

# Multi-Attribute Multi-Grained Adaptation of Pre-Trained Language Models for Text Understanding from Bayesian Perspective

You Zhang<sup>1</sup>, Jin Wang<sup>1\*</sup>, Liang-Chih Yu<sup>2\*</sup>, Dan Xu<sup>1</sup>, Xuejie Zhang<sup>1</sup>

<sup>1</sup>School of Information Science and Engineering, Yunnan University, Yunnan, P.R.China

<sup>2</sup>Department of Information Management, Yuan Ze University, Taiwan  
{yzhang0202, wangjin}@ynu.edu.cn, lcyu@saturn.yzu.edu.tw

## Abstract

Current neural networks often employ multi-domain-learning or attribute-injecting mechanisms to incorporate non-independent and identically distributed (non-IID) information for text understanding tasks by capturing individual characteristics and the relationships among samples. However, the extent of the impact of non-IID information and how these methods affect pre-trained language models (PLMs) remains unclear. This study revisits the assumption that non-IID information enhances PLMs to achieve performance improvements from a Bayesian perspective, which unearths and integrates non-IID and IID features. Furthermore, we proposed a multi-attribute multi-grained framework for PLM adaptations (M2A), which combines multi-attribute and multi-grained views to mitigate uncertainty in a lightweight manner. We evaluate M2A through prevalent text-understanding datasets and demonstrate its superior performance, mainly when data are implicitly non-IID, and PLMs scale larger.

**Code** — <https://github.com/yoyo-yun/M2A>.

## Introduction

Modern neural networks, especially pre-trained language models (PLMs), for specific tasks require large amounts of data and computational resources (Kong, Wang, and Zhang 2021; Han et al. 2021); therefore, collecting as much data as possible from diverse sources to achieve satisfactory task performance is desirable (Yuan, Zhao, and Qin 2022; Katsarou, Jeney, and Stefanidis 2023; Liu, Zhang, and Liu 2018; Yang and Hospedales 2015). For instance, sentiment analysis data can be gathered from different users, categories, and platforms, where corresponding golden labels are annotated by different users and intended to build sufficient samples leveraged for robust sentiment model training (Chen and Chao 2021; Zhang et al. 2023). These networks unanimously assume that all samples in a data pool (for a specific downstream task) are independent and identically distributed (IID) in a coarse-grained view (Pang and Lee 2006; Zhang, Wang, and Liu 2018). However, data distributions from different resources are not typically IID, which needs

fine-grained view to deal with the data heterogeneity problem. Without the fine-grained view, the coarse-grained IID assumption may impede the optimal model convergences (uncertainty), degrading ideal performance when data are non-IID, either explicitly or implicitly, even though carrying out sufficient data (Zhang et al. 2023; Yao et al. 2024).

One intuitive solution to address the issue is to abstract the heterogeneities in data-sufficient while resource-diverse tasks and guide models adapted to diverse individual scenarios. This brings two significant challenges in recent adaptations of PLMs: 1) **representing the heterogeneities among data samples** and 2) **making PLMs capable of recognizing them effectively and efficiently in the training and inference phase**. Recently, many efforts have been committed to introducing multi-domain-learning (Yuan, Zhao, and Qin 2022; Katsarou, Jeney, and Stefanidis 2023), and attribute-injecting mechanisms (Zhang et al. 2021; Amplayo 2019) for investigating data heterogeneities. Multi-domain learning aims to learn domain-shared features and combine them with separate domain-specific features for multi-domain text classifications. Attribute injecting treats concrete domain information as attribute knowledge and injects it into coarse neural networks for fine-grained performance. However, these high-performance methods do not systematically develop the model design and optimization approach, constraining the ability to represent robust non-IID and IID features as well as their internal relatedness. Moreover, traditional full-model fine-tuning (FFT) and sophisticated structure modification render limited scalability for larger PLMs (Min et al. 2024).

We argue that most text-understanding data contain multiple attributes, i.e., user, category, and platform. From a data distribution perspective, each attribute covers coarse-grained and fine-grained views, as shown in Figure 1(a). The coarse-grained view with one domain, which treats all samples as IID, ignores data heterogeneities. This approach is advantageous for accommodating large number of samples, as shown in Figure 1(b). The fine-grained view, e.g., the category, considers all samples divided into numerous fine-grained domains. Within each domain, samples are IID, whereas samples across different domains are non-IID. This perspective allows models to manage diverse samples by employing individual modules for each domain, as shown in Figure 1(c).

\*Corresponding authors.

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

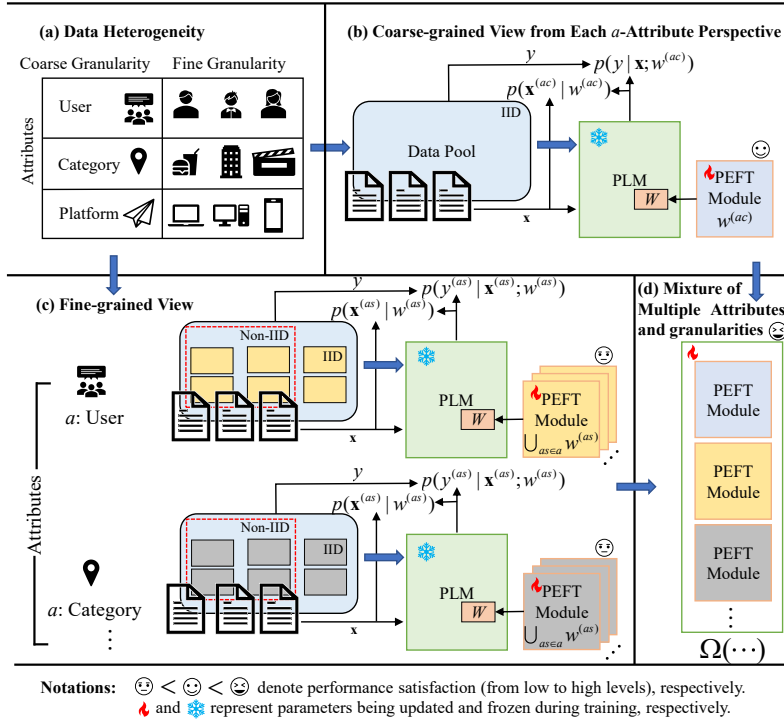


Figure 1: A conception of the proposed method.

To integrate IID and non-IID information for complementary benefits, we rethink the current multi-domain learning and attribute injecting mechanism and propose a multi-attribute, multi-grained adaptation (M2A) framework for mitigating the uncertainty of PLMs’ adaptation, as shown in Figure 1(d). Theoretically, we adopt a Bayesian inference to analyze the relatedness between coarse-grained and fine-grained views and a Bayesian neural network (BNN) to represent robust data heterogeneities (Magris and Iosifidis 2023; Jospin et al. 2022) (see §Preliminaries). Through a combination of multiple attributes and granularities from the Bayesian inference perspective, we found that the proposed M2A framework could facilitate PLMs to be effectively adapted to non-IID tasks. Moreover, a parameter-efficient fine-tuning (PEFT) module (Houlsby et al. 2019) and a joint learning strategy are proposed to facilitate the scalability of M2A for diverse PLMs (Zhang et al. 2023). In experiments, we utilized various PLMs to evaluate M2A on multi-domain and personalized sentiment analysis tasks, which contain sufficient data for task adaptation while involving data heterogeneities.

Our key contributions in this paper are threefold.

- We rethink the current multi-domain-learning and attribute-injecting mechanism from a Bayesian perspective and provide a robust M2A framework for text understanding.
- Based on PEFT and joint learning methods, M2A makes PLMs effectively and efficiently adapted for diverse scales and broader scenarios.
- Extensive experiments were conducted on several text-

understanding tasks and provided empirical analysis.

## Preliminaries

### Background: Text Understanding

Regarding a text understanding task with a training dataset  $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_n^N$  with  $N$  IID and label-annotated samples, a neural network  $\text{NN}_w(\cdot)$  parameterized with weights  $w$  is optimized to estimate the likelihood  $p(\mathcal{D}_y | \mathcal{D}_x; w)$  that links the inputs and the outputs via:

$$p(\mathcal{D}_y | \mathcal{D}_x; w) = \arg \max_w \prod_{n=1}^N \text{Loss}(y_n, \hat{y}_n) \quad (1)$$

where  $\mathcal{D}_x = \{\mathbf{x}_n\}_{n=1}^N$  and  $\mathcal{D}_y = \{y_n\}_{n=1}^N$  represents the input data and the output labels;  $\hat{y}_n = \text{NN}_w(\mathbf{x}_n)$  denotes the network output for input  $\mathbf{x}_n$ ;  $\text{Loss}(\cdot)$  the loss function for model optimizations.

Regarding non-IID heterogeneities in a specific attribute  $a \in \mathcal{A}$ , the dataset  $\mathcal{D} = \cup_{as \in \mathcal{A}} \mathcal{D}^{(as)}$  can be divided into  $|\mathcal{A}|_f$  fine-grained domains (sub-datasets) where each domain contains IID samples, denoted as  $\mathcal{D}^{(as)} = \{(\mathbf{x}_n^{(as)}, y_n^{(as)})\}_n^{|\mathcal{D}^{(as)}|}$ . For each domain  $as$ , it typically requires a neural network  $\text{NN}_{w^{(as)}}(\cdot)$  to estimate the domain-specific sub-datasets. Notably, only one IID domain  $ac$  is covered in coarse-grained view, where  $\mathcal{D} = \mathcal{D}^{(ac)}$ , and the corresponding model is  $\text{NN}_{w^{(ac)}}(\cdot)$ .

### A Bayesian Perspective

**Bayesian Inference.** We propose a Bayesian perspective to investigate diversity and relatedness among multiple at-

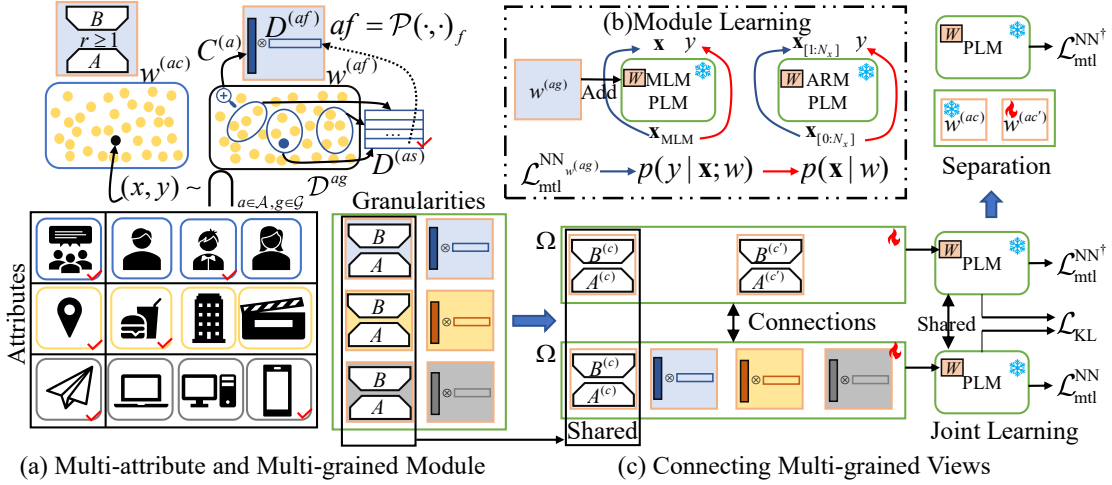


Figure 2: Overview of the M2A Framework.

tributes and granularities, respectively. In Bayesian learning (de Freitas, Niranjan, and Gee 2000),  $p(w|\mathcal{D})$  is the posterior distribution of the model being learned. Therefore, the predictive probabilities can be defined as:

$$p(y|\mathbf{x}; \mathcal{D}) = \int p(y|\mathbf{x}; w)p(w|\mathcal{D})d_w \quad (2)$$

which can be approximated by Monte Carlo method (Kroese et al. 2014) and aggregated from  $|S|$  models:

$$p(y|\mathbf{x}; \mathcal{D}) \approx \frac{1}{|S|} \sum_{s \in S} p(y|\mathbf{x}; w^{(s)}) \quad (3)$$

where  $w^{(s)} \sim p(w^{(s)}|\mathcal{D})$ .

**Bayesian Neural Network.** It targets the estimation of posterior distribution  $p(w|\mathcal{D})$  based on a Bayesian theorem via (Jospin et al. 2022):

$$p(w|\mathcal{D}_x, \mathcal{D}_y) = \frac{p(\mathcal{D}_y|\mathcal{D}_x; w) \cdot p(\mathcal{D}_x|w) \cdot p(w)}{p(\mathcal{D}_x, \mathcal{D}_y)} \quad (4)$$

$$\propto p(\mathcal{D}_y|\mathcal{D}_x; w) \cdot p(\mathcal{D}_x|w)$$

where  $p(\mathcal{D}_x|w)$  is the likelihood of the model that predicts input data;  $p(w)$  is the prior distribution of the model;  $p(\mathcal{D}) = p(\mathcal{D}_x, \mathcal{D}_y)$  is the evidence.

To sum up, the question of our proposed method is: how to (1) integrate multiple attributes and granularities for robust text representations, (2) estimate  $p(w|\mathcal{D})$  in each domain for capturing the non-IID and IID features, and (3) connect non-IID and IID information for modeling the internal relatedness.

## M2A Framework

Figure 2 illustrates the conceptual diagram of the proposed M2A method, which consists of three steps from a Bayesian perspective. The first step views each data from different attribute and granularity perspectives, utilizing M2A modules to integrate non-IID and IID features; the second step optimizes M2A modules through multitask learning; and the

third step connects multi-grained views and introduces a joint learning strategy. Additionally, our framework can be decomposed to improve parameter efficiency.

## Multi-attribute and Multi-grained Module in Bayesian Inference

Based on Bayesian inference in Eq. (3), we introduce a multi-attribute and multi-grained module to represent and integrate non-IID and IID information for robust hidden representation.

Each data sample  $(\mathbf{x}, y)$  can be viewed from different attributes  $a \in \mathcal{A}$  and multi-grained views  $g \in \mathcal{G} = \{c, f\}$  and denoted by  $(\mathbf{x}, y) \sim \bigcap_{a \in \mathcal{A}, g \in \mathcal{G}} \mathcal{D}^{(ag)}$ , where  $af$  represents  $(\mathbf{x}, y)$  in the fine-grained domain  $f$  from  $a$ -attribute perspective<sup>1</sup> via the formulation of  $af = \mathcal{P}((\mathbf{x}, y), a)_f$ . In contrast to Monte Carlo sampling, we sample  $w^{(s)}$  from diverse data distributions, i.e.,  $\mathcal{D}^{(ag)}$  with  $a \in \mathcal{A}$  and  $g \in \mathcal{G}$ ; therefore, the predictive probabilistic of  $(\mathbf{x}, y)$  can be formulated as:

$$p(y|\mathbf{x}; \mathcal{D}) \approx \frac{1}{|\mathcal{A}| |\mathcal{G}|} \sum_{a \in \mathcal{A}, g \in \mathcal{G}} p(y|\mathbf{x}; w^{(ag)}) \quad (5)$$

where  $w^{(ag)} \sim p(w^{(ag)}|\mathcal{D}^{(ag)})p(\mathcal{D}^{(ag)}|\mathcal{D})$ . Since these  $w^{(ag)}$ s have complementary and individual semantic information,  $w^{(ag)}$ s can be fused to a unified model in current models (Zhang et al. 2024):

$$p(y|\mathbf{x}; \mathcal{D}) \approx p(y|\mathbf{x}; \Omega(\bigcup_{ag} w^{(ag)})) \quad (6)$$

where  $\Omega(\cdot)$  represents module integrating operations.

Inspired by PEFT methods that introduce a lightweight module to enable PLMs to be adapted to downstream tasks, we utilize the LoRA family to implement our M2A framework for multi-objective adaptations efficiently, facilitating further research when PLMs scale large (Hu et al. 2021).

<sup>1</sup>For simplification, we argue each data sample  $(\mathbf{x}, y)$  belongs to only one fine-grained domain  $f \in a$  in each attribute view  $a$ .

Each module  $w^{(ag)}$  is formulated as  $A^{(ag)}B^{(ag)}$  with low-rank instinct features of  $r^{(ag)} \ll \min(d_{in}, d_{out})$ , which are then added into PLM matrix parameters  $W \in \mathbb{R}^{d_{in} \times d_{out}}$  for model calibration:

$$W = W + w^{(ag)} \quad (7)$$

Due to the cumulative and plug-and-play characteristics of LoRA (Wu, Huang, and Wei 2023), we instantiate modules integrating operations as  $\Omega(\bigcup_{a \in \mathcal{A}, g \in \mathcal{G}} w^{(ag)}) = \frac{1}{|\mathcal{A}| |\mathcal{G}|} \sum_{ag} w^{(ag)}$ . After that, as shown in Figure 2, our M2A module is simplified and denoted as:

$$W = W + \frac{1}{|\mathcal{A}| |\mathcal{G}|} \sum_{a \in \mathcal{A}, g \in \mathcal{G}} w^{(ag)} \quad (8)$$

## Module Learning in BNN

The essential factor in optimizing the module  $w$  locates how to maximize  $p(w|\mathcal{D}_x, \mathcal{D}_y)$  based on PLM backbones.

Although PLMs have provided general knowledge for text generation or cloze tasks, they could not effectively distinguish the data heterogeneities. For example, due to different pragmatic writing styles, PLMs make it hard to determine word prediction for different writers with the same given texts. Based on Eq. (4) with the Bayesian diagram, we introduce a multitask learning to estimate  $p(w|\mathcal{D})$  via maximizing  $p(\mathcal{D}_y|\mathcal{D}_x; w)$  and  $p(\mathcal{D}_x|w)$ , simultaneously:

$$\mathcal{L}_{\text{mtl}}^{\text{NN}w} = \text{Loss}(y_n, \hat{y}_n) + \alpha \text{Loss}(\mathbf{x}_n, \hat{\mathbf{x}}_n) \quad (9)$$

where  $\alpha$  is the balance factor;  $\hat{\mathbf{x}}_n$  denote reconstructed text tokens under  $p(\mathcal{D}_x|w)$  estimation and  $\text{NN}_w(\cdot)$ . Empirically, from the writer-specific perspective (i.e.,  $a$  denotes writer attribute),  $p(\mathcal{D}_y|\mathcal{D}_x; w^{(as)})$  aims to facilitate PLMs to learn *how a specific writer (i.e.,  $s$ ) would influence the annotations of the texts*. Meanwhile,  $p(\mathcal{D}_x|w^{(as)})$  is introduced to optimize PLMs to learn *What pragmatic contents a specific writer (i.e.,  $s$ ) would generate*.

To implement  $p(\mathcal{D}_x|w)$  based on masked language modeling (MLM) (Devlin et al. 2019) and autoregressive modeling (ARM) (Radford et al. 2019), respectively, we utilized different strategies. For MLM, we randomly mask 15% of input tokens as [MASK] and get unified models to predict original tokens of texts and annotated labels based on [CLS] or <s> tokens. For ARM-based PLMs, ARM tasks are used to predict both shift-right text labels and downstream task labels.

## Connecting Multi-grained Views

Acknowledging that identifying exact non-IID information locations could facilitate neural network models to mitigate predictive uncertainty, it is difficult to collect the locations due to extensive labor in retrieving sensitive information and bypassing privacy concerns (Zhang et al. 2023). To further investigate internal relatedness among coarse-grained and fine-grained views, Bayesian inference in Eq. (3) guides the M2A framework with a joint learning strategy for improving the model performance and broadening practical scenarios. A more detailed analysis can be found in the technical appendices. To generalize fine-grained domains, we can

get a multi-attribute coarse-grained M2A module (dubbed as M2A $\dagger$ ), denoted by  $W + \Omega(\bigcup_{a \in \mathcal{A}, g \in \{c, c'\}} w^{(ag)})$ , where  $w^{(ac')}$  is a special coarse-view module used to align diverse fine-grained modules.

To align with M2A and M2A $\dagger$ , we resort to a joint learning strategy with knowledge distillation for final optimizations (Aguilar et al. 2020), formally:

$$\mathcal{L} = \prod_{(\mathbf{x}, y) \in \mathcal{D}} \mathcal{L}_{\text{mtl}}^{\text{NN}} + \mathcal{L}_{\text{mtl}}^{\text{NN}\dagger} + \text{KL}(\text{NN}(\mathbf{x}), \text{NN}\dagger(\mathbf{x})) \quad (10)$$

where  $\text{NN} = \text{NN}_{W + \Omega(\bigcup_{a \in \mathcal{A}, g \in \{c, f\}} w^{(ag)})}$  and  $\text{NN}\dagger = \text{NN}_{W + \Omega(\bigcup_{a \in \mathcal{A}, g \in \{c, c'\}} w^{(ag)})}$ . During the training phase, only parameters of  $\bigcup_{a \in \mathcal{A}, g \in \{c, c', f\}} w^{(ag)}$  are simultaneously updated. Moreover, we have found that NN converged faster than NN $\dagger$  (see §Experiments); therefore, we introduce a module separation strategy. That means when the learning procedure of NN is optimally stopped during the training phase, NN $\dagger$  will continue learning with only lightweight  $\bigcup_{a \in \mathcal{A}, g \in \{c'\}} w^{(ag)}$  to being updated until its optima convergence.

## Model Decomposition

Due to uncertainties in model capacity, the number of domains employed, and module redundancy, we further improve the M2A framework in parameter efficiency via model decomposition (see Figure 2).

**Fine-grained domain parameters reduction.** Regarding model capacity, the fine-grained module ( $w^{(af)}$ ) covers much fewer data samples than the coarse-grained module ( $w^{(ac)}$  or  $w^{(ac')}$ ); thus, we introduce a KronA module (Edalati et al. 2022), which only contains one-rank weight parameters and provides feasible decompositions, to replace the original LoRA for fine-grained modules, i.e.,  $w^{(af)} = C^{(af)} \otimes D^{(af)}$  where  $\otimes$  is Kronecker product,  $C^{(af)} \in \mathbb{R}^{(d_{in}/r') \times r'}$ , and  $D^{(af)} \in \mathbb{R}^{r' \times (d_{out}/r')}$  with  $1 \leq r' \leq \min(d_{in}, d_{out})$ .

**Sharing attribute information among fine-grained domains.** From a fine-grained perspective, global data can be divided into many domains, which linearly scale model parameters and memory budgets. Here, we share one  $C^{(a)}$  with all  $\bigcup_{a, s \in a} C^{(as)}$ .

**Sharing coarse-grained attribute information across attributes.** Across multiple attributes, the coarse-grained views cover the same global dataset, i.e.,  $\forall a \in \mathcal{A} \mathcal{D}^{(ac)} = \mathcal{D}$ . To reduce redundancy information, we share one  $w^{(c)}$  with all  $\bigcup_{a \in \mathcal{A}} w^{(ac)}$ .

**Efficiency analysis.** Without mode decomposition, the external complexity (ignoring  $W$  in PLMs) is  $\sum_a (|a|_f + 2) \cdot r \cdot (d_{in} + d_{out})$ , which is larger than that of our decomposition version with  $(\sum_a (d_{in} + |a|_f \cdot d_{out})) + 2 \cdot r \cdot (d_{in} + d_{out})$ .

## Experiments

### Datasets and Evaluations

We evaluate the proposed M2A frameworks on two types of prevalent datasets. (1) **Multi-domain sentiment analysis** includes **FDU-MTL**, which contains product and movie

Domains	BERT	DAEA	BERTMasker	KCL-KB	B-MTL	B-M2A‡	G-M2A‡	R-M2A‡
Books	87.00	89.00	93.00	93.08	94.75	94.75	93.50	<b>97.25</b>
Electronics	88.30	91.80	93.25	94.92	94.00	95.50	93.75	<b>96.00</b>
DVD	85.60	88.30	89.25	89.92	90.75	93.25	91.50	<b>93.50</b>
Kitchen	91.00	90.30	90.75	92.50	92.00	93.25	95.25	<b>96.00</b>
Apparel	90.00	89.00	92.25	92.67	91.25	92.50	93.25	<b>94.50</b>
Camera	90.00	92.00	92.75	93.67	94.75	95.00	93.00	<b>95.25</b>
Health	88.30	89.80	95.25	95.67	94.25	96.50	95.50	<b>97.50</b>
Music	86.80	88.00	89.50	90.42	90.75	91.75	92.00	<b>93.75</b>
Toys	91.30	91.80	93.75	93.33	92.75	93.50	92.25	<b>94.25</b>
Video	88.00	92.30	91.25	91.67	92.00	<b>94.50</b>	91.75	94.00
Baby	91.50	92.30	92.75	94.58	95.75	<b>96.25</b>	94.00	<b>96.25</b>
Magazines	92.80	<b>96.50</b>	94.50	94.17	94.25	94.75	94.00	<b>96.50</b>
Software	89.30	92.80	93.00	94.33	95.75	<b>96.50</b>	96.25	<b>96.50</b>
Sports	90.80	90.80	92.50	94.42	94.25	<b>95.00</b>	94.50	94.75
IMDB	85.80	90.80	86.00	90.83	93.00	93.25	92.25	<b>94.50</b>
MR	74.80	77.00	83.75	85.58	84.25	<b>85.75</b>	84.00	85.50
Avg.	88.16	90.16	91.47	92.62	92.78	93.86	92.92	<b>94.70</b>

Table 1: Comparative test Acc (%) results for multi-domain sentiment analysis (the category attribute only).

Models	IMDB			Yelp-2013			Yelp-2014		
	Acc	RMSE	F1	Acc	RMSE	F1	Acc	RMSE	F1
non-IID free or Coarse-grained view (IID) only									
BERT	52.2	1.163	49.3	67.7	0.628	65.5	67.7	0.615	65.6
GPT2	51.5	1.222	47.5	67.6	0.622	64.5	68.0	0.614	65.5
RoBERTa	53.0	1.147	49.4	69.2	0.590	65.1	69.0	0.601	66.5
R-M2A†	<b>54.2</b>	<b>1.100</b>	<b>50.3</b>	<b>70.7</b>	<b>0.574</b>	<b>69.4</b>	<b>70.5</b>	<b>0.578</b>	<b>68.5</b>
Multiple granularities (IID+non-IID)									
B-IUPC	53.8	1.151	-	70.5	0.589	-	71.2	0.592	-
B-MAA	57.3	1.042	-	70.3	0.588	-	71.4	0.573	-
B-GS	57.2	1.042	54.5	70.2	0.593	68.3	71.1	0.585	68.5
R-GNNLM	54.4	1.102	-	72.2	0.573	-	72.1	0.568	-
B-M2A‡	58.7	1.021	54.8	72.0	0.569	69.7	72.4	0.560	69.1
G-M2A‡	58.1	1.074	53.6	70.4	0.598	67.4	71.8	0.569	68.4
R-M2A‡	60.3	<b>0.954</b>	<b>56.7</b>	<b>73.7</b>	0.548	<b>70.7</b>	<b>74.2</b>	<b>0.535</b>	<b>71.4</b>
R-M2A	<b>60.6</b>	0.960	56.6	73.4	<b>0.543</b>	70.2	73.6	0.545	70.9

Table 2: Comparative test results for three personalized sentiment classification datasets (user and item attributes). The underline and backbone scores respectively meant the best scores in baselines and all models in each group. Our M2A models outperformed the previous works significantly ( $p < 0.05$ ) in Acc.

reviews across 16 domains from the category-attribute perspective (Liu, Qiu, and Huang 2017). (2) **Personalized sentiment analysis** includes **IMDB**, **Yelp-2013**, and **Yelp-2014**. These datasets exhibit data heterogeneities from user and item-attribute perspectives (Tang, Qin, and Liu 2015). Regarding annotated labels, FDU-MTL, IMDB, and Yelps are treated as a 2-class, 10-class, and 5-class classification problems, respectively.

For evaluation metrics, we adopt **Accuracy (Acc)** to measure the effectiveness of models for multi-domain sentiment analysis. For personalized sentiment analysis, we used **Acc** (primarily), **Rooted-Mean Square Error (RMSE)**, and **Macro-F1 (F1)** scores (Yuan, Zhao, and Qin 2022; Zhang et al. 2023).

## Experimental Settings

We compared our M2A framework with previous state-of-the-art models across different datasets.

(1) For multi-domain datasets, the baselines include **BERT** (Devlin et al. 2019), **DAEA** (Cai and Wan 2019), **BERTMaker** (Yuan, Zhao, and Qin 2022), **KCL-KB** (Yuan et al. 2023), and **MTL** (Thrun 1995).

(2) For personalized datasets, the baselines include **variants of PLMs** in a coarse-grained view, as well as **IPUC** (Lyu, Foster, and Graham 2020), **MAA** (Zhang et al. 2021), **GS** (Zhang et al. 2023), and **GNNLM** (Kertkeidkachorn and Shirai 2023) from multi-grained perspectives.

To investigate the effect of our M2A, we introduced diverse PLMs as backbones, including BERT (B) (Devlin et al. 2019), RoBERTa (R) (Liu et al. 2019), and GPT2 (G) (Radford et al. 2019). We adopted FFT and PEFT for optimization, denoted as **M2A‡** and **M2A**, respectively. Here, M2A‡ provided a fair comparison with previous works also using FFT. **M2A†** represented the coarse-grained version joint learned with M2A, handling non-IID-free scenarios, as discussed in §M2A Framework. All results were averaged over

five runs. Detailed model configurations of hyperparameters can be found in the technical appendices.

### Comparative Results and Analysis

Tables 1 and 2 reported the comparative results of the proposed M2A framework against previous works on multi-domain and personalized sentiment analysis datasets.

From Table 1, when individually trained for each domain, BERT achieved relatively the worst Acc scores. This phenomenon suggested that although PLMs could gain high performance in sentiment analysis, they continued encountering challenges due to data scarcity within each domain. When all domain datasets were gathered, multi-domain models achieved higher performance than BERT, indicating that more extensive and aggregated data could facilitate data-hungry neural networks to achieve generalized performance. B-MTL, which utilized multitask learning to handle all domains, achieved competitive results. However, it treated all samples as IID, disregarding data heterogeneities, which led to inferior performance compared to our B-M2A. The proposed M2A, from a Bayesian perspective, outperformed other baselines by ensemble multiple views and encoding robust domain features. Moreover, the results showed that PLMs using the same transformer structure exhibited varying performance with M2A, emphasizing the importance of selecting appropriate prior distributions ( $p(w)$  and  $p(x|w)$ ). The proposed M2A framework, which integrates both IID and non-IID features, consistently outperformed other models by effectively addressing data heterogeneities and leveraging domain-specific information.

As shown in Table 2, models integrating multiple granularities (from user and item-attribute perspectives) outperformed those considering only coarse-grained views. This demonstrated the importance of leveraging data heterogeneities collected from diverse sources during the downstream task adaptation of PLMs. Compared with GNNLM, which introduced a graph neural work to encode robust attribute or domain representations for improved performance, our M2A provided a multitask learning method based on BNN for the same purpose, resulting in competitive performance across all three datasets.

Furthermore, we compared a general version of M2A (R-M2A†) against RoBERTa. The superior performance validated that our Bayesian learning approach, combined with joint learning, endowed M2A with a generalized capability to handle scenarios with unclear data heterogeneities. Consequently, R-M2A‡ obtained the best results on all three datasets across all three metrics, suggesting the advantage of our Bayesian learning-based M2A framework, especially as the backbone PLMs became more powerful.

### Ablation Study

The test results of the ablation study on both multi-domain and personalized sentiment analysis are reported in Table 3. We used BERT and GPT-2 as backbones for these investigations.

First, we gradually eliminated attribute representations from coarse to fine-grained views. As different granularities were removed, the performance of both B-M2A‡ and

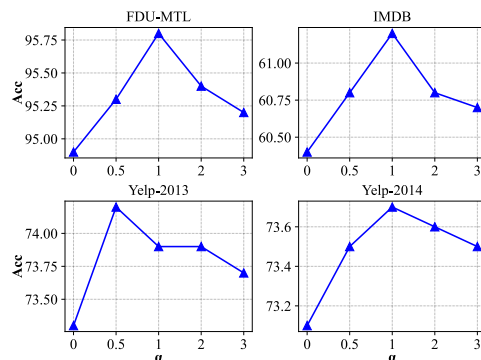


Figure 3: Dev Acc of R-M2A with different multi-task.

G-M2A‡ declined to different extents. This indicated integrating multi-grained information enhanced final performance by mitigating the uncertainty of models. In personalized scenarios, eliminating user non-IID information had a more significant impact than item non-IID information, consistent with previous findings (Wu et al. 2018). This was because user preferences tend to create more distinct non-IID distributions among texts than item characteristics, revealing the diversity among different attributes.

Next, we removed text generation tasks in the module learning strategy to investigate the effect of the prior likelihood estimator  $p(x|w)$ . The results for both B-M2A‡ and G-M2A‡ decreased with the removal of  $p(x|w)$ , indicating its significance. To explore the improvements of B-M2A‡ derived from  $p(x|w)$  or 15% tokens masking, we removed  $p(x|w)$  while pretraining the MLM targets and found that B-M2A‡ without  $p(x|w)$  was worse than the above settings. This phenomenon further demonstrated the effect of  $p(x|w)$  through BNN in Eq. (4).

### The Effect of Bayesian Learning

An extensive experimental analysis was conducted to investigate how Bayesian learning theoretically facilitated our framework for text understanding.

**Mixture of granularities.** We initially employed different LoRA methods to evaluate the Bayesian learning-based mixture of multiple views. From Table 4, it was evident that the ascendancy of LoRA over KronA had a more pronounced effect on the final performance of R-M2A in coarse views than in fine-grained views. Consequently, LoRA was adopted to sample coarse models. Since KronA had much fewer parameters than LoRA and fine-grained views typically encompass many domains, we utilized KronA to instantiate domain modules in fine-grained views.

**Text generation tasks.** We varied the balance factor  $\alpha$  to explore the impact of text generation tasks in BNN assumptions. As shown in Figure 3, with the introduction of text generation tasks for optimizations, the final dev performance was improved, demonstrating the effect of text generation through Bayesian learning. The performance varied with different values of  $\alpha$ ; models with either smaller or larger values  $\alpha$  exhibited relatively lower Acc compared to those with an appropriate selection. This finding suggests caution when

Models		FDU-MTL	IMDB	Yelp-2013	Yelp-2014
B-M2A‡		93.86	58.70	72.00	72.44
M2 Mixture	- coarse view	90.98	-	-	-
	- fine-grained view (user)	-	52.80	68.57	68.85
	- fine-grained view (item)	-	58.43	71.42	71.58
	- fine-grained view (all)	92.78	52.21	67.65	67.65
Module Learning	- text generation task	92.70	57.42	69.36	70.29
	- text generation task (except randomly masking)	93.56	57.41	69.80	71.02
G-M2A‡		92.92	58.10	70.41	71.79
M2 Mixture	- coarse view	91.84	-	-	-
	- fine-grained view (user)	-	52.35	68.04	68.41
	- fine-grained view (item)	-	57.25	70.35	71.49
	- fine-grained view (all)	92.86	51.46	67.63	67.97
Module Learning	- text generation task	92.87	56.43	70.18	70.99

Table 3: The ablative test Acc on multi-domain and personalized sentiment analysis in terms of the mixture of multiple attributes and granularities (guided by Bayesian Inference) and module learning strategies (guided by BNN), respectively.

Methods		IMDB	Yelp-2013	Yelp-2014
CGV	FGV			
FFT	LoRA	OOM	OOM	OOM
	KronA	58.60	72.68	73.52
LoRA	KronA	59.35	71.18	72.22
KronA		57.63	68.03	67.38

Table 4: Dev Acc of R-M2A with different PEFTs for {coarse}+{fine-grained} views (denoted by CGV and FGV). LoRA has a low rank of 64, and KronA estimates a low-rank matrix using rank-one parameters. OOM represents out-of-memory in our settings due to large external parameters.

B-M2A‡	FDU-MTL
w/o text generation	92.70
w/ labeled data	93.86
w/ unlabeled data	93.65

Table 5: Test Acc of B-M2A‡ performing different text generalization tasks.

configuring the application in practice.

To further investigate the impact of text generation tasks, we introduced unlabeled data in each fine-grained domain to facilitate models to learn domain modules. As shown in Table 5, B-M2A‡ w/o text generation tasks represented models optimized with  $p(y|x, w)$  on labeled data, and B-M2A‡ w/ unlabeled data represented models optimized with  $p(y|x, w)$  on labeled data and  $p(x|w)$  on unlabeled data. B-M2A‡ w/ unlabeled data model achieved better results than B-M2A‡ w/o generation tasks, indicating that unlabeled domain-specific data could also leverage domain-specific (non-IID) knowledge for fine-grained adaptations.

**Connection of granularities.** We compared our joint learning strategy with other ensemble methods for integrating general view representation. From Table 6, R-M2A† achieved the best performance among all methods, demonstrating the effect of our proposed joint learning strategy guided by the Bayesian inference assumption. With the re-

Connections	IMDB	Yelp-2013	Yelp-2014
RoBERTa	53.0	69.2	69.0
R-M2A w/ Avg	52.2	70.3	69.9
R-M2A w/ Rand	47.3	66.7	66.2
R-M2A†	<b>54.2</b>	<b>70.7</b>	<b>70.5</b>
R-M2A† w/o Sep	52.8	69.9	70.0
R-M2A† w/o KL	53.8	70.5	70.4
R-M2A† w/o $\mathcal{L}_{mtl}^{NN†}$	53.8	70.3	70.1

Table 6: Test Acc of R-M2A with strategies of connecting granularities for IID-only scenarios. Avg means averaging all fine-grained modules for a general one. Rand randomly selected a fine-grained module for input data. Sep represents separation operation.

moval of KL or coarse-grained losses  $\mathcal{L}_{mtl}^{NN†}$ , the performance of R-M2A† degraded, revealing these losses could facilitate our framework to connect granularities. Moreover, R-M2A† outperformed R-M2A† w/o Sep, validating that  $w^{(ac')}$  was required to continue being updated for a better generalization performance for IID-only scenarios.

## Conclusions

This study proposes an M2A framework to extract data heterogeneities from multi-source data for fine-grained adaptation. Our approach introduced a Bayesian analysis to rethink previous multi-domain-learning and attribute-injecting methods and provided a PEFT and joint learning strategy for facilitating PLMs' adaptations. Experimental findings from multi-domain and personalized sentiment analysis showed that the proposed method could integrate models sampled from multiple attributes and granularities to eliminate data uncertainty. Moreover, the proposed method utilized BNN paradigms to leverage domain modules for performance improvements.

Future works intend to collect a multi-view dataset that contains more kinds of sources for further analysis and build an automatic data heterogeneity detector to verify the effectiveness of our methods.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant Nos. 61966038, 62266051, and 62162068, the National Science and Technology Council (NSTC), Taiwan, ROC, under Grant No. 113-2221-E-155-046-MY3, and the Caiyun Postdoctoral Science Foundation under Grant No. C615300504090. The authors would like to thank the anonymous reviewers for their constructive comments.

## References

- Aguilar, G.; Ling, Y.; Zhang, Y.; Yao, B.; Fan, X.; and Guo, C. 2020. Knowledge distillation from internal representations. In *Proceedings of the AAAI conference on artificial intelligence (AAAI-2020)*, 7350–7357.
- Amplayo, R. K. 2019. Rethinking Attribute Representation and Injection for Sentiment Classification. 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-2019)*, 5601–5612.
- Cai, Y.; and Wan, X. 2019. Multi-Domain Sentiment Classification Based on Domain-Aware Embedding and Attention. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-2019)*, 4904–4910.
- Chen, H.-Y.; and Chao, W.-L. 2021. FedBE: Making Bayesian Model Ensemble Applicable to Federated Learning. In *The Ninth International Conference on Learning Representations (ICLR-2021)*.
- de Freitas, J. F. G.; Niranjan, M.; and Gee, A. H. 2000. Hierarchical Bayesian models for regularization in sequential learning. *Neural computation*, 12(4): 933–953.
- Devlin, J.; Chang, M. W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT-2019)*, 4171–4186.
- Edalati, A.; Tahaei, M.; Kobzyev, I.; Nia, V. P.; Clark, J. J.; and Rezagholizadeh, M. 2022. Krona: Parameter efficient tuning with kronecker adapter. *arXiv preprint arXiv:2212.10650*.
- Han, X.; Zhang, Z.; Ding, N.; Gu, Y.; Liu, X.; Huo, Y.; Qiu, J.; Yao, Y.; Zhang, A.; Zhang, L.; Han, W.; Huang, M.; Jin, Q.; Lan, Y.; Liu, Y.; Liu, Z.; Lu, Z.; Qiu, X.; Song, R.; Tang, J.; Wen, J.-R.; Yuan, J.; Zhao, W. X.; and Zhu, J. 2021. Pre-trained models: Past, present and future. *AI Open*, 2: 225–250.
- Houlsby, N.; Giurghi, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-Efficient Transfer Learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning (ICML-2019)*, volume 97, 2790–2799.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations (ICLR-2021)*.
- Jospin, L. V.; Laga, H.; Boussaid, F.; Buntine, W.; and Bennamoun, M. 2022. Hands-On Bayesian Neural Networks—A Tutorial for Deep Learning Users. *IEEE Computational Intelligence Magazine*, 17(2): 29–48.
- Katsarou, K.; Jeney, R.; and Stefanidis, K. 2023. MUTUAL: Multi-Domain Sentiment Classification via Uncertainty Sampling. In *Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing (SAC-2023)*, 331–339.
- Kertkeidkachorn, N.; and Shirai, K. 2023. Sentiment Analysis using the Relationship between Users and Products. In *Findings of the Association for Computational Linguistics (ACL-2023)*, 8611–8618.
- Kong, J.; Wang, J.; and Zhang, X. 2021. Accelerating Pretrained Language Model Inference Using Weighted Ensemble Self-distillation. In *CCF International Conference on Natural Language Processing and Chinese Computing (NLPCC-2021)*, 224–235.
- Kroese, D. P.; Brereton, T.; Taimre, T.; and Botev, Z. I. 2014. Why the Monte Carlo method is so important today. *Wiley Interdisciplinary Reviews: Computational Statistics*, 6(6): 386–392.
- Liu, P.; Qiu, X.; and Huang, X. 2017. Adversarial Multi-task Learning for Text Classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL-2017)*, 1–10.
- Liu, Q.; Zhang, Y.; and Liu, J. 2018. Learning Domain Representation for Multi-Domain Sentiment Classification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-2018)*, 541–550.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Lyu, C.; Foster, J.; and Graham, Y. 2020. Improving Document-Level Sentiment Analysis with User and Product Context. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING-2020)*, 6724–6729.
- Magris, M.; and Iosifidis, A. 2023. Bayesian learning for neural networks: an algorithmic survey. *Artificial Intelligence Review*, 56(10): 11773–11823.
- Min, B.; Ross, H.; Sulem, E.; Veyseh, A. P. B.; Nguyen, T. H.; Sainz, O.; Agirre, E.; Heintz, I.; and Roth, D. 2024. Recent Advances in Natural Language Processing via Large Pre-trained Language Models: A Survey. *ACM Computing Surveys*, 56(2): 1–40.
- Pang, B.; and Lee, L. 2006. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 1(2): 91–231.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners.

- Tang, D.; Qin, B.; and Liu, T. 2015. Learning Semantic Representations of Users and Products for Document Level Sentiment Classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-2015)*, 1014–1023.
- Thrun, S. 1995. Multitask Learning. *Economic Outlook*, 19(4): 46–47.
- Wu, X.; Huang, S.; and Wei, F. 2023. MoLE: Mixture of LoRA Experts. In *The Twelfth International Conference on Learning Representations (ICLR-2023)*.
- Wu, Z.; Dai, X. Y.; Yin, C.; Huang, S.; and Chen, J. 2018. Improving review representations with user attention and product attention for sentiment classification. In *Proceedings of The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, 5989–5996.
- Yang, Y.; and Hospedales, T. 2015. A Unified Perspective on Multi-Domain and Multi-Task Learning. In *3rd International Conference on Learning Representations (ICLR-2015)*.
- Yao, D.; Pan, W.; Dai, Y.; Wan, Y.; Ding, X.; Yu, C.; Jin, H.; Xu, Z.; and Sun, L. 2024. FedGKD: Toward Heterogeneous Federated Learning via Global Knowledge Distillation. *IEEE Transactions on Computers*, 73(1): 3–17.
- Yuan, J.; Zhao, Y.; and Qin, B. 2022. Learning to share by masking the non-shared for multi-domain sentiment classification. *International Journal of Machine Learning and Cybernetics*, 13(9): 2711–2724.
- Yuan, Z.; Li, Y.; Li, Y.; Zheng, H.-T.; He, Y.; Liu, W.; Huang, D.; and Wu, B. 2023. A Curriculum Learning Approach for Multi-Domain Text Classification Using Keyword Weight Ranking. *Electronics*, 12(14): 3040.
- Zhang, L.; Wang, S.; and Liu, B. 2018. Deep learning for sentiment analysis: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4).
- Zhang, Y.; Wang, J.; Yu, L.-C.; Xu, D.; and Zhang, X. 2023. Domain Generalization via Switch Knowledge Distillation for Robust Review Representation. In *Findings of the Association for Computational Linguistics (ACL-2023)*, 12812–12826.
- Zhang, Y.; Wang, J.; Yu, L.-C.; Xu, D.; and Zhang, X. 2024. Personalized LoRA for Human-Centered Text Understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-2024)*, 19588–19596.
- Zhang, Y.; Wang, J.; Yu, L.-C.; and Zhang, X. 2021. MA-BERT: Learning Representation by Incorporating Multi-Attribute Knowledge in Transformers. In *Findings of the Association for Computational Linguistics (ACL-IJCNLP 2021)*, 2338–2343.