

Intention-Adaptive LLM Fine-Tuning for Text Revision Generation

Zhexiong Liu, Diane Litman

Department of Computer Science, Learning Research & Development Center
University of Pittsburgh, Pittsburgh, Pennsylvania, USA 15260
zhexiong@cs.pitt.edu, dlitman@pitt.edu

Abstract

Large Language Models (LLMs) have achieved impressive capabilities in various context-based text generation tasks, such as summarization and reasoning; however, their applications in intention-based generation tasks remain under-explored. One such example is revision generation, which requires the generated text to explicitly reflect the writer’s actual intentions. Identifying intentions and generating desirable revisions are challenging due to their complex and diverse nature. Although prior work has employed LLMs to generate revisions with few-shot learning, they struggle with handling entangled multi-intent scenarios. While fine-tuning LLMs using intention-based instructions appears promising, it demands large amounts of annotated data, which is expensive and scarce in the revision community. To address these challenges, we propose Intention-Tuning, an intention-adaptive layer-wise LLM fine-tuning framework that dynamically selects a subset of LLM layers to learn the intentions and subsequently transfers their representations to revision generation. Experimental results suggest that Intention-Tuning is effective and efficient on small revision corpora, outperforming several PEFT baselines.

1 Introduction

Text revision has been regarded as an essential part of writing as it typically improves the final written work through multiple rounds of editing (Sommers, 1980). However, a writer’s actual intentions, which motivate the revision, are difficult to capture due to their entangled and multi-intent nature (Fitzgerald, 1987). Figure 1 shows the same text edited based on single-intent and multi-intent revisions. One involves a *meaning-changed* revision, while the other focuses on *clarity* and *fluency* improvement. These examples demonstrate that text revisions are often driven by diverse intentions, whether single or multiple, which makes it challenging to generate

Single-Intent Revision: We introduce TLDR generation for scientific papers, a new automatic summarization task with high source compression, requiring expert background knowledge and ~~complex language understanding~~ understanding of complex domain-specific language [meaning change].

Multi-Intent Revision: We introduce TLDR generation for scientific papers [clarity], a new automatic summarization task with high source compression, requiring and requires [fluency] expert background knowledge and complex language understanding.

Figure 1: Revision examples: the original text is revised differently based on a single (meaning change) intention and multiple (clarity and fluency) intentions. The examples are from the ITERATER corpus (Du et al., 2022b).

revisions using computational methods. In the literature, prior work primarily employs sequence-to-sequence language models, such as BART (Lewis et al., 2019), to generate revisions while keeping intentions as inputs (Du et al., 2022a,b), or utilizes prefix-tuning to learn intention representations for revision generation (Chong et al., 2023). However, these models often struggle with multiple entangled intentions, due to their limited capacity (Skitalinskaya and Wachsmuth, 2023). Although Ziegenbein et al. (2024) utilize reinforcement learning policies to prompt large language models (LLMs) for argument rewriting, no specific intentions are used for fine-tuning. While Shu et al. (2024) develop LLMs for text rewriting, their fine-tuning is based on implicit writing instructions rather than explicit intentions. We argue that LLMs fine-tuned on dedicated intention-based revision corpora are critically needed for addressing revision tasks.

LLMs have achieved impressive success in various NLP tasks, including summarization (Takeshita et al., 2024), reasoning (Li et al., 2024), and question answering (Peng et al., 2024). However, their application in revision tasks remains underexplored. This might be because text revision requires an iterative editing process, involving additions, deletions, and modifications, each driven by specific inten-

tions; nevertheless, LLMs are mostly pre-trained to generate just final texts. Although [Shu et al. \(2024\)](#); [Ruan et al. \(2024a\)](#) explore revision LLMs using parameter-efficient fine-tuning (PEFT) methods, these approaches require large-scale training data. In contrast, most revision corpora are small and contain limited annotations, making such methods unsustainable in the revision community.

Increasing attention has recently been given to layer-wise PEFT, since prior work ([Zhang et al., 2023](#); [Elhoushi et al., 2024](#)) suggests fine-tuning LLM layers that contribute more, while freezing those that contribute less, could benefit downstream tasks. For instance, [Yao et al. \(2024\)](#) introduce an importance-aware sparse tuning (IST) to sample a subset of LLM layers for PEFT. While effective, it suffers from over-sampling or under-sampling issues, resulting in suboptimal results ([Liu and Litman, 2025](#)). In addition, current layer-wise PEFT methods ([Pan et al., 2024](#); [Zhu et al., 2024](#); [Yao et al., 2024](#); [Wei et al., 2025](#)) have limited exploration in revision generation tasks, since they are primarily fine-tuned on context-based, which relies heavily on contextualized understanding, rather than intention-based generation, which requires the revisions to specifically reflect the writer’s actual intentions. Therefore, we propose Intention-Tuning¹, a novel intention-adaptive PEFT framework dedicated to learning revision intentions in one task and fine-tuning revision generation in another, all through selected LLM layers. Although Intention-Tuning introduces an additional task to facilitate the transfer of learned representations, it remains efficient as the two tasks are fine-tuned in a sequential manner rather than a joint optimization.

To assess the Intention-Tuning for revision generation, we study three research questions: **RQ1**: Can Intention-Tuning align intention prediction and revision generation tasks through selected LLM layers? **RQ2**: Can Intention-Tuning generate effective revisions using small annotated corpora? **RQ3**: Can Intention-Tuning be generally efficient across different PEFT adapters? In particular, we make the following contributions:

- We are the first work to use an intention-adaptive layer-wise PEFT method for revision generation.
- We develop a framework to align revision intention and revision generation through LLM layers.
- We demonstrate the feasibility and generalizability of the framework on small revision corpora.

¹<https://github.com/ZhexiongLiu/Intention-Tuning>

2 Related Work

Revision Generation Text revision primarily focuses on analyzing human edits to identify revision intentions and generating revisions based on the specific intentions ([Skitalinskaya and Wachsmuth, 2023](#); [Skitalinskaya et al., 2023](#); [Mita et al., 2024](#); [Jourdan et al., 2024](#); [Ruan et al., 2024a](#)). Although an essential NLP task, revision generation remains challenging, mainly due to the scarcity of human annotated corpora ([Anthonio et al., 2020](#); [Spangher et al., 2022](#); [Du et al., 2022b](#); [D’Arcy et al., 2024](#); [Liu et al., 2025](#)). Prior work ([Afrin and Litman, 2018](#); [Kashefi et al., 2022](#); [Afrin et al., 2020](#); [Afrin and Litman, 2023](#)) develops feature-based approaches to predict revision intentions but not revision generation. While [Jiang et al. \(2022\)](#); [Chong et al. \(2023\)](#) utilize sequence-to-sequence language models, such as BART ([Lewis et al., 2019](#)), to study revision generation, these models struggle to capture complex revision patterns due to their limited capacity. More recent efforts have been made to prompt LLMs using instructional inputs ([Ruan et al., 2024b](#)) and fine-tune LLMs ([Ziegenbein et al., 2024](#); [Shu et al., 2024](#)) on a large amount of data. In contrast, we explore more efficient layer-wise PEFT with small corpora.

Layer-wise PEFT PEFT provides efficient solutions for LLM fine-tuning, as it employs adapter-based ([Hu et al., 2023](#)), low-rank adaptive ([Hu et al., 2022](#)), and prompt-based tuning ([Zhao et al., 2024](#)) to minimize expense. However, these methods apply an identical PEFT strategy across all LLM layers, which cannot leverage distinct layer-wise contributions to downstream tasks. To address it, [Kaplun et al. \(2023\)](#) introduce a greedy search to select informative layers, [Pan et al. \(2024\)](#) instead adapt random layer selection, and [Zhu et al. \(2024\)](#) use directional heuristics for the layer-wise fine-tuning. Despite their effectiveness, these methods either incur high computational costs or rely on simple heuristics incompatible with complex tasks. In addition, [Yao et al. \(2024\)](#); [Wei et al. \(2025\)](#) propose sampling important layers based on scoring metrics; however, their methods risk sampling too many redundant layers due to their fixed number of layer selections. Although [Liu and Litman \(2025\)](#) address this issue using a dynamic method, they focus on intention prediction rather than revision generation. Recently, PEFT has been integrated with multi-task learning. [Liu et al. \(2024a\)](#) fine-tune LLMs by leveraging LoRA ([Hu](#)

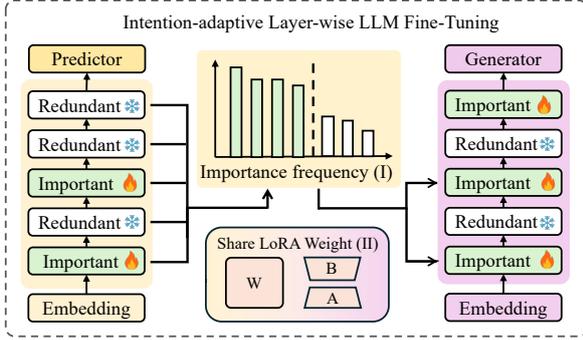


Figure 2: The Intention-Tuning framework. In the intention prediction task (predictor), the important LLM layers are fine-tuned while the redundant LLM layers are frozen. Upon completion, (I) its layer-wise importance frequency is used to finalize the important layers for the revision generation task (generator); (II) the learned layer-wise representations (LoRA weights) in the predictor are shared with the generator.

et al., 2022) and QLoRA (Dettmers et al., 2023) to optimize a balanced multi-task approach, while Agiza et al. (2024) employ task-agnostic and task-specific PEFT adapters for joint fine-tuning. In addition, Baek et al. (2025) develop task-aware feature adaptation by considering task-specific input contexts, whereas Cheng et al. (2025) mitigate task-wise gradient conflicts at individual model layers. Although helpful, these methods primarily focused on either shared adapter modules or non-LLM layers. In contrast, we build on prior work (Liu and Litman, 2025) to dynamically select LLM layers for intention-adaptive multi-task PEFT.

3 Methods

3.1 Preliminaries

We use the notations based on the previous text revision tasks (Liu et al., 2023a). We denote \mathcal{X} , \mathcal{Y} , and \mathcal{I} as original text, revised text, and revision intentions, respectively. We formulate revision intention prediction task $\mathcal{T}_{pre} = f(\mathcal{I}; \mathcal{X}, \mathcal{Y})$ as classifying a pair of revision $\{\mathcal{X}, \mathcal{Y}\}$ into an intention label \mathcal{I} , and revision generation task $\mathcal{T}_{gen} = f(\mathcal{Y}; \mathcal{X}, \mathcal{I})$ as generating \mathcal{Y} given a pair of the original text and its revision intention $\{\mathcal{X}, \mathcal{I}\}$. We propose an intention-adaptive layer-wise PEFT framework that learns intentions in one task and generates revisions in another, all through selected LLM layers. Figure 2 shows the Intention-Tuning framework.

3.2 LLM Layer Selection

Inspired by work (Zhang et al., 2023; Elhoushi et al., 2024; Wei et al., 2025) that suggests fine-tuning partial LLM layers is a feasible PEFT strat-

egy, we use a two-step method to implement this idea, which first probes important layers in \mathcal{T}_{pre} , and then uses the importance frequency based on the probed layers to finalize the important layers in \mathcal{T}_{gen} . Note that the used layers in both tasks have high overlap but might not be exactly the same.

Probing Important Layer Given an LLM $\mathcal{M} = \{m_i\}_{i=1}^{\ell}$ consisting of ℓ transformer layer m . The objective is to split ℓ layers into an important subset \mathcal{S} and a redundant subset $\bar{\mathcal{S}}$, where the layers in \mathcal{S} are used for fine-tuning, and the layers in $\bar{\mathcal{S}}$ are frozen. We use the method in Liu and Litman (2025) to probe important layers based on the layer-wise importance scores (gradient norms) in \mathcal{M} for the task \mathcal{T}_{pre} , as they suggest that layer parameters with high gradient norms make key contributions to the rapid update of the loss function, thus facilitating efficient gradient descent. Additionally, layers with large gradient norms carry information relevant to downstream tasks, making LLM layers informative during fine-tuning. The detailed probing method is described in Appendix A.1.

Finalizing Important Layer Suppose the task \mathcal{T}_{pre} obtains an important layer subset \mathcal{S}_{pre} and redundant layer subset $\bar{\mathcal{S}}_{pre}$ of \mathcal{M} in a single fine-tuning step, its layer-wise importance (updating) frequency $\mathcal{G} = \{p_i\}_{i=1}^{\ell}$ can be obtained, where p_i is the frequency of transformer layer m_i being selected throughout a k -step fine-tuning. Here, $p_i = \sum_{j=1}^k q_{i,j}$, where

$$q_{i,j} = \begin{cases} 1 & \text{if layer } m_i \text{ is selected at step } j, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

We argue that frequently selected layers in task \mathcal{T}_{pre} are well-trained to capture revision intentions, which can be reused for the task of revision generation \mathcal{T}_{gen} . Specifically, layer-wise updating frequency \mathcal{G} from task \mathcal{T}_{pre} is again used to split ℓ layers of \mathcal{M} into important subset \mathcal{S}_{gen} and redundant subset $\bar{\mathcal{S}}_{gen}$ for the task of \mathcal{T}_{gen} , by solving Equation 10, 11 in the Appendix. This design ensures learned representations in the revision intention task can be transferred to the revision generation task through selected important LLM layers. To measure the layer alignment between the two tasks, we define the layer alignment ratio r :

$$r = \frac{|\mathcal{S}_{pre} \cap \mathcal{S}_{gen}|}{|\mathcal{S}_{gen}|}, \quad (2)$$

which measures the percentage of the shared LLM layers across the two tasks.

3.3 Intention-Adaptive LLM Fine-Tuning

Given that text revisions are commonly motivated by the writer’s diverse intentions, we develop an intention-adaptive multi-task PEFT framework to fine-tune LLMs on the tasks of intention prediction and revision generation. Note that we fine-tune the two tasks sequentially, rather than simultaneously, to alleviate computational burdens. The two tasks share weights across LLM layers and maintain separate prediction and generation headers.

Single-Intent Revision Objective Given a pair of revisions $\{\mathcal{X}, \mathcal{Y}\}$ and intention labels $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$, an intention prediction header aims to learn a label distribution through a softmax function and a cross-entropy loss \mathcal{L}_{pre} :

$$\mathcal{P}(\hat{\mathcal{I}}_i) = \frac{\exp(\hat{\mathcal{I}}_i)}{\sum_{i=1}^n \exp(\hat{\mathcal{I}}_i)}, \quad (3)$$

$$\mathcal{L}_{pre} = - \sum_{i=1}^n \mathcal{I}_i \log(\hat{\mathcal{I}}_i), \quad (4)$$

where $\hat{\mathcal{I}}_i$ is a predicted intention. In contrast, a revision generation header aims to generate revised text $\hat{\mathcal{Y}}$, given the original text \mathcal{X} and revision intention \mathcal{I} . Although the generation task also uses cross-entropy loss \mathcal{L}_{gen} , it is optimized to maximize the probability of the next token y_t rather than the intention label:

$$\mathcal{P}(\hat{\mathcal{Y}} | \mathcal{X}, \mathcal{I}) = \prod_{t=1}^d \mathcal{P}(y_t | \mathcal{X}, \mathcal{I}, y_{<t}), \quad (5)$$

$$\mathcal{L}_{gen} = - \sum_{t=1}^d \log \mathcal{P}(y_t | \mathcal{X}, \mathcal{I}, y_{<t}), \quad (6)$$

where $y_{<t}$ are the tokens before y_t and d is the number of tokens in the sequence.

Multi-Intent Revision Objective In the cases that text revisions are driven by multiple intentions, the revised sequences need to simultaneously fulfill all the revision purposes $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$. Thus, an intention prediction header aims to learn multiple labels in \mathcal{I} . Here, a multi-label cross-entropy loss \mathcal{L}'_{pre} is formulated as:

$$\mathcal{L}'_{pre} = - \sum_i^n \left[\mathcal{I}_i \log(\hat{\mathcal{I}}_i) + (1 - \mathcal{I}_i) \log(1 - \hat{\mathcal{I}}_i) \right], \quad (7)$$

where the predicted label $\hat{\mathcal{I}}_i$ is passed to a sigmoid function for $\hat{\mathcal{I}}_i = \sigma(z_i) = \frac{1}{1+e^{-z_i}}$. Similarly, the

Algorithm 1 Intention-adaptive layer-wise PEFT

Output: An optimized LLM \mathcal{M}'

Input: Datasets $\mathcal{D}_{pre}, \mathcal{D}_{gen}$ for Task \mathcal{T}_{pre} and \mathcal{T}_{gen} , Fine-tuning steps \mathcal{K} , LLM \mathcal{M} with PEFT adapters

Parameter: Important layer $\mathcal{S} = \emptyset$, redundant layer $\bar{\mathcal{S}} = \emptyset$

- 1: **for** $k = 1$ to \mathcal{K} **do**
 - 2: Compute fine-tuning loss \mathcal{L}_{pre} and layer-wise gradient norm \mathcal{N}_k on \mathcal{T}_{pre}
 - 3: Use \mathcal{N}_k to obtain important and redundant layer sets $\mathcal{S}_k, \bar{\mathcal{S}}_k$ by Eq. 10, 11
 - 4: Activate LLM layer $m \in \mathcal{S}_k$
 - 5: freeze LLM layer $m' \in \bar{\mathcal{S}}_k$
 - 6: **end for**
 - 7: Compute layer-wise importance frequency \mathcal{G} by Eq. 1 and Task \mathcal{T}_{pre}
 - 8: Use \mathcal{G} to obtain another important and redundant layer sets $\mathcal{S}, \bar{\mathcal{S}}$ by Eq. 10, 11
 - 9: **for** $k = 1$ to \mathcal{K} **do**
 - 10: Activate LLM layer $m \in \mathcal{S}$
 - 11: Freeze LLM layer $m' \in \bar{\mathcal{S}}$
 - 12: Compute fine-tuning loss \mathcal{L}_{gen} on \mathcal{D}_{gen}
 - 13: **end for**
-

optimized probability of the next tokens in the revised text is formulated as:

$$\mathcal{P}'(\hat{\mathcal{Y}} | \mathcal{X}, \mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_k) = \prod_{t=1}^d \mathcal{P}(y_t | \mathcal{X}, \mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n, y_{<t}), \quad (8)$$

$$\mathcal{L}'_{gen} = - \sum_{t=1}^d \log \mathcal{P}(y_t | \mathcal{X}, \mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n, y_{<t}), \quad (9)$$

where d is the number of tokens in the sequence.

3.4 Layer Update with PEFT

While standard PEFT methods have reduced computational overhead, their efficiency can be further enhanced by selectively fine-tuning a subset of layers. In particular, we integrate Low-Rank Adaptation (LoRA) (Hu et al., 2022) into important LLM layers, i.e., important LoRA layers, for efficient low-resource tuning, while freezing those of redundant layers. As shown in Figure 2, both intention prediction and revision generation tasks are fine-tuned based on shared adapter-based PEFT (LoRA). The pseudo-algorithm is described in Algorithm 1.

4 Corpora

ArgRevision Text revision is rarely annotated because its annotation is a costly process; thus,

Intention	Relevant	Irrelevant	Repeated	LCE	not LCE	Commentary	Others	Total
Add	1,774	397	225	1,069	317	425	435	4,642
Delete	598	102	45	318	121	194	126	1,504
Modify	138	26	6	105	23	41	108	447
Total	2,510	525	276	1,492	461	660	669	6,593

Table 1: Statistics of revision intentions in ArgRevision.

Intention	Clarity	Fluency	Coherence	Style	Meaning	Others	Total
Add	94	208	37	2	342	6	689
Delete	282	111	180	20	11	11	615
Modify	1,225	623	176	106	543	41	2,714
Total	1,601	942	393	128	896	58	4,018

Table 2: Statistics of revision intention in ITERATER.

we use a previously collected and annotated argument revision corpus named ArgRevision (Liu and Litman, 2025), which contains 660 pairs of essay drafts written by students in grades four to eight. The ArgRevision corpus includes three essay drafts, where students write initial drafts, and revise their essays in the second drafts based on provided feedback, and further revise their essays in the third drafts based on another round of feedback. The pairs of drafts (e.g., draft1-draft2 and draft2-draft3) are annotated with sentence-level revisions, using their dedicated intention labels: *relevant*, *irrelevant*, *repeated*, *linked claim-evidence (LCE)*, *not LCE*, *commentary* and *others*. Detailed label taxonomy and examples are described in Appendix A.2. The intention label statistics are shown in Table 1.

ITERATER We utilize a publicly available revision corpus named ITERATER (Du et al., 2022b), which annotates 4,018 sentence-level text revisions from Wikipedia, ArXiv, and news articles. The corpus contains six intention labels: *clarity*, *fluency*, *coherence*, *style*, *meaning-changed*, and *others*. Detailed label taxonomy and annotation examples are provided in Appendix A.3. The label statistics are shown in Table 2.

Data Preprocessing We preprocess the two corpora for sentence-level and document-level revision generation. In ArgRevision, the original sentence is empty in *adding*, the revised sentence is empty in *deleting*; only *modifying* revisions are used, in which, however, only 447 examples are collected (see Table 1). Hence, we excluded ArgRevision from sentence-level revision generation; instead, we used it for document-level generation, where a document is revised based on multiple intentions. To simplify the task, we added an <edit> tag before and after an edited sentence to note a text revision in the ArgRevision dataset. Regarding the ITERATER corpus, we utilize its sentence-level dataset, ITERATER-sent, which is single-intent revisions, and its document-level dataset, ITERATER-doc, which contains multi-intent revisions (see Figure 1).

Corpus	ITERATER-sent	ITERATER-doc	ArgRevision
Train/Val/Test	3,215 / 385 / 360	481 / 27 / 51	528 / 66 / 66

Table 3: Data splits across different corpora.

We exclude *others* label in both corpora, yielding six intentions in ArgRevision and five intentions in ITERATER. The data split for the three datasets is summarized in Table 3. The following instruction prompts are used for the intention prediction and revision generation tasks, respectively:

- **Intention Prediction:** Identify the intention of the revision between the original text and the revised text. The possible intentions include: \mathcal{I} . Original Text: \mathcal{X} . Revised Text: \mathcal{Y} .
- **Revision Generation:** Revise the original text based on the intention \mathcal{I} . Original Text: \mathcal{X} . Revised Text: \mathcal{Y} .

5 Experiments

We use cutting-edge LLMs, i.e., upgraded Mistral-7B-Instruct-v0.3 (Mistral-7B) known for efficient inference (Jiang et al., 2023), Llama3.1-8B-Instruct (Llama3.1-8B) good for general NLP tasks (Grattafiori et al., 2024), and Qwen2.5-14B with relatively larger parameters and strong multi-lingual capability (Yang et al., 2024). We compare our Intention-Tuning to the following baselines:

- **Copy-Baseline:** We copy the original text as revised text for a non-edit baseline.
- **BART-Baseline:** We train BART-large (Lewis et al., 2019) as a small language model baseline.
- **ICL-Baseline:** We use in-context learning as a non-finetuned LLM baseline. The used prompt is based on Raheja et al. (2023) and Ziegenbein et al. (2024), as described in Appendix A.4.
- **CoT-Baseline:** We use in-context learning with a chain-of-thought prompt (Wei et al., 2022) as another non-finetuned baseline in Appendix A.4.
- **LISA-Baseline:** We fine-tune randomly selected four layers based on Pan et al. (2024). We apply PEFT (LoRA) to the LLM layers.
- **IST-Baseline:** We compute layer-wise importance scores based on Yao et al. (2024) and sample the top eight LLM layers for PEFT (LoRA).
- **IR-Baseline:** We use gradient norms based on Liu and Litman (2025) to dynamically select a subset of LLM layers for PEFT (LoRA).
- **Full-Finetuning:** We fine-tune all LLM layers with PEFT (LoRA) as a strong baseline.

In the implementation, we build the framework using PyTorch and HuggingFace, and optimize task losses using the Adam optimizer on a Nvidia A100

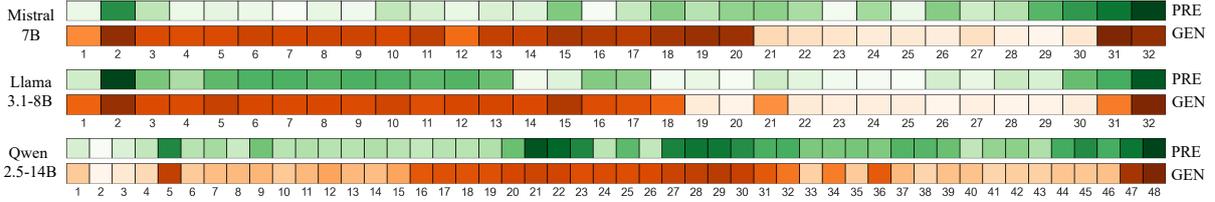


Figure 3: The layer-wise importance score alignment between the intention prediction (green boxes) and the revision generation (brown boxes) tasks while fine-tuning LLMs on the ITERATER-sent. All PEFT uses LoRA. The dark colors indicate high scores (important layers) while the light colors indicate low scores (redundant layers).

GPU. We set the batch size and max length as 16 and 256 for ITERATER-sent, four and 1,024 for ITERATER-doc, two and 1,536 for ArgRevision, respectively. We use a learning rate of $2e-4$, set the maximum epochs to two, select important and redundant layers at every step, and log the loss every 20 steps. We fine-tune LLMs on the training sets, tune hyperparameters on the validation sets, and report **SARI** (Xu et al., 2016), **GLEU** (Napoles et al., 2015), and **Update-ROUGE** (**Update-R**) (Iv et al., 2022) scores on the test sets. The SARI is designed to evaluate addition, deletion, and modification in n-grams; GLEU is customized to penalize changed n-grams; Update-R focuses on updated sentences rather than full text (Shu et al., 2024). Detailed metric implementation and hyperparameter settings during fine-tuning, generation, and evaluation are described in Appendix A.5.

6 Results

6.1 LLM Layer Alignment

To investigate layer-wise alignment between the intention prediction and revision generation tasks, Figure 3 visualizes layer-wise importance scores, i.e., average gradient norms throughout fine-tuning, for the two tasks. Regarding Mistral-7B, a few top, middle, and bottom layers in the prediction task (green) have high importance scores, e.g., dark green boxes in layers 2, 17 to 22, and 29 to 32, which are mostly selected for fine-tuning. In contrast, the top and bottom layers in the generation task have high importance scores, e.g., dark brown boxes in layers 1 to 20 and 31 to 32, which suggests that these layers are frequently selected during revision generation. Also, the layers in the lower middle have low importance scores, e.g., light brown boxes in layers 21 to 30, which suggests these layers are regarded as redundant and are mostly frozen. In terms of Llama3.1-8B, the top to middle and a few bottom layers are both identified as important layers in the two tasks. For Qwen2.5-14B, the middle and bottom layers are mainly used for the

prediction task, while similar middle and bottom layers with slightly shifted locations are mostly used for the generation task. Additionally, several important and redundant layers are located in consecutive positions, suggesting that neighboring layers may share similar fine-tuning patterns.

Despite differences across three LLMs, the important and redundant layers are mostly aligned between the prediction and generation tasks, e.g., dark green and dark brown boxes are mostly aligned², and particularly aligned for layers 2 and 32 in Mistral-7B and Llama3.1-8B, as well as layer 5 in Qwen2.5-14B. Similar patterns are observed in Figure 6 and Figure 7 in the Appendix regarding the ITERATER-doc and ArgRevision datasets. These findings reveal that despite the unique designs across different LLMs, their layers exhibit similar patterns. This consistency demonstrates the feasibility of aligning the intention prediction and revision generation tasks through selected LLM layers for Intention-Tuning, which answers **RQ1**.

6.2 Performance Comparison

We evaluate Intention-Tuning in Table 4. The Copy-Baseline that copies the original text as revised text has low SARI and Update-R scores, but relatively high GLEU scores. This suggests that the original and revised texts have a high overlap; SARI and Updated-R are key metrics for evaluating revision quality. Also, the BART-Baseline exhibits low performance on ITERATER and poor results on ArgRevision due to its limited capability to generate complex revisions (e.g., argument essay revisions), which implies the need for LLM solutions. In ITERATER-sent, Intention-Tuning achieves the best results on Llama3.1-8B, and the best Update-R and Average scores on Qwen2.5-14B. Although not competitive regard-

²Although the colored boxes indicate the LLM layer-wise importance scores, they do not exactly reflect the binary important and redundant layers as well as their alignment across tasks, since there is a threshold to split layer-wise importance scores. Their layer alignment ratio is shown in Table 5.

Models	Methods	ITERATER-sent				ITERATER-doc				ArgRevision			
		SARI	GLEU	Update-R	Average	SARI	GLEU	Update-R	Average	SARI	GLEU	Update-R	Average
Copy-Baseline	N/A	29.46	68.86	2.50	33.61	27.17	53.49	0.00	26.89	30.07	80.60	3.75	38.14
	BART-Baseline	31.74	66.85	35.57	44.72	41.68	47.57	48.22	45.82	16.78	2.34	15.31	11.48
Mistral-7B	Full-Finetuning	37.29	64.26	56.88	52.81	42.85	53.71	32.86	43.14	36.00	62.48	14.19	37.56
	ICL-Baseline	31.43	13.10	61.99*	35.51	37.93	23.90	53.69*	38.51	28.13	26.97	18.79*	24.63
	CoT-Baseline	31.53	12.48	61.83*	35.28	36.74	20.42	52.45*	36.54	24.36	11.24	22.07*	19.22
	LISA-Baseline	36.68	66.33*	56.71	53.24*	35.86	53.91*	35.66*	41.81	34.14	62.23	11.28	35.88
	IST-Baseline	39.42*	64.04	67.95*	57.14*	38.03	54.53*	30.90	41.15	34.29	62.06	9.80	35.38
	IR-Baseline	38.60*	63.79	59.51*	53.97*	45.03*	54.81*	34.21*	44.68*	36.55*	61.77	13.58	37.30
	Intention-Tuning	39.33*	63.72	65.71*	56.25*	45.40*	54.58*	35.15*	45.04*	36.78*	62.07	12.79	37.21
Llama3.1-8B	Full-Finetuning	42.52	63.64	77.26	61.14	47.61	54.79	51.76	51.39	37.80	62.79	15.63	38.74
	ICL-Baseline	38.96	27.93	69.12	45.34	39.03	29.58	54.01*	40.87	35.66	52.24	19.08*	35.66
	CoT-Baseline	37.77	25.64	66.75	43.39	39.24	18.45	47.66	35.12	33.09	48.42	15.54	32.35
	LISA-Baseline	36.86	62.63	74.63	58.04	38.66	54.03	46.11	46.27	34.38	64.43*	12.04	36.95
	IST-Baseline	39.28	63.16	76.82	59.75	41.68	54.76	48.98	48.47	34.38	64.38*	12.20	36.99
	IR-Baseline	41.77	63.88*	76.73	60.79	46.94	54.74	48.79	50.16	37.19	62.48	16.83*	38.83*
	Intention-Tuning	42.10	64.08*	77.99*	61.39*	48.77*	54.86*	52.08*	51.90*	37.37	63.72*	16.58*	39.22*
Qwen2.5-14B	Full-Finetuning	41.95	60.89	76.08	59.64	50.13	55.25	38.60	47.99	41.38	60.96	25.35	42.56
	ICL-Baseline	34.00	29.80	61.61	41.80	38.28	27.86	51.85*	39.33	33.82	55.90	15.02	34.91
	CoT-Baseline	35.72	18.52	65.34	39.86	38.09	19.00	52.63*	36.57	33.07	50.75	17.81	33.88
	LISA-Baseline	39.82	63.13*	74.44	59.13	37.63	53.89	32.76	41.43	34.14	63.65*	13.40	37.06
	IST-Baseline	41.12	63.09*	75.71	59.97*	41.06	53.84	38.93*	44.61	36.22	63.81*	14.29	38.11
	IR-Baseline	41.79	60.60	76.88*	59.76*	45.07	51.80	39.51*	45.46	37.51	60.91	21.74	40.05
	Intention-Tuning	41.13	62.94*	77.50*	60.52*	46.14	54.53	40.67*	47.11	38.08	63.09*	19.94	40.37

Table 4: The performance of Intention-Tuning and baselines on ITERATER-sent, ITERATER-doc, and ArgRevision datasets. The bold numbers represent the best results, and the asterisks (*) indicate that the results are better than those of Full-Finetuning.

ing Mistral-7B, Intention-Tuning achieves close performance to the best SARI score, and generally outperforms Full-Finetuning. In ITERATER-doc, Intention-Tuning largely outperforms the baselines, and sometimes better than Full-Finetuning, regarding Llama3.1-8B and Qwen2.5-14B. Although ICL and CoT baselines achieve relatively high Update-R scores, they exhibit low GLEU scores, indicating that while in-context learning produces revision-like outputs, the edits do not align well with the reference revisions. In ArgRevision, while Intention-Tuning mostly performs better on average, the key metrics are lower than those in the ITERATER datasets, which again suggests that the argument revisions are more complicated than the scientific and news revisions. In general, Llama3.1-8B and Qwen2.5-14B with Intention-Tuning achieve better Average scores than the IR-Baseline across three datasets, which highlights the advantages of transferring layer representations from intention prediction to revision generation.

We argue that the degree of the LLM layer alignment between the two revision tasks could impact its performance, given that both Llama3.1-8B and Qwen2.5-14B exhibit higher layer-wise alignment and higher performance, compared to Mistral-7B, as shown in Figure 3 and Table 4. In particular, we employ the layer alignment ratio metric, as defined in Equation 2, to measure the layer-wise alignment between the two tasks. Table 5 shows that more than half of the layers could be aligned using Intention-Tuning, of which Llama3.1-

LLMs	ITERATER-sent	ITERATER-doc	ArgRevision
Mistral-7B	50.00 %	68.18 %	66.67 %
Llama3.1-8B	83.33 %	56.25 %	72.73 %
Qwen2.5-14B	63.64 %	52.38 %	56.67 %

Table 5: The layer alignment ratio (percentage) between the intention prediction and revision generation tasks.

Methods	Clarity	Fluency	Coherence	Style	Meaning
Full-Finetuning	57.25	75.78	64.39	60.36	44.15
ICL-Baseline	37.96	63.55	44.93	26.20	35.69
CoT-Baseline	35.53	55.68	41.14	41.00	32.92
LISA-Baseline	55.68	69.43	65.05*	46.81	40.88
IST-Baseline	56.53	73.29	63.29	55.74	43.79
IR-Baseline	58.69*	74.48	60.99	54.02	43.23
Intention-Tuning	58.46*	74.72	64.70*	54.86	45.52*

Table 6: The average scores on ITERATER-sent for the Llama3.1-8B model. The bold numbers represent the best results, and the asterisks indicate that the results are better than Full-Finetuning.

8B generally has high alignment, and Mistral-7B performs better in ITERATER-doc than the other two datasets. These observations are in line with the performance in Table 4. Hence, we reason that Intention-Tuning would work best for task-agnostic LLM layers where the layer-wise performance remains consistent across tasks (e.g., layers 2 to 13, 16, 17, 31, and 32 in Llama3.1-8B, as shown in Figure 3), which will be studied in future work.

We evaluate the contribution of individual intention to the revision generation task on ITERATER-sent. Table 6 shows Intention-Tuning achieves the best on *fluency* and *meaning-changed*, and close to IR-Baseline on *clarity*, but lower than the PEFT baselines on *coherence* and *style*. This might be because the two intentions have the smallest number of annotations used for the fine-tuning, as shown

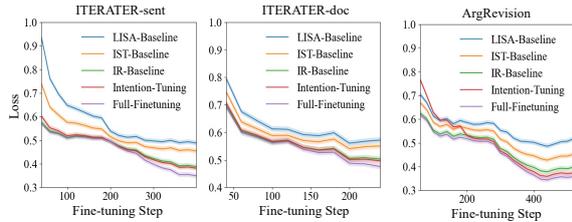


Figure 4: Llama3.1-8B fine-tuning loss on the training set of the datasets. All PEFT uses LoRA.

in Table 2. Although lower than the PEFT results, Llama3.1-8B outperforms Full-Finetuning on *clarity*, *coherence*, and *meaning-changed*, as well as ICL and CoT on all intentions. These observations suggest Intention-Tuning is effective across various revision corpora, which answers **RQ2**.

6.3 Efficiency Evaluation

To evaluate Intention-Tuning efficiency, we visualized the loss convergence of Llama3.1-8B, as it generally exhibits high performance in Table 4 and high layer-wise alignment in Table 5. Figure 4 shows that Intention-Tuning with LoRA consistently achieves faster convergence than the other baselines and is even close to Full-Finetuning. Similar patterns are observed while using other PEFT adapters, i.e., the recent DoRA (Liu et al., 2024b) and classic Bottleneck (Houlsby et al., 2019), as described in Appendix A.7. These results demonstrate that Intention-Tuning remains efficient during fine-tuning. Additionally, we measured peak GPU memory allocations while fine-tuning Llama3.1-8B using standard PEFT versus Intention-Tuning with a batch size of one. Table 7 shows that Intention-Tuning uses less GPU memory than standard PEFT methods across all adapters and three datasets. Among the methods, Intention-Tuning generally saves 5% to 14%. We further compare Intention-Tuning performance using different adapters in Table 8. The results suggest that LoRA stands out for its minimal memory allocation and best performance on the ITERATER datasets. While Bottleneck achieves the highest performance on ArgRevision, it requires 10% more resources than LoRA, and LoRA’s performance remains competitive. These findings indicate that Intention-Tuning is efficient in both fine-tuning convergence and GPU memory allocations, and the findings generalize across different PEFT adapters and datasets, which answers **RQ3**.

7 Qualitative Analysis

We show revision examples in Figure 5, comparing those edited by humans and generated by Llama3.1-

Adapter	ITERATER-sent			ITERATER-doc			ArgRevision		
	Std	IntT	Save	Std	IntT	Save	Std	IntT	Save
LoRA	19	18	5%	26	24	8%	30	28	7%
DoRA	20	19	5%	31	25	19%	36	31	14%
Bottleneck	23	20	13%	30	28	7%	33	31	6%

Table 7: Llama3.1-8B GPU allocations (Gigabytes) between standard (Std) PEFT and Intention-Tuning (IntT).

Adapter	ITERATER-sent	ITERATER-doc	ArgRevision
LoRA	61.39	51.90	39.22
DoRA	59.95	51.28	37.97
Bottleneck	59.58	49.83	39.42

Table 8: Llama3.1-8B with Intention-Tuning. The results are the average scores of the metrics.

8B using Intention-Tuning on the ITERATER-sent corpus. The human revision identifies a *clarity* issue and replaces the phrase “less than optimal” with “inefficient”, to explain computational inefficiency, while the LLM replaces it with “suboptimal”, which improves the *clarity* differently. Regarding *fluency*, the human emphasizes the “optimizing” while the LLM makes a concise edit. In *coherence*, the human uses “and” to connect two actions, and the LLM employs the word “achieving” to link the outcome of the annotation, which achieves a more cohesive sentence flow. For *style*, those two have exactly the same revisions. In terms of *meaning-changed*, the human adds additional content to draw out the implications of the finding; however, the LLM only makes a superficial change, which is less informative. Although LLM revisions do not entirely match those of humans, they are typically effective in reflecting the writer’s actual intentions. Regarding multi-intent revisions, LLMs attempt to perform minimal edits, but sometimes generate hallucinated revisions that do not accurately reflect the writer’s intentions. Such examples (e.g., Figure 10 and 11 in the Appendix) and their error analysis are described in Appendix A.8. These case studies suggest the task difficulty and potential room for future improvement.

8 Conclusion

We present Intention-Tuning, an intention-adaptive layer-wise LLM fine-tuning framework, which transfers the learned representations from the intention prediction to the revision generation through LLM layers. Experiments suggest that it achieves strong performance across LLMs and diverse corpora, while maintaining fast convergence and low GPU allocation. These results reveal its potential and generalizability for complex revision tasks.

[Clarity] Human Revision: We motivate the development of our framework with several cautionary tales of previous research, which has developed numerous less-than-optimal inefficient algorithms for computing expectations...
[Clarity] LLM Revision: We motivate the development of our framework with several cautionary tales of previous research, which has developed numerous less-than-optimal suboptimal algorithms for computing expectations...
[Fluency] Human Revision: Instead of the optimization of optimizing one-step-ahead prediction in the traditional sequence-to-sequence model...
[Fluency] LLM Revision: Instead of the optimization of one-step ahead prediction in traditional sequence-to-sequence model...
[Coherence] Human Revision: Each comment was annotated by three different annotators, which and achieved high inter-annotator agreement.
[Coherence] LLM Revision: Each comment was annotated by three different annotators, which achieved achieving high inter-annotator agreement.
[Style] Human Revision: There were few Muslims there with whom they could engage in dialogue, Noack said claimed .
[Style] LLM Revision: There were few Muslims there with whom they could engage in dialogue, Noack said claimed .
[Meaning-Changed] Human Revision: Finally, we show that the examined models substantially under-perform humans in retrieval, highlighting substantial room for future progress .
[Meaning-Changed] LLM Revision: Finally, we show that the examined models substantially under-perform humans in retrieval tasks .

Figure 5: The example of the human and generated revisions by Llama3.1-8B with Intention-Tuning on the ITERATER-sent dataset. The yellow and green colors denote additions, and the purple denotes deletions.

Ethics Statement

We utilize existing human-annotated revision corpora, which are collected and annotated for academic purposes, and do not pose any ethical implications, such as those related to identity, gender, and race. The corpora include text revisions from both skilled and less skilled writers, which, instead, increases the diversity of the revision research. Furthermore, we acknowledge the existing issues with LLM-generated revisions, such as inaccuracies and hallucinations, which need to be addressed in future work. Lastly, our work on text revision, particularly for argumentative essay revisions, is critically needed for the development of automated student writing evaluation, which would benefit both the education and NLP communities.

Limitations

Hypothesis Our method is grounded in a well-established revision theory that text revisions are primarily motivated by specific intentions. Specifically, we design a two-step method to model the human revision process, in which writers formulate intentions to revise in one step (intention prediction) and implement intention-based revisions in another (revision generation). Also, based on the theory, the revision process is one-directional from intention to generation, but not the opposite. In other words, optimizing the two tasks simulta-

neously is not theoretically justifiable in the text revision domain. However, our framework could be extended to support both tasks mutually.

Algorithm Our algorithm performs multi-task transfer learning on the selected LLM layers, hypothesizing that the important layers for revision generation are aligned with those for intention prediction. Although the findings in Section 6.2 suggest layer-wise alignment could be task-agnostic, there exist task-specific layers that exhibit different patterns in different tasks, and the alignment could vary based on different intentions and LLMs. Thus, further investigation into LLM layer alignments across tasks and intentions would be helpful. Additionally, the algorithm demonstrates the feasibility of LLM fine-tuning with small corpora, which is particularly advantageous for low-resource NLP tasks. The method’s utilization in broader generation tasks will require corpora annotated with specific intentions. Moreover, the algorithm transfers the learned intention representations to the revision generation, which initially results in a high fine-tuning loss and later decreases to normal levels. Nevertheless, it does not introduce significant computational overhead compared to standard PEFT, based on our pilot studies.

LLMs Our evaluation is constrained by computational limitations with relatively lightweight models. Although larger LLMs such as Llama3.3-70B provide stronger contextual understanding, they exceed the scope of this study. In addition, our experiments rely on small annotated corpora, which limits the generalizability of the results to diverse real-world revisions. Also, we follow the literature that commonly reports single-run results due to the high expenses of the LLM fine-tuning; we will perform a significance test based on multi-run results in future work. Furthermore, we focus on layer-wise PEFT methods and do not investigate other PEFT baselines, such as prompt-tuning. Thus, further investigation into alternative PEFT methods and instruction designs could achieve more robust task-specific fine-tuning results. In addition, we randomly select four layers for the LISA-Baseline and the top eight layers for the IST-Baseline because these methods do not provide optimal solutions for determining the number of layers to use. Since the selection of sampling LLM layers is underexplored, we use the optimal numbers based on prior work (Liu and Litman, 2025) to achieve a compromise between performance and efficiency.

Metrics We use the standard text revision generation metrics that focus on the edit-edit alignment rather than text-text alignment. Although embedding similarity is helpful for text-text comparison, it might not reflect the actual revision quality, given most of the text in a document is unchanged, and thus has high embedding similarity. To the best of our knowledge, LLM-as-a-Judge is not commonly used for text revision generation tasks because it requires criteria tailored to each revision issue. It is interesting to explore, but it needs human evaluation on the LLM judges themselves before it can be implemented for assessing the revision quality, which will be our future work.

Acknowledgments

This research was supported by National Science Foundation Award #2202347 and by a supplement made in response to NSF DCL 24-093. We also acknowledge the National Artificial Intelligence Research Resource (NAIRR) Pilot and Voltage Park for contributing to our research results. The opinions expressed are those of the authors and do not represent the views of the institutes. The authors would like to thank the anonymous reviewers and the Pitt PETAL and NLP groups for their valuable feedback on this work.

References

- Tazin Afrin and Diane Litman. 2018. Annotation and classification of sentence-level revision improvement. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 240–246, New Orleans, Louisiana.
- Tazin Afrin and Diane Litman. 2023. [Predicting desirable revisions of evidence and reasoning in argumentative writing](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2550–2561, Dubrovnik, Croatia. Association for Computational Linguistics.
- Tazin Afrin, Elaine Lin Wang, Diane Litman, Lindsay Clare Matsumura, and Richard Correnti. 2020. Annotation and classification of evidence and reasoning revisions in argumentative writing. In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, Seattle, Washington, USA (Remote).
- Ahmed Agiza, Marina Neseem, and Sherief Reda. 2024. Mtlora: Low-rank adaptation approach for efficient multi-task learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16196–16205.
- Talita Anthonio, Irshad Bhat, and Michael Roth. 2020. [wikiHowToImprove: A resource and analyses on edits in instructional texts](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5721–5729, Marseille, France. European Language Resources Association.
- Seungmin Baek, Soyul Lee, Hayeon Jo, Hyesong Choi, and Dongbo Min. 2025. Tadformer: Task-adaptive dynamic transformer for efficient multi-task learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14858–14868.
- Tiancong Cheng, Ying Zhang, Rajiv Ratn Shah, Roger Zimmermann, Zhiwen Yu, and Bin Guo. 2025. Compmtl: Layer-wise competitive multi-task learning. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Ruining Chong, Cunliang Kong, Liu Wu, Zhenghao Liu, Ziye Jin, Liner Yang, Yange Fan, Hanghang Fan, and Erhong Yang. 2023. [Leveraging prefix transfer for multi-intent text revision](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1219–1228, Toronto, Canada. Association for Computational Linguistics.
- Mike D’Arcy, Alexis Ross, Erin Bransom, Bailey Kuehl, Jonathan Bragg, Tom Hope, and Doug Downey. 2024. [ARIES: A corpus of scientific paper edits made in response to peer reviews](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6985–7001, Bangkok, Thailand. Association for Computational Linguistics.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: efficient finetuning of quantized llms. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022a. [Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision](#). In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, pages 96–108, Dublin, Ireland. Association for Computational Linguistics.
- Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, and Dongyeop Kang. 2022b. [Understanding iterative revision from human-written text](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3573–3590, Dublin, Ireland. Association for Computational Linguistics.
- Jane Dwivedi-Yu, Timo Schick, Zhengbao Jiang, Maria Lomeli, Patrick Lewis, Gautier Izacard, Edouard Grave, Sebastian Riedel, and Fabio Petroni. 2024. [EditEval: An instruction-based benchmark for text](#)

- improvements. In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 69–83, Miami, FL, USA. Association for Computational Linguistics.
- Mostafa Elhoushi, Akshat Shrivastava, Diana Liskovich, Basil Hosmer, Bram Wasti, Liangzhen Lai, Anas Mahmoud, Bilge Acun, Saurabh Agarwal, Ahmed Roman, Ahmed Aly, Beidi Chen, and Carole-Jean Wu. 2024. **LayerSkip: Enabling early exit inference and self-speculative decoding**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12622–12642, Bangkok, Thailand. Association for Computational Linguistics.
- Jill Fitzgerald. 1987. Research on revision in writing. *Review of educational research*, 57(4):481–506.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. **The llama 3 herd of models**. *Preprint*, arXiv:2407.21783.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Lee. 2023. **LLM-adapters: An adapter family for parameter-efficient fine-tuning of large language models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5254–5276, Singapore. Association for Computational Linguistics.
- Robert Iv, Alexandre Passos, Sameer Singh, and Ming-Wei Chang. 2022. **FRUIT: Faithfully reflecting updated information in text**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3670–3686, Seattle, United States. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. **Mistral 7b**. *Preprint*, arXiv:2310.06825.
- Chao Jiang, Wei Xu, and Samuel Stevens. 2022. **arXivEdits: Understanding the human revision process in scientific writing**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9420–9435, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Herv   J  gou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- L  ane Jourdan, Florian Boudin, Nicolas Hernandez, and Richard Dufour. 2024. **CASIMIR: A corpus of scientific articles enhanced with multiple author-integrated revisions**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2883–2892, Torino, Italia. ELRA and ICCL.
- Gal Kaplun, Andrey Gurevich, Tal Swisa, Mazor David, Shai Shalev-Shwartz, and Eran Malach. 2023. Less is more: Selective layer finetuning with sub-tuning. *arXiv preprint arXiv:2302.06354*.
- Omid Kashefi, Tazin Afrin, Meghan Dale, Christopher Olshefski, Amanda Godley, Diane Litman, and Rebecca Hwa. 2022. Argrewrite v. 2: an annotated argumentative revisions corpus. *Language Resources and Evaluation*, pages 1–35.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Nianqi Li, Jingping Liu, Sihang Jiang, Haiyun Jiang, Yanghua Xiao, Jiaqing Liang, Zujie Liang, Feng Wei, Jinglei Chen, Zhenghong Hao, and 1 others. 2024. Cr-llm: A dataset and optimization for concept reasoning of large language models. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13737–13747.
- Bingchang Liu, Chaoyu Chen, Zi Gong, Cong Liao, Huan Wang, Zhichao Lei, Ming Liang, Dajun Chen, Min Shen, Hailian Zhou, and 1 others. 2024a. Mft-coder: Boosting code llms with multitask fine-tuning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5430–5441.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024b. **DoRA: Weight-decomposed low-rank adaptation**. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 32100–32121. PMLR.
- Zhexiong Liu and Diane Litman. 2025. **Efficient layer-wise LLM fine-tuning for revision intention prediction**. In *Findings of the Association for Compu-*

- tational Linguistics: EMNLP 2025*, pages 15319–15334, Suzhou, China. Association for Computational Linguistics.
- Zhexiong Liu, Diane Litman, Elaine Wang, Lindsay Matsumura, and Richard Correnti. 2023a. [Predicting the quality of revisions in argumentative writing](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 275–287, Toronto, Canada. Association for Computational Linguistics.
- Zhexiong Liu, Diane Litman, Elaine L Wang, Tianwen Li, Mason Gobat, Lindsay Clare Matsumura, and Richard Correnti. 2025. [eRevise+RF: A writing evaluation system for assessing student essay revisions and providing formative feedback](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 173–190, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zhexiong Liu, Licheng Liu, Yiqun Xie, Zhenong Jin, and Xiaowei Jia. 2023b. [Task-adaptive meta-learning framework for advancing spatial generalizability](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12):14365–14373.
- Masato Mita, Keisuke Sakaguchi, Masato Hagiwara, Tomoya Mizumoto, Jun Suzuki, and Kentaro Inui. 2024. [Towards automated document revision: Grammatical error correction, fluency edits, and beyond](#). In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 251–265, Mexico City, Mexico. Association for Computational Linguistics.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. [Ground truth for grammatical error correction metrics](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 588–593, Beijing, China. Association for Computational Linguistics.
- Rui Pan, Xiang Liu, Shizhe Diao, Renjie Pi, Jipeng Zhang, Chi Han, and Tong Zhang. 2024. [Lisa: layer-wise importance sampling for memory-efficient large language model fine-tuning](#). *Advances in Neural Information Processing Systems*, 37:57018–57049.
- Yixing Peng, Quan Wang, Licheng Zhang, Yi Liu, and Zhendong Mao. 2024. [Chain-of-question: A progressive question decomposition approach for complex knowledge base question answering](#). In *ACL (Findings)*.
- Vipul Raheja, Dhruv Kumar, Ryan Koo, and Dongyeop Kang. 2023. [CoEdit: Text editing by task-specific instruction tuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5274–5291, Singapore. Association for Computational Linguistics.
- Qian Ruan, Iliia Kuznetsov, and Iryna Gurevych. 2024a. [Are large language models good classifiers? a study on edit intent classification in scientific document revisions](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15049–15067, Miami, Florida, USA. Association for Computational Linguistics.
- Qian Ruan, Iliia Kuznetsov, and Iryna Gurevych. 2024b. [Re3: A holistic framework and dataset for modeling collaborative document revision](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4635–4655, Bangkok, Thailand. Association for Computational Linguistics.
- Lei Shu, Liangchen Luo, Jayakumar Hoskere, Yun Zhu, Yinxiao Liu, Simon Tong, Jindong Chen, and Lei Meng. 2024. [RewritelM: An instruction-tuned large language model for text rewriting](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18970–18980.
- Gabriella Skitalinskaya, Maximilian Spliethöver, and Henning Wachsmuth. 2023. [Claim optimization in computational argumentation](#). In *Proceedings of the 16th International Natural Language Generation Conference*, pages 134–152, Prague, Czechia. Association for Computational Linguistics.
- Gabriella Skitalinskaya and Henning Wachsmuth. 2023. [To revise or not to revise: Learning to detect improvable claims for argumentative writing support](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15799–15816, Toronto, Canada. Association for Computational Linguistics.
- Nancy Sommers. 1980. Revision strategies of student writers and experienced adult writers. *College Composition & Communication*, 31(4):378–388.
- Alexander Spangher, Xiang Ren, Jonathan May, and Nanyun Peng. 2022. [NewsEdits: A news article revision dataset and a novel document-level reasoning challenge](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 127–157, Seattle, United States. Association for Computational Linguistics.
- Sotaro Takeshita, Tommaso Green, Ines Reinig, Kai Eckert, and Simone Ponzetto. 2024. [ACLSum: A new dataset for aspect-based summarization of scientific publications](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6660–6675, Mexico City, Mexico. Association for Computational Linguistics.
- Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. 2021. [Augmented SBERT: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks](#). In *Proceedings of*

the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 296–310, Online. Association for Computational Linguistics.

Chenxing Wei, Yao Shu, Ying Tiffany He, and Fei Yu. 2025. Flexora: Flexible low-rank adaptation for large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14643–14682, Vienna, Austria. Association for Computational Linguistics.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA. Curran Associates Inc.

Yiqun Xie, Erhu He, Xiaowei Jia, Han Bao, Xun Zhou, Rahul Ghosh, and Praveen Ravirathinam. 2021. A statistically-guided deep network transformation and moderation framework for data with spatial heterogeneity. In *2021 IEEE International Conference on Data Mining (ICDM)*, pages 767–776. IEEE.

Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Kai Yao, Penglei Gao, Lichun Li, Yuan Zhao, Xiaofeng Wang, Wei Wang, and Jianke Zhu. 2024. Layer-wise importance matters: Less memory for better performance in parameter-efficient fine-tuning of large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1977–1992, Miami, Florida, USA. Association for Computational Linguistics.

Kaiyan Zhang, Ning Ding, Biqing Qi, Xuekai Zhu, Xinwei Long, and Bowen Zhou. 2023. CRaSh: Clustering, removing, and sharing enhance fine-tuning without full large language model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9612–9637, Singapore. Association for Computational Linguistics.

Zheng Zhao, Yftah Ziser, and Shay B Cohen. 2024. Layer by layer: Uncovering where multi-task learning happens in instruction-tuned large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15195–15214, Miami, Florida, USA. Association for Computational Linguistics.

Ligeng Zhu, Lanxiang Hu, Ji Lin, and Song Han. 2024. Lift: Efficient layer-wise fine-tuning for large model models.

Timon Ziegenbein, Gabriella Skitalinskaya, Alireza Bayat Makou, and Henning Wachsmuth. 2024. LLM-based rewriting of inappropriate argumentation using reinforcement learning from machine feedback. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4455–4476, Bangkok, Thailand. Association for Computational Linguistics.

A Appendix

A.1 Probing Important Layer

Given an LLM $\mathcal{M} = \{m_i\}_{i=1}^{\ell}$ consisting of ℓ transformer layer m , its layer-wise gradient norms are denoted as $\mathcal{N} = \{a_i\}_{i=1}^{\ell}$. The objective is to split ℓ layers into an important subset \mathcal{S} and a redundant subset $\bar{\mathcal{S}}$, where $\mathcal{S} = \{m_i \mid a_i > \gamma\}$, $\bar{\mathcal{S}} = \{m_i \mid a_i \leq \gamma\}$, and γ is a threshold to obtain. Based on prior work (Liu et al., 2023b; Liu and Litman, 2025), the layer splitting task can be formulated as a distribution divergence problem, arguing that important and redundant layers are drawn from different distributions. Thus, this problem can be solved as

$$\mathcal{S}^* = \operatorname{argmax}_{\mathcal{N}} \log \frac{\text{Likelihood}(\mathcal{H}_1 \mid \mathcal{N})}{\text{Likelihood}(\mathcal{H}_0)}, \quad (10)$$

where \mathcal{S}^* contains selected important layers, \mathcal{H}_0 is a null hypothesis that the importance scores of layers in \mathcal{M} follows a single distribution, and \mathcal{H}_1 is an alternative hypothesis that there exists a subset \mathcal{S} of \mathcal{M} , where importance scores of layers in \mathcal{S} follow a different distribution from those of the remaining layers in $\bar{\mathcal{S}}$. Xie et al. (2021) suggest the likelihood of the alternative hypothesis can be optimized by minimizing the variances of the importance scores in $\mathcal{N}_{\mathcal{S}} = \{a_i \mid m_i \in \mathcal{S}\}$ and $\mathcal{N}_{\bar{\mathcal{S}}} = \{a_i \mid m_i \in \bar{\mathcal{S}}\}$. Hence, an optimal threshold γ^* is obtained by minimizing the sum of variances $\text{Var}(\mathcal{N}_{\mathcal{S}})$ and $\text{Var}(\mathcal{N}_{\bar{\mathcal{S}}})$ such that

$$\text{Var}(\mathcal{N}_{\mathcal{S}}) + \text{Var}(\mathcal{N}_{\bar{\mathcal{S}}}) \leq \text{Var}(\mathcal{N}). \quad (11)$$

A.2 ArgRevision Examples

We use the argument essays corpus previously described in Liu and Litman (2025). The corpus addresses a crucial need for argument revision in automated student writing assessment, which offers valuable contributions to both the education and NLP communities. The annotation taxonomy (Afrin et al., 2020) includes:

- **Relevant:** relevant revision is about examples or details that are relevant to the claim.
- **Irrelevant:** irrelevant revision is about examples or details that are appropriate and unnecessary, impertinent to, or disconnected from claims.
- **Repeated:** repeat evidence revision is about examples or details that already exist.
- **LCE:** linked claim-evidence revision is about the explanation that connects evidence with claims.
- **Not LCE:** non-linked claim-evidence revision is about the explanation that does not connect evidence with claims.
- **Commentary:** commentary revision is about the explanation that is unrelated to claims or source text; instead, it is the writer’s personal experience.
- **Other:** other argument revisions that are not included in the above intentions.

We show annotated sentence-level revision examples in Table 14, where the original sentence is empty in adding, and the revised sentence is empty in deleting.

A.3 ITERATER Examples

We used a publicly available annotated corpus named ITERATER (Du et al. 2022b), which contains revisions annotated from Wikipedia, ArXiv, and news articles. Wikipedia revisions typically aim to enhance clarity and structure. ArXiv edits are made by researchers who revise content, such as hypotheses, experimental findings, and interpretations. News article edits are performed by editors who focus on improving clarity and readability. The annotation taxonomy includes:

- **Fluency:** the revision to fix grammatical errors in the text.
- **Coherence:** the revision to make the text more cohesive, logically linked, and consistent.
- **Clarity:** the revision to make the text more formal, concise, readable, and understandable.
- **Style:** the revision to convey the writer’s writing preferences, including emotions, tone, voice, etc.
- **Meaning-Changed:** the revision to update or add new information to the text.
- **Other:** the revisions that are not recognizable and do not belong to the above intentions.

The annotated ITERATER corpus includes ITERATER-sent and ITERATER-doc datasets. We include annotated sentence-level revision examples in Table 15.

```
### Instruction: revise the original text based on the intention(s).
```

```
Relevant: relevant revision is about examples or details that are relevant to the claim.
Irrelevant: irrelevant revision is about examples or details that are appropriate and unnecessary, impertinent to, or disconnected from claims.
Repeated: repeat evidence revision is about examples or details that already exist.
LCE: linked claim-evidence revision is about the explanation that connects evidence with claims.
Not LCE: non-linked claim-evidence revision is about the explanation that does not connect evidence with claims.
Commentary: commentary revision is about the explanation that is unrelated to claims or source text; instead, it is the writer’s personal experience.
```

```
Here is the example of the revision: {example}
```

```
### Original text: {original text}
### Revised text:
```

Table 9: The in-context learning prompt regarding six intentions in ArgRevision.

```
### Instruction: revise the original text based on the intention(s).
```

```
Relevant: relevant revision is about examples or details that are relevant to the claim.
Irrelevant: irrelevant revision is about examples or details that are appropriate and unnecessary, impertinent to, or disconnected from claims.
Repeated: repeat evidence revision is about examples or details that already exist.
LCE: linked claim-evidence revision is about the explanation that connects evidence with claims.
Not LCE: non-linked claim-evidence revision is about the explanation that does not connect evidence with claims.
Commentary: commentary revision is about the explanation that is unrelated to claims or source text; instead, it is the writer’s personal experience.
```

```
When you revise the text, please break down the revision process into multiple steps, and check if the revision fulfills the intention requirement in each step. Only output the final revised text. Do not output the intermediate steps.
```

```
Here is the example of the revision: {example}
```

```
### Original text: {original text}
### Revised text:
```

Table 10: The in-context learning with chain-of-thought prompt regarding six intentions in ArgRevision.

A.4 In-Context Learning Prompts

We design an in-context learning prompt based on Raheja et al. (2023) and Ziegenbein et al. (2024) as a non-finetuned LLM baseline. The prompt includes the intention taxonomies described in Appendices A.2 and A.3 to indicate the revision objectives³. The prompt also uses a one-shot example retrieved from the training set to demonstrate the

³The intention taxonomies included in a prompt are determined by the intention labels associated with each revision. Not all the taxonomies are used if a prompt does not require the revisions to achieve all the intentions.

Instruction: revise the original text based on the intention(s).

Fluency: the revision to fix grammatical errors in the text.
Coherence: the revision to make the text more cohesive, logically linked, and consistent.

Clarity: the revision to make the text more formal, concise, readable, and understandable.

Style: the revision to convey the writer’s writing preferences, including emotions, tone, voice, etc.

Meaning-Changed: the revision to update or add new information to the text.

Here is the example of the revision: {example}

Original text: {original text}

Revised text:

Table 11: The in-context learning prompt regarding five intentions in ITERATER.

Instruction: revise the original text based on the intention(s).

Fluency: the revision to fix grammatical errors in the text.
Coherence: the revision to make the text more cohesive, logically linked, and consistent.

Clarity: the revision to make the text more formal, concise, readable, and understandable.

Style: the revision to convey the writer’s writing preferences, including emotions, tone, voice, etc.

Meaning-Changed: the revision to update or add new information to the text.

When you revise the text, please break down the revision process into multiple steps, and check if the revision fulfills the intention requirement in each step. Only output the final revised text. Do not output the intermediate steps.

Here is the example of the revision: {example}

Original text: {original text}

Revised text:

Table 12: The in-context learning with chain-of-thought prompt regarding five intentions in ITERATER.

revisions. Specifically, the original texts in the test and training sets are encoded with Sentence Transformer (Thakur et al., 2021). The top similar example is retrieved using Faiss (Johnson et al., 2019). The retrieved original text paired with its revised text is used as a one-shot example. In addition, we use a chain-of-thought prompt as another baseline, which adds the text “*when you revise the text, please break down the revision process into multiple steps, and check if the revision fulfills the intention requirement in each step. Only output the final revised text. Do not output the intermediate steps*” to the prompt. The detailed prompts are shown in Tables 9 and 10, and Tables 11 and 12 for ArgRevision and ITERATER datasets, respectively.

A.5 Hyperparameters and Metrics

We implement the framework using Python v3.9.20, PyTorch v2.5.1, and HuggingFace Transformers

Hyperparameter	ITERATER-sent	ITERATER-doc	ArgRevision
rank r		32	
α		64	
dropout		0.05	
where		{Q, K, V, Up, Down}	
fine-tuning batch size	16	4	2
inference batch size	32	8	4
max length	256	1024	1536
max new tokens	128	512	768
optimizer		AdamW	
learning rate		2e-4	
lr scheduler		Warmup Steps	
weight decay		0	
warmup steps		100	
epochs		2	
sample		true	
beam num		4	
top p		0.75	
top k		40	
temperature		0.2	

Table 13: Hyperparameters for PEFT with Intention-Tuning on ITERATER-sent, ITERATER-doc, and ArgRevision.

v4.46.0, and utilize Weights & Biases v0.19.1 for logging. The used fine-tuning hyperparameters⁴ are shown in Table 13. We use SARI⁵, GLEU⁶, and Update-R⁶ as revision generation metrics. Here, SARI and Update-R are crucial metrics for text revision, given that the original and revised texts often exhibit high overlap in words. We use evaluation packages adopted from prior work (Dwivedi-Yu et al., 2024; Xu et al., 2016).

A.6 Layer Alignment

We visualize the LLM layer alignment across three LLMs in Figure 6 and Figure 7, specifically for the multi-intent revision datasets. Despite architectural differences, the important and redundant layers show observable alignment between the intention prediction and the revision generation tasks. Regarding ITERATER-doc, the dark green and dark brown boxes coincide, especially for layers 2 and 32 in Mistral-7B and Llama3.1-8B, and layers 5 to 12 and 28 to 31 in Qwen2.5-14B. In terms of ArgRevision, the alignments are relatively higher on Llama3.1-8B than Mistral-7B and Qwen2.5-14B (see Table 5). Although layer alignment varies across LLMs, the important layers used for revision generation remain consistent across datasets, e.g., layers 5 to 31 in Qwen2.5-14B. This observation highlights that specific LLM layers are consistently important in revision generation tasks across dif-

⁴The values of the *max length* parameter are doubled in the ICL and CoT baselines (i.e., setting *max length* as 512, 2,048, 3,072) since their prompts include one-shot examples that are long documents exceeding the limits in PEFT settings.

⁵<https://github.com/cocoxu/simplification>

⁶<https://github.com/facebookresearch/EditEval>

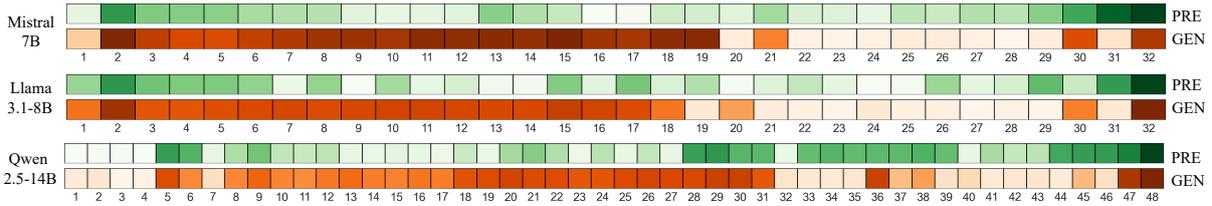


Figure 6: The layer-wise importance score alignment between the intention prediction (green boxes) and the revision generation (brown boxes) tasks while fine-tuning LLMs on the ITERATER-doc. All PEFT uses LoRA. The dark colors indicate high scores (important layers) while the light colors indicate low scores (redundant layers).

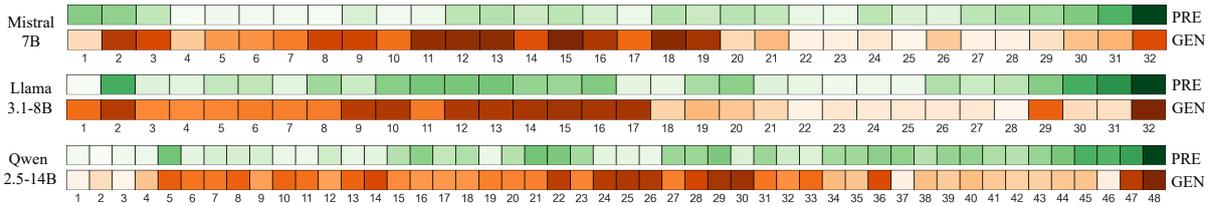


Figure 7: The layer-wise importance score alignment between the intention prediction (green boxes) and the revision generation (brown boxes) tasks while fine-tuning LLMs on the ArgRevision. All PEFT uses LoRA. The dark colors indicate high scores (important layers) while the light colors indicate low scores (redundant layers).

ferent datasets; however, this does not hold true for intention prediction, which might be because LLMs are largely pre-trained for generation rather than classification. This observation will be evaluated in future work.

A.7 Convergence on Adapters

We conduct additional experiments to visualize the loss convergence of Llama3.1-8B while fine-tuning with DoRA (Liu et al., 2024b) and Bottleneck (Houlsby et al., 2019) adapters. Figure 8 and 9 show that Intention-Tuning has fast convergence across different adapters and datasets. Although the loss convergence in IR-Baseline is more rapid than Intention-Tuning, they are close. These observations suggest that Intention-Tuning is generally efficient during fine-tuning.

A.8 Error Analysis

We analyze multi-intent revision generation on the ITERATER-doc and ArgRevision datasets. Figure 10 shows *clarity* and *fluency* revisions mostly involve wording and surface issues. The human replaces “was broken” with “broke”, and “closed” with “shut”; however, the LLM does not make these revisions. Instead, it corrects the typo “facory” with “factory”, and replaces “causing” with “caused” to improve *clarity*, which is in line with human edits. This example illustrates that LLMs can generate desirable revisions to align with the writer’s intentions, but have limited capability to identify the text spans that have such issues, resulting in minimal revisions. Similarly, another example shows that

human involves more content-level revisions, e.g., adding additional information to indicate the exhibition “opened on May 15”, and revises sentence orders to improve *coherence*. In contrast, the LLM is limited in handling such revisions. Although LLMs with Intention-Tuning achieve relatively higher performance than the baselines, they are more conservative than humans in making revisions, which could contribute to counter-intuitive high GLEU and Average scores, similar to the Copy-Baseline performance in Table 4.

Figure 11 shows argument revisions in ArgRevision. The human adds relevant examples to explain that the reason why they need “funding space exploration” is not only because it helps the “scientist,” but also because it “helps others.” Although the LLM revision provides evidence to support the claim, its content does not align with the human revision. These observations illustrate the challenges in argument revision generation, which requires providing relevant evidence linked to the main argument. In addition, another example shows the LLM makes minimal changes, despite being requested for *relevant*, *LCE*, and *not LCE* revisions. This again demonstrates that LLMs are conservative in providing challenging revisions, possibly because they struggle to identify the text spans (e.g., sentences that require revision) and are uncertain about revision actions (e.g., adding, deleting, or modifying sentences). These examples offer insights into enhancing revision quality by improving the identification of revision spans and actions in future work.

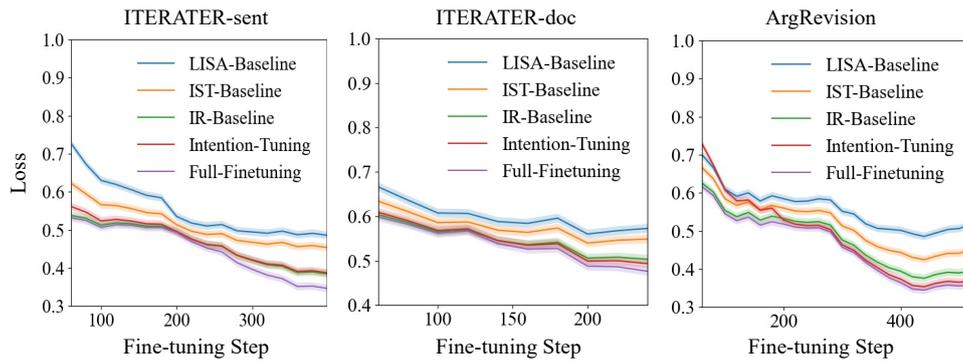


Figure 8: Llama3.1-8B fine-tuning loss on the training set of the three datasets. All PEFT uses DoRA.

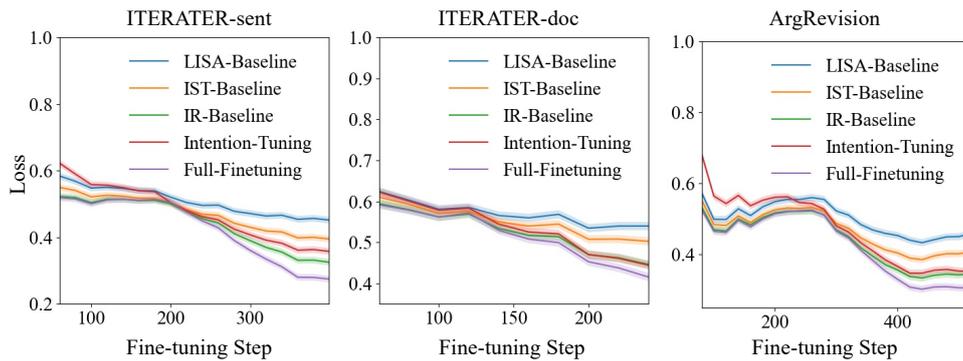


Figure 9: Llama3.1-8B fine-tuning loss on the training set of the three datasets. All PEFT uses Bottleneck.

[Clarity and Fluency] Human Revision:
 Nine people are confirmed dead in a fire in a leather bag ~~factory~~ factory in Kolkata's South 24 Pargana district in Topsia area. Eighteen people were seriously injured and admitted to ~~the~~ National Medical College and Hospital. Inspector General of Police (Law and Order) Raj Kanojia said that the fire ~~was broken broke~~ out at about a.m. ~~local time~~, when the victims were inside the factory and all the doors were closed. Director General of Fire Brigade Gopal Bhattacharjee said that it was an illegal factory in Kolkata located at third floor of the building. At the time of fire, all the ~~premises exits~~ were ~~closed shut~~ and workers could not come out from the factory. The cause of the fire ~~was is~~ yet to be ascertained. Kolkata Mayor Bikash Ranjan Bhattacharya ~~predicted suggested~~ that a short circuit or an unattended cigarette ~~causing may have caused the~~ fire ~~in the factory~~.

[Clarity and Fluency] LLM Revision:
 Nine people are confirmed dead in a fire in a leather bag ~~factory~~ factory in Kolkata's South 24 Pargana district in Topsia area. Eighteen people were seriously injured and admitted to ~~the~~ National Medical College and Hospital. Inspector General of Police (Law and Order) Raj Kanojia said that the fire was broken out at about 3 am when the victims were inside the factory and all the doors were closed. Director General of Fire Brigade Gopal Bhattacharjee said that it was an illegal factory in Kolkata located at third floor of the building. At the time of fire, all the premises were closed and workers could not come out from the factory. The cause of the fire was yet to be ascertained. Kolkata Mayor Bikash Ranjan Bhattacharya predicted ~~a~~ short circuit or an unattended cigarette ~~causing caused the~~ fire ~~in the factory~~.

[Clarity, Coherence, and Meaning-Changed] Human Revision:
 noicon]Launch of the Opening Ceremony. 2008 The 27th Young Designers' Exhibition, ~~opened on May 15 at~~ ~~famed as the largest show from students' creations, recognized by~~ International Council of Societies of Industrial Design (ICSID) since last year, ~~started at the~~ Taipei World Trade Center yesterday and will close at Sunday [May 18]. It features ~~with~~ ~~participations from 87 academical units~~ ~~participation by 87 academic groups~~ in Taiwan and 20 ~~units~~ ~~groups~~ from United States, United Kingdom, Italy, Netherlands, New Zealand, and Australia to showcase ~~varied~~ ~~various~~ achievements ~~from~~ in industrial design. It is ~~recognized by the International Council of Societies of Industrial Design (ICSID) as the largest show of student creations~~ ~~designing industry~~.

[Clarity, Coherence, and Meaning-Changed] LLM Revision:
 noicon]Launch of the Opening Ceremony. 2008 The 27th Young Designers' Exhibition, famed as the largest show from students' creations, recognized by International Council of Societies of Industrial Design (ICSID) since last year, started at Taipei World Trade Center yesterday and will close at Sunday (May 18) with participations from 87 ~~academical~~ academic units in Taiwan and 20 units from United States, United Kingdom, Italy, Netherlands, New Zealand, and Australia to showcase varied achievements from ~~the design industry~~ ~~designing industry~~.

Figure 10: The example of the human and generated revisions by Llama3.1-8B with Intention-Tuning on the ITERATER-doc dataset. The yellow and green colors denote additions, and the purple denotes deletions.

[Relevant] Human Revision:
 this story is very convincing about the topic, "the importance of space exploration". yes, the author convinced me that "space exploration leads to long-term benefits" justifies the cost by telling me that the money spent could be used differently to improve people's lives...we can use this information to improve the way we produce and distribute food. this will enable us to provide more food at a lower price to people who need it." when we explore space, we're not just doing it just for scientist, we are also helping to solve serious problems on earth. ~~funding space exploration doesn't just help scientist tho it helps others in need.~~ 46.2 million americans (15%) live in poverty. ~~funding space exploration can help those in need.~~

[Relevant] LLM Revision:
 this story is very convincing about the topic, "the importance of space exploration". yes, the author convinced me that "space exploration leads to long-term benefits" justifies the cost by telling me that the money spent could be used differently to improve people's lives... we can use this information to improve the way we produce and distribute food. this will enable us to provide more food at a lower price to people who need it." when we explore space, we're not just doing it just for scientist, we are also helping to solve serious problems on earth. ~~in the text it states,~~ ~~"even the problems of hunger and poverty can be tackled by space exploration. satellites that circle the earth can monitor land and the atmosphere. they can track and measure the conditions of crops, soil, and rainfall."~~

[Relevant, LCE, Not LCE] Human Revision:
 no, because 20,000 people die from malaria every day. they only make 1\$ a day and the bed nets are 5\$. if they ~~can't~~ ~~can't~~ afford that than there is no way they can do it in a lifetime ~~because if they buy just one bed net that is five days of work spent, and there are a lot of other things that humans need to survive, for example water, food shelter, clothing and so much more.~~ they can't afford school supplies either and the school's themselves don't have enough books and supplies for everyone. ~~if they can't afford that then they won't be able to have any sports teams, so no one will be able to play sports. all of the kids run around with minimum clothing and bare feet, it can be very hard to run around with bare feet,~~ another reason is that hospitals are in very bad shape.

[Relevant, LCE, Not LCE] LLM Revision:
 no, because 20,000 people die from malaria every day. they only make 1\$ a day and the bed nets are 5\$. if they ~~can't~~ ~~can't~~ afford that than there is no way they can do it in a lifetime. ~~another reason is that~~ they can't afford school supplies either and the school's themselves don't have enough books and supplies for everyone. another reason is that hospitals are in very bad shape.

Figure 11: The example of the human and generated revisions by Llama3.1-8B with Intention-Tuning on the ArgRevision dataset. The yellow and green colors denote additions, and the purple denotes deletions.

ID	Sentence in original text	Sentence in revised text	Edit	Intention
1	I think that we should still help nasa but keep in mind that there are lots of kids and adults out there that need help.	I think that we should still help nasa but keep in mind that there are lots of kids and adults out there that need help.	N/A	N/A
2	
3		Nineteen billion was only 1.2% of the total national budget.	Add	Irrelevant
4	So i think that we should keep giving money but keep thinking of others.		Delete	not LCE
5		It is very important to keep thinking of others.	Add	LCE
6	Because it is not just people that need help.	because it is not just people that need help.	N/A	N/A
7	
8	A program to develop clean energy could be viewed as a worthy investment.	A program to develop clean energy could be viewed as a worthy investment.	N/A	N/A
9	Maybe exploring space should not be a priority when there is so much that needs to be done on earth.	Maybe exploring space should not be a priority when there is so much that needs to be done on earth.	N/A	N/A
10	In 2012, united states spent 19 billion dollars for space exploration.	In 2012, united states spent 19 billion dollars for space exploration.	N/A	N/A
11	Some people think that money should be spent to help heal the people and the earth.	Some people think that money should be spent to help heal the people and the earth.	N/A	N/A
12		The arguments against space exploration stem from a belief that the money spent could be used differently – to improve people’s lives.	Add	Relevant
13		In 1953, president eisenhower captured this viewpoint.	Add	Relevant
14	
15	Also nineteen billion was only 1.2% of the total national budget.		Delete	Irrelevant

Table 14: Example of revision intention annotation for an essay in ArgRevision.

ID	Sentence in original text	Sentence in revised text	Intention
1	Bidirectional Encoder Representations from Transformers (BERT) models for biomedical specialties such as BioBERT and clinical-BERT have significantly improved in biomedical text-mining tasks and enabled us to extract valuable information from biomedical literature .	Bidirectional Encoder Representations from Transformers (BERT) models for medical specialties, such as BioBERT and clinical-BERT have significantly improved in biomedical text-mining tasks and enabled us to extract valuable information from biomedical literature .	Style
2	We introduce the method to train a BERT model on a small medical corpus both in English and Japanese, respectively, and then we evaluate each of them in terms of the biomedical language understanding evaluation (BLUE) benchmark and the medical-document-classification task in Japanese, respectively.	We introduce the method to train a BERT model on a small medical corpus both in English and Japanese, respectively, and then we evaluate each of them in terms of the biomedical language understanding evaluation (BLUE) benchmark and the medical document classification task in Japanese, respectively.	Clarity
3	Therefore, we propose a method that realizes a high-performance BERT model by using a small corpus.	Therefore, we propose a method that realizes a high-performance BERT model using a small corpus.	Coherence
4	After confirming their satisfactory performances, we apply our method to develop a model that outperforms the pre-existing models.	After confirming their satisfactory performances, we applied our method to develop a model that outperforms the pre-existing models.	Fluency
5	The total score is 1.0 points above that of BioBERT .	The total score is 1.0 points above that of BioBERT and 0.3 points above that of the ablated model trained without our proposed method. This proposed technique is an effective approach to develop localized medical BERT models and to enhance domain-specific models in the biomedical domain .	Meaning-Changed

Table 15: Example of revision intention annotation for an article in ITERATER.