Bosma, Zongwei Zhou, Tao Wang, Yu Emma Wang, Kellie Webster, Marie Pellat, Kevin Robinson, Kathleen Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V Le, Yonghui Wu, Zhifeng Chen, and Claire Cui. 2022. [Glam: Efficient scaling of language models with mixture-of-experts](#). *Preprint*, arXiv:2112.06905.
- Jean-Antoine Désidéri. 2012. [Multiple-gradient descent algorithm \(mgda\) for multiobjective optimization](#). *Comptes Rendus Mathématique*, 350(5):313–318.
- Simin Fan, Matteo Pagliardini, and Martin Jaggi. 2023. [DOGE: Domain reweighting with generalization estimation](#). In *Second Agent Learning in Open-Endedness Workshop*.
- Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. 2017. [Forward and reverse gradient-based hyperparameter optimization](#). *Preprint*, arXiv:1703.01785.
- Samir Yitzhak Gadre, Georgios Smyrnis, Vaishaal Shankar, Suchin Gururangan, Mitchell Wortsman, Rulin Shao, Jean Mercat, Alex Fang, Jeffrey Li, Sedrick Keh, Rui Xin, Marianna Nezhurina, Igor Vasiljevic, Jenia Jitsev, Luca Soldaini, Alexandros G. Dimakis, Gabriel Ilharco, Pang Wei Koh, Shuran Song, Thomas Kollar, Yair Carmon, Achal Dave, Reinhard Heckel, Niklas Muennighoff, and Ludwig Schmidt. 2024. [Language models scale reliably with over-training and on downstream tasks](#). *Preprint*, arXiv:2403.08540.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The pile: An 800gb dataset of diverse text for language modeling](#). *Preprint*, arXiv:2101.00027.
- Ce Ge, Zhijian Ma, Daoyuan Chen, Yaliang Li, and Bolin Ding. 2024. [Bimix: Bivariate data mixing law for language model pretraining](#). *Preprint*, arXiv:2405.14908.
- Suchin Gururangan, Margaret Li, Mike Lewis, Weijia Shi, Tim Althoff, Noah A. Smith, and Luke Zettlemoyer. 2023. [Scaling expert language models with unsupervised domain discovery](#). *Preprint*, arXiv:2303.14177.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#). *Preprint*, arXiv:2203.15556.
- Yiding Jiang, Allan Zhou, Zhili Feng, Sadhika Malladi, and J. Zico Kolter. 2024. [Adaptive data optimization: Dynamic sample selection with scaling laws](#). *Preprint*, arXiv:2410.11820.
- Feiyang Kang, Yifan Sun, Bingbing Wen, Si Chen, Dawn Song, Rafid Mahmood, and Ruoxi Jia. 2025. [Autoscale: Scale-aware data mixing for pre-training llms](#). *Preprint*, arXiv:2407.20177.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *Preprint*, arXiv:2001.08361.
- Jeongyeol Kwon, Dohyun Kwon, Stephen Wright, and Robert Nowak. 2023. [A fully first-order method for stochastic bilevel optimization](#). *Preprint*, arXiv:2301.10945.
- Tianyi Lin, Chi Jin, and Michael Jordan. 2020. [On gradient descent ascent for nonconvex-concave minimax problems](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6083–6093. PMLR.
- Bingchang Liu, Chaoyu Chen, Cong Liao, Zi Gong, Huan Wang, Zhichao Lei, Ming Liang, Dajun Chen, Min Shen, Hailian Zhou, Hang Yu, and Jianguo Li. 2023a. [Mftcoder: Boosting code llms with multitask fine-tuning](#). *Preprint*, arXiv:2311.02303.
- Bo Liu, Yihao Feng, Peter Stone, and Qiang Liu. 2023b. [Famo: Fast adaptive multitask optimization](#). *Preprint*, arXiv:2306.03792.
- Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. 2021a. [Conflict-averse gradient descent for multi-task learning](#). *Advances in Neural Information Processing Systems*, 34.
- Dong C Liu and Jorge Nocedal. 1989. [On the limited memory bfgs method for large scale optimization](#). *Mathematical programming*, 45(1):503–528.
- Liyang Liu, Yi Li, Zhanghui Kuang, Jing-Hao Xue, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang. 2021b. [Towards impartial multi-task learning](#). In *International Conference on Learning Representations*.
- Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. 2024. [Regmix: Data mixture as regression for language model pre-training](#). *arXiv preprint arXiv:2407.01492*.
- Jonathan Lorraine, Paul Vicol, and David Duvenaud. 2019. [Optimizing millions of hyperparameters by implicit differentiation](#). *Preprint*, arXiv:1911.02590.

- Zheheng Luo, Xin Zhang, Xiao Liu, Haoling Li, Yeyun Gong, Chen Qi, and Peng Cheng. 2024. [Velocitune: A velocity-based dynamic domain reweighting method for continual pre-training](#). *Preprint*, arXiv:2411.14318.
- Sachin Mehta, Mohammad Hossein Sekhavat, Qingqing Cao, Maxwell Horton, Yanzi Jin, Chenfan Sun, Iman Mirzadeh, Mahyar Najibi, Dmitry Belenko, Peter Zatloukal, and Mohammad Rastegari. 2024. [Openlm: An efficient language model family with open training and inference framework](#). *Preprint*, arXiv:2404.14619.
- Yonatan Oren, Shiori Sagawa, Tatsunori B. Hashimoto, and Percy Liang. 2019. [Distributionally robust language modeling](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4227–4237, Hong Kong, China. Association for Computational Linguistics.
- Rui Pan, Xiang Liu, Shizhe Diao, Renjie Pi, Jipeng Zhang, Chi Han, and Tong Zhang. 2024a. [Lisa: Layerwise importance sampling for memory-efficient large language model fine-tuning](#). *Preprint*, arXiv:2403.17919.
- Rui Pan, Jipeng Zhang, Xingyuan Pan, Renjie Pi, Xiaoyu Wang, and Tong Zhang. 2024b. [Scalebio: Scalable bilevel optimization for llm data reweighting](#). *arXiv preprint arXiv:2406.19976*.
- Haoran Que, Jiaheng Liu, Ge Zhang, Chenchen Zhang, Xingwei Qu, Yinghao Ma, Feiyu Duan, Zhiqi Bai, Jikai Wang, Yuanxing Zhang, Xu Tan, Jie Fu, Wenbo Su, Jiamang Wang, Lin Qu, and Bo Zheng. 2024. [D-cpt law: Domain-specific continual pre-training scaling law for large language models](#). *Preprint*, arXiv:2406.01375.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Meisam Razaviyayn, Tianjian Huang, Songtao Lu, Maher Nouiehed, Maziar Sanjabi, and Mingyi Hong. 2020. [Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances](#). *IEEE Signal Processing Magazine*, 37(5):55–66.
- Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. 2020. [Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization](#). *Preprint*, arXiv:1911.08731.
- Zhiqiang Shen, Tianhua Tao, Liqun Ma, Willie Neiswanger, Zhengzhong Liu, Hongyi Wang, Bowen Tan, Joel Hestness, Natalia Vassilieva, Daria Soboleva, and Eric Xing. 2024. [Sлимпajama-dc: Understanding data combinations for llm training](#). *Preprint*, arXiv:2309.10818.
- Xingyou Song, Qiuyi Zhang, Chansoo Lee, Emily Fertig, Tzu-Kuo Huang, Lior Belenki, Greg Kochanski, Setareh Ariaifar, Srinivas Vasudevan, Sagi Perel, and Daniel Golovin. 2024. [The vizier gaussian process bandit algorithm](#). *Preprint*, arXiv:2408.11527.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert-Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silterra, James T Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-Baeza. 2020. [Scipy 1.0: fundamental algorithms for scientific computing in python](#). *Nature Methods*, 17(3):261–272.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#). *Preprint*, arXiv:2109.01652.
- Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. 2024. [Sheared llama: Accelerating language model pre-training via structured pruning](#). *Preprint*, arXiv:2310.06694.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le, Tengyu Ma, and Adams Wei Yu. 2023. [Doremi: Optimizing data mixtures speeds up language model pretraining](#). *arXiv preprint arXiv:2305.10429*.

- Yongxin Yang and Timothy M. Hospedales. 2015. [A unified perspective on multi-domain and multi-task learning](#). *Preprint*, arXiv:1412.7489.
- Jiasheng Ye, Peiju Liu, Tianxiang Sun, Yunhua Zhou, Jun Zhan, and Xipeng Qiu. 2024. [Data mixing laws: Optimizing data mixtures by predicting language modeling performance](#). *arXiv preprint arXiv:2403.16952*.
- Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. [Gradient surgery for multi-task learning](#). *arXiv preprint arXiv:2001.06782*.
- Jiawei Zhang, Peijun Xiao, Ruoyu Sun, and Zhi-Quan Luo. 2022. [A single-loop smoothed gradient descent-ascent algorithm for nonconvex-concave min-max problems](#). *Preprint*, arXiv:2010.15768.
- Yihua Zhang, Prashant Khanduri, Ioannis Tsaknakis, Yuguang Yao, Mingyi Hong, and Sijia Liu. 2023. [An introduction to bi-level optimization: Foundations and applications in signal processing and machine learning](#). *Preprint*, arXiv:2308.00788.

A Static Weights vs. Dynamic Weights

Static weights. Static weights are primarily used to combine the training data $\mathcal{D}_{\text{train}}$ by assigning fixed sampling probabilities to each domain during model training. These weights remain unchanged throughout the entire training process. In offline methods, the optimal domain weights \mathbf{w}^* used to train the target model are considered static weights.

Dynamic weights. Dynamic weights are adaptively updated during model training, in contrast to static weights that remain fixed. However, directly adjusting the data mixing ratio for all domains within each mini-batch during training is not straightforward. A more common approach is to apply uniform sampling to sample all domains in $\mathcal{D}_{\text{train}}$ and alter the data contribution through dynamic weighting of the loss function, which can be expressed as: $\sum_{i=1}^k w_i^t \ell_{\text{train},i}(\boldsymbol{\theta}^t)$. Here $\ell_{\text{train},i}(\boldsymbol{\theta}^t)$ represents the average loss of $\boldsymbol{\theta}^t$ on the i -th domain of the mini-batch data at each training step t . The domain weights \mathbf{w}^t used in online methods are typically dynamic weights, which will be updated at each training iteration or a fixed interval t_{update} .

B Offline Methods

B.1 Heuristic Methods

Heuristic-based methods, such as uniform sampling and proportional sampling, determine domain weights based on simple heuristics.

Uniform sampling defines the optimal static domain weights $\mathbf{w}^* = [\frac{1}{k}, \dots, \frac{1}{k}]$ when mixing k training domains $\mathcal{D}_{\text{train}} = \{D_{\text{train},1}, \dots, D_{\text{train},k}\}$. Proportional sampling is commonly used in large-scale training, where domain weights are set in proportion to token counts. Specifically, the best domain weights \mathbf{w}^* are defined as $w_i^* = \frac{|D_{\text{train},i}|}{\sum_{j=1}^k |D_{\text{train},j}|}$, $i \in [k]$, where $|D_{\text{train},i}|$ denotes the number of tokens in each training domain data $D_{\text{train},i}$.

The domain weights in these methods remain fixed rather than being optimized and thus they do not rely on proxy models. Considering that they can be implemented by omitting line 3 from Framework 1, and for the sake of completeness in our framework, we treat them as degenerate cases of offline methods.

B.2 Algorithm-based Methods

B.2.1 Bi-level Optimization

$$\begin{aligned} \mathbf{w}^* \in \arg \min_{\mathbf{w} \in \Delta^k} \ell_{\text{val}}(\phi^*(\mathbf{w})) \\ \text{s.t. } \phi^*(\mathbf{w}) \in \arg \min_{\phi} \sum_{i=1}^k w_i \ell_{\text{train},i}(\phi). \end{aligned} \quad (4)$$

There are three primary optimization frameworks (Zhang et al., 2023) used to solve the bi-level optimization problems: the Implicit Function (IF)-based approach, the Gradient Unrolling (GU)-based approach, and the Value Function (VF)-based approach. The primary obstacle hindering the scalability of bi-level optimization stems from the interdependence between the outer-level and inner-level problems. This mutual dependency introduces significant computational challenges, such as the need to compute the Hessian and Jacobian, especially when handling large-scale problem instances. The first two classes both leverage (some approximated version of) the Implicit Gradient (IG). The key difference is how the IG approximation is conducted: One directly assumes that a given procedure can provide a high-quality solution to the inner-level problem, while the other approximates the inner-level solution by unrolling a given algorithm for a fixed number of steps. The Value Function (VF)-based approach reformulates the bi-level optimization as a single-level regularized optimization problem, which offers flexibility in handling inner-level constraints.

B.3 Function Fitting-based Methods

B.3.1 The D-CPT law

Que et al. (2024) study mixing laws for Domain-Specific Continual Pre-Training (D-CPT) by following the Chinchilla scaling law (Hoffmann et al., 2022). In the simplest scenario, where the training data consists of a general corpus $\mathcal{D}_{\text{train},0}$ and a domain-specific corpus $\mathcal{D}_{\text{train},1}$, they model the mixing law $f_j, j \in \{0, 1\}$ as:

$$\ell_{\text{val},j}(w_j, D, N) = E^j + \frac{A^j}{N^{\alpha^j}} + \frac{B^j \cdot w_j^\eta}{D^{\beta^j}} + \frac{C^j}{w_j^\gamma}, \quad (5)$$

where D is dataset size, N is model size, $w_j' = w_j + \epsilon$, and ϵ is a small constant used to ensure the stability of $\ell_{\text{val},j}(N, D, w_j)$ when w_j approaches zero. The set of parameters $\mathbf{E}, \mathbf{A}, \mathbf{B}, \mathbf{C}, \alpha, \beta, \eta \in \mathbb{R}^2$ is learned by fitting the model using the L-BFGS optimization algorithm (Liu and Nocedal, 1989).

They propose three designs for the aggregation function g to obtain the optimal static domain weights \mathbf{w}^* . Among them, the most widely adopted strategy typically considers a trade-off between the general and domain-specific abilities of the model after Continued Pre-Training. In this approach, the optimal proportion for the domain corpus is determined as $w_1^* = \arg \min_{w_1} \ell_{\text{val},1}(N = N_0, D = D_0, w_1)$

s.t. $\frac{\ell_{\text{val},0} - \ell_{\text{val},0}^0}{\ell_{\text{val},0}^0} < T$, where T denotes a threshold that constrains the degradation in the model's general ability. Therefore, the optimal domain weights is $\mathbf{w}^* = (1 - w_1^*, w_1^*)$.

C Online Methods

C.1 Online Min-Max Optimization Methods

C.1.1 Gradient Manipulation Methods

A class of methods in multi-task learning, known as gradient manipulation methods, solves Equation 6 by finding a new update d_t at each step in an online fashion to prevent task optimization unbalance:

$$\min_{\theta} \max_{\mathbf{w}^t \in \Delta^k} \sum_{i=1}^k w_i^t f_i^t(\theta), \quad (6)$$

where $f_i^t(\theta)$ is related with $\ell_{\text{train},i}^t(\theta)$. In a gradient descent style iterative update: $\theta^{t+1} = \theta^t - \eta d_t$, η is the learning rate, $d_t = \sum_i w_i^t \nabla \ell_{\text{train},i}^t(\theta)$. There are several approaches (Désidéri, 2012; Liu et al., 2021a,b) targeting to improve the "worst-case improvement" among all tasks, where $f_i^t(\theta) = \ell_{\text{train},i}^t(\theta) - \ell_{\text{train},i+1}^t(\theta)$. Yu et al. (2020) propose a form of gradient surgery to avoid detrimental gradient interference, also helping improve the "worst improvement" across all tasks. Liu et al. (2023b) propose a method called FAMO to optimize the "largest worst-case improvement rate" across all tasks to achieve better task balance, $f_i^t(\theta) = (\ell_{\text{train},i}^t(\theta) - \ell_{\text{train},i+1}^t(\theta)) / \ell_{\text{train},i}^t(\theta)$.

C.1.2 Comparison of Online Min-Max Optimization Methods

DoReMi_{online} jointly optimizes both the target model θ and the domain weights \mathbf{w} during training. Following the framework of Algorithm 2, the target model θ is updated as

$$\theta^t = \theta^{t-1} - \eta \sum_{i=1}^k w_i^{t-1} \nabla \ell_{\text{train},i}(\theta^{t-1}),$$

while the dynamic domain weights w^t are updated according to

$$w^t = \frac{\alpha^t}{\sum_{i=1}^k \alpha_i^t}, \text{ where } \alpha^t = \alpha^{t-1} \exp(\zeta \max\{\ell_{\text{train},i}(\theta^t) - \ell_{\text{train},i}(\theta_{\text{ref}}), 0\}),$$

where ζ is a step size, and θ_{ref} is a pre-trained reference model, typically of the same size as the target model θ . The reference model’s training loss $\ell_{\text{train},i}(\theta_{\text{ref}})$ serves as a baseline for evaluating the target model’s current performance. More specifically, the excess loss $\ell_{\text{train},i}(\theta^t) - \ell_{\text{train},i}(\theta_{\text{ref}})$ measures how much the target model has yet to learn from domain i . For ease of comparison with Velocitune (Luo et al., 2024), we also denote this quantity as: $V_i^t = \ell_{\text{train},i}(\theta^t) - \ell_{\text{train},i}(\theta_{\text{ref}})$.

Training a reference model θ_{ref} to obtain baseline losses for evaluating the learning progress of the target model on each domain can be prohibitively expensive. To reduce this cost, ShearedLLaMa (Xia et al., 2024) improves upon DoReMi_{online} by using losses predicted by a fitted scaling law for each domain as the reference losses. Additionally, it replaces the training loss $\ell_{\text{train},i}(\theta^t)$ with the validation loss $\ell_{\text{val},i}(\theta^t)$ when dynamically updating the domain weights w^t to optimize the generalization ability of the target model (i.e., $V_i^t = \ell_{\text{val},i}(\theta^t) - \ell_{\text{ref},i}$).

Velocitune (Luo et al., 2024) introduces a more precise metric to measure how much the target model has yet to learn from each domain, defined as $V_i^t = \frac{\ell_{\text{val},i}(\theta^t) - \ell_{\text{ref},i}}{\ell_{\text{init},i} - \ell_{\text{ref},i}}$, $i \in [k]$, where $\ell_{\text{init},i}$ denotes the initial validation loss for domain i , and $\ell_{\text{ref},i}$ denotes the reference loss. Similar to ShearedLLaMa, $\ell_{\text{ref},i}$ is estimated using a fitted scaling law on each domain i , providing a cost-efficient approximation of the desired learning target for each domain.

C.2 Others

C.2.1 Online Data Mixing (ODM)

Albalak et al. (2023) propose an efficient method, Online Data Mixing (ODM), which leverages multi-armed bandit (MAB) algorithms by treating each data domain as an arm. In ODM, the principle of updating w^t at each training step t is to increase the mixing ratio for domains that provide the most valuable information for model training. Drawing from information theory that perplexity can be interpreted as the expected information gain from learning the next token, they use the training loss per domain as a reward signal. The ODM method

is ultimately modified based on the Exponential-weight algorithm for Exploration and Exploitation (Exp3) (Auer et al., 2002), with the key difference being that ODM employs a moving average estimated reward instead of a cumulative estimated reward, paying more attention to the recent rewards.

C.2.2 Adaptive Data Optimization (ADO)

Jiang et al. (2024) propose Adaptive Data Optimization (ADO), where they adjust the dynamic weights w^t at a fixed interval t_{update} by fitting a domain scaling law for each domain. The domain scaling law predicts the i -th domain’s training loss after training on n samples: $\widehat{\ell}_{\text{train},i}(n) = \widehat{\ell}(n; \alpha_i, \beta_i, \varepsilon_i) = \varepsilon_i + \beta_i n^{-\alpha_i}$, ε_i corresponds to the estimated irreducible loss of domain i . The derivative of the loss to the number of samples can be interpreted as follows:

$$\begin{aligned} \frac{d\widehat{\ell}_{\text{train},i}(n)}{dn} &= \frac{-\alpha_i \beta_i n^{-\alpha_i}}{n} \\ &= -\frac{1}{n} \underbrace{\alpha_i}_{\text{Learning speed}} \underbrace{(\widehat{\ell}_{\text{train},i}(n) - \varepsilon_i)}_{\text{Reducible loss}} \end{aligned}$$

To determine w^t during model training, they define a preference distribution

$$\rho_i^t \propto -\mu_i \frac{\partial}{\partial n} \widehat{\ell}_{\text{train},i}(n) \lambda_i^t, \quad (7)$$

here μ_i is the prior domain weights, and $\frac{\partial}{\partial n} \widehat{\ell}_{\text{train},i}(n)$ is the learning speed forecast by a scaling law. λ_i^t is a credit assignment score for domain i , which is a real positive-valued function that indicates how much data from the i -th domain contributed to recent changes in the loss \mathcal{L}_i . The intuition behind ADO is to prioritize sampling from a domain i if the model is decreasing its loss \mathcal{L}_i quickly, but only if that decrease can be attributed to data from domain i . The dynamic weights w^t is a linear combination of the moving average of preference policy ρ^t at every step.

D Differences between the previous survey paper and our paper

Albalak et al. (2024) is a comprehensive survey on data selection, with data mixture being a small part of it. They categorize existing data mixture methods into two main types: offline methods and online methods. For each category, they discuss a limited number of methods, mainly because there were not many related works available at that time.

For our survey paper, we propose a fine-grained categorization of existing data mixture methods, extending beyond their offline and online classification to include and contextualize a broader range of recent and subsequent developments. Specifically, we divide offline methods into three subcategories: heuristic-based, algorithm-based, and function fitting-based methods. For algorithm-based methods, we systematically analyze their connections to min-max optimization and bi-level optimization. For function fitting-based methods, we identify three key components and present a formal problem formulation. We further compare min-max optimization methods, bi-level optimization methods, and function fitting-based approaches, and highlight their interrelationships. In particular, we point out that the key distinction between min-max and bi-level optimization lies in how the optimal domain weights are defined, while function fitting-based methods convert the bi-level optimization problem into a single-level problem via function fitting.

We further draw connections between online methods and the optimization frameworks underlying offline methods by interpreting most online approaches as online variants of offline techniques. Based on this perspective, we categorize online methods into three groups: online min-max optimization, online mixing law, and other methods. We also provide a systematic comparison and summary of online min-max optimization methods and online mixing law methods, highlighting their differences in objective design, update strategies.

Finally, we survey these methods primarily from an algorithmic perspective, aiming to provide deeper understanding and valuable insights into optimization strategies for data mixture, as we believe they are applicable to any stage of model training.