

A Reasoning Paradigm for Named Entity Recognition

Hui Huang^{1,2}, Yanping Chen^{1,2*}, Ruizhang Huang^{1,2}, Chuan Lin^{1,2}, Yongbin Qin^{1,2*}

¹Text Computing & Cognitive Intelligence Engineering Research Center of National Education Ministry, College of Computer Science and Technology, Guizhou University, Guiyang, China

²Big Data Laboratory, Guizhou University, Guiyang, China
{gs.hhuang21,ypench,rzhuang,clin,ybqin}@gzu.edu.cn

Abstract

Generative LLMs typically improve Named Entity Recognition (NER) performance through instruction tuning. They excel at generating entities by semantic pattern matching but lack an explicit, verifiable reasoning mechanism. This “cognitive shortcutting” leads to suboptimal performance and weak generalization, especially in zero-shot and low-resource scenarios where reasoning from limited contextual cues is crucial. To address this issue, a reasoning framework is proposed for NER, which shifts the extraction paradigm from implicit pattern matching to explicit reasoning. This framework consists of three stages: Chain of Thought (CoT) generation, CoT tuning, and reasoning enhancement. First, a dataset annotated with NER-oriented CoTs is generated, which contain task-relevant reasoning chains. Then, they are used to tune the NER model to generate coherent rationales before deriving the final answer. Finally, a reasoning enhancement stage is implemented to optimize the reasoning process using a comprehensive reward signal. This stage ensures explicit and verifiable extractions. Experiments show that ReasoningNER demonstrates impressive cognitive ability in the NER task, achieving competitive performance. In zero-shot settings, it achieves SoTA performance, outperforming GPT-4 by 12.3 percentage points on the F1 score. Analytical results demonstrate its great potential to advance research in reasoning-oriented information extraction.

Code — <https://github.com/HuiResearch/ReasoningIE>

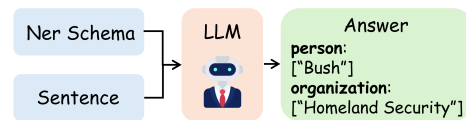
Extended version —

<https://www.arxiv.org/abs/2511.11978>

Introduction

Named entity recognition (NER) is a key task in information extraction, which identifies named entities from unstructured text. Because named entities are phrases that refer to objects of the world, precisely identifying them is very important for understanding sentence semantics. Traditionally, this task is implemented as a discriminative paradigm, which applies a supervised learning model to classify each token or span to indicate its semantic type. In this paradigm, discriminative language models are widely used to make full use of contextual features and semantic dependencies of

(a) Instruction-tuning based NER



(b) ReasoningNER

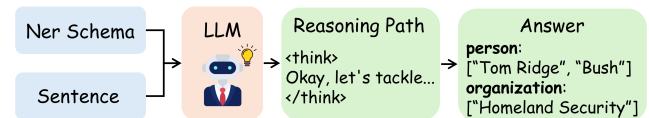


Figure 1: Comparison between Matching and Reasoning

a sentence. They usually make a prediction based on statistical models. Therefore, they are also known as “conditional models”, which are prone to overfitting the training data and often falter when applied to new domains or unseen entity types (Liu et al. 2021).

This limitation has spurred interest in another paradigm for NER: the generative paradigm. In this approach, generative LLMs are employed to leverage their broad knowledge, which is effective to advance the generalization ability of NER models, exhibiting impressive performance in zero-shot settings. This paradigm explores the use of in-context learning to guide models in performing the NER task (Wang et al. 2023a). However, as the model parameters are not specifically tuned for the NER task, their performance still falls significantly short of specialized models (Wang et al. 2023b). To address this gap, researchers focus on instruction-tuning LLMs. For example, InstructUIE (Wang et al. 2023b) fine-tunes a T5-11B model on a diverse collection of IE tasks using natural language prompts. A further approach incorporates structured prompts or coding-inspired schemas. For instance, KnowCoder (Li et al. 2024) represents extraction schemas in a code format and uses a two-phase learning process to imbue LLMs with structured knowledge.

Despite this progress, a fundamental limitation persists in both the discriminative and generative paradigms: existing NER models operate primarily through semantic pattern matching rather than explicit reasoning. The difference between the two paradigms is illustrated in Figure 1.

As shown in Figure 1(a), instruction-tuning prompts an

*Corresponding author.

LLM to directly generate entity labels in an opaque, single step. The NER task heavily depends on semantic patterns in the training dataset. The performance often degrades when these models are applied to unseen entities, especially in zero-shot and low-resource scenarios. In contrast, in the reasoning workflow, the extraction process is guided by an explicit, intermediate reasoning path, where a series of reasoning steps are designed to make use of contextual cues to facilitate the NER task. This paradigm incorporates an explicit reasoning mechanism before generating the final answer and can benefit from the vast knowledge and latent reasoning ability of generative LLMs.

Motivated by the above discussions, we propose a reasoning paradigm to support the NER task, named ReasoningNER. This paradigm is composed of three stages: CoT Generation, CoT Tuning, and Reasoning Enhancement. In the first stage, a NER-CoT dataset is constructed, which contains named entities annotated with step-by-step reasoning traces. In the second stage, the NER model is optimized with a CoT analysis, verbalizing how it identifies entities based on contextual clues before outputting the final answer. This paradigm shifts the task from direct prediction to a reasoned extraction process. In the third stage, a Group Relative Policy Optimization (GRPO) algorithm is employed to further refine the model’s policy. Two task-specific reward functions are designed to score the final predicted entities against the ground truth and to verify that the model’s output conforms to a predefined schema. The reasoning paradigm enables the extraction process to be guided by contextual clues, semantic hints, and logical constraints, rather than relying purely on semantic patterns.

Extensive experiments are conducted to evaluate the performance and generalization of ReasoningNER. The results show that ReasoningNER demonstrates impressive cognitive ability, attaining competitive performance on multiple NER evaluation datasets. In zero-shot and low-resource settings, ReasoningNER achieves state-of-the-art performance, outperforming GPT-4 by 12.3% in F1 score. It demonstrates consistent superiority across all target domains and showcases a more robust transfer capability to unseen entity types. Our main contributions are as follows:

- A reasoning paradigm is proposed to support NER that formulates the task as a reasoning process. It enables reasoning-oriented information extraction.
- A generative NER model (ReasoningNER) is constructed to evaluate the effectiveness of the reasoning paradigm, which shows the potential to support other NLP tasks.
- A NER-CoT corpus is built in this paper in which every entity is accompanied by an explicit CoT trace. The corpus is helpful to promote research in this field.

Related Work

Recent studies have significantly advanced the capabilities of LLMs on NER by reformulating the task within a generative, prompt-driven paradigm. Similar works like UIE (Lu et al. 2022), Flan-T5 (Chung et al. 2024), and InstructUIE (Wang et al. 2023b) have demonstrated that multi-task instruction tuning on a diverse set of IE tasks enables models

to generalize to novel tasks, achieving performance comparable to BERT-based models in specific scenarios. To further enhance generalization, subsequent research have explored methods for injecting explicit structured knowledge into these models. One line of work focuses on innovative schema representations. For instance, KnowCoder (Li et al. 2024) encodes entity types as Python class definitions, while GoLLIE (Sainz et al. 2024) directly trains the model to follow human-written annotation guidelines.

Another stream of research concentrates on dynamically providing contextual information at inference time rather than relying on the model’s static, internalized knowledge. RUIE (Liao et al. 2025) employs a retrieval-augmented mechanism to furnish input samples with relevant demonstrations, whereas GPT-NER (Wang et al. 2023a) leverages few-shot prompting and a self-verification step to guide the model. Although these approaches have achieved considerable success in improving the flexibility of NER, their mechanism remains a direct mapping from instructions to labels. They lack an explicit reasoning process, which limits their reliability in complex and ambiguous scenarios that demand deep contextual understanding or logical inference.

CoT reasoning, which prompts LLMs to generate a series of intermediate steps to solve complex problems, unlocks the potential of LLMs in tasks such as arithmetic and commonsense reasoning (Wei et al. 2022; Wang et al. 2022; Ma, Li, and Zhang 2023; Agarwal, Joshi, and Rojkova 2025). The central idea is to transform the implicit, single-step “black-box” prediction into an explicit, decomposable reasoning trajectory, thereby enhancing model accuracy and interpretability. However, the application of CoT to IE remains nascent. A few pioneering studies have begun to integrate reasoning into IE. For example, PromptNER (Ashok and Lipton 2023) decomposes NER into “judgment” and “extraction” steps, offering initial validation for the efficacy of simple reasoning prompts. Similarly, ERA-CoT (Liu et al. 2024) addresses multi-entity relation problems by first extracting all entities and then applying CoT to reason over this structured information.

These studies suggest that guiding a model to “think” about its extraction decisions is beneficial. However, current models primarily depend on prompts at inference time. They fail to supervise the generation of a reasoning chain as a core part of the model’s functionality. In the information extraction field, there is a lack of a systematic framework for training models to generate and utilize a verifiable reasoning process. By explicitly supervising the reasoning steps, our proposed method combines the advantages of CoT with the generalization power of instruction tuning to achieve a more robust and transparent reasoning-driven approach.

Methodology

The architecture of the ReasoningNER model is composed of three stages: CoT Generation (CG), CoT Tuning (CT) and Reasoning Enhancement (RE). The architecture is illustrated in Figure 2. The CG stage constructs a semi-supervised NER-CoT dataset for training the model. In the CT stage, a supervised fine-tuning is applied to establish the model reasoning capability. The RE stage enhances the reasoning

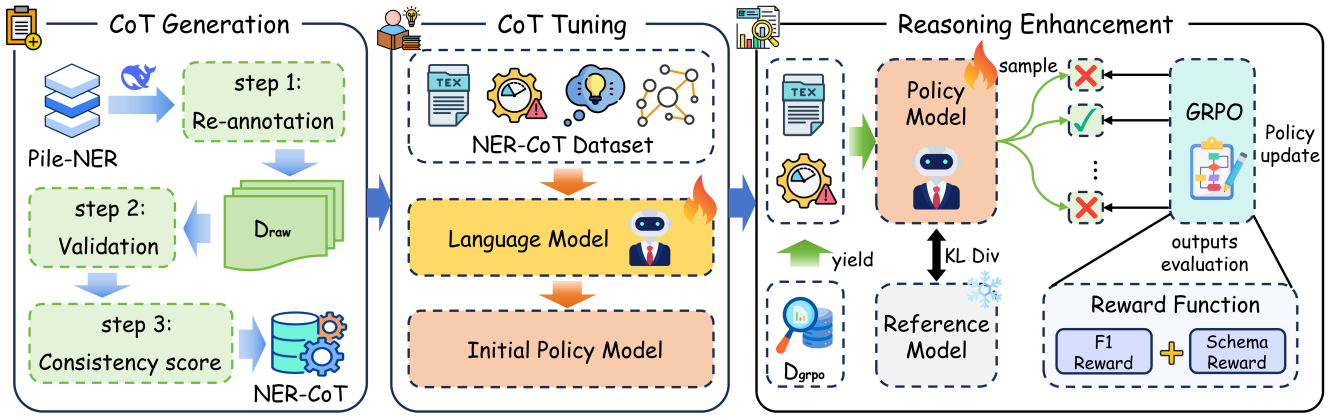


Figure 2: The architecture of the ReasoningNER model

ability through a reinforcement learning. In the sections that follow, we will first provide a problem formulation, then describe each stage in details.

Problem Formulation

Let $X = [w_1, w_2, \dots, w_n]$ be an input text composed of n tokens, $\mathcal{S} = \{\ell_1, \ell_2, \dots, \ell_m\}$ is an entity schema, where ℓ_i denotes to a predefined entity type.

All named entities in X can be refer as a set:

$$E = \{(s_i, \ell_j) | s_i \subset X, \ell_j \in \mathcal{S}\}, \quad (1)$$

where s_i is a contiguous subsequence $[w_p, \dots, w_q]$ ($1 \leq p \leq q \leq n$), $\ell_j \in \mathcal{S}$ is a entity type of s_i .

A CoT $\mathcal{C} = [c_1, c_2, \dots, c_L]$ is a depiction composed of L tokens, that provides the reasoning path for identifying and classifying entities in X .

The objective of reasoning NER is to produce an output $Y = (\mathcal{C}, E)$ that encompasses a CoT and entity set belonging to X .

CoT Generation

Before generating the NER-CoT dataset, a NER dataset should be adopted to produce not only extracted entities but also the explicit reasoning path for each sentence. In our work, the Pile-NER corpus (Zhou et al. 2023) is used, which contains a collection of sentences annotated with fine-grained entity schemas.

Based on this dataset, as illustrated in the CG stage on the left of Figure 2, the construction contains three steps: *Re-Annotation*, *Validation* and *Consistency*.

Re-annotation Step: Based on sentences and entity schemas in Pile-NER, a specialized prompt template is designed to guide a LLM (DeepSeek-R1 (Guo et al. 2025)) to produce not only entities but also explicit reasoning paths. This step yields an initial, raw collection of CoT annotations, denoted as \mathcal{D}_{raw} . Each instance in \mathcal{D}_{raw} is a tuple $(X, \mathcal{S}, \mathcal{C}_{\text{raw}}, E_{\text{raw}})$, where \mathcal{C}_{raw} is the raw reasoning chain and E_{raw} is the list of extracted entities.

Validation Step: because the outputs \mathcal{D}_{raw} may contain logical fallacies or structural errors, we propose a strict structural integrity to validate the initial outputs. In this validation step, a sample is retained only if its reasoning trace

\mathcal{C} explicitly justifies every entity in the list E , and the list E is well-formatted and compliant with the predefined entity schema \mathcal{S} . Samples failing these consistency checks are discarded.

Consistency Step: in this step, we move beyond structural checks to evaluate the semantic quality of the reasoning path. An auxiliary LLM (Qwen3 32B) is employed to assess the internal coherence, logical soundness, and factual accuracy of the reasoning process \mathcal{C} relevant to the input X and the extracted entities E . This evaluation yields a quantitative consistency score on a scale of 0 to 10. By setting a high threshold of 9, we ensure that only the most coherent and logically sound samples are kept.

This CoT generation process builds the final NER-CoT dataset, denoted as \mathcal{D}_{cot} , which is used for the subsequent CoT Tuning and Reasoning Enhancement stages.

CoT Tuning

In this stage, the NER-CoT dataset is used to increase the reasoning capability of a LLM.

Each sample in the NER-CoT dataset is a quadruple $(X, \mathcal{S}, \mathcal{C}, E)$, where X is the input sentence, \mathcal{S} is the entity schema, \mathcal{C} is the CoT reasoning trace, and E is the final list of entities. The output of the model is $Y = (\mathcal{C}, E)$. It can be represented as a token sequence: $Y = \mathcal{C} \oplus E = [y_1, y_2, \dots, y_T]$ formed by concatenating the reasoning trace \mathcal{C} and the entity list E .

The training objective is to maximize the conditional likelihood of this target sequence given the input sentence X and schema \mathcal{S} . This is achieved by minimizing the negative log-likelihood loss with respect to the model parameters θ :

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{(X, \mathcal{S}, Y) \sim \mathcal{D}_{\text{CoT}}} \left[\sum_{t=1}^T \log \pi_{\theta}(y_t | X, \mathcal{S}, y_{<t}) \right], \quad (2)$$

where π_{θ} denotes the language model with parameters θ , and y_t is the t -th token in the target sequence Y . This CoT Tuning process yields the Initial Policy Model, referred as π_{θ} , as shown in Figure 2.

Reasoning Enhancement

In this stage, the Initial policy model (π_θ) obtained from CoT Tuning stage is tuned by a GRPO algorithm (Guo et al. 2025). To prevent the model from deviating too far from the learned CoT behavior, a reference model (π_{ref}) is built to regularize the policy updates. We initialize both the policy model, which is the policy to be optimized, and the reference model with the weights of the initial model.

A dataset for GRPO training, denoted $\mathcal{D}_{\text{grpo}}$, is constructed by stratified sampling without replacement from the InstructUIE dataset. It generates a fixed-size subset of 4,703 samples. For each query $q = (X, \mathcal{S})$ in $\mathcal{D}_{\text{grpo}}$, a group of G candidate outputs $\{o_1, o_2, \dots, o_G\}$ is sampled from the old policy $\pi_{\theta_{\text{old}}}$. Each output o_i comprises a reasoning chain \mathcal{C}_i and a final list of entities E_i , i.e., $o_i = (\mathcal{C}_i, E_i)$. Subsequently, each output is evaluated based on a reward function.

As shown in Figure 2, generated outputs are sent to the ‘‘outputs evaluation’’ step and scored by a Reward Function. To guide the model towards optimal NER performance, a composite reward is designed to assess both task accuracy and schema adherence. First, a span-level F1 reward, R_{F1} , quantifies the match between the extracted entities and the ground truth. It is computed as the micro-averaged F1 score, linearly scaled to the range $[0, 1]$, providing a direct incentive for precision and recall. Second, a structural consistency reward, R_{schema} , evaluates adherence to the predefined output format and entity type constraints. Outputs with any schema violations (e.g., malformed structures or unpermitted entity types) receive $R_{\text{schema}} = 0$, whereas compliant outputs receive $R_{\text{schema}} = 1$.

The total reward $R(o_i)$ is a weighted combination:

$$R(o_i) = \lambda_{F1} R_{F1}(o_i) + \lambda_{\text{schema}} R_{\text{schema}}(o_i). \quad (3)$$

In our experiments, the coefficients λ_{F1} and λ_{schema} are set to 10 and 1, respectively, to prioritize entity correctness while strongly enforcing structural validity.

Finally, the computed rewards guide the GRPO Policy update. The average reward of the samples within the group is utilized as a baseline to compute the relative advantage. Let \bar{R} be a mean reward within the group, computed as:

$$\bar{R} = \frac{1}{G} \sum_{j=1}^G R(o_j). \quad (4)$$

Then, the advantage A_i for o_i is defined as:

$$A_i = R(o_i) - \bar{R}. \quad (5)$$

This advantage A_i is used for policy updates. The policy parameters θ are optimized by maximizing the following objective function:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q, \{o_i\} \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{G} \sum_{i=1}^G (\mathcal{L}_i - \beta D_{\text{KL}}) \right] \quad (6)$$

$$D_{\text{KL}} = \mathbb{D}_{\text{KL}}(\pi_\theta || \pi_{\text{ref}}) \quad (7)$$

$$\mathcal{L}_i = \min(r_i A_i, \text{clip}(r_i, 1 - \varepsilon, 1 + \varepsilon) A_i) \quad (8)$$

$$r_i = \frac{\pi_\theta(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}, \quad (9)$$

where $\pi_{\theta_{\text{old}}}$ is the old policy, π_θ is the current policy being optimized, and ε is a clipping hyperparameter. The \mathbb{D}_{KL} term represents KL divergence regularization, penalizing deviations of the current policy π_θ from a reference policy π_{ref} .

Experiments

Datasets

Our framework utilizes data across two training stages. First, for the CoT Tuning stage, we use our proposed NER-CoT dataset, which contains 45,787 training samples. Second, for the Reasoning Enhancement stage, we constructed a 4,703-sample dataset by applying stratified sampling to the InstructUIE (Wang et al. 2023b). This collection aggregates 20 diverse NER datasets. To ensure a balanced representation and prevent large-scale datasets from dominating the sampling proportions, we imposed a cap, limiting the size of any dataset to 10,000 when calculating its sampling ratio. The 4,703 samples were then drawn proportionally from all 20 datasets based on these adjusted, capped sizes.

To evaluate model performance, we use a suite of seven datasets. To assess out-of-domain generalization, we employ five datasets from the Cross-NER (Liu et al. 2021): artificial intelligence, literature, music, politics, and science. To evaluate in-domain performance, we use the MIT-Movie and MIT-Restaurant datasets (Liu et al. 2021). These two datasets are considered in-domain because their training data is part of the InstructUIE collection used in our RE sampling pool. Consistent with previous research (Wadden et al. 2019; Lin et al. 2020; Van Nguyen, Lai, and Nguyen 2021; Lu et al. 2022; Li et al. 2024), we use span-based Micro-F1 as the evaluation metric.

Implementation Details

For our experiments, the qwen3-8B-Base (Yang et al. 2025a) model is optimized with AdamW. The initial supervised fine-tuning is conducted for five epochs with a sequence length of 8192, a learning rate of 2×10^{-5} using a cosine scheduler, and a batch size of 256. Subsequently, we performed reinforcement learning using GRPO for one epoch, with a clipping threshold ε of 0.2 and a KL coefficient β of 0.04, maintaining an effective batch size of 384. To enhance computational efficiency, we integrated bfloat16 mixed-precision, gradient checkpointing, FlashAttention-2 (Dao 2023), and the Liger-kernel (Hsu et al. 2024).

Main Results

The efficacy of ReasoningNER is substantiated by its performance on the CrossNER benchmark, as detailed in Table 1. In the zero-shot cross-domain setting, ReasoningNER achieves new state-of-the-art results, with ReasoningNER 8B attaining an average F1 score of 72.4 across all seven domains. Notably, even our smaller ReasoningNER 1.7B model demonstrates competitive performance, achieving an average F1 of 65.1 and surpassing many larger models.

Comparison with Non-Reasoning models. In the cross-domain setting, ReasoningNER significantly outperforms all non-reasoning baselines. As shown in Table 1, ReasoningNER 8B achieves a new state-of-the-art average F1 score

Model	#Num	Movie.	Rest.	AI	Litera.	Music	Politics	Science	Avg
<i>Non Reasoning Model</i>									
ChatGPT-3.5(Ouyang et al. 2022)	-	5.3	32.8	52.4	39.8	66.6	68.5	67.0	47.5
UniNER 7B(Zhou et al. 2023)	45,884	42.4	31.7	53.5	59.4	65.0	60.8	61.1	53.4
UniNER 13B(Zhou et al. 2023)	45,884	48.7	36.2	54.2	60.9	64.5	61.4	63.5	55.6
InstructUIE 11B(Wang et al. 2023b)	225,930	-	-	48.4	48.8	54.4	49.9	49.4	-
GoLLIE 7B(Sainz et al. 2024)	165,163	63.0	43.4	59.1	62.7	67.8	57.2	55.5	58.4
GoLLIE 13B(Sainz et al. 2024)	165,163	62.5	49.8	56.7	59.7	65.5	54.4	56.2	57.8
KnowCoder 7B(Li et al. 2024)	4,590,403	50.0	48.2	60.3	61.1	70.0	72.2	59.1	60.1
GPT-4(Achiam et al. 2023)	-	60.4	59.7	50.0	55.2	69.2	63.4	63.2	60.1
GLiNER-L(Zaratiana et al. 2024)	44,889	57.2	42.9	57.2	64.4	69.6	72.6	62.6	60.9
B ² NER 7B♣(Yang et al. 2025b)	70,292	67.6	53.3	59.0	63.7	68.6	67.8	72.0	64.6
<i>Reasoning Model</i>									
deepseek-r1 8B(Guo et al. 2025)	-	60.1	50.6	47.6	57.4	53.3	58.1	55.3	54.6
deepseek-r1 32B(Guo et al. 2025)	-	70.4	57.5	60.4	52.3	70.4	71.1	65.9	64.0
qwen3 4B(Yang et al. 2025a)	-	69.6	53.4	59.2	56.8	62.3	68.4	65.4	62.1
qwen3 8B(Yang et al. 2025a)	-	70.1	57.4	61.2	58.0	71.0	71.9	68.6	65.4
ReasoningNER 1.7B(Ours)	50,787	70.2	52.4	63.6	59.1	71.6	68.6	69.9	65.1
ReasoningNER 8B(Ours)	50,787	76.3	56.8	71.0	69.4	78.7	78.8	75.8	72.4
ReasoningNER 8B♣(Ours)	50,787	79.3	67.7	72.2	77.1	84.0	79.8	81.4	77.3

Table 1: Out-of-domain evaluation results. “#Num” denotes the number of Information Extraction related samples used for model training, and “Avg” denotes the average performance across the seven datasets. ♣ denotes the results evaluated on the test set processed by B²NER(Yang et al. 2025b).

of 72.4 on the CrossNER benchmark, surpassing B²NER 7B by 12.7 points and the larger GPT-4 by 12.3 points. This highlights a fundamental limitation of conventional models. The origin of this issue lies in the instruction-tuning paradigm, where models are taught a direct mapping from input to final entity labels, omitting explicit analytical steps. This training objective incentivizes the model to find the shortest path to the answer, a behavior we identify as “cognitive shortcutting.” While efficient at recognizing patterns in training data, this shortcut is the primary cause of their brittle generalization, leading to precipitous performance degradation when faced with misleading surface-level patterns in unseen domains. In contrast, ReasoningNER leverages explicit reasoning guided by CoT to comprehend contextual logic, enabling it to accurately identify entities in novel domains and effectively overcome the limitations of methods that depend on direct input-output mappings.

Comparison with Reasoning Models. ReasoningNER 8B also demonstrates a clear advantage over models with innate reasoning capabilities, outperforming Qwen3 8B by 7.0 average F1 points. This validates the efficacy of our proposed methods, which imbue the base model with NER-specific reasoning abilities that are more potent than its native, general-purpose architecture. Moreover, ReasoningNER 8B surpasses the much larger DeepSeek-R1 32B by 8.4 F1 points, while our smaller 1.7B model substantially outperforms DeepSeek-R1 8B by 10.5 points. This proves our two-stage targeted optimization (SFT and GRPO) creates a critical task-specific alignment, enabling ReasoningNER to achieve superior performance over models that possess general reasoning abilities but lack focused optimization for re-

solving ambiguity in complex NER contexts.

Ablation Study

In this section, we analyze the impact of the NER-CoT dataset, the introduction of CoT during SFT, and the subsequent RE phase. As shown in Table 2, our analysis yields the following observations: (1) SFT on the Pile-NER dataset boosts the F1 score of the Qwen3-1.7B-Base model to 40.3, an increase of 8.4 points. This improvement underscores the foundational importance of supervised fine-tuning on task-specific data. (2) Notably, training solely on our constructed NER-CoT annotations, even without CoT reasoning, achieves a substantial 20.3-point F1 improvement over training on the standard Pile-NER dataset. This result highlights the superior quality of our reconstructed annotations. (3) Incorporating CoT reasoning during the SFT process further increases the F1 score by 2.4 points, confirming that explicitly modeling the reasoning process enables the model to capture more complex entity patterns. (4) Finally, the RE phase provides an additional 2.1-point F1 gain, demonstrating that fine-tuning the reasoning pathways with a task-aligned reward signal via GRPO effectively refines the model’s predictions and improves overall performance.

Analysis

Analysis of Supervised Evaluation

For in-domain supervised evaluation, we address the absence of CoT annotations across the 20 NER benchmarks by creating a mixed-inference training set (Yang et al. 2025a). We first construct a targeted subset by sampling 50 examples per category, along with 50 negative examples, from each

Model	F1
Qwen3-1.7B-Base	31.9
+ SFT on Pile-NER	40.3
+ SFT on NER-CoT (w/o cot)	60.6
+ SFT on NER-CoT (w/ cot)	63.0
+ SFT on NER-CoT (w/ cot) & RE	65.1

Table 2: Ablation study showing the average F1 score across the CrossNER and MIT datasets. We start with the Qwen3-1.7B-Base model and incrementally add: (1) SFT on Pile-NER; (2) SFT on our NER-CoT dataset annotations (w/o CoT); (3) SFT with CoT reasoning; and (4) the final Reasoning Enhancement (RE) phase.

of the 20 datasets. On this curated subset, we apply GRPO reinforcement learning for 3 epochs to train a specialized policy model capable of generating reasoning paths. This model performs inference on the full training set to generate synthetic CoT. Through a filtering process, we retain these reasoning paths only for samples where the model’s predictions are verifiably correct, yielding a high-fidelity, CoT-augmented dataset. This is then merged with the remaining non-CoT samples to form a hybrid corpus, enabling our final ReasoningNER model to be fine-tuned with a combination of explicit reasoning and standard supervision.

The results of this strategy, presented in Table 3, demonstrate its efficacy. Our ReasoningNER model achieves a new state-of-the-art average F1 score of 85.2, outperforming strong baselines like B²NER by 1.3 percentage points and securing top performance on 11 of the 20 datasets. Remarkably, this significant improvement is achieved even though the CoT-augmented data constitutes only a small fraction of the training corpus. This indicates that even a limited exposure to explicit reasoning provides a powerful inductive bias, enhancing the model’s learning dynamics and sample efficiency, while the larger set of standard examples ensures robust generalization. We plan to explore methods for creating more comprehensive CoT data in future work to unlock greater performance.

Analysis of Zero-Shot Performance We conducted a zero-shot performance evaluation on these same 20 datasets. The results, as shown in Table 4, are directly from the ReasoningNER 8B model that has only undergone the C-T stage. This model was not trained on any of these 20 datasets. ReasoningNER establishes a new SoTA in zero-shot NER, achieving an average F1 score of 56.8 across 20 benchmark datasets. This performance represents a significant leap over existing methods, outperforming the next-best model, GLiNER-L, by 9 F1 points and demonstrating a 11.1-point improvement over UniNER. The model’s consistently superior performance across all 20 diverse datasets, without any task-specific training, substantiates the efficacy of our proposed reasoning-centric learning paradigm. By first internalizing explicit reasoning steps through supervised fine-tuning on the NER-CoT corpus and subsequently refining these reasoning pathways via GRPO, ReasoningNER acquires a generalizable, procedural approach to entity

Dataset	InstructUIE	UniNER	GLiNER	B ² NER	Ours
ACE05	79.9	86.7	82.8	83.0	86.6
AnatEM	88.5	88.7	88.9	89.2	89.9
bc2gm	80.7	82.4	83.7	82.0	84.5
bc4chemd	87.6	89.2	87.9	89.0	91.6
bc5cdr	89.0	89.3	88.7	88.5	90.1
Broad Twitter	80.3	81.3	82.5	82.2	81.2
CoNLL03	91.5	93.3	92.6	92.6	94.4
FabNER	78.4	81.9	77.8	78.8	82.4
FindVehicle	87.6	98.3	95.7	97.9	98.4
GENIA	75.7	77.5	78.9	76.4	78.3
HarveyNER	74.7	74.2	68.6	73.7	74.2
Movie.	89.6	90.2	87.9	90.8	89.6
Restaurant.	82.6	82.4	83.6	83.7	81.4
MultiNERD	90.3	93.7	93.8	94.0	94.5
ncbi	86.2	87.0	87.8	84.8	87.9
OntoNotes	88.6	89.9	89.0	84.3	89.1
PolyglotNER	53.3	65.7	61.5	62.0	65.8
TweetNER7	66.0	65.8	54.4	66.3	65.8
WikiANN	64.5	84.9	83.7	85.1	85.9
wikiNeural	88.3	93.3	91.3	93.0	92.8
Average	81.2	84.8	82.9	83.9	85.2

Table 3: Results of in-domain supervised evaluation on 20 NER benchmarks.

identification. This methodology moves beyond superficial pattern matching, enabling the model to develop a more robust and adaptable understanding of entity concepts, which translates into its zero-shot generalization capabilities.

Analysis of Few-Shot Performance To evaluate the model’s sample efficiency in low-resource settings, we followed the experimental setup of UIE: the CoNLL03 training set was randomly downsampled to 1%, 5%, and 10% of its original size, and these data subsets were used to continue training the ReasoningNER 8B model. We compare it against two baseline models, UIE-base and KnowCoder-7B, with the results shown in Figure 3. Across all data proportions, the performance of ReasoningNER 8B surpasses both baseline models. This performance advantage is most significant in the most data-scarce setting (1%), where our model achieves an F1 score of 87.1%, outperforming UIE-base by 4.3 points and KnowCoder-7B by 7.9 points. When the amount of training data increases to 5% and 10%, ReasoningNER 8B maintains the best performance, achieving F1 scores of 91.0% and 92.9%, respectively. These results indicate that while large-scale instruction tuning contribute to improving model generalization, the fine-grained logical reasoning supervision employed in our model is crucial for maximizing sample efficiency in low-resource scenarios.

Analysis of the Impact of Backbone LLMs To ascertain that the superior performance of ReasoningNER is not contingent upon a specific LLM, we conducted a comprehensive evaluation across a suite of mainstream models with varying architectures and scales. For this analysis, all models were uniformly fine-tuned exclusively during the SFT phase, thereby circumventing the computationally intensive GRPO process. As detailed in Table 5, when trained on a mod-

Dataset	ChatGPT	UniNER	GLiNER-L	Ours
ACE05	26.6	36.9	27.3	33.8
AnatEM	30.7	25.1	33.3	49.5
bc2gm	40.2	46.2	47.9	43.3
bc4chemd	35.5	47.9	43.1	53.1
bc5cdr	52.4	68.0	66.4	74.8
Broad Twitter	61.8	67.9	61.2	76.0
CoNLL03	52.5	72.2	64.6	79.0
FabNER	15.3	24.8	23.6	26.0
FindVehicle	10.5	22.2	41.9	57.4
GENIA	41.6	54.1	55.5	58.9
HarveyNER	11.6	18.2	22.7	26.8
MIT Movie	5.3	42.4	57.2	76.0
MIT Restaurant	32.8	31.7	42.9	57.6
MultiNERD	58.1	59.3	59.7	65.5
ncbi	42.1	60.4	61.9	69.9
OntoNotes	29.7	27.8	32.2	44.2
PolyglotNER	33.6	41.8	42.9	49.4
TweetNER7	40.1	42.7	41.4	52.5
WikiANN	52.0	55.4	58.9	58.6
wikiNeural	57.7	69.2	71.8	83.4
Average	36.5	45.7	47.8	56.8

Table 4: Zero-shot performance of ReasoningNER 8B compared against other models on 20 NER benchmarks. All results are F1 scores, and the best results are bolded. Results of ChatGPT are reported from UniNER (Zhou et al. 2023).

est 45,787 samples from our NER-CoT dataset, all ReasoningNER variants demonstrate a significant performance advantage over three competitive baselines. Employing the same InternLM2-7B model and test set, our method achieves an F1 score of 74.9, surpassing B²NER’s 64.6. Similarly, the Llama2-7B variant reaches 63.5 F1, surpassing KnowCoder (60.1 F1) despite the latter being trained on a dataset two orders of magnitude larger. Furthermore, we observe a monotonic improvement that correlates with model capacity: the Qwen3 series scales from 63.0 F1 (1.7B) to 69.0 F1 (4B), ultimately achieving a peak performance of 70.2 F1 with the 8B model. This result surpasses both InternLM2-7B (69.1 F1) and Llama3.1-8B (68.9 F1). These findings suggest that incorporating explicit chains of reasoning is an effective strategy for unlocking the full potential of various LLMs and demonstrates substantially greater data efficiency than large-scale instruction tuning. Given its superior performance in this comparative analysis, we have selected Qwen3-8B as the backbone for our main experiments, where it is further optimized with GRPO to establish a new SoTA.

Conclusion

This paper introduces ReasoningNER, an NER framework grounded in a reasoning paradigm. It reframes the task from implicit pattern matching to explicit, step-by-step logical inference. This paradigmatic shift enhances the generalization and data efficiency of large language models. To operationalize this paradigm, we present a three-stage ap-

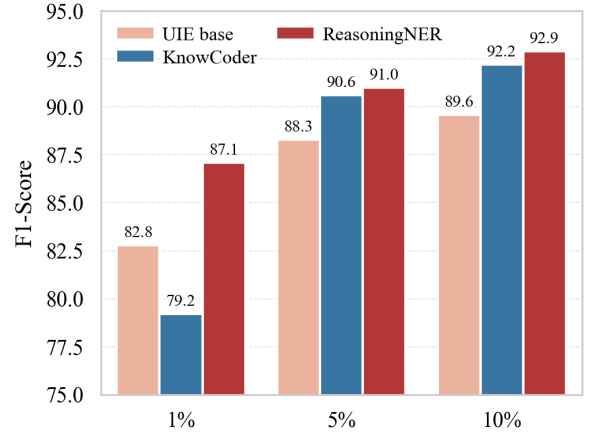


Figure 3: Low-resource results on CoNLL03 dataset.

Model	#Num	F1
<i>SFT on NER-CoT</i>		
Qwen3-8B	45,787	70.2
Qwen3-8B [*]	45,787	75.9
Qwen3-4B	45,787	69.0
Qwen3-1.7B	45,787	63.0
InternLM2-7B	45,787	69.1
InternLM2-7B [*]	45,787	74.9
Llama3.1-8B	45,787	68.9
Qwen2.5-7B	45,787	68.6
Llama2-7B	45,787	63.5
<i>Baselines</i>		
GPT-4 (Zero-shot)	-	60.1
KnowCoder (Llama2-7B)	4,590,403	60.1
B ² NER [*] (InternLM2-7B)	70,292	64.6

Table 5: Performance Comparison of Supervised Fine-Tuning on the NER-CoT Dataset across Various LLMs, Measured by the Average F1-Score on Seven Datasets from CrossNER and MIT.

proach. First, we construct NER-CoT, the first corpus for NER where entities are annotated with fine-grained Chain-of-Thought rationales. Second, we employ CoT Tuning to supervise the model in generating structured reasoning paths before making a prediction. Finally, a Reasoning Enhancement stage utilizes the GRPO algorithm to optimize the reasoning policy, guided by a composite reward function that ensures both task accuracy and structural-format consistency. Empirically, ReasoningNER achieves SoTA results, significantly outperforming existing methods across various benchmarks and demonstrating superior robustness in zero-shot, cross-domain, and low-resource scenarios. Despite its efficacy, a primary limitation is the increased inference latency from generating longer reasoning chains. Future work will explore hybrid CoT strategies to strike a better balance between model performance and inference efficiency.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Agarwal, B.; Joshi, I.; and Rojkova, V. 2025. Think Inside the JSON: Reinforcement Strategy for Strict LLM Schema Adherence. *arXiv preprint arXiv:2502.14905*.
- Ashok, D.; and Lipton, Z. C. 2023. Promptner: Prompting for named entity recognition. *arXiv preprint arXiv:2305.15444*.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Dao, T. 2023. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hsu, P.-L.; Dai, Y.; Kothapalli, V.; Song, Q.; Tang, S.; Zhu, S.; Shimizu, S.; Sahni, S.; Ning, H.; and Chen, Y. 2024. Liger Kernel: Efficient Triton Kernels for LLM Training. *arXiv preprint arXiv:2410.10989*.
- Li, Z.; Zeng, Y.; Zuo, Y.; Ren, W.; Liu, W.; Su, M.; Guo, Y.; Liu, Y.; Lixiang, L.; Hu, Z.; et al. 2024. KnowCoder: Coding Structured Knowledge into LLMs for Universal Information Extraction. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8758–8779.
- Liao, X.; Duan, J.; Huang, Y.; and Wang, J. 2025. RUIE: Retrieval-based Unified Information Extraction using Large Language Model. In *Proceedings of the 31st International Conference on Computational Linguistics*, 9640–9655.
- Lin, Y.; Ji, H.; Huang, F.; and Wu, L. 2020. A joint neural model for information extraction with global features. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, 7999–8009.
- Liu, Y.; Peng, X.; Du, T.; Yin, J.; Liu, W.; and Zhang, X. 2024. ERA-CoT: Improving Chain-of-Thought through Entity Relationship Analysis. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8780–8794.
- Liu, Z.; Xu, Y.; Yu, T.; Dai, W.; Ji, Z.; Cahyawijaya, S.; Madotto, A.; and Fung, P. 2021. Crossner: Evaluating cross-domain named entity recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 13452–13460.
- Lu, Y.; Liu, Q.; Dai, D.; Xiao, X.; Lin, H.; Han, X.; Sun, L.; and Wu, H. 2022. Unified Structure Generation for Universal Information Extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5755–5772.
- Ma, X.; Li, J.; and Zhang, M. 2023. Chain of Thought with Explicit Evidence Reasoning for Few-shot Relation Extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2334–2352.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Sainz, O.; García-Ferrero, I.; Agerri, R.; de Lacalle, O. L.; Rigau, G.; and Agirre, E. 2024. GoLLIE: Annotation Guidelines improve Zero-Shot Information-Extraction. In *ICLR*.
- Van Nguyen, M.; Lai, V. D.; and Nguyen, T. H. 2021. Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 27–38.
- Wadden, D.; Wennberg, U.; Luan, Y.; and Hajishirzi, H. 2019. Entity, Relation, and Event Extraction with Contextualized Span Representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5784–5789.
- Wang, S.; Sun, X.; Li, X.; Ouyang, R.; Wu, F.; Zhang, T.; Li, J.; and Wang, G. 2023a. Gpt-ner: Named entity recognition via large language models. *arXiv preprint arXiv:2304.10428*.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Wang, X.; Zhou, W.; Zu, C.; Xia, H.; Chen, T.; Zhang, Y.; Zheng, R.; Ye, J.; Zhang, Q.; Gui, T.; et al. 2023b. InstructTIE: Multi-task instruction tuning for unified information extraction. *arXiv preprint arXiv:2304.08085*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yang, Y.; Zhao, W.; Huang, C.; Ye, J.; Wang, X.; Zheng, H.; Nan, Y.; Wang, Y.; Xu, X.; Huang, K.; et al. 2025b. Beyond Boundaries: Learning a Universal Entity Taxonomy across Datasets and Languages for Open Named Entity Recognition. In *Proceedings of the 31st International Conference on Computational Linguistics*, 10902–10923.
- Zaratiana, U.; Tomeh, N.; Holat, P.; and Charnois, T. 2024. GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer. In *Proceedings of the 2024*

Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), 5364–5376.

Zhou, W.; Zhang, S.; Gu, Y.; Chen, M.; and Poon, H. 2023. Universalner: Targeted distillation from large language models for open named entity recognition. *arXiv preprint arXiv:2308.03279*.