

# SOI Matters: Analyzing Multi-Setting Training Dynamics in Pretrained Language Models via Subsets of Interest

**Shayan Vassef\***      **Amirhossein Dabiriaghdam\***      **Mohammadreza Bakhtiari\***  
University of Illinois Chicago      University of British Columbia      Stony Brook University  
svass@uic.edu      amirhossein@ece.ubc.ca      mohammadreza.bakhtiari@stonybrook.edu

**Yadollah Yaghoobzadeh**  
University of Tehran  
y.yaghoobzadeh@ut.ac.ir

## Abstract

This work investigates the impact of multi-task, multi-lingual, and multi-source learning approaches on the robustness and performance of pretrained language models. To enhance this analysis, we introduce Subsets of Interest (SOI), a novel categorization framework that identifies six distinct learning behavior patterns during training, including forgettable examples, unlearned examples, and always correct examples. Through SOI transition heatmaps and dataset cartography visualization, we analyze how examples shift between these categories when transitioning from single-setting to multi-setting<sup>1</sup> configurations. We perform comprehensive experiments across three parallel comparisons: multi-task vs. single-task learning using English tasks (entailment, paraphrase, sentiment), multi-source vs. single-source learning using sentiment analysis datasets, and multi-lingual vs. single-lingual learning using intent classification in French, English, and Persian. Our results demonstrate that multi-source learning consistently improves out-of-distribution performance by up to 7%, while multi-task learning shows mixed results with notable gains in similar task combinations. We further introduce a two-stage fine-tuning approach where the second stage leverages SOI-based subset selection to achieve additional performance improvements. These findings provide new insights into training dynamics and offer practical approaches for optimizing multi-setting language model performance.

## 1 Introduction

Deep learning has revolutionized natural language processing (NLP), with Transformer-based models (Vaswani et al., 2017) achieving remarkable

success across various tasks. These architectures primarily fall into two categories: decoder-only models, such as GPT-2 (Radford et al., 2019), and encoder-only models, including BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). Recently, decoder-only large language models (LLMs) have gained prominence, notably following the success of ChatGPT. While newer open-source LLMs, such as the Llama3 family (Dubey et al., 2024), facilitate human-friendly interactions across diverse tasks, they do not consistently outperform traditional models (Abaskohi\* et al., 2024; Ghaffarzadeh-Esfahani et al., 2024). Furthermore, current benchmarks for evaluating LLMs emphasize general capabilities such as comprehension and reasoning, frequently neglecting specialized NLP tasks like text classification and named entity recognition. Recent research by Yu et al. (2023) indicates that smaller, fine-tuned encoder-only pretrained language models (PLMs), such as RoBERTa, can match or exceed the performance of larger LLMs across various specialized datasets. Although closed-source LLMs like GPT-4o (Hurst et al., 2024) can occasionally surpass PLMs with extensive prompt engineering, smaller open-source models offer substantial advantages regarding speed, cost-effectiveness, and transparency. Therefore, systematic analyses of PLM training dynamics remain crucial, even as LLMs increasingly dominate NLP research.

Motivated by this, we systematically investigate the impact of multi-task, multi-lingual, and multi-source learning approaches on the robustness and performance of PLMs. Multi-task learning, which leverages shared knowledge across related tasks, has shown considerable promise for enhancing model generalization and robustness, particularly under constraints of limited data and computational resources. Concurrently, multi-source learning exploits diverse data origins to provide models with a broader understanding of target problems,

\* Equal Contribution

<sup>1</sup>For brevity, we use "multi-setting" to refer to multi-task, multi-source, or multi-lingual learning, and "single-setting" to refer to single-task, single-source, or single-lingual learning, respectively.

while multi-lingual learning enables the acquisition of language-agnostic knowledge, significantly improving cross-lingual transfer and performance on low-resource languages.

Despite these advances, a key challenge in training PLMs is handling *Forgettable Examples*, samples that are difficult or out-of-scope, leading models to frequently oscillate between correct and incorrect predictions during training. Although fine-tuning on these challenging examples has proven beneficial in enhancing model robustness (Yaghoobzadeh et al., 2021), systematic analysis of their underlying learning patterns is currently lacking.

To address this gap, we introduce *Subsets of Interest* (SOI), a novel framework for categorizing dataset samples based on distinct learning behaviors observed during training. Specifically, SOI consists of six categories: *Unlearned Examples* (UNE), *Always Correct Examples* (ACE), *1-time Forgettable Examples* ( $1t$ -FRGE), *At least 2-times Forgettable Examples* ( $\geq 2t$ -FRGE), *Early-Learned Examples* (ELE), and *Late-Learned Examples* (LLE). Collectively, these subsets enable detailed insights into the dynamics of model learning behaviors under single- or multi-setting configurations, spanning different tasks, languages, and sources. Furthermore, we investigate the potential of SOI subsets to enhance out-of-distribution performance through second-stage fine-tuning strategies based on various SOI combinations.

The key contributions of this work are as follows: First, we introduce the SOI framework, systematically classifying training samples into distinct learning behavior subsets (Section 4.1). Second, we visualize model learning dynamics via dataset cartography and SOI transition heatmaps, offering intuitive insights into sample-level training behaviors (Sections 4.2 and 4.3). Third, we provide a comprehensive comparative analysis of multi-task, multi-lingual, and multi-source learning methods, evaluating their impacts on both in-distribution (ID) and out-of-distribution (OOD) performances (Subsection 5.1). Lastly, we extend our OOD evaluations through second-stage fine-tuning on strategically chosen subsets derived from SOI analyses, demonstrating additional performance gains (Subsection 5.2).<sup>2</sup>

---

<sup>2</sup>Our code is publicly available at [this GitHub repository](#). It builds upon the implementation provided [here](#), adapting Hugging Face’s transformers library for multi-setting training.

## 2 Related Work

In recent years, extensive research has focused on developing multilingual models as well as models capable of performing multiple NLP tasks simultaneously. Multi-task learning leverages shared representations to jointly optimize model performance across various related tasks, enhancing model generalization, robustness, and computational efficiency. Early foundational work by Collobert and Weston (2008) introduced multi-task learning concepts to NLP, illustrating that training multiple tasks concurrently could lead to better feature generalization and more robust representations. Subsequent studies have widely adopted transfer learning techniques (Howard and Ruder, 2018), demonstrating how pretrained language model knowledge can significantly enhance performance on various downstream NLP tasks. Multi-lingual learning, another promising direction, enables models to gain language-agnostic knowledge to understand, generate, and generalize textual information across multiple languages. Conneau et al. (2020) introduced XLM-R, a robust cross-lingual PLM trained on diverse multilingual data, significantly improving performance on low-resource languages and facilitating effective cross-lingual transfer.

Understanding model behavior at the individual example level represents another critical aspect in training language models. The phenomenon of example forgetting, instances where models oscillate between correct and incorrect predictions during training, has been thoroughly investigated by Yaghoobzadeh et al. (2021). Their work demonstrated that fine-tuning models specifically on these challenging, forgettable examples can significantly enhance model robustness and generalization on task-specific OOD datasets. Complementary to this perspective, Swayamdipta et al. (2020) proposed dataset cartography, a visualization technique characterizing training samples based on prediction confidence and variability metrics. Their method categorizes data into easy-to-learn, hard-to-learn, and ambiguous regions, providing intuitive insights into model behavior throughout training. They conclude that training the model from scratch on the ambiguous region achieves the best ID and OOD performances compared to other scenario cases, including training on hard-to-learn and forgetting examples.

Inspired by these foundational works, our study introduces the Subsets of Interest (SOI) framework,

extending beyond previous categorizations with a finer-grained, analytical perspective. Instead of limiting analysis to three regions, SOI systematically classifies training examples into six distinct learning subsets based on their dynamic behaviors during training. Our comprehensive categorization enriches existing analytical tools, offering nuanced insights into model OOD generalization capability across various multi-task, multi-lingual, and multi-source training scenarios.

### 3 Experiments Setup

In this section we introduce three parallel experimental comparisons: multi-task vs. single-task learning, multi-source vs. single-source learning, and multi-lingual vs. single-lingual learning. For each comparison, we conducted similar experiments to evaluate both performance and generalizability. Our experimental framework encompasses various tasks, languages, datasets, and a unified model architecture detailed below.

#### 3.1 Tasks, Languages and Sources

Our experimental framework spans across multiple dimensions of learning. In the multi-task learning, we utilize three English tasks: entailment (E), paraphrase (P), and sentiment (S). Entailment and paraphrase tasks require binary decisions on semantic relationships between two textual inputs, while sentiment analysis processes single inputs, allowing us to explore combinations of similar tasks (P & E) versus dissimilar ones (S & P, S & E). For the multi-source learning, we focus on sentiment analysis across different data distributions using English datasets, isolating the effects of data source variation from task variation. In our multi-lingual experiments, we conduct intent classification across French (Fr), English (En), and Persian (Fa). This language selection enables us to examine the impact of script and linguistic similarities, as English and French share common features while Persian differs significantly in both script and structure.

#### 3.2 Datasets

We employed several benchmark datasets tailored to different learning settings. For each setting, such as multi-task learning, we construct three pairs of datasets, where each pair includes one in-distribution (ID) and one out-of-distribution (OOD) dataset. Each ID dataset is divided into training, validation, and test splits, whereas the corresponding OOD dataset is treated as a single

evaluation set without internal splits. In the following subsections, we detail the specific datasets chosen for each setting.

##### 3.2.1 Multi-task Learning

For entailment, SciTail (Khot et al., 2018) serves as the ID dataset, comprising 23,097 training examples. The OOD counterpart is the RTE training set from the GLUE benchmark (Wang et al., 2018), comprising 2,490 samples. In the paraphrase detection task, we use the Microsoft Research Paraphrase Corpus (MRPC) (Dolan and Brockett, 2005) as the ID dataset, which includes 3,668 training instances. The OOD dataset in this case is a reduced version of the Quora Question Pairs (QQP) training set from the GLUE (Wang et al., 2018), subsampled to 4,000 examples.

For sentiment classification, we utilize a modified version of the Twitter US Airline Sentiment dataset (Rane and Kumar, 2018), containing 8,078 samples after removing "neutral" labels to enforce binary sentiment polarity. For the corresponding OOD dataset, we adopt a reduced version of the Stanford Sentiment Treebank (SST-2) (Socher et al., 2013) dataset, limited to 4,000 examples to maintain balance across tasks.

##### 3.2.2 Multi-source Learning

For multi-source experiments, we use three sentiment analysis datasets as our ID datasets, each containing 50,000 examples sampled from the full dataset with an 80-10-10 train-eval-test split, resulting in 40,000 training instances per dataset. The IMDB movie reviews dataset (Rudra and Gopalakrishnan, 2023) serves as our first source, the Yelp Reviews dataset (Hemalatha and Ramathmika, 2019) comprises business reviews with binary sentiment labels, and Sentiment140 (Habib and Sultani, 2021) provides sentiment-labeled Twitter content for social media analysis.

As the OOD dataset, we use the Stanford Sentiment Treebank (SST-2) (Socher et al., 2013), comprising 5,000 examples. Since all sources share the same task, we use the same OOD dataset for all three sources.

##### 3.2.3 Multi-lingual learning

For our multilingual experiments, we adopted three intent classification datasets as the ID datasets: Persian subset of MASSIVE (FitzGerald et al., 2023) (11,514 training examples), *Small* subset of CLINC150 (Larson et al., 2019) for English (7,600




Figure 1: Unified Architecture for Our Multi-Setting Learning Experiments.

training samples), and LORIA subset of MIAM (Colombo et al., 2021) for French (8,465 training samples).

Since the number of intent classes can differ across datasets in the intent classification task, we translated each ID dataset into Burmese, a very low-resource language, and treated these translated versions as the OOD datasets. For translation, we employed the No Language Left Behind machine translation model<sup>3</sup> (NLLB Team et al., 2022).

### 3.3 Architecture

Our experiments employ a unified architecture (see Figure 1) that leverages shared knowledge through a common encoder. For multi-task and multi-source experiments, we use BERT-base<sup>4</sup> (Devlin et al., 2019), while multi-lingual experiments utilize the multilingual XLM-R<sup>5</sup> (Conneau et al., 2020) model. Each setting maintains specialized classification heads (task-specific, source-specific, or language-specific) attached to the shared encoder. For multi-setting experiments, training occurs with pairs of tasks/sources/languages: sentiment-entailment (SE), sentiment-paraphrase (SP), and paraphrase-entailment (PE) for tasks; IMDB-Yelp (IY), Sentiment140-Yelp (SY), and IMDB-Sentiment140 (IS) for sources; and finally, English-Persian (En-Fa), French-English (Fr-En), and French-Persian (Fr-Fa) for languages.

## 4 Subsets of Interest

In this section, we present a comprehensive framework for analyzing deep learning models through the lens of training dynamics. We introduce the concept of *Subsets of Interest* (SOI), a novel categorization system that partitions training examples based on their unique learning patterns observed

during the training process. Our analysis unfolds in three complementary parts: first, we formally define the six distinct SOI categories and their characteristics; second, we employ dataset cartography to visualize how these subsets manifest in the confidence-variability space; and third, we introduce transition heatmaps to track how examples migrate between SOI categories under different training configurations. Together, these components provide a systematic approach to understanding and analyzing the complex dynamics of neural network training.

### 4.1 SOI Framework and Definitions

In this section, we introduce a novel approach to analyzing deep learning models by extracting specific samples from the training set, based on unique learning patterns observed during training. Based on these patterns, the training set spans six distinct subsets, which we call *Subsets of Interest* (SOI): 1. *Unlearned Examples* (UNE), 2. *Always Correct Examples* (ACE), 3. *1-time Forgettable Examples* ( $1t$ -FRGE), 4. *At least 2-times Forgettable Examples* ( $\geq 2t$ -FRGE), 5. *Early-Learned Examples* (ELE), and 6. *Late-Learned Examples* (LLE).

UNE refers to samples that show no sign of learning from a certain point onward in the training process. A representative prediction pattern over ten epochs of fine-tuning, assuming the true label is 1, might be  $[1, 0, 0, 0, 0, 0, 0, 0, 0, 0]$ . ACE denotes samples that the model finds particularly easy to learn, exhibiting consistently correct predictions across all epochs, such as  $[1, 1, 1, 1, 1, 1, 1, 1, 1, 1]$ .

$1t$ -FRGE and  $\geq 2t$ -FRGE represent samples that undergo forgetting events, inspired by Toneva et al.’s (2019) work on forgetting dynamics during training. In that framework, our UNE is interpreted as a subset of forgettable examples exhibiting an infinite number of forgetting events, denoted by  $\infty t$ -FRGE. A more recent study by Yaghoobzadeh et al. (2021) defined forgettable examples as those that experience at least one forgetting event (i.e.,  $\geq 1t$ -FRGE), or are never learned at all (i.e., UNE).

In our framework, a forgettable example is defined as one that exhibits at least one forgetting and one recollecting event. A *forgetting event* occurs when a previously correct prediction becomes incorrect in a subsequent epoch, while a *recollecting event* is the reverse, an incorrect prediction followed by a correct one. This distinction ensures that FRGE includes dynamic behav-

<sup>3</sup>We used “nllb-200-3.3B” model.

<sup>4</sup>We used “bert-base-uncased” model.

<sup>5</sup>We used “xlm-roberta-base” model.

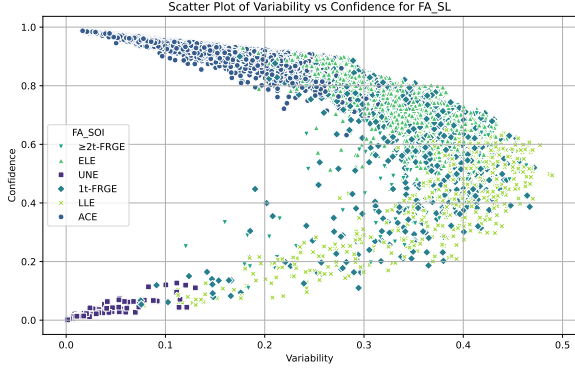


Figure 2: Dataset cartography Map for Single-Lingual Learning in Persian (Fa), showing confidence (average highest prediction probabilities) vs. variability (standard deviation across training epochs).

ior, separating it from UNE, which lacks recollection. For example, a prediction pattern such as  $[0, 1, 0, 0, 0, 1, 0, 1, 0, 0]$  contains three forgetting and two recollecting events, and would be categorized as  $\geq 2t$ -FRGE.

**ELE** and **LLE** refer to samples that initially elude correct classification but eventually reach a point of consistent accuracy. If the first correct prediction occurs on or before epoch 5, the sample is considered ELE; otherwise, it is categorized as LLE, reflecting late-stage learning. For instance, prediction patterns such as  $[0, 0, 1, 1, 1, 1, 1, 1, 1, 1]$  and  $[0, 0, 0, 0, 0, 0, 0, 0, 1, 1]$  represent ELE and LLE, respectively. All training was conducted over 10 epochs, meaning the ELE/LLE classification is influenced by this hyperparameter. However, the broader notion of early- vs. late-stage learning generalizes across training durations.

## 4.2 SOI Visualization via Dataset Cartography

To better illustrate our definition of SOI, we use dataset cartography analysis, following the approach of Swayamdipta et al.’s (2020), to visualize how the model learns over time. This method maps training examples onto a two-dimensional space based on two metrics: *confidence* (the average of the model’s highest prediction probabilities) and *variability* (the standard deviation of these predictions across training epochs). This mapping helps us understand how the model behaves with different examples during training.

With cartography we divide the examples into three main regions: (1) *easy-to-learn*, with high confidence and low variability; (2) *hard-to-learn*, with low confidence and low variability; and (3)




Figure 3: SOI Transition Heatmap: Tracking Training Example Migrations from Single-task (E) to Multi-task (SE). Each cell  $H_{i,j}$  shows the number of examples transitioning from SOI category  $i$  in single-task to category  $j$  in multi-task learning, with the final row and column representing total sums.

*ambiguous*, with high variability.

Figure 2 shows the cartography plot for single-lingual learning in Persian. In this plot, the UNE category mainly appears in the hard-to-learn region, while ACE is mostly found in the easy-to-learn region. The LLE class spreads across the hard-to-learn and ambiguous regions, showing a wide range of variability but generally low confidence. On the other hand, ELE stretches from the ambiguous to the easy-to-learn region, suggesting higher confidence even when variability differs. Both  $1t$ -FRGE and  $\geq 2t$ -FRGE appear in all three regions, with more examples found in the ambiguous area, which suggests less stable learning behavior. A full version of our cartography visualization is provided in Appendix B.

## 4.3 SOI Transitions through Heatmaps

To analyze how training dynamics evolve under different configurations, we introduce *SOI transition heatmaps*, which capture how examples shift between learning behavior categories when moving from a *single-setting* (e.g., single-task) to a *multi-setting* configuration (e.g., multi-task).

Each transition is represented as a one-to-one mapping from a subset of training examples in a given SOI category under the single-setting to a potentially different category under the multi-setting. This mapping highlights how the training setup influences the model’s ability to learn or forget

certain examples.

To visualize these transitions, we construct  $7 \times 7$  heatmaps. The first 6 rows and columns correspond to the defined SOI categories, while the final row and column represent the *row sums* and *column sums*, respectively. Each cell  $H_{i,j}$  in the heatmap records the number of training examples that transitioned from category  $i$  (row, under the single-setting) to category  $j$  (column, under the multi-setting). For instance in Figure 3, if 24 examples labeled as LLE in single-task learning become ELE in multi-task learning, then the corresponding cell is  $H_{\text{LLE,ELE}} = 24$ .

## 5 Results & Analyses

In this section, we present our experimental results and analyses, focusing on the impact of different multi-setting configurations on ID and OOD performance, relative to their corresponding single-setting baselines.

### 5.1 First-Stage Fine-Tuning

The goal of first-stage fine-tuning is to adapt PLMs introduced in Subsection 3.3 to the in-distribution (ID) training sets under both single-setting and multi-setting configurations. Evaluation is then performed using the corresponding **ID test sets** and **OOD test sets**.

#### 5.1.1 In-Distribution Performance

Overall, we observed no substantial improvements in **ID** performance when moving from single-setting to multi-setting fine-tuning. A notable exception was the entailment task, which exhibited a performance gain of 2.6% when trained jointly with the sentiment task. Additionally, five cases showed marginal improvements (between 0.5% and 1.0%), mostly attributed to multi-source learning. The remaining configurations either showed negligible changes ( $\leq 0.5\%$ ) or experienced slight performance degradation (see middle columns of Tables 1, 2, and 3).

#### 5.1.2 Out-of-Distribution Performance

OOD performance across both single-setting and multi-setting configurations was consistently lower than ID performance. However, the comparative OOD performance between the two configurations revealed insightful trends.

Referencing Table 1, we observed similar OOD behavior for French and Persian, while English exhibited a distinct pattern. Specifically, French

showed one performance decline (in the French-English pair) and one improvement (French-Persian), while Persian followed a symmetric trend with a drop in the Persian-French case and an increase in Persian-English. In contrast, English experienced performance drops in both of its OOD pairings. These results suggest that multilingual OOD behavior cannot be easily generalized from single-lingual learning. Notably, the positive impact of one language on another’s OOD performance (e.g., Persian improving French) does not imply reciprocal benefit (i.e., French may not enhance Persian).

Turning to Table 2, the multi-source learning configuration demonstrated consistent OOD improvements across all six evaluated cases. For Sentiment140, we observed the most significant gain, with a 7% improvement in OOD accuracy. Other datasets exhibited improvements exceeding 3%, confirming the effectiveness of multi-source learning in enhancing generalization beyond the training distribution.

Finally, Table 3 echoes the patterns seen in Table 1, with one key distinction: in multi-task learning, when one task enhances another’s OOD performance, the improvement is typically mutual. This is evident in the Paraphrase-Entailment configuration (similar tasks), where OOD performance increased by 1.8% for Entailment and 6.9% for Paraphrase. In contrast, dissimilar task combinations such as Sentiment-Paraphrase led to performance drops in both tasks under OOD evaluation.

Overall, OOD performance improves most when we hold the task and language fixed (e.g., English sentiment analysis) and vary only the data sources (multi-source). With a fixed task but varying languages (multilingual intent classification), the effect is language-dependent—some language pairs improve while others regress. When tasks differ (multitask), gains are conditional and appear primarily when the tasks are closely related (e.g., entailment and paraphrase).

### 5.2 Second-Stage Fine-Tuning

In the second-stage fine-tuning, we investigate whether the multi-setting models fine-tuned in Subsection 5.1 can be further improved to enhance OOD performance. The fine-tuning sets for this stage are selected based on the heatmaps introduced in Section 4.3, which reveal SOI transitions for a given task, language, or source. For instance, referring to Figure 3, we can subsample the En-

tailment training set by extracting all *ELE* examples from the single-task configuration (e.g., all entries along row  $H_{LLE,-}$ ). Using this approach, we experimented with multiple subsampled sets, each defined by a specific heatmap-based criterion, and selected the one that achieved the best average OOD performance across the three multi-task combinations. The selected strategy was then applied to the multi-source and multi-lingual setups as well (see below). Fine-tuning was conducted for 4 epochs, and evaluation was performed solely on OOD sets under multi-setting conditions.

### 5.2.1 Heatmap-Based Fine-Tuning Set Selection

We evaluated several fine-tuning set selection strategies based on the transition patterns identified from the heatmaps: **I.** Transitions representing shifts from more favorable to less favorable learning behaviors (9 out of 36 heatmap transitions:  $[ACE, ELE, LLE] \rightarrow 1t-FRGE$ ,  $[LLE, ELE, ACE, 1t-FRGE] \rightarrow \geq 2t-FRGE$ , and  $[ACE, ELE] \rightarrow LLE$ ); **II.** Diagonal entries excluding both  $ACE \rightarrow ACE$  and  $ELE \rightarrow ELE$ ; **III.** Diagonal entries excluding only  $ACE \rightarrow ACE$ ; **IV.** All forgettable examples identified in single-task learning; **V.** All forgettable examples identified in multi-task learning; and **VI.** The entire training set. Among these strategies, method **III** produced the highest average out-of-distribution (OOD) performance across various multi-task configurations.

### 5.2.2 Out-of-Distribution Performance

We compare the second-stage results against first-stage OOD performance (Subsection 5.1.2). In multi-lingual learning (Table 1), English-French continued to decline, while English-Persian and French-Persian each showed marginal improvements of about 0.3%. Here, one possible explanation for the overall limited improvements lies in the nature of the OOD test language—*Burmese*, a low-resource language that XLM-R may struggle to represent effectively. As a result, improvements made through training on English, French, or Persian datasets may not transfer well to Burmese, regardless of the fine-tuning strategy. In multi-source learning (Table 2), no further gains were observed, likely because the first-stage fine-tuning had already maximized performance. In multi-task learning (Table 3), each combination showed a clear improvement for one task and a slight decline

for the other. These improvements often occurred where first-stage fine-tuning had previously led to performance drops (e.g., Paraphrase dropped from 62.7% to 57.3% in the first stage, then improved to 58.8%).

Based on our analysis, we found that second-stage fine-tuning was most beneficial in the multi-task setting, had limited or no effect in the multi-source setting, and largely preserved performance in the multi-lingual setting. These results suggest that optimizing the fine-tune set selection with the help of SOI transitions heatmaps is a promising direction for improving OOD robustness in multi-setting configurations.

## 6 Conclusion & Future Work

In this work, we conducted a comprehensive investigation into the effects of multi-task, multi-source, and multi-lingual training on PLMs, emphasizing the learning dynamics through the introduction of SOI. By leveraging SOI transition heatmaps and dataset cartography, we provided novel insights into how different training configurations influence both ID and OOD performance. Our results reveal that multi-source learning consistently enhances OOD generalization, while multi-task and multi-lingual learning exhibit more nuanced behavior, offering benefits primarily when task or language similarities exist. The proposed two-stage fine-tuning approach, particularly when guided by SOI-based sample selection, showed further gains in OOD performance, especially in multi-task settings. To sum up, our work highlights the potential of multi-setting configurations in creating more adaptable, robust PLMs capable of generalizing across tasks, languages, and sources.

While our study focused on encoder-based PLMs, future work could apply the SOI framework to large decoder-based language models, such as GPT-style models, to gain insights into their training behaviors and generalization capabilities. Additionally, expanding beyond pairwise combinations to train models on multiple (three or more) tasks, sources, or languages simultaneously could provide a deeper understanding of scaling trends in multi-setting learning. Another direction involves investigating curriculum learning strategies where training is staged according to SOI categories.

Tables 1, 2, and 3 summarize experimental results from initial (first stage) and SOI-guided fine-tuning (second stage). The **initial fine-tuning** trains PLMs separately (single-setting) or jointly (multi-setting) on ID datasets; ID columns report in-distribution evaluations, while OOD columns show out-of-distribution performance. The **SOI-guided fine-tuning** (second stage) further optimizes multi-setting models using targeted subsets strategically selected via SOI transition heatmaps (Section 5.2.1), with improvements measured under the second-stage OOD columns. To interpret these tables, first compare single-setting to multi-setting performances from the initial fine-tuning, then evaluate the additional gains obtained from the subsequent SOI-guided (second stage) fine-tuning.

Table 1: Single and Multi-lingual learning performances. For the multi-lingual setting, we translate each ID dataset (in English, French, or Farsi) into Burmese, and treat the translated samples as OOD evaluation set.

Model Type	Language	First stage fine-tuning		Second stage fine-tuning
		ID	OOD	OOD
Single-lingual	English	84.5	52.8	-
	French	88.5	49	-
	Persian	87.4	62.9	-
Multi-lingual (En-Fr)	English	84.4	51.9	51.8
	French	88.7	41.6	40.9
Multi-lingual (En-Fa)	English	84.7	48	48.1
	Persian	87.4	63.3	63.6
Multi-lingual (Fr-Fa)	French	89.4	52.2	52.2
	Persian	87.2	61	61.4

Table 2: Single and Multi-source learning performances. OOD dataset: SST-2 is used for all three sources.

Model Type	Dataset	First stage fine-tuning		Second stage fine-tuning
		ID	OOD	OOD
Single-source	IMDB	89.4	79.4	-
	Yelp	93.8	79.6	-
	Sentiment140	82.7	76	-
Multi-source (IY)	IMDB	90.2	83.9	83.6
	Yelp	94.1	84.3	84.1
Multi-source (SY)	Sentiment140	83.6	79	79.4
	Yelp	93.7	83.2	82.7
Multi-source (IS)	IMDB	90.2	85.5	84.9
	Sentiment140	83.5	83	83.1

Table 3: Single and Multi-task learning performances. OOD dataset: RTE for entailment, QQP for paraphrase, and SST-2 for sentiment.

Model Type	Task	First stage fine-tuning		Second stage fine-tuning
		ID	OOD	OOD
Single-task	Entailment	89.3	43.9	-
	Sentiment	94.6	76.7	-
	Paraphrase	81.7	62.7	-
Multi-task (SP)	Sentiment	95	75.3	74.4
	Paraphrase	80.3	57.3	58.8
Multi-task (SE)	Sentiment	95.1	62.7	64.9
	Entailment	91.9	38.6	38.2
Multi-task (PE)	Paraphrase	79.3	69.6	70
	Entailment	89.6	45.7	45.1



## Limitations

We outline the known limitations of our current implementation: (1) The subsets of interest (SOI) categorize dataset samples by aggregating their learning pattern—the binary status of learned vs. not learned per epoch—over 10 epochs, yielding six learning categories. However, SOI (i) does not account for the per-epoch class probability distributions and therefore does not fully capture the training dynamics of the samples, and (ii) we did not validate the optimal number of epochs (here, 10), defined as the epoch at which samples within each SOI category exhibit the least shift to another category when the epoch increases by one. (2) The SOI concept is limited to discriminative tasks (e.g., deciding whether a sentence entails another sentence) that require ground-truth labels; for generative models trained on unlabeled text, it does not generalize. (3) The multi-task/multi-source/multilingual experiments were conducted on two datasets; exploring a larger number of datasets remains unexplored. (4) For computational efficiency, we applied the multi-task heatmap-based subsampling (which defines the second-stage fine-tuning set) to both the multi-source and multilingual configurations. However, a single subsampling policy may not generalize across distinct multi-setting configurations.

## References

- Amirhossein Abaskohi\*, Amirhossein Dabiriaghdam\*, Lele Wang, and Giuseppe Carenini. 2024. Bcamirs at semeval-2024 task 4: Beyond words: A multimodal and multilingual exploration of persuasion in memes. *arXiv preprint arXiv:2404.03022*.
- Ronan Collobert and Jason Weston. 2008. [A unified architecture for natural language processing: Deep neural networks with multitask learning](#). In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 160–167, New York, NY, USA. Association for Computing Machinery.
- Pierre Colombo, Emile Chapuis, Matthieu Labeau, and Chloé Clavel. 2021. [Code-switched inspired losses for spoken dialog representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8320–8337, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, volume 1 (long and short papers)*, pages 4171–4186.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Nataraajan. 2023. [MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4277–4302, Toronto, Canada. Association for Computational Linguistics.
- Mohammadreza Ghaffarzadeh-Esfahani, Mahdi Ghaffarzadeh-Esfahani, Arian Salahi-Niri, Hossein Toreyhi, Zahra Atf, Amirali Mohsenzadeh-Kermani, Mahshad Sarikhani, Zohreh Tajabadi, Fatemeh Shojaeian, Mohammad Hassan Bagheri, and 1 others. 2024. Large language models versus classical machine learning: Performance in covid-19 mortality prediction using high-dimensional tabular data. *arXiv preprint arXiv:2409.02136*.
- Mohammad W Habib and Zainab N Sultani. 2021. Twitter sentiment analysis using different machine learning and feature extraction techniques. *Al-Nahrain Journal of Science*, 24(3):50–54.
- S Hemalatha and Ramathmika Ramathmika. 2019. Sentiment analysis of yelp reviews by machine learning. In *2019 International Conference on Intelligent Computing and Control Systems (ICCS)*, pages 700–704. IEEE.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow,

- Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. [Scitail: A textual entailment dataset from science question answering](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An evaluation dataset for intent classification and out-of-scope prediction](#). 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)*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia-Gonzalez, Prangthip Hansanti, and 20 others. 2022. No language left behind: Scaling human-centered machine translation.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Ankita Rane and Anand Kumar. 2018. [Sentiment classification system of twitter data for us airline service analysis](#). In *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, volume 01, pages 769–773.
- Rizwanul Islam Rudra and Anilkumar Kothalil Gopalakrishnan. 2023. Sentiment analysis of consumer reviews using machine learning approach. In *2023 20th International Joint Conference on Computer Science and Software Engineering (JCSSE)*, pages 49–54. IEEE.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. 2019. [An empirical study of example forgetting during deep neural network learning](#). In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Yadollah Yaghoobzadeh, Soroush Mehri, Remi Tachet des Combes, T. J. Hazen, and Alessandro Sordani. 2021. [Increasing robustness to spurious correlations using forgettable examples](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3319–3332, Online. Association for Computational Linguistics.
- Hao Yu, Zachary Yang, Kellin Pelrine, Jean Francois Godbout, and Reihaneh Rabbany. 2023. Open, closed, or small language models for text classification? *arXiv preprint arXiv:2308.10092*.

## A Experiments Environment

All of the experiments were conducted on the Google Colab virtual systems with around 12.7GB of available RAM and an Nvidia T4 GPU with around 15GB available VRAM.

## B Complete Dataset Cartography Visualizations

To complement our dataset cartography analysis in Section 5, we provide here the complete set of cartography visualizations across all learning configurations: Figures 4, 5, and 6 present the confidence-variability distributions for single-setting learning across tasks, sources, and languages, respectively.

## C Complete Heatmap Visualizations

As part of our experimental analysis, we generated 18 transition heatmaps - six for each learning mode (multi-task, multi-source, and multi-lingual). While Section 6 presents a detailed analysis of these transitions, here we provide the complete set




Figure 4: Single-Task Learning (ST) cartography showing the distribution of examples for Entailment, Sentiment, and Paraphrase tasks.




Figure 5: Single-Source Learning (SS) cartography showing the distribution of examples for IMDB, Sentiment140, and Yelp sources.




Figure 6: Single-Lingual Learning (SL) cartography showing the distribution of examples for English, French, and Farsi languages.

of heatmaps for reference: Figures 7, 8, and 9 show how samples transition happen between different SOI categories when moving from single-mode to multi-mode learning. Each cell indicates the number of samples that moved from one category to another, with rows representing the initial (single-mode) categories and columns showing the final (multi-mode) categories.

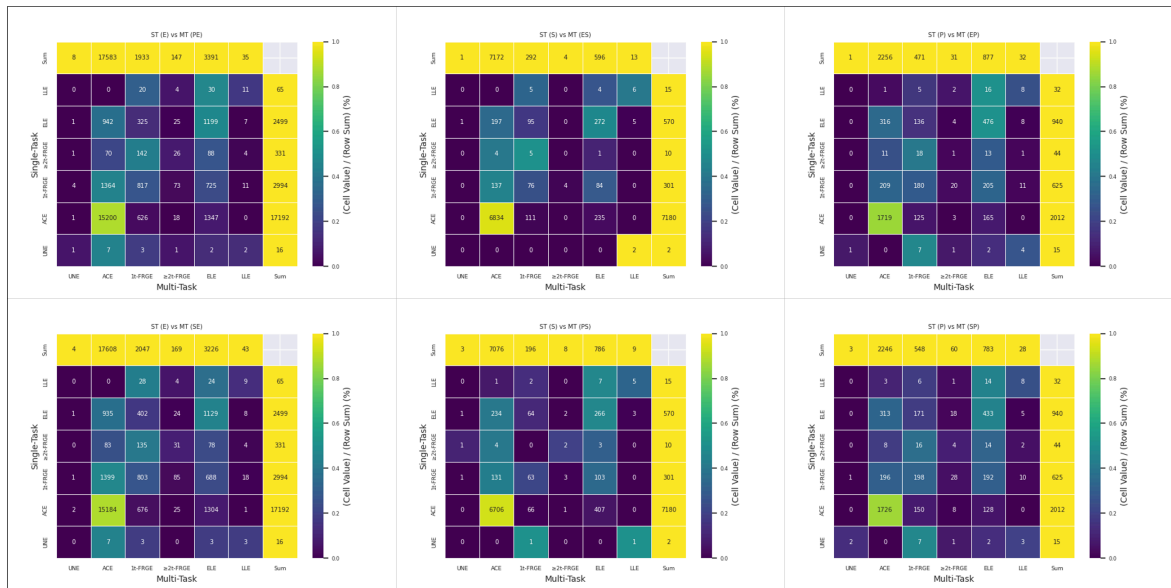


Figure 7: Multi-Task Learning (MT) transition heatmaps showing SOI transitions for all task combinations.

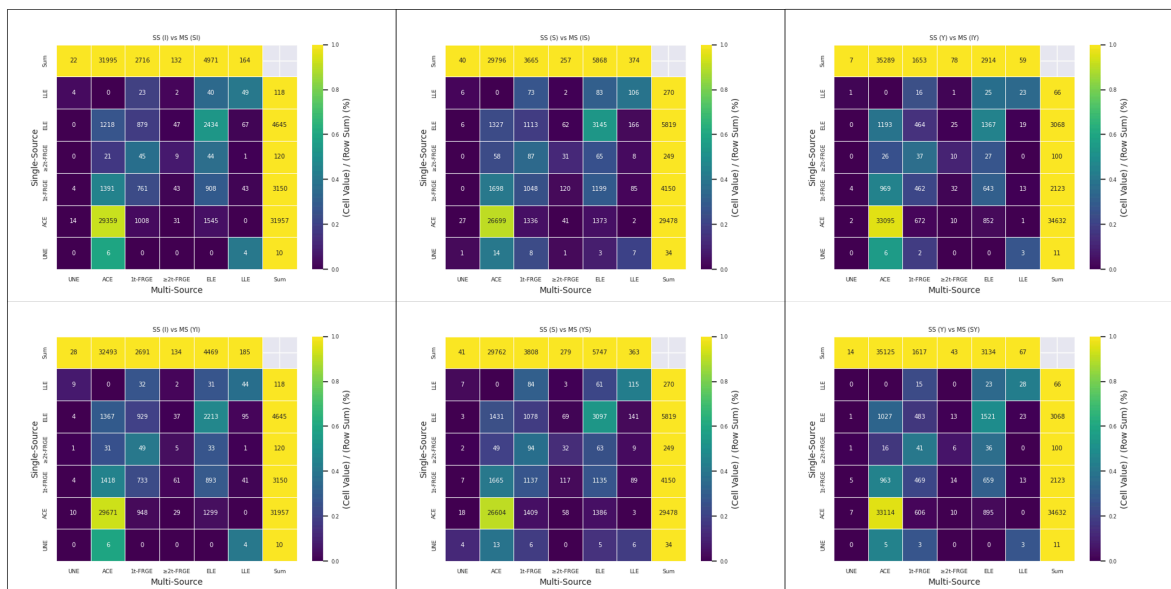


Figure 8: Multi-Source Learning (MS) transition heatmaps showing SOI transitions for all source combinations.

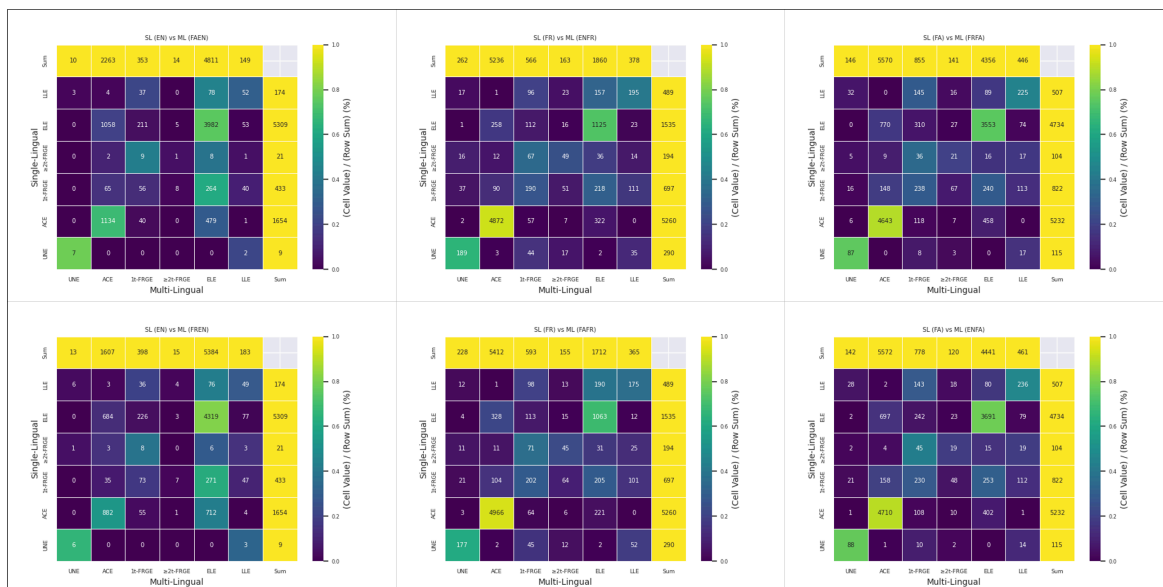


Figure 9: Multi-Lingual Learning (ML) transition heatmaps showing SOI transitions for all language combinations.