

AnchorAlign: A Novel Anchor Alignment-enhanced Generative Method for Joint Named Entity Recognition and Relation Extraction

Anonymous ACL submission

Abstract

Named Entity Recognition (NER) and Relation Extraction (RE) are two fundamental and interdependent tasks in information extraction (IE), aiming to identify entities and relations from unstructured text. Recently, generative methods have become mainstream instead of discriminative methods for IE, especially joint multi-task IE, due to their promising performance and flexibility. For joint NER and RE, existing methods suffer from misalignment between entities and relations, as well as misalignment among relations. To address these issues, we propose AnchorAlign, a novel generative method enhanced by anchor alignment. Specifically, we first introduce an anchor entity selection mechanism to identify key entities in the text as anchor points, which serve as semantic pivots to bridge the two tasks. Then, we design a dual-level anchor alignment module: at the semantic level, we construct a cross-task semantic alignment space to align the semantic representations of anchor entities and their associated relations; at the generation level, we introduce an anchor-guided generation constraint to guide the model to generate entities and relations with strict alignment based on the anchor points. Extensive experiments on five benchmark datasets show that AnchorAlign outperforms SOTA baselines, demonstrating its effectiveness. Our work provides a new perspective for optimizing the joint modeling of NER and RE, and has potential to be extended to more complex multi-task IE such as NER and Event Extraction (EE).

1 Introduction

Information Extraction (IE) is a fundamental technology for converting unstructured text into structured knowledge. Traditional approaches decompose IE into isolated sub-tasks such as Named Entity Recognition (NER), Relation Extraction (RE), and Event Extraction (EE), while recent advancements have shifted to joint extraction frameworks,

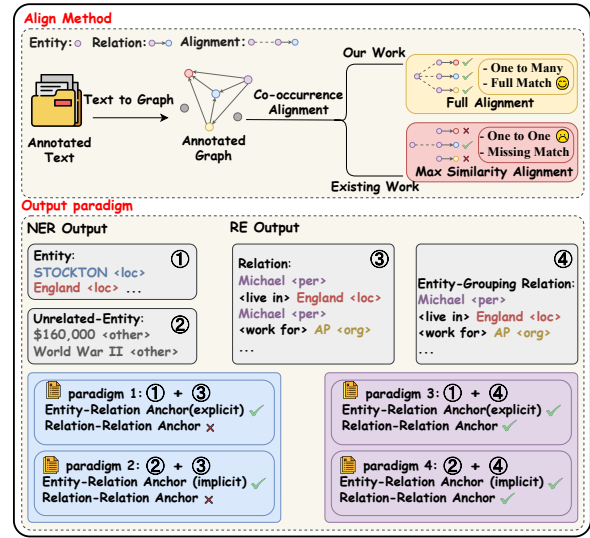


Figure 1: Comparative Analysis of Structured Output Paradigms and Alignment Mechanisms in generative methods for joint NER and RE.

particularly generative methods. These joint multi-task IE frameworks based on generative methods, including Universal Information Extraction (UIE) (Lu et al., 2022; Lou et al., 2023; Wang et al., 2023; Xiao et al., 2023; Li et al., 2024; Sainz et al., 2023), have demonstrated promising performance with high flexibility.

In the case of typical IE tasks, NER and RE, existing joint generative methods still have the following limitations in structured output paradigms and multi-task alignment.

First, although some effective structured output paradigms (e.g., paradigm 1 (Asada and Miwa, 2025) and 2 (Lu et al., 2022) in Figure 1) have been proposed for generative models, there is no study to compare and unify them.

Second, the alignment mechanisms in most existing generative methods only focus on inter-task dependencies (i.e., dependencies between NER and RE), ignoring intra-task dependencies (e.g., dependencies among entity relations). In the case of the

065 alignment mechanisms for dependencies between
066 NER and RE, auxiliary objectives are specially de-
067 signed to align entities in NER with entities in RE.
068 The existing alignment mechanisms (Zheng et al.,
069 2021; He and Tang, 2023) usually align one entity
070 in NER to the closet entity in one relation in RE no
071 matter how many relation the same entity in. They
072 may lead to some biases as shown in Figure 1.

073 To address the limitations mentioned above, we
074 propose AnchorAlign, a novel generative method
075 enhanced by anchor alignment with the following
076 advantages:

- 077 • We introduce anchor entity to unify the exist-
078 ing output paradigms of generative methods,
079 and two new paradigms (i.e., paradigms 3 and
080 4 in Figure 1) are derived them in the unified
081 way.
- 082 • Based on anchor entities, we design new se-
083 mantic alignment mechanism, which not only
084 considers intra-task (RE) dependencies, but
085 also avoids biases caused by inter-task (NER
086 and RE) dependencies.
- 087 • Extensive experiments on five benchmark
088 datasets demonstrate the effectiveness of our
089 proposed method, outperforming other state-
090 of-the-art (SOTA) methods for comparison.
091 The introduced new output paradigms as well
092 as semantic alignment mechanism bring sig-
093 nificant improvement.

094 2 Related Work

095 2.1 Generative Information Extraction

096 Traditional IE typically relies on discriminative
097 architectures to classify text spans into predefined
098 schemas, often necessitating complex, task-specific
099 engineering. In contrast, generative IE flattens
100 structured information into linear sequences for
101 direct output generation.

102 Generative paradigms have recently eclipsed
103 discriminative methods to become the dominant
104 framework for IE. The landscape of generative IE
105 is largely defined by three primary paradigms for
106 formulating the output, differentiated by how they
107 represent the target structured information:

- 108 • **Structured Language Generation** This
109 paradigm represents the inherent structure of
110 the information by using special symbolic to-
111 kens or by introducing new, semantically ini-
112 tialized token embeddings. This approach

113 is exemplified by models such as BART-
114 NER (Yan et al., 2021), REBEL (Cabot and
115 Navigli, 2021), UIE (Lu et al., 2022), and
116 USM (Lou et al., 2023).

- 117 • **Natural Language Generation** This
118 paradigm represents the target structured
119 information as a natural language sequence.
120 It typically leverages prompts, often including
121 extraction guidelines or annotations, to
122 guide a general-purpose Language Model
123 toward the correct task interpretation. This
124 method is employed by a wide range of
125 models, including TANL (Paolini et al.,
126 2021), DEEPSTRUCT (Wang et al., 2022),
127 GenIE (Josifoski et al., 2022), and Instruc-
128 tUIE (Wang et al., 2023).
- 129 • **Code Style Generation** This paradigm is
130 distinguished by its use of a Code-LLM as
131 the base model. It reframes the extraction
132 task by converting the schema into code con-
133 structs (e.g., classes) and represents the struc-
134 tured output using code-style syntax. Notable
135 examples in this category are CODEIE (Li
136 et al., 2023), GoLLIE (Sainz et al., 2023), and
137 KnowCoder (Li et al., 2024).

138 Our work builds upon the generative architecture
139 of BARTNER (Yan et al., 2021), extending it with
140 a structured generation language designed for both
141 NER and RE, thereby situating our approach within
142 the paradigm of structured language generation.

143 2.2 Joint Extraction

144 Joint extraction aims to unify multiple IE sub-tasks
145 within a single model. However, realizing in dis-
146 criminative architectures remains difficult due to
147 the rigid coupling of task-specific modules. Design-
148 ing specialized decoding structures for heteroge-
149 neous tasks often leads to combinatorial complexity
150 and conflicting optimization objectives, making it
151 challenging to extend these models beyond basic
152 tasks. Consequently, discriminative joint extrac-
153 tion research (Ye et al., 2022; He and Tang, 2023)
154 has predominantly converged on NER and RE, ex-
155 ploiting the strong correlation where NER acts as a
156 natural and tractable precursor to RE.

157 Generative architectures, conversely, offer a
158 more flexible paradigm by unifying any IE sub-
159 task into a text generation problem. This inher-
160 ent flexibility significantly lowers the barrier to
161 integrating diverse sub-tasks, which has led to the

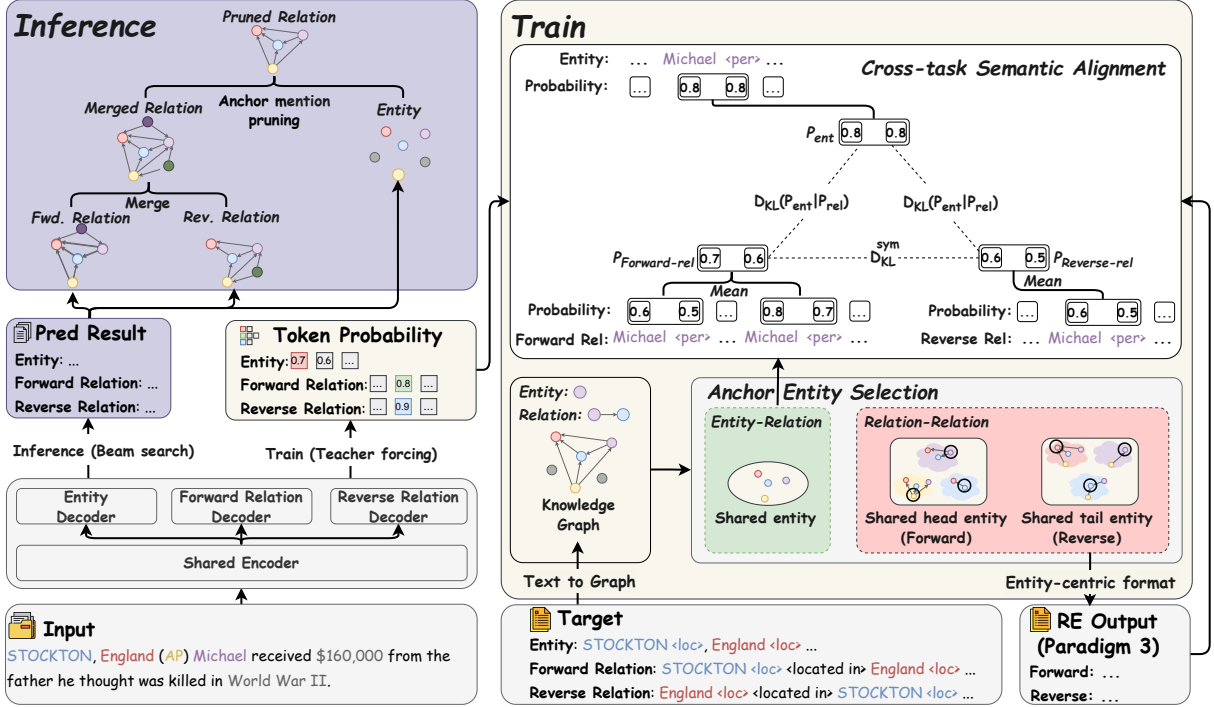


Figure 2: Overview of the AnchorAlign Framework

development of numerous Universal Information Extraction (UIE) models capable of concurrently handling NER, RE, and EE, such as UIE (Lu et al., 2022), InstructUIE (Wang et al., 2023), and KnowCoder (Li et al., 2024). Our work is situated within the generative paradigm for joint NER and RE. It yields significant performance gains on the relation extraction task, driven by a strategic enhancement of inter-task consistency.

3 Methodology

This section details the methodology of AnchorAlign. We first describe the Anchor Entity Selection Mechanism, followed by the model’s output paradigm. Finally, we introduce a dual-level anchor alignment mechanism, comprising cross-task semantic alignment for training and anchor-guided generation constraints for inference. These components are integrated into a unified framework. The overall architecture is illustrated in Figure 2.

3.1 Anchor Entity Selection Mechanism

Alignment in joint extraction encompasses two distinct scenarios: entity-relation alignment and relation-relation alignment.

Both alignment types are facilitated by key entities that link disparate task components. We define the entities shared between NER and RE as **anchors for entity-relation**, while entities shared

among multiple relations function as **anchors for relation-relation**.

In the following sections, we elaborate on the detailed designs of the output paradigm, training objectives, and inference strategies, all of which are centered around the aforementioned anchors.

3.2 Output Paradigm

AnchorAlign represents labels as unique special tokens (e.g., <PER>, <WORK FOR>) integrated into the vocabulary. To leverage semantic priors, these tokens are initialized by mean-pooling the embeddings of their original label text:

$$E_{\text{label}} = \text{Mean}(\text{TokenEmbed}(\text{text}_{\text{label}})) \quad (1)$$

This single-token format eliminates boundary markers (e.g., “()”) and shortens sequences, which are then generated by two specialized decoders:

NER Decoder This decoder generates a linearized sequence of entities, where each entity (e_i) is represented as a pair of its textual mention (m_i) and its corresponding single-token type ($\langle t_i \rangle$):

$$O_{\text{NER}} = \underbrace{(m_1, \langle t_1 \rangle)}_{e_1}, \dots, \underbrace{(m_n, \langle t_n \rangle)}_{e_n} \quad (2)$$

RE Decoder To align **anchors for relation-relation**, this decoder employs a subject-centric,

factorized format. All relations for a given subject entity (e_{s_i}) are grouped into a single block. The subject is generated once, followed by all its corresponding object-relation pairs ($e_{o_{i,j}}, \langle r_{i,j} \rangle$):

$$O_{RE} = \left(\underbrace{(e_{s_1}, e_{o_{1,1}}, \langle r_{1,1} \rangle, e_{o_{1,2}}, \langle r_{1,2} \rangle)}_{\text{Block for } e_{s_1}}, \underbrace{(e_{s_2}, \dots)}_{\text{Block for } e_{s_2 \sim n}} \right) \quad (3)$$

Similarly, the reverse RE decoder adopts an object-centric strategy.

This output paradigm simultaneously optimizes generation efficiency and ensures the alignment of **anchors for relation-relation**.

3.3 Cross-task Semantic Alignment

3.3.1 Positive Alignment

Our core design principle is to enforce alignment within the output probability space rather than the internal feature space. This strategy ensures consistency at the decision level while accommodating inherent disparities in internal representation capacities across different task components.

Positive Sampling Alignment comprises two components: NER-RE alignment and Forward-Reverse RE alignment. Both utilize Kullback-Leibler divergence (D_{KL}) as the objective.

NER-RE Alignment To facilitate the semantic alignment of **anchors for entity-relation**, we align both the forward ($d = \text{fwd-RE}$) and reverse ($d = \text{rev-RE}$) RE decoders with the NER target distributions in the output probability space:

$$\mathcal{L}_{NER,d} = \frac{1}{|\mathcal{E}_d|} \sum_{e \in \mathcal{E}_d} D_{KL}(p_{NER}(\cdot|e') \| p_d(\cdot|e)) \quad (4)$$

where \mathcal{E}_d is the set of entities from the target sequence of decoder d , $e' \in \mathcal{E}_{NER}$ is the corresponding entity for e , and $p_d(\cdot|e)$ is the predictive distribution from decoder d .

The final NER-RE alignment loss is the average of these two components:

$$\mathcal{L}_{NER-RE} = \frac{1}{2} (\mathcal{L}_{NER, \text{fwd-RE}} + \mathcal{L}_{NER, \text{rev-RE}}) \quad (5)$$

Forward-Reverse RE Alignment To mitigate biases inherent in the generation order and indirectly reinforce **anchors for relation-relation** alignment, we introduce a symmetric Kullback-Leibler (KL) divergence loss to enforce consistency between the forward and reverse RE decoders for each triplet:

$$D_{KL}^{\text{sym}}(p_1 \| p_2) = \frac{1}{2} (D_{KL}(p_1 \| p_2) + D_{KL}(p_2 \| p_1)) \quad (6)$$

The final loss is the average symmetric divergence over all ground-truth triplets $\tau \in \mathcal{T}$:

$$\mathcal{L}_{RE-Align} = \frac{1}{|\mathcal{T}|} \sum_{\tau \in \mathcal{T}} D_{KL}^{\text{sym}}(p_{\text{fwd}}(\tau) \| p_{\text{rev}}(\tau)) \quad (7)$$

where $p_{\text{fwd}}(\tau)$ and $p_{\text{rev}}(\tau)$ are the sequences of predictive distributions for the triplet τ .

3.3.2 Negative Suppression

Although positive alignment ensures consistency for positive samples, it lacks explicit penalties for valid but incorrect tokens. To address this, we introduce a negative suppression mechanism to complement the positive alignment by strictly discouraging non-target generation.

Formulation of the Negative Set Let \mathcal{V}_{NER} and \mathcal{V}_{RE} represent the ground-truth token sets. The negative set \mathcal{V}_{neg} is defined as the restricted vocabulary $\mathcal{V}_{\text{restricted}}$ (comprising tokens of input, label and special symbols) excluding these targets:

$$\mathcal{V}_{\text{neg}} = \{v \in \mathcal{V}_{\text{restricted}} \mid v \notin \mathcal{V}_{NER} \wedge v \notin \mathcal{V}_{RE}\} \quad (8)$$

Suppression Loss We apply the suppression objective independently to both RE decoders to minimize the probability mass assigned to the negative set \mathcal{N} . For a decoder $d \in \{\text{fwd-RE}, \text{rev-RE}\}$ with target sequence length T_d , the loss is formulated as:

$$\mathcal{L}_{\text{suppress}}^d = -\frac{1}{T_d} \sum_{t=1}^{T_d} \frac{1}{|\mathcal{N}|} \sum_{v \in \mathcal{N}} \log(1 - p_d(v|\mathbf{h}_t)) \quad (9)$$

where $p_d(v|\mathbf{h}_t)$ is the predicted probability of a negative token v at step t . The final suppression loss is the average over both directions:

$$\mathcal{L}_{\text{suppress}} = \frac{1}{2} (\mathcal{L}_{\text{suppress}}^{\text{fwd-RE}} + \mathcal{L}_{\text{suppress}}^{\text{rev-RE}}) \quad (10)$$

This objective effectively sharpens the predictive distribution by penalizing the generation of irrelevant tokens for the given instance.

3.3.3 Overall Training Objective

The overall training objective integrates standard generative extraction tasks with our proposed anchor alignment module. We first define the total generation loss, \mathcal{L}_{Gen} , as the aggregation of cross-entropy losses over the decoder set $\mathcal{D} = \{\text{NER}, \text{fwd-RE}, \text{rev-RE}\}$:

$$\mathcal{L}_{\text{Gen}} = \sum_{d \in \mathcal{D}} \mathcal{L}_{CE}^d \quad (11)$$

The final objective \mathcal{L} combines this generation loss with the auxiliary alignment and suppression terms:

$$\mathcal{L} = \mathcal{L}_{\text{Gen}} + \alpha \mathcal{L}_{\text{NER-RE}} + \beta \mathcal{L}_{\text{RE-Align}} + \gamma \mathcal{L}_{\text{suppress}} \quad (12)$$

where α, β, γ are hyperparameters weighing the NER-RE alignment, forward-reverse RE alignment, and negative suppression terms, respectively.

3.4 Anchor-guided Generation Constraint

While aggregating forward and reverse RE outputs maximizes recall, it inevitably introduces false positives. To mitigate this, we employ an inference-time **Anchor-guided constraint filter**.

Principle We designate mentions of **anchors for entity-relation** as generation constraints to guide the model to generate relations with strict alignment based on the anchor points.

Specifically, by treating the NER output as a high-confidence reference, we cross-validate the entity mentions within the predicted relation triplets. Any triplet containing an entity mention unrecognized by the NER module is classified as low-confidence and subsequently discarded.

Formulation Let $\hat{\mathcal{E}}_{\text{NER}}$, $\hat{\mathcal{T}}_{\text{fwd}}$, and $\hat{\mathcal{T}}_{\text{rev}}$ denote the prediction sets from the NER and the two RE decoders, respectively. We first form a candidate pool by aggregating the RE outputs: $\hat{\mathcal{T}}_{\text{cand}} = \hat{\mathcal{T}}_{\text{fwd}} \cup \hat{\mathcal{T}}_{\text{rev}}$. We then apply the generation constraint to derive the final set $\hat{\mathcal{T}}_{\text{final}}$:

$$\hat{\mathcal{T}}_{\text{final}} = \{(s, r, o) \in \hat{\mathcal{T}}_{\text{cand}} \mid \{s_{\text{mention}}, o_{\text{mention}}\} \subseteq \hat{\mathcal{E}}_{\text{NER}}\} \quad (13)$$

This constraint guarantees high precision while preserving recall gains.

3.5 Overall Training Framework

Drawing inspiration from prior works (Lu et al., 2022; Lou et al., 2023; Li et al., 2024), we adopt a two-stage training framework: task-adaptive pre-training followed by supervised fine-tuning (SFT). Specifically, the pre-training phase utilizes large-scale LLM-augmented corpora to equip the backbone with an understanding of structural paradigms and common IE capabilities.

3.5.1 Data Augmentation

Inspired by guidance-based prompting for NER (Kang et al., 2024), we devise a data

augmentation strategy tailored for joint NER and RE extraction. We employ DeepSeek-V3 to generate high-quality pseudo-labeled instances by prompting it with the original training exemplars. Specific details are provided in Appendix A.

3.5.2 Pre-training Stage

Following augmentation, the base model \mathcal{M} undergoes a two-step pre-training phase on the generated dataset to internalize joint extraction capabilities.

- 1. Entity Type Understanding** We replace entity mentions in the augmented data with type-specific mask tokens ($\langle \text{mask}_i \rangle$). Training on this masked data forces the model to learn context-type associations rather than memorizing surface forms.
- 2. Progressive Objective Activation** Training on the unmasked dataset proceeds in two sequential phases. Initially, standard NER and RE objectives are employed to secure fundamental capabilities. Once the model produces semantically valid outputs, the cross-task semantic alignment objectives are activated to explicitly inject cross-task consistency.

Upon completion, the pre-trained backbone is termed \mathcal{M}' .

3.5.3 Supervised Fine-Tuning (SFT) Stage

The SFT stage adapts \mathcal{M}' to the downstream ground-truth dataset $\mathcal{D}_{\text{train}}$ using a curriculum learning paradigm. We partition $\mathcal{D}_{\text{train}}$ based on triplet density, defining a simple-sample subset $\mathcal{D}_{\text{simple}}$ with a triplet count threshold K :

$$\mathcal{D}_{\text{simple}} = \{x \in \mathcal{D}_{\text{train}} \mid |\mathcal{T}_x| \leq K\} \quad (14)$$

The optimization proceeds in two phases:

- 1. Curriculum Warm-up.** Training initiates on $\mathcal{D}_{\text{simple}}$ using only standard generation objectives. This facilitates smooth adaptation to the target distribution’s easy samples.
- 2. Full Data Training.** We then transition to the complete dataset $\mathcal{D}_{\text{train}}$. In this phase, we rigorously apply the **Progressive Objective Activation** protocol (Sec. 3.5.2).

4 EXPERIMENTS

4.1 Experiment Setups

Datasets. We conduct experiments on five widely used benchmark datasets, including

ACE05 (Walker and Consortium, 2005), SciERC (Luan et al., 2018), ADE (Gurulingappa et al., 2012), NYT (Riedel et al., 2010), and Text2DT (Zhu et al., 2023).

Evaluation Metrics. Following prior works, we adopt two metrics to evaluate RE performance:

1. **Relation Strict F1.** A predicted relation is correct if the mentions and types of both the subject and object entities, as well as the relation type, are all correctly identified.
2. **Relation Boundary F1.** A predicted relation is correct if the mentions of both the subject and object entities and the relation type are correctly identified, regardless of the correctness of entity types.

Hyperparameter Configuration. On each dataset, we manually tuned hyper-parameters and retained the configuration that yielded the highest relation F1 on the development set (10-fold cross-validation for ADE). Please refer to Appendix D for the detailed settings.

4.2 BASELINES

We select several SOTA models for relation extraction, which can be categorized into generative and discriminative paradigms.

Generative Models. We compare our method against the following SOTA generative baselines: **UIE** (Lu et al., 2022) unifies various IE tasks by reformulating them into a text-to-structure generation format. **REBEL** (Cabot and Navigli, 2021) treats relation extraction as a sequence-to-sequence task, fine-tuned on the BART-Large architecture. **YAYI-UIE** (Xiao et al., 2023) is an end-to-end model based on Baichuan2-13B that enhances performance by co-training on both dialogue and IE data. **InstructUIE** (Wang et al., 2023) fine-tunes FlanT5-11B using natural language instructions to guide the generation process. Finally, **Know-Coder** (Li et al., 2024), based on LLaMA2-7B, reframes IE tasks as a Python code generation problem within a two-stage fine-tuning framework.

Discriminative Models. We compare against competitive discriminative baselines: **PL-Marker** (Ye et al., 2022) models relations between spans via a packed levitated marker strategy. **BiSPN** (He and Tang, 2023) employs set prediction for parallel entity and relation extraction to

mitigate error propagation. **USM** (Lou et al., 2023) decomposes IE into token-token and label-token linking subtasks within a unified pipeline.

4.3 Main Results

Table 1 summarizes the comprehensive performance comparison across five public datasets, reporting both Relation Strict F1 and Relation Boundary F1.

In terms of Relation Strict F1, AnchorAlign achieves new SOTA results on SciERC and ADE, surpassing previous bests by **0.61%** and **0.14%**, respectively. On ACE05, it remains highly competitive, ranking second only to the discriminative model USM by a narrow margin of **0.65%**.

Regarding Relation Boundary F1, our model establishes new SOTA performance on four out of five datasets, achieving substantial gains on ACE05 (**+0.72%**), SciERC (**+1.00%**), ADE (**+0.88%**), and Text2DT (**+0.12%**). Although slightly trailing the discriminative baseline USM on NYT, AnchorAlign surpasses the prior SOTA average by **0.46%** overall.

Two baseline discrepancies warrant notice: (1) PL-Marker’s single-label constraint is incompatible with NYT’s multi-label nature (allowing multiple relations per entity pair). (2) Second, UIE’s lower score on Text2DT stems from the limited capacity of the `uie-char-small` variant compared to its larger English counterparts, likely constrains its representation capacity for this task.

These results demonstrate that AnchorAlign effectively surpasses prior methods, particularly in comparison to generative peers, validating the efficacy of our structural design.

4.4 Output Paradigm Design Comparison

As illustrated in Figure 1, we propose two output configurations for NER and RE tasks, respectively:

- **NER Outputs (Circles 1 & 2):** Circle 1 generates a full concatenation of all entities, while Circle 2 provides a simplified version containing only non-relational entities.
- **RE Outputs (Circles 3 & 4):** Circle 3 outputs a concatenation of all relation triplets. Circle 4 implements an entity-centric factorization, where each unique subject/object is produced followed by all associated object/subject-relation pairs.

By combining these configurations, we derive four paradigms distinguished by anchor utilization:

Metric	Architecture	Method	ACE05	SciERC	ADE	NYT	Text2DT	AVG
Strict	Discriminative	BiSPN (He and Tang, 2023)	63.74	36.65	83.70	–	–	61.36
		PL-Marker (Ye et al., 2022)	62.80	36.93	83.20	–	–	60.98
		USM (Lou et al., 2023)	67.88	37.36	–	–	–	–
	Generative	UIE (Lu et al., 2022)	66.06	36.53	82.10	–	–	61.56
		REBEL (Cabot and Navigli, 2021)	–	–	82.21	–	–	–
	–	SoTa	67.88	37.36	83.70	–	–	62.98
Generative	Ours	67.23	37.97	83.84	–	–	63.01	
Boundary	Discriminative	BiSPN (He and Tang, 2023)	66.29	46.40	84.42	92.56	94.90	76.91
		PL-Marker (Ye et al., 2022)	65.72	47.81	84.23	80.64	93.30	74.34
		USM (Lou et al., 2023)	–	–	–	94.07	–	–
	Generative	UIE (Lu et al., 2022)	68.58	48.30	82.10	93.40	83.31	75.14
		REBEL (Cabot and Navigli, 2021)	–	–	–	91.96	–	–
		InstructUIE (Wang et al., 2023)	–	45.15	82.31	90.47	–	–
		YAYI-UIE (Xiao et al., 2023)	–	40.94	84.14	89.97	–	–
		KnowCoder (Li et al., 2024)	64.50	40.00	84.30	93.70	–	–
	–	SoTa	68.58	48.30	84.42	94.07	94.90	78.05
	Generative	Ours	69.30	49.30	85.30	93.64	95.02	78.51

Table 1: Performance comparison on five public datasets. We report **Relation Strict F1** and **Relation Boundary F1**. Consistent with prior research, we evaluate NYT and Text2dt solely on Relation Boundary F1.

- **Paradigms 1 and 2:** Paradigm 1 explicitly includes **anchors for entity-relation** within the NER output. Conversely, Paradigm 2 leverages them implicitly by decoupling them from the NER component. Notably, both paradigms lack a mechanism to leverage **anchors for relation-relation**.
- **Paradigms 3 (winner) and 4:** These extend the previous two by incorporating entity-centric factorization (Circle 4). By leveraging **anchors for relation-relation**, they effectively bridge relations that share common entities.

Detailed comparative experiments and analyses regarding the impact of shared entity density are provided in Appendix B.

4.5 Ablation Study

Table 2 presents an ablation study based on Relation Strict F1, distinguishing between training-related and inference-related components:

Training-Related Components. 1) w/o Pre-training: Removing the pre-training stage leads to a 1.17% decrease in F1. This demonstrates that pre-training effectively injects prior IE knowledge into the model and enhances its comprehension of structured output paradigms. **2) w/o Semantic Alignment:** Excluding the cross-task semantic alignment objective causes the most significant performance drop (2.00%), underscoring its criticality.

Appendix C visualizes its impact on co-occurring entity consistency by comparing constraint loss on ACE05 (*w/* vs. *w/o*).

Inference-Related Strategies. 1) w/o Decoder Aggregation: Removing the bidirectional merge mechanism causes a 1.66% decline. This indicates that the recall gain from recovering true positives outweighs the precision loss incurred by potential false positives, leading to an overall performance improvement. **2) w/o Generation Constraint:** The removal of the generation constraint results in a 1.72% drop. This confirms that the benefits of pruning false positives significantly outweigh the risk of erroneously filtering valid predictions. **3) w/o Aggregation & Constraint:** The simultaneous removal of both aggregation and filtering leads to the most severe degradation of 2.72%. Since these two strategies optimize performance from complementary directions (enhancing recall vs. precision) without conflict, their combined application achieves optimal results.

4.6 Case Study

Figure 3 illustrates two representative cases from the Text2DT and SciERC datasets.

Valid Case: The baseline suffers from a span nesting issue, predicting an erroneous entity that encapsulates the correct one. This ambiguity triggers a cascading error, leading to the inevitable misallocation of the corresponding tail entity. Although the model trained with semantic alignment

Input	Ours		Baseline	
	Entity	Relation	Entity	Relation
<p>Valid case</p> <p>原发性纤毛运动障碍综合征患者@急性肺部感染时, 需及时使用有效抗生素抗感染治疗; 对于肺部反复感染者, 可以使用化痰药物 (如N-乙酰半胱氨酸、福多司坦、羧甲司坦) 治疗</p>	<p>True Pred: (急性肺部感染, <症状>) (抗生素, <药物>) (福多司坦, <药物>) (N-乙酰半胱氨酸, <药物>) (化痰药物, <药物>) (羧甲司坦, <药物>) (原发性纤毛运动障碍综合征患者, <患者, <病患>) (肺部反复感染, <症状>)</p> <p>Wrong Pred: (抗感染治疗, <治疗>)</p> <p>Missing Relation: ()</p>	<p>True Pred: (原发性纤毛运动障碍综合征患者, <治疗药物>, 抗生素) (原发性纤毛运动障碍综合征患者, <治疗药物>, 福多司坦) (原发性纤毛运动障碍综合征患者, <治疗药物>, N-乙酰半胱氨酸) (原发性纤毛运动障碍综合征患者, <治疗药物>, 化痰药物) (原发性纤毛运动障碍综合征患者, <治疗药物>, 羧甲司坦) (原发性纤毛运动障碍综合征患者, <治疗药物>, 福多司坦) (原发性纤毛运动障碍综合征患者, <临床表现>, 肺部反复感染) (原发性纤毛运动障碍综合征患者, <临床表现>, 急性肺部感染)</p> <p>Wrong Pred: (抗感染治疗, <治疗>)</p> <p>Missing Relation: ()</p>	<p>True Pred: (急性肺部感染, <症状>) (抗生素, <药物>) (福多司坦, <药物>) (N-乙酰半胱氨酸, <药物>) (化痰药物, <药物>) (羧甲司坦, <药物>) (原发性纤毛运动障碍综合征患者, <患者, <病患>) (肺部反复感染, <症状>)</p> <p>Wrong Pred: (抗生素抗感染治疗, <治疗>)</p> <p>Missing Relation: ()</p>	<p>True Pred: (原发性纤毛运动障碍综合征患者, <治疗药物>, 福多司坦) (原发性纤毛运动障碍综合征患者, <治疗药物>, N-乙酰半胱氨酸) (原发性纤毛运动障碍综合征患者, <治疗药物>, 化痰药物) (原发性纤毛运动障碍综合征患者, <治疗药物>, 羧甲司坦) (原发性纤毛运动障碍综合征患者, <临床表现>, 肺部反复感染) (原发性纤毛运动障碍综合征患者, <临床表现>, 急性肺部感染)</p> <p>Wrong Pred: (原发性纤毛运动障碍综合征患者, <治疗方案>, 抗生素抗感染治疗)</p> <p>Missing Relation: (原发性纤毛运动障碍综合征患者, <治疗药物>, 抗生素)</p>
<p>Failure case</p> <p>It has also been studied in the framework of Japanese information extraction -LRB- -LSB- 3 -RSB- -RRB- in recent years .</p>	<p>True Pred: (It, <generic>) (Japanese information extraction, <task>)</p> <p>Wrong Pred: ()</p> <p>Missing Relation: ()</p>	<p>True Pred: ()</p> <p>Wrong Pred: (It, <used-for>, Japanese information extraction) (Japanese information extraction, <part-of>, It)</p> <p>Missing Relation: (Japanese information extraction, <used-for>, It)</p>	<p>True Pred: (It, <generic>) (Japanese information extraction, <task>)</p> <p>Wrong Pred: ()</p> <p>Missing Relation: ()</p>	<p>True Pred: ()</p> <p>Wrong Pred: (It, <used-for>, Japanese) (Japanese, <part-of>, It)</p> <p>Missing Relation: (Japanese information extraction, <used-for>, It)</p>

Figure 3: Cases from Text2DT and SciERC.

Method	Strict F1	Δ F1
AnchorAlign (Full Model)	67.23	-
<i>Impact of Training Components</i>		
w/o Pre-training	66.06	-1.17
w/o Semantic Alignment	65.23	-2.00
<i>Impact of Inference Strategies</i>		
w/o Generation Constraint	65.51	-1.72
w/o Decoder Aggregation	65.57	-1.66
w/o Aggregation & Constraint	64.51	-2.72

Table 2: Ablation study of different components on the ACE05 dataset. We report the Relation Strict F1 and the performance drop (Δ) when removing each component.

does not fully rectify the entity prediction error, it enforces clearer boundary distinctions between overlapping entities. Consequently, the tail entity is correctly identified during the relation extraction phase, demonstrating the efficacy of our method.

Failure Case: Although the proposed semantic alignment successfully rectifies the span boundary for “Japanese information extraction”, it fails to recover the relation triplet due to type inconsistency and head-tail reversal. This suggests that structural improvements alone cannot compensate for the lack of intrinsic domain knowledge needed for complex reasoning.

5 Conclusion

In this work, we systematically explore structured output paradigms and introduce AnchorAlign to address the critical misalignment in joint extraction. Specifically, we propose the concept of the anchor entity to unify existing output paradigms, which theoretically enables the derivation of two novel output paradigms. Building on this foundation, we design a semantic alignment mechanism that leverages anchor entities to rigorously model intra-task RE dependencies while effectively mitigating biases arising from inter-task (NER-RE) dependencies. Extensive empirical results across five benchmark datasets confirm that our approach yields significant improvements over competitive baselines, achieving new SOTA performance.

566 Limitations

567 While the semantic alignment objective effectively
568 sharpens entity boundaries and facilitates the cor-
569 rect assignment of relation arguments, it offers lim-
570 ited gains in rectifying the semantic correctness of
571 predicted entities, as evidenced by persistent en-
572 tity encapsulation errors. Furthermore, the model’s
573 performance remains constrained by its inherent
574 knowledge base. As demonstrated by the Text2DT
575 case, accurate entity extraction does not automati-
576 cally translate into accurate relation extraction. The
577 model continues to exhibit misalignments in rela-
578 tion classification and directionality judgment (e.g.,
579 head-tail reversal), indicating that alignment con-
580 straints cannot fully compensate for deficiencies in
581 the model’s reasoning knowledge. In future work,
582 we will attempt to incorporate more dependency
583 information among IE subtasks and explore new
584 avenues for injecting more comprehensive domain
585 knowledge to bolster the model’s reasoning capa-
586 bilities.

587 References

588 Masaki Asada and Makoto Miwa. 2025. [Improving](#)
589 [relation extraction by sequence-to-sequence-based](#)
590 [dependency parsing pre-training](#). In *Proceedings of*
591 *the 31st International Conference on Computational*
592 *Linguistics*, pages 7099–7105, Abu Dhabi, UAE. As-
593 sociation for Computational Linguistics.

594 Pere-Lluís Huguet Cabot and Roberto Navigli. 2021.
595 Rebel: Relation extraction by end-to-end language
596 generation. In *Findings of the association for compu-*
597 *tational linguistics: emnlp 2021*, pages 2370–2381.

598 Harsha Gurulingappa, Abdul Mateen Rajput, Angus
599 Roberts, Juliane Fluck, Martin Hofmann-Apitius, and
600 Luca Toldo. 2012. [Development of a benchmark](#)
601 [corpus to support the automatic extraction of drug-](#)
602 [related adverse effects from medical case reports](#).
603 *Journal of Biomedical Informatics*, 45(5):885–892.
604 Text Mining and Natural Language Processing in
605 Pharmacogenomics.

606 Yuxin He and Buzhou Tang. 2023. Bispn: Generating
607 entity set and relation set coherently in one pass. In
608 *Findings of the Association for Computational Lin-*
609 *guistics: EMNLP 2023*, pages 2066–2077.

610 Martin Josifoski, Nicola De Cao, Maxime Peyrard,
611 Fabio Petroni, and Robert West. 2022. [GenIE: Gen-](#)
612 [erative information extraction](#). In *Proceedings of*
613 *the 2022 Conference of the North American Chap-*
614 *ter of the Association for Computational Linguistics:*
615 *Human Language Technologies*, pages 4626–4643,
616 Seattle, United States. Association for Computational
617 Linguistics.

Hyeonseok Kang, Hyein Seo, Jeesu Jung, Sangkeun
Jung, Du-Seong Chang, and Riwoo Chung. 2024. [Guidance-based prompt data augmentation in spe-](#)
618 [cialized domains for named entity recognition](#). In
619 *Proceedings of the 62nd Annual Meeting of the As-*
620 *sociation for Computational Linguistics (Volume 2:*
621 *Short Papers)*, pages 665–672, Bangkok, Thailand.
622 Association for Computational Linguistics. 625

Peng Li, Tianxiang Sun, Qiong Tang, Hang Yan, Yuan-
bin Wu, Xuanjing Huang, and Xipeng Qiu. 2023. [CodeIE: Large code generation models are better](#)
626 [few-shot information extractors](#). In *Proceedings*
627 *of the 61st Annual Meeting of the Association for*
628 *Computational Linguistics (Volume 1: Long Papers)*,
629 pages 15339–15353, Toronto, Canada. Association
630 for Computational Linguistics. 633

Zixuan Li, Yutao Zeng, Yuxin Zuo, Weicheng Ren,
Wenxuan Liu, Miao Su, Yucan Guo, Yantao Liu,
Xiang Li, Zhilei Hu, and 1 others. 2024. Know-
634 coder: Coding structured knowledge into llms for
635 universal information extraction. *arXiv preprint*
636 *arXiv:2403.07969*. 639

Jie Lou, Yaojie Lu, Dai Dai, Wei Jia, Hongyu Lin, Xi-
anpei Han, Le Sun, and Hua Wu. 2023. Universal
640 information extraction as unified semantic matching.
641 In *Proceedings of the AAAI conference on Artificial*
642 *Intelligence*, volume 37, pages 13318–13326. 644

Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu
Lin, Xianpei Han, Le Sun, and Hua Wu. 2022. [Uni-](#)
645 [fied structure generation for universal information](#)
646 [extraction](#). In *Proceedings of the 60th Annual Meet-*
647 *ing of the Association for Computational Linguistics*
648 *(Volume 1: Long Papers)*, pages 5755–5772, Dublin,
649 Ireland. Association for Computational Linguistics. 651

Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh
Hajishirzi. 2018. [Multi-task identification of entities,](#)
652 [relations, and coreference for scientific knowledge](#)
653 [graph construction](#). In *Proceedings of the 2018 Con-*
654 *ference on Empirical Methods in Natural Language*
655 *Processing*. 657

Giovanni Paolini, Ben Athiwaratkun, Jason Krone,
Jie Ma, Alessandro Achille, Rishita Anubhai, Ci-
cero Nogueira dos Santos, Bing Xiang, and Stefano
658 Soatto. 2021. Structured prediction as translation be-
659 tween augmented natural languages. *arXiv preprint*
660 *arXiv:2101.05779*. 663

Sebastian Riedel, Limin Yao, and Andrew McCallum.
2010. [Modeling Relations and Their Mentions with-](#)
664 [out Labeled Text](#), page 148–163. 666

Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri,
Oier Lopez de Lacalle, German Rigau, and Eneko
667 Agirre. 2023. Gollie: Annotation guidelines improve
668 zero-shot information-extraction. *arXiv preprint*
669 *arXiv:2310.03668*. 671

C. Walker and Linguistic Data Consortium. 2005. [ACE](#)
672 [2005 Multilingual Training Corpus](#). LDC corpora.
673 Linguistic Data Consortium. 674

Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, and Dawn Song. 2022. [DeepStruct: Pre-training of language models for structure prediction](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 803–823, Dublin, Ireland. Association for Computational Linguistics.

Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, and 1 others. 2023. [Instructuie: Multi-task instruction tuning for unified information extraction](#). *arXiv preprint arXiv:2304.08085*.

Xinglin Xiao, Yijie Wang, Nan Xu, Yuqi Wang, Hanxuan Yang, Minzheng Wang, Yin Luo, Lei Wang, Wenji Mao, and Daniel Zeng. 2023. [Yayi-uie: A chat-enhanced instruction tuning framework for universal information extraction](#). *arXiv preprint arXiv:2312.15548*.

Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. [A unified generative framework for various NER subtasks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5808–5822, Online. Association for Computational Linguistics.

Deming Ye, Yankai Lin, Peng Li, and Maosong Sun. 2022. [Packed levitated marker for entity and relation extraction](#). In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 4904–4917.

Hengyi Zheng, Rui Wen, Xi Chen, Yifan Yang, Yunyan Zhang, Ziheng Zhang, Ningyu Zhang, Bin Qin, Xu Ming, and Yefeng Zheng. 2021. [PRGC: Potential relation and global correspondence based joint relational triple extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6225–6235, Online. Association for Computational Linguistics.

Wei Zhu, Wenfeng Li, Xiaoling Wang, Wendi Ji, Yuanbin Wu, Jin Chen, Liang Chen, and Buzhou Tang. 2023. [Extracting Decision Trees from Medical Texts: An Overview of the Text2DT Track in CHIP2022](#), page 89–102.

A Data Augmentation

The data augmentation process is decomposed into two distinct phases: entity mention rewriting and context rewriting, as illustrated in Figures 7 and 8. The expanded corpus is subsequently synthesized by cross-matching the generated mention variants with the context variants. A critical component of this process is our index-based consistency constraint: we assign a unique index to each entity and strictly enforce that both the entity type associated

with a specific index and the existing relations between indexed entities remain invariant throughout the rewriting process. This mechanism ensures the automated generation of high-quality labeled data tailored for joint entity and relation extraction.

Finally, we performed data augmentation on each dataset, capped at a maximum of 10,000 samples. The augmented data was processed via an automated annotation script, retaining only valid instances to construct the final pre-training corpus. Table 3 illustrates the detailed statistics of the original versus augmented data counts.

Dataset	labels	Samples			Augment train
		Train	Val	Test	
ACE05	5	10051	2420	2050	9845
ADE(10-fold avg.)	1	3845	427	-	9782
NYT	24	56196	5000	5000	9767
SciERC	7	1861	275	551	9658
Text2DT	7	400	100	-	3144

Table 3: Statistics of Relation Extraction (RE) Datasets

B Output Paradigm Design Comparison

According to the comparative results in Table 4, Paradigm 1 exhibits the lowest average ranking. Conversely, Paradigm 3 outperforms other competitive variants (Paradigms 2 and 4) by offering a more robust balance of stability and accuracy. Driven by these observations, we prioritize Paradigm 3 as our default paradigm.

Viewing Paradigm 3 as an optimized variant of Paradigm 1, we further investigate the correlation between its performance gains and the density of co-occurring entities in the datasets (detailed in Table 5). We measure the density of shared entities using both per-sample and per-entity metrics, which yield consistent dataset rankings.

A distinct positive correlation emerges: datasets with higher shared density correspond to greater F1 improvements for Paradigm 3 over Paradigm 1. This finding validates our design rationale, confirming that the entity-grouping strategy effectively reduces output redundancy, particularly in scenarios with high shared density.

C Visual Validation of Semantic Alignment Effectiveness

To visually demonstrate the enhancement in anchor consistency brought by cross-task semantic alignment training, we compare the consistency loss trajectories of models trained with ($w/$) and

Output Paradigm	ACE05	SciERC	ADE
Paradigm1	66.95	47.22	84.86
Paradigm2	68.76	47.24	85.10
Paradigm3	67.05	47.62	85.24
Paradigm4	67.21	47.01	85.13

Table 4: Relation Boundary F1 performance comparison of different output paradigms on three public datasets.

Metric	ACE05	SciERC	ADE
<i>Sample-level Shared Density</i>			
Shared Head Ratio	4.30%	21.59%	13.76%
Shared Tail Ratio	2.30%	23.74%	21.79%
Average Ratio	3.30%	22.67%	17.78%
<i>Entity-level Shared Density</i>			
Shared Head Ratio	20.89%	35.41%	27.65%
Shared Tail Ratio	11.11%	37.37%	42.02%
Average Ratio	16.00%	36.39%	34.84%
Δ F1 (vs. Paradigm 1)	+0.10	+0.40	+0.38

Table 5: Impact of Shared Entity Density (Sample- and Entity-level) on the Performance Gain of Paradigm 3 over Paradigm 1.

without (*w/o*) this objective. This comparative analysis is conducted across the training, validation, and test splits of the ACE05 dataset, providing concrete evidence that substantiates the efficacy of our proposed method.

We logged the step-wise loss values for both the baseline and the semantic alignment enhanced models throughout the training process. During post-processing, invalid entries—specifically samples devoid of relation annotations which yield zero values—were filtered out. To mitigate high-frequency stochastic fluctuations and highlight underlying optimization trends, we applied a Simple Moving Average (SMA) smoothing technique with a window size of $w = 7$.

As shown in Figure 4, since both models have achieved near-complete convergence on the seen data (training set), the disparity in consistency loss is negligible. However, a comparison of Figures 5 and 6 reveals that the model trained with semantic alignment exhibits significantly lower consistency loss on unseen samples (validation and test sets). This demonstrates that the semantic alignment objective promotes greater consistency and robustness in the outputs of the two tasks, effectively validating the generalization capability of our approach.

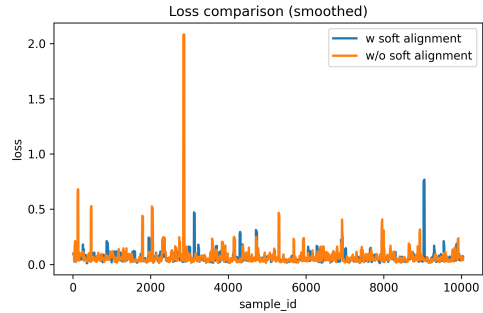


Figure 4: Semantic Alignment Loss (Train)

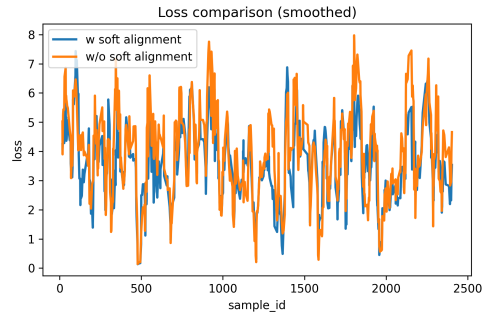


Figure 5: Semantic Alignment Loss (Dev)

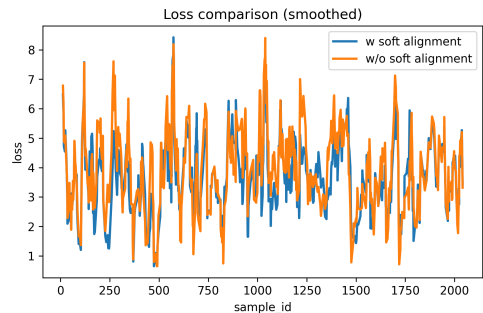


Figure 6: Semantic Alignment Loss (Test)

D Hyperparameter Configuration

For the pre-training stage, we report the exact epochs spent in **Entity Type Understanding** and **Progressive Objective Activation** (both without alignment and within alignment training) while keeping the learning rate fixed.

For the SFT stage, we list the epochs and corresponding learning rates used in **Curriculum Warm-up** and **Full Data Training** (again, both without alignment and within alignment training). The final hyper-parameter choices for every dataset are provided in Table 6.

Entity Mention Rewriting


Your task is to rewrite the entity mentions represented in dictionary format {"entity_type": entity_mention} within the input sentence, ensuring that the total number of entities remains unchanged while preserving each entity's original order and category. Only the entity mentions themselves should be rewritten—the rewritten mentions must align with the contextual environment and maintain consistent relationships with other entity mentions, while keeping all other parts of the sentence structure intact. Output each version on a separate line.

Input:
 Meanwhile, {"peop": "Shi Liming"} at the {"org": "Institute of Zoology"} of {"loc": "Kunming"} found that pandas lack variety in their protein heredity , which may serve as one of the major reasons for pandas 'near extinction .

Output:

1. Meanwhile, {"peop": "a leading researcher"} at the {"org": "Zoological Research Center"} of {"loc": "Yunnan"} found that pandas lack variety in their protein heredity, which may serve as one of the major reasons for pandas' near extinction.
2. Meanwhile, {"peop": "Dr. Shi"} at the {"org": "Chinese Academy of Sciences"} of {"loc": "Kunming City"} found that pandas lack variety in their protein heredity, which may serve as one of the major reasons for pandas' near extinction.
3. Meanwhile, {"peop": "a scientist"} at the {"org": "Wildlife Conservation Institute"} of {"loc": "southwestern China"} found that pandas lack variety in their protein heredity, which may serve as one of the major reasons for pandas' near extinction.

Input:
 "It is my sincere hope that we will reignite our united purpose , " Senator {"person": "Mitch McConnell"} of {"location": "Kentucky"} , the Republican whip , said on the Senate floor before the ceremony .

 1. "It is my sincere hope that we will reignite our united purpose," Senator {"person": "the GOP leader"} of {"location": "the Bluegrass State"} , the Republican whip, said on the Senate floor before the ceremony.

2. "It is my sincere hope that we will reignite our united purpose," Senator {"person": "McConnell"} of {"location": "KY"} , the Republican whip, said on the Senate floor before the ceremony.

3. "It is my sincere hope that we will reignite our united purpose," Senator {"person": "the Senate Minority Leader"} of {"location": "his home state"} , the Republican whip, said on the Senate floor before the ceremony.

Figure 7: Entity mention rewriting

Context Rewriting


Your task is to rewrite the given sentence into three distinct versions, ensuring that each unique <mask_id> represents a masked entity. The rewritten versions must retain the same number of <mask_id> placeholders, with the contextual environment remaining consistent with the original <mask_id> while preserving the implied relationships between them. Output each version on a separate line.

Input:
 Meanwhile , <mask0> at the <mask1> of <mask2> found that pandas lack variety in their protein heredity , which may serve as one of the major reasons for pandas 'near extinction .

Output:

1. Meanwhile, <mask0> at the <mask1> of <mask2> discovered that pandas have limited protein heredity diversity, potentially contributing to their near-extinction status.
2. At the <mask1> of <mask2>, <mask0> found that pandas exhibit a lack of protein heredity variation, which could be a key factor in their near extinction.
3. Pandas' near extinction may be partly due to their low protein heredity variety, as revealed by <mask0> at the <mask1> of <mask2>.

Input:
 "It is my sincere hope that we will reignite our united purpose , " Senator <mask0> of <mask1> , the Republican whip , said on the Senate floor before the ceremony .

 1. "It is my sincere hope that we will reignite our united purpose," declared Senator <mask0> of <mask1>, the Republican whip, during a speech on the Senate floor prior to the ceremony.

2. Before the ceremony, Senator <mask0> of <mask1>, the Republican whip, expressed on the Senate floor, "It is my sincere hope that we will reignite our united purpose."

3. Speaking on the Senate floor ahead of the ceremony, Senator <mask0> of <mask1>, the Republican whip, stated, "It is my sincere hope that we will reignite our united purpose."

Figure 8: Context rewriting

Dataset	Epochs		Learning Rate ($\times 10^{-5}$)		Batch Size	
	Pre-training	SFT	Pre-training	SFT	Pre-training	SFT
ACE05	[4, 6, 10]	[10, 35, 55]	3	[2, 2, 2]	25	25
SciERC	[2, 3, 10]	[30, 70, 100]	3	[3, 3, 3]	25	25
ADE	[2, 3, 5]	[10, 35, 55]	2	[2, 2, 2]	30	30
NYT	[2, 3, 15]	[10, 35, 55]	3	[3, 3, 2]	25	25
Text2DT	[2, 3, 10]	[10, 70, 100]	3	[3, 3, 2]	10	15

Table 6: Detailed hyper-parameter settings across Pre-training and Supervised Fine-tuning (SFT).