

Zero-Shot Relation Classification with Large Language Models: Binary Prompting and Relevance Elicitation*

Sheila Alemany, Danielle Sullivan-Pao, Olga Simek

MIT Lincoln Laboratory

Lexington, MA, USA

Abstract

Relation classification (RC) transforms unstructured text into structured knowledge by identifying semantic relationships between entities. As natural language data volumes rapidly increase, manual knowledge extraction has become impractical. Existing zero-shot LLM-based approaches face limitations due to sparse RC task representation in instruction-tuning datasets and inherent biases toward certain relation types. To address these challenges, we propose Binary Relevance Elicitation (BRE), a true/false prompting framework that systematically queries an LLM for each candidate relation between two entities from a closed set of relation types. BRE incorporates relevance scoring to identify the most likely relation among those classified as “true,” reducing false positive predictions. We evaluate BRE using multiple open-source LLMs and achieve state-of-the-art zero-shot RC accuracy across the TACRED family of datasets. Our results demonstrate performance improvements with newer model releases that are independent of parameter count, suggesting information representation advances within LLMs drive zero-shot RC capabilities.

1 Introduction

Social network analysis reveals how entities (people, organizations, locations) connect and how

*DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited. This material is based upon work supported by the Combatant Commands under Air Force Contract No. FA8702-15-D-0001 or FA8702-25-D-B002. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Combatant Commands.

© 2026 Massachusetts Institute of Technology. Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.




Figure 1: Zero-shot BRE prompting pipeline using the relevance elicitation measure (ρ) to select the relation for a given context (c) and entity pair. The predicted relation is the one with the highest relevance measure from the set of possible relations ($r_{pred} \in r_{potential} \subseteq R$).

these connections influence social behavior and outcomes. Relation classification (RC) identifies semantic relationships between entities in sentences, converting unstructured text into structured data. These triplets enhance knowledge classification, improve information retrieval and question answering, and generate knowledge graphs (Zhao et al., 2024a; Nayak et al., 2021).

RC has achieved state-of-the-art performance by leveraging pre-trained language models (PLMs) to extract relation triplets (Park and Kim, 2021; Han et al., 2022; Zhou and Chen, 2022; Huguet Cabot et al., 2021). However, a fundamental misalignment persists: top-performing PLMs rely on pre-training objectives like masked language modeling that fail to capture the structured semantic relationships between predefined entity pairs (Devlin, 2018; Liu, 2019). While some efforts have reframed RC as textual entailment or natural lan-

guage inference (Obamuyide and Vlachos, 2018; Chanthran et al., 2026), PLM-based approaches remain critically dependent on large volumes of high-quality labeled training data (Huguet Cabot et al., 2021; Elsahar et al., 2018; Zhao et al., 2024a). Language ambiguity makes annotating semantic relations inherently challenging, producing incomplete or inconsistent ground-truth data that varies by annotator perspective and dataset domain. The absence of standardized, closed-set training and testing datasets represents a fundamental obstacle to RC advancement, making zero-shot approaches not merely desirable but essential for domain-specific applications.

Large language models (LLMs) have shown promise with minimal training data (Hang et al., 2025), yet they falter in critical scenarios: capturing implied knowledge in long-tail cases where entity relations must be globally inferred (Zhao et al., 2024b), and handling complex multi-relation sentences without handcrafted templates for each relation type (Li et al., 2024; Geroimenko, 2025; Aguda et al., 2025). We introduce a zero-shot approach to RC centered on *relevance elicitation*, a quantitative ranking mechanism that identifies optimal relation types and directly addresses LLMs’ systematic over-prediction bias. While our method uses binary true/false queries as a lightweight foundation which avoids manual template customization, the performance gains stem primarily from aggregating these queries into relevance scores rather than treating them as final decisions. This shift toward ranked, relevance-based selection is the key driver of our results, mirroring the binary relevance filtering approach of Zhuang et al. (2024) for document ranking. Our method achieves state-of-the-art zero-shot RC performance across the TACRED family of datasets (Zhang et al., 2017) using various open-source LLMs.

2 Related Work

2.1 From Supervised to Zero-Shot Closed RC

RC methods vary in supervision requirements and output structures. Traditional supervised RC approaches using deep neural networks require substantial annotated training data, limiting their applicability to new domains or relation types. Unsupervised relation extraction methods eliminate this annotation burden, primarily employing clustering-based techniques that extract entities via named-entity taggers and group entity pairs based on con-

textual similarity (Detroja et al., 2023). However, these methods produce open-ended clusters rather than assignments to predefined relation labels, making interpretation ambiguous and challenging to evaluate when sentences contain entity pairs with multiple relations. Zero-shot closed RC addresses these limitations by enabling users to specify a predefined set of relations of interest without requiring annotated training examples. This approach combines the flexibility of unsupervised methods (no annotation requirement) with the interpretability of supervised methods (explicit, user-defined label assignments), making it particularly suitable for domain-specific applications where relation types are known but labeled data is scarce.

2.2 Pre-Trained Language Models (PLMs)

PLMs have advanced RC by eliciting rich knowledge from large pre-trained models (Huguet Cabot et al., 2021; Tang et al., 2022; Sui et al., 2023). Models like BERT (Devlin, 2018), RoBERTa (Liu, 2019), and ELMo (Peters et al., 2018) produce contextualized representations where token meaning depends on surrounding context, enabling nuanced understanding of entity relationships within sentences. PLM-based RC approaches employ fine-tuning to adapt these models for specific RC tasks. However, standard pre-training objectives (e.g., masked language modeling, next-token prediction) do not naturally capture structured semantic relationships, creating reliance on task-specific training data (Hang et al., 2025). This creates two key limitations that zero-shot closed RC addresses: significant computational requirements for training and inference, and tendency to overfit on domain-specific datasets (Detroja et al., 2023).

2.3 Large Language Models (LLMs)

LLMs have transformed zero-shot RC through two primary approaches: direct RC and synthetic data generation for training downstream models. For direct classification, prompt engineering approaches typically require no additional fine-tuning, simplifying deployment and eliminating the need for annotated few-shot data (Chia et al., 2022; Li et al., 2024). Alternatively, LLMs can generate synthetic training data as supplementary samples for unseen relations (Xu et al., 2023). However, effective prompt engineering for direct classification requires considerable subject matter expertise specific to the RC-employed domain (Geroimenko, 2025), and out-of-the-box LLM performance is

limited by general language understanding capabilities and training biases (Aguda et al., 2025; Dai et al., 2024).

Recent efforts show LLM performance varies significantly with prompt templates (Liu et al., 2023; Zhang et al., 2023a; Xu et al., 2024). To address this, approaches have incorporated additional context (Wan et al., 2023), entity type information (Zhang et al., 2023a), multi-step prompting for chain-of-thought control (Layegh et al., 2024; Hannah et al., 2025), and logical searches for discrete prompts to avoid manual design (Xu et al., 2024). A common observation is that task simplification correlates with improved LLM performance (Zhou et al.; Khot et al.). Our approach uniquely simplifies RC to a binary decision and adapts the LLM-as-a-judge framework (Gu et al., 2024) to eliminate domain-specific prompt-tuning while improving performance.

3 Methodology

BRE, described in Algorithm 1, simplifies RC to a series of binary decisions while enabling principled relation selection through relevance scoring. For each entity pair (e_h, e_t) in context c , BRE iterates through type-compatible relations from a predefined set R , prompting the LLM to evaluate whether each relation r_i holds and rate its relevance ρ to the context on a 0-10 scale, given the relation definition d_i and the statement “ e_h and e_t are somehow connected through the relationship of ‘ r_i .’”

Pairing binary classification with a relevance score is critical: the binary decision identifies plausible relations, while ρ enables ranking among multiple true relations. BRE returns the highest-scoring relation among all ‘true’ classifications¹, or ‘no relation’ if none qualify. This mechanism directly addresses the multi-relation challenge in RC, providing a principled, interpretable basis for relation selection without domain-specific prompt engineering or few-shot examples. See Appendix A.1 for the prompt template.

3.1 Closed RC Definition

The RC task considers a given sentence c , a set of entities $E \subseteq c$, entity pairs $(e_h, e_t) \in E \times E$ where $e_h \neq e_t$, and a set of relation types R . The task identifies the semantic relations between entity

¹If two relations are assigned the same relevance score, the selected relation is randomly chosen for this evaluation. Additional details are provided in Appendix A.5

Algorithm 1 Binary Relevance Elicitation (BRE)

Require: sentence c , entity pair (e_h, e_t) , entity types (e_h^{type}, e_t^{type}) , relation set R , valid type pairs T_r , definitions D , LLM θ
 $R_{valid} \leftarrow \{r \in R \mid (e_h^{type}, e_t^{type}) \in T_r\}$
 $r_{pot} \leftarrow \emptyset$ $\triangleright$ Store LLM-selected relations
for $r_i \in R_{valid}$ **do**
 $prompt \leftarrow \text{construct_prompt}(c, e_h, r_i, e_t, D(r_i))$
 $(y_i, \rho_i) \leftarrow \theta(prompt)$ $\triangleright y_i \in \{\text{TRUE}, \text{FALSE}\}, \rho_i \in [0, 10]$
 if $y_i = \text{TRUE}$ **then**
 $r_{pot} \leftarrow r_{pot} \cup \{(r_i, \rho_i)\}$
 end if
end for
return $|r_{pot}| = 0 ? \text{‘no relation’} : \text{argmax}_{(r, \rho) \in r_{pot}} \rho$

pairs in c . The RC output is a set of sentence-level triplets $\langle e_h, r_{pred}, e_t \rangle$ where $r_{pred} \in R \cup \{\text{‘no relation’}\}$. Zero-shot RC employs no fine-tuning on the target relations, and no examples are included in the prompt. The performance varies substantially across LLMs even under identical task framing. By holding the dataset and prompt fixed, this work isolates and evaluates these cross-model differences. Figure 1 shows the flow to identify the predicted relation for a valid entity pair.

3.2 Entity-Type Aware Querying

Entity type constraints critically reduce the search space of potential relation triplets. By defining valid entity type pairs (e_h^{type}, e_t^{type}) for each relation, we systematically filter candidate relations before LLM inference. Specifically,

$$R_{valid} = \{r \in R \mid (e_h^{type}, e_t^{type}) \in T_r\}$$

where T_r is the set of valid type pairs for relation r as defined by TACRED dataset annotation guidelines (Ellis et al., 2015). This leverages domain knowledge to eliminate semantically invalid combinations (e.g., testing “founded by” for two location entities) (Zhang et al., 2023a). This optional type-aware filtering improves computational efficiency and extraction precision by focusing on semantically plausible candidates, yielding more reliable and cost-effective RC systems.

3.3 Relation Definition Context

Including relation definitions in prompts serves as a disambiguation mechanism (Zhou et al., 2024). Natural language descriptions establish clear semantic boundaries for interpreting relation types, which is useful for ambiguous or overlapping relations (e.g., distinguishing “employed by” from “affiliated with”). Grounding reasoning in explicit

Model	Huggingface Model ID	Parameter Size (B)	Release Date
Llama-3.2-3B	meta-llama/Llama-3.2-3B-Instruct	3	09/25/2024
Llama-3.1-8B	meta-llama/Llama-3.1-8B-Instruct	8	07/23/2024
Llama-3.1-70B	meta-llama/Llama-3.1-70B-Instruct	70	07/23/2024
Llama-3.3-70B	meta-llama/Llama-3.3-70B-Instruct	70	12/06/2024
Llama-4-Scout-17B-16E	meta-llama/Llama-4-Scout-17B-16E-Instruct	109	04/05/2025
Mistral-7B	mistralai/Mistral-7B-Instruct-v0.2	7	12/25/2023
Mixtral-8x7B	mistralai/Mixtral-8x7B-Instruct-v0.1	56	12/10/2023
Mixtral-8x22B	mistralai/Mixtral-8x22B-Instruct-v0.1	176	04/17/2024
Qwen2.5-VL-32B	Qwen/Qwen2.5-VL-32B-Instruct	32	02/19/2025
DeepSeek-R1-Distill-Qwen-32B	deepseek-ai/DeepSeek-R1-Distill-Qwen-32B	32	02/23/2025

Table 1: Evaluated LLM version details including their corresponding Huggingface model IDs, their parameter sizes, and release dates in the format MM/DD/YYYY

definitions reduces semantic confusion and ensures consistent interpretation across contexts, rather than relying on implicit, context-dependent associations. Specifically, we define $D = \{r_i : d_i\}$ such that d_i is the definition for r_i as described in the TACRED family dataset annotation guidelines (Ellis et al., 2015; Stoica et al., 2021).

4 Evaluation

4.1 Experimental Setup

LLM Models We evaluate open-source LLMs from multiple families (Llama, Mistral, Qwen, and DeepSeek) detailed in Table 1, spanning varying parameter counts to examine correlations between model size, organization, release date, and RC performance. The same prompt template was used across all models.

Computational Usage We implemented our approach using VLLM on NVIDIA H100 GPUs, scaling resources by model size (up to four H100s for the largest model). We performed five inference runs per model and report averaged results to account for stochastic variation.

NoTA Evaluation LLM-based RC systems exhibit a critical bias: they over-assign named relations rather than predicting “no relation” or the none-of-the-above (NoTA) class, even when no valid relationship exists (Wan et al., 2023; Ma et al.; Zhang et al., 2023b). This bias is amplified by the asymmetric nature of many relations, requiring evaluation of numerous ordered entity pairs with potentially no predefined relationship (Qin et al., 2021; Baek and Choi, 2022), resulting in severe class imbalance that increases false positives and reduces precision. Our relevance elicitation method mitigates this by ranking potential false positives through comparative relevance scoring

across candidate relations. Since Gutierrez et al. (Gutierrez et al., 2022) attributed early LLM under-performance on RC largely to poor NoTA handling, we report both NoTA-included and NoTA-excluded metrics for comparison with prior work.

Parsing LLM Outputs All evaluated LLMs followed a specified format starting with “true” or “false” followed by the numerical relevance score. Llama-4-Scout-17B-16E consistently placed this output at the end of its response and required a modified parsing script. Complete parsing scripts are provided in Appendix A.2. Additional analysis regarding invalid outputs and no-score outputs are included in Appendix A.5.

4.2 Baselines

Few-Shot Baselines ODDA (Zhong et al., 2025) employs data augmentation to iteratively generate few-shot examples based on the RE model’s predictions for target relations, then produces diverse training data using attribute-based constraints. ConsistRE (Zheng et al., 2024) uses relation-specific keywords as generation constraints to encourage relation consistency. However, it requires access to an external corpus for keyword extraction, an assumption that may not be feasible in all scenarios. Lastly, both approaches use DeepSeek-v3 (671B parameters) along with its PLM backbone.

Zero-Shot Baselines QA4RE (Zhang et al., 2023a) reformulates RC as multiple-choice question-answering by incorporating valid entity pairs into pre-defined relation templates as answer options. SUMASK (Li et al., 2023) takes a different approach, combining LLMs with natural language inference and chain-of-thought reasoning to sequentially query whether each candidate relation exists between entity pairs through iterative decomposition. REA (Layegh et al., 2024)

Dataset	Model	macro-average			weighted-average			micro- F_1 Accuracy
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	
Re-TACRED	Llama-3.2-3B	17.5 $\pm$ 0.29	41.4 $\pm$ 0.91	20.9 $\pm$ 0.4	58.0 $\pm$ 0.37	46.8 $\pm$ 0.74	49.1 $\pm$ 0.59	42.3 $\pm$ 0.56
	Llama-3.1-8B	21.8 $\pm$ 0.3	53.5 $\pm$ 1.02	27.7 $\pm$ 0.24	64.0 $\pm$ 0.22	51.0 $\pm$ 0.88	50.1 $\pm$ 1.0	43.6 $\pm$ 1.02
	Llama-3.1-70B	61.4 $\pm$ 0.49	53.6 $\pm$ 1.07	54.4 $\pm$ 0.76	80.7 $\pm$ 0.2	76.0 $\pm$ 0.17	75.8 $\pm$ 0.11	79.1 $\pm$ 0.25
	Llama-3.3-70B	57.9 $\pm$ 1.92	62.9 $\pm$ 3.21	57.8 $\pm$ 1.53	73.2 $\pm$ 8.32	72.4 $\pm$ 2.93	71.4 $\pm$ 4.47	79.5 $\pm$ 0.4
	Llama-4-Scout-17B-16E	57.8 $\pm$ 0.37	73.3 $\pm$ 0.62	62.1 $\pm$ 0.47	86.9 $\pm$ 0.33	83.9 $\pm$ 0.33	84.6 $\pm$ 0.29	83.5 $\pm$ 0.29
	Mistral-7B	39.1 $\pm$ 0.89	55.8 $\pm$ 2.84	40.6 $\pm$ 2.31	67.1 $\pm$ 4.91	63.7 $\pm$ 3.35	63.4 $\pm$ 3.44	66.1 $\pm$ 1.55
	Mixtral-8x7B	48.2 $\pm$ 0.96	58.0 $\pm$ 1.02	47.2 $\pm$ 0.47	75.4 $\pm$ 0.41	72.7 $\pm$ 0.57	72.4 $\pm$ 0.39	72.1 $\pm$ 0.61
	Mixtral-8x22B	50.7 $\pm$ 1.28	61.3 $\pm$ 1.16	52.6 $\pm$ 1.25	81.7 $\pm$ 0.35	77.7 $\pm$ 0.34	78.2 $\pm$ 0.33	79.5 $\pm$ 0.5
	Qwen2.5-VL-32B	66.2 $\pm$ 0.16	64.4 $\pm$ 0.27	61.9 $\pm$ 0.22	86.5 $\pm$ 0.01	82.5 $\pm$ 0.03	83.2 $\pm$ 0.03	85.4 $\pm$ 0.03
TACREV	DeepSeek-R1-Distill-Qwen-32B	44.4 $\pm$ 0.29	60.7 $\pm$ 2.36	48.4 $\pm$ 0.78	76.2 $\pm$ 0.38	74.2 $\pm$ 0.36	74.2 $\pm$ 0.28	74.4 $\pm$ 0.45
	Llama-3.2-3B	21.2 $\pm$ 0.24	47.6 $\pm$ 0.82	24.2 $\pm$ 0.39	75.7 $\pm$ 0.23	54.5 $\pm$ 0.41	61.3 $\pm$ 0.19	49.6 $\pm$ 0.36
	Llama-3.1-8B	22.0 $\pm$ 0.52	65.0 $\pm$ 1.64	28.8 $\pm$ 0.67	79.2 $\pm$ 0.32	52.1 $\pm$ 1.05	58.9 $\pm$ 1.02	47.6 $\pm$ 0.99
	Llama-3.1-70B	56.8 $\pm$ 0.55	60.1 $\pm$ 2.03	54.6 $\pm$ 1.16	85.7 $\pm$ 0.33	87.5 $\pm$ 0.06	86.1 $\pm$ 0.18	85.8 $\pm$ 0.13
	Llama-3.3-70B	55.7 $\pm$ 0.2	70.5 $\pm$ 0.01	59.6 $\pm$ 0.31	87.9 $\pm$ 0.37	88.2 $\pm$ 0.26	87.8 $\pm$ 0.3	86.4 $\pm$ 0.22
	Llama-4-Scout-17B-16E	52.2 $\pm$ 0.52	82.2 $\pm$ 1.21	60.6 $\pm$ 0.53	90.9 $\pm$ 0.11	88.3 $\pm$ 0.14	88.9 $\pm$ 0.11	86.7 $\pm$ 0.1
	Mistral-7B	37.5 $\pm$ 0.13	69.1 $\pm$ 0.73	44.5 $\pm$ 0.15	84.6 $\pm$ 0.13	74.3 $\pm$ 0.21	77.8 $\pm$ 0.15	73.3 $\pm$ 0.22
	Mixtral-8x7B	46.1 $\pm$ 0.79	70.6 $\pm$ 1.4	52.5 $\pm$ 0.48	86.2 $\pm$ 0.24	80.2 $\pm$ 0.4	82.3 $\pm$ 0.25	78.6 $\pm$ 0.39
	Mixtral-8x22B	46.8 $\pm$ 0.43	72.1 $\pm$ 0.79	53.9 $\pm$ 0.44	88.0 $\pm$ 0.11	87.1 $\pm$ 0.14	87.2 $\pm$ 0.11	85.3 $\pm$ 0.11
TACRED	Qwen2.5-VL-32B	62.6 $\pm$ 0.6	76.7 $\pm$ 0.25	65.0 $\pm$ 0.1	91.8 $\pm$ 0.19	90.6 $\pm$ 0.02	90.5 $\pm$ 0.01	90.0 $\pm$ 0.01
	DeepSeek-R1-Distill-Qwen-32B	42.3 $\pm$ 0.6	70.9 $\pm$ 0.6	49.2 $\pm$ 0.47	85.6 $\pm$ 0.17	83.5 $\pm$ 0.45	84.0 $\pm$ 0.23	80.7 $\pm$ 0.47
	Llama-3.2-3B	22.9 $\pm$ 0.39	46.0 $\pm$ 1.06	25.4 $\pm$ 0.34	75.2 $\pm$ 0.24	54.2 $\pm$ 0.36	61.1 $\pm$ 0.28	49.8 $\pm$ 0.34
	Llama-3.1-8B	23.6 $\pm$ 0.26	59.8 $\pm$ 0.96	30.2 $\pm$ 0.4	78.1 $\pm$ 0.05	51.7 $\pm$ 0.98	58.6 $\pm$ 0.83	48.0 $\pm$ 0.87
	Llama-3.1-70B	54.7 $\pm$ 1.36	50.7 $\pm$ 1.34	50.5 $\pm$ 1.11	84.0 $\pm$ 0.19	84.8 $\pm$ 0.19	83.9 $\pm$ 0.18	83.4 $\pm$ 0.3
	Llama-3.3-70B	52.3 $\pm$ 0.38	62.1 $\pm$ 1.55	55.1 $\pm$ 1.0	86.8 $\pm$ 0.29	84.9 $\pm$ 0.04	85.6 $\pm$ 0.15	84.1 $\pm$ 0.07
	Llama-4-Scout-17B-16E	52.0 $\pm$ 0.31	70.0 $\pm$ 0.96	57.6 $\pm$ 0.38	88.5 $\pm$ 0.11	85.3 $\pm$ 0.15	86.5 $\pm$ 0.13	84.6 $\pm$ 0.14
	Mistral-7B	39.0 $\pm$ 0.82	60.1 $\pm$ 0.48	43.2 $\pm$ 0.38	83.0 $\pm$ 0.07	72.9 $\pm$ 0.32	76.5 $\pm$ 0.23	72.3 $\pm$ 0.32
	Mixtral-8x7B	47.7 $\pm$ 0.39	60.1 $\pm$ 0.63	49.8 $\pm$ 0.42	84.4 $\pm$ 0.09	78.7 $\pm$ 0.31	80.9 $\pm$ 0.2	77.4 $\pm$ 0.25
TACRED	Mixtral-8x22B	48.0 $\pm$ 0.2	63.6 $\pm$ 0.18	52.6 $\pm$ 0.16	86.1 $\pm$ 0.44	84.2 $\pm$ 0.15	84.9 $\pm$ 0.26	83.1 $\pm$ 0.16
	Qwen2.5-VL-32B	60.6 $\pm$ 0.53	63.4 $\pm$ 0.2	58.4 $\pm$ 0.12	89.0 $\pm$ 0.21	86.8 $\pm$ 0.04	87.3 $\pm$ 0.04	86.9 $\pm$ 0.04
	DeepSeek-R1-Distill-Qwen-32B	43.3 $\pm$ 0.84	61.0 $\pm$ 1.02	47.7 $\pm$ 0.71	83.8 $\pm$ 0.26	81.4 $\pm$ 0.32	82.2 $\pm$ 0.21	79.2 $\pm$ 0.37

Table 2: NoTA-included zero-shot RC performance on the TACRED family of datasets. The gray scores indicate highest performing score of the category. Performance and standard deviations are represented as percentages (%).

employs a more elaborate four-stage prompting strategy: identifying entity types, filtering relation candidates by entity type compatibility, generating hypothesis sentences with confidence scores for each candidate, and selecting the final relation based on accumulated evidence.

4.3 Datasets

TACRED (Zhang et al., 2017) is a widely-used dataset for sentence-level relation extraction with 42 relation types, containing annotated news articles and web documents with entities and predefined relations between persons and organizations for knowledge base construction. Updated versions include TACREV (Alt et al., 2020), which corrects systematic labeling errors (particularly false “no relation” labels), and Re-TACRED (Stoica et al., 2021), which employed improved crowd-sourced annotation and relabeled 23.9% of the original dataset for enhanced accuracy. Evaluating all three datasets (entire test sets) demonstrates our methodology’s robustness to annotation inconsistencies.

5 Results

5.1 Relation Classification Performance

Qwen2.5-VL-32B consistently shows high RC performance likely due to its advanced attention mechanisms and multimodal pre-training which enhance structured relationship understanding (Wang et al., 2024). MoE architectures (e.g., Llama-4-Scout and Mixtral) with specialized tuning showed high accuracy with balanced classification performance (Shen et al., 2023). Both represent architectural innovations beyond standard decoder-only transformers, specifically beneficial for tasks requiring precise semantic relationship extraction.

Observation: DeepSeek-R1-Distill-Qwen-32B showed inferior expected performance suggesting distillation disproportionately affected semantic relationship identification.

5.2 Baseline Comparisons

BRE prompting demonstrates substantial improvements over existing zero-shot baselines across all

	Re-TACRED	TACREV	TACRED
Qwen2.5-VL-32B	74.7	82.4	73.9
Llama-4-Scout-17B-16E	<u>78.8</u>	<u>86.3</u>	<u>78.8</u>
Llama-3.3-70B	65.6	78.9	73.9
Mixtral-8x22B	69.0	79.4	70.7
Llama-3.1-70B	62.1	68.3	62.7
DeepSeek-R1-Distill-Qwen-32B	66.7	75.7	68.5
Mixtral-8x7B	65.9	71.3	68.0
REA + Mixtral-8x7B	-	-	58.0
REA + GPT-3.5-Turbo	-	-	63.5
QA4RE + FLAN-T5 XXL	67.3	54.3	54.5
QA4RE + GPT-3.5-Turbo	61.2	59.4	59.4
SUMASK + GPT-3.5-Turbo	-	-	57.7

Table 3: NoTA-excluded comparison to state-of-the-art zero-shot RC methods. The gray scores indicate the top two micro- F_1 scores and the underlined score is the highest.

	Re-TACRED	TACREV	TACRED
Qwen2.5-VL-32B	<u>85.6</u>	<u>90.0</u>	<u>86.8</u>
Llama-4-Scout-17B-16E	83.2	86.7	84.4
Llama-3.3-70B	79.5	86.4	83.9
Mixtral-8x22B	79.4	85.9	83.0
Llama-3.1-70B	79.1	85.8	83.4
DeepSeek-R1-Distill-Qwen-32B	74.5	80.8	79.3
Mixtral-8x7B	72.4	78.7	77.7
SUMASK + GPT-3.5-Turbo	73.8	75.1	79.6
ODDA/TYP Marker (8-shot)	76.8	-	75.8
ODDA/KnowPrompt (8-shot)	<u>84.1</u>	-	81.0
ConsistRE/TYP Marker (8-shot)	71.8	-	72.6
ConsistRE/KnowPrompt (8-shot)	75.8	-	80.1

Table 4: NoTA-included comparison to state-of-the-art zero- and few-shot RC methods. All approaches in this table are zero-shot unless otherwise specified. The gray scores indicate the top two micro- F_1 scores and the underlined score is the highest.

evaluated datasets². Most notably, BRE outperforms REA (Layegh et al., 2024) even when both methods employ the same foundational Mixtral 8x7B model, achieving 17% relative increased accuracy on TACRED. The consistent gains across datasets highlights that our approach effectively addresses the over-prediction bias inherent in LLM-based RC.

Using Qwen2.5-VL-32B and Llama-4-Scout, BRE matches or surpasses the performance of state-of-the-art few-shot baselines which require labeled

²Reported baseline performance numbers are from original efforts. Missing reported values were not included in the original work.




Figure 2: NoTA-included micro- F_1 comparison to state-of-the-art zero-shot RC methods on TACRED dataset




Figure 3: NoTA-excluded micro- F_1 comparison to state-of-the-art zero-shot RC methods on TACRED dataset

data or iterative data augmentation. However, analysis of NoTA-excluded metrics reveals an important distinction between our top-performing models: Llama-4-Scout outperforms Qwen2.5-VL-32B in NoTA-excluded scenarios.

Observation: Qwen2.5-VL-32B’s high overall scores stem largely from strong NoTA classification, while Llama-4-Scout maintains more balanced performance.

In Figures 2 and 3, circles (○) represent BRE, and diamonds (◇) represent the baseline methods. The

size of the shape indicates the parameter size in the foundational model.³ The baseline release dates included are those of the research publication. While the figures show an upward trend in RC task performance over time, our statistical analysis indicates that this trend is primarily attributable to parameter size rather than release date itself: a regression controlling for both factors found parameter size to be a significant predictor of performance ($p = 0.006$), while release date was not ($p = 0.667$), even after excluding performance inflation from NoTA relations in Figure 3.

Observation: Statistical analysis highlights performance gains over time are largely explained by increases in model parameter size rather than by time-dependent improvements in natural language representation alone.

5.3 Relation-Specific Results

Qwen2.5-VL-32B and Llama-4-Scout achieved the highest accuracy across all three datasets when paired with BRE prompting. However, the discrepancy between weighted-average and macro-average scores indicates that performance varies significantly across different relation types. Figure 4 presents per-relation F_1 scores on the TACRED dataset, revealing that all models exhibit consistently low performance on the same subset of relations regardless of overall accuracy differences. Similar heatmaps for the TACREV and ReTACRED dataset are included in Appendix A.3.

Relations such as ‘members of’, ‘members,’ and ‘shareholders’ prove particularly difficult, likely due to semantic overlap between relation definitions. When multiple relations share similar meanings, the relevance elicitation process may assign comparable scores across overlapping categories for a given entity pair, making correct distinction difficult. Notably, BRE’s relevance elicitation identifies these challenging cases by providing interpretable scores, enabling more transparent failure analysis compared to baseline approaches.

Observation: Relevance elicitation enables transparent identification of challenging relation types, particularly those with semantic overlap that confound classification.

³Figure 2 baseline parameter sizes are not scaled since ODDA and ConsistRE with the DeepSeek-v3 (671B parameters) backbones were too large to visually scale.




Figure 4: F_1 scores for the varying relation types in the TACRED dataset. The relations are sorted based on their F_1 and the models are sorted based on their macro- F_1 scores.

5.4 Ablation Study

Ablation results exploring the impact of: (1) only including the binary decision in the methodology, and (2) only using the relevance elicitation without the binary step are included in Appendix A.4.

Impact of Relation Definitions The impact of relation-specific challenges are further evidenced in our ablation study, where removing relation definitions resulted in an average relative accuracy degradation of 4% across the TACRED family of datasets ($p\text{-value} = 0.01 < 0.05$). Including relation definitions serves as a semantic anchor grounds the binary classification task and provides explicit criteria for what constitutes a valid relationship beyond mere label matching. It also contributes more context regarding what a potential relation may include in scenarios where the sentence or relation

Dataset	Model	macro-average			weighted-average			micro- F_1 Accuracy
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	
Re-TACRED	Qwen2.5-VL-32B	65.1 $\pm$ 0.24	54.9 $\pm$ 0.4	54.9 $\pm$ 0.27	84.9 $\pm$ 0.09	77.6 $\pm$ 0.08	78.4 $\pm$ 0.1	82.7 $\pm$ 0.09
	Llama-4-Scout-17B-16E	50.0 $\pm$ 0.79	63.2 $\pm$ 0.3	51.4 $\pm$ 0.45	83.3 $\pm$ 0.39	77.3 $\pm$ 0.36	78.9 $\pm$ 0.36	77.4 $\pm$ 0.31
	Llama-3.3-70B	54.3 $\pm$ 0.86	64.7 $\pm$ 1.88	54.1 $\pm$ 0.92	82.8 $\pm$ 0.42	73.0 $\pm$ 0.34	74.8 $\pm$ 0.12	76.8 $\pm$ 0.33
	Mixtral-8x22B	45.0 $\pm$ 0.11	60.1 $\pm$ 0.89	47.8 $\pm$ 0.4	78.0 $\pm$ 0.2	72.2 $\pm$ 0.07	73.0 $\pm$ 0.05	74.9 $\pm$ 0.06
	Llama-3.1-70B	45.9 $\pm$ 0.85	59.5 $\pm$ 2.32	47.8 $\pm$ 0.95	79.6 $\pm$ 0.34	73.4 $\pm$ 0.3	74.5 $\pm$ 0.28	74.5 $\pm$ 0.54
TACREV	Qwen2.5-VL-32B	63.0 $\pm$ 0.14	70.9 $\pm$ 0.28	61.6 $\pm$ 0.19	91.0 $\pm$ 0.04	89.7 $\pm$ 0.05	89.5 $\pm$ 0.04	89.1 $\pm$ 0.05
	Llama-4-Scout-17B-16E	47.9 $\pm$ 0.91	78.0 $\pm$ 0.81	55.8 $\pm$ 0.83	89.7 $\pm$ 0.05	85.1 $\pm$ 0.1	86.4 $\pm$ 0.08	83.0 $\pm$ 0.04
	Llama-3.3-70B	51.1 $\pm$ 0.55	72.8 $\pm$ 0.25	56.6 $\pm$ 0.36	88.1 $\pm$ 0.25	85.9 $\pm$ 0.07	86.4 $\pm$ 0.14	84.2 $\pm$ 0.1
	Mixtral-8x22B	41.7 $\pm$ 0.0	71.4 $\pm$ 0.0	49.7 $\pm$ 0.0	88.1 $\pm$ 0.0	83.4 $\pm$ 0.0	84.8 $\pm$ 0.0	81.6 $\pm$ 0.0
	Llama-3.1-70B	44.5 $\pm$ 0.58	70.6 $\pm$ 3.75	50.7 $\pm$ 0.74	85.5 $\pm$ 0.82	83.8 $\pm$ 0.81	84.0 $\pm$ 0.46	80.9 $\pm$ 0.79
TACRED	Qwen2.5-VL-32B	60.9 $\pm$ 0.22	57.7 $\pm$ 0.15	55.1 $\pm$ 0.1	88.3 $\pm$ 0.03	86.2 $\pm$ 0.03	86.2 $\pm$ 0.03	86.2 $\pm$ 0.04
	Llama-4-Scout-17B-16E	47.3 $\pm$ 0.59	66.0 $\pm$ 0.62	52.3 $\pm$ 0.41	87.4 $\pm$ 0.19	82.3 $\pm$ 0.35	84.1 $\pm$ 0.3	81.0 $\pm$ 0.36
	Llama-3.3-70B	50.9 $\pm$ 0.81	65.4 $\pm$ 1.6	54.7 $\pm$ 0.5	87.2 $\pm$ 0.66	82.9 $\pm$ 0.73	84.5 $\pm$ 0.2	82.0 $\pm$ 0.51
	Mixtral-8x22B	43.8 $\pm$ 0.0	64.7 $\pm$ 0.0	49.8 $\pm$ 0.0	86.1 $\pm$ 0.0	81.8 $\pm$ 0.0	83.4 $\pm$ 0.0	80.5 $\pm$ 0.0
	Llama-3.1-70B	45.0 $\pm$ 1.2	62.3 $\pm$ 1.69	49.4 $\pm$ 0.35	84.5 $\pm$ 0.43	81.4 $\pm$ 0.76	82.3 $\pm$ 0.37	79.1 $\pm$ 0.72
Average Δ		-5.29	-0.31	-4.33	-0.21	-2.95	-2.05	-3.29

Table 5: (Ablation 1) Excluding relation definitions during prompting

Dataset	Model	macro-average			weighted-average			micro- F_1 Accuracy
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	
Re-TACRED	Qwen2.5-VL-32B	46.5 $\pm$ 6.46	8.6 $\pm$ 0.86	12.5 $\pm$ 1.46	68.9 $\pm$ 0.74	61.1 $\pm$ 0.18	55.0 $\pm$ 0.28	71.2 $\pm$ 0.15
	Llama-4-Scout-17B-16E	32.8 $\pm$ 3.52	8.7 $\pm$ 0.14	11.6 $\pm$ 0.4	62.5 $\pm$ 1.77	60.3 $\pm$ 0.16	54.2 $\pm$ 0.29	69.6 $\pm$ 0.12
	Llama-3.3-70B	39.2 $\pm$ 5.13	8.8 $\pm$ 0.48	12.0 $\pm$ 0.72	64.0 $\pm$ 1.04	59.7 $\pm$ 0.08	53.1 $\pm$ 0.15	69.9 $\pm$ 0.12
	Mixtral-8x22B	19.6 $\pm$ 1.88	6.4 $\pm$ 0.86	8.0 $\pm$ 1.23	53.0 $\pm$ 0.94	58.8 $\pm$ 0.23	51.7 $\pm$ 0.33	68.8 $\pm$ 0.16
	Llama-3.1-70B	34.6 $\pm$ 1.86	6.7 $\pm$ 0.39	8.9 $\pm$ 0.47	61.4 $\pm$ 1.24	59.6 $\pm$ 0.12	52.4 $\pm$ 0.17	70.0 $\pm$ 0.14
TACREV	Qwen2.5-VL-32B	35.6 $\pm$ 3.99	8.5 $\pm$ 1.0	11.6 $\pm$ 1.35	75.8 $\pm$ 0.52	80.8 $\pm$ 0.14	74.9 $\pm$ 0.29	81.9 $\pm$ 0.13
	Llama-4-Scout-17B-16E	25.6 $\pm$ 1.63	8.1 $\pm$ 1.03	10.2 $\pm$ 0.95	73.2 $\pm$ 0.42	79.6 $\pm$ 0.16	74.0 $\pm$ 0.17	80.4 $\pm$ 0.19
	Llama-3.3-70B	29.6 $\pm$ 4.31	7.8 $\pm$ 0.58	10.2 $\pm$ 0.76	75.6 $\pm$ 0.87	80.2 $\pm$ 0.14	74.6 $\pm$ 0.09	81.2 $\pm$ 0.25
	Mixtral-8x22B	15.5 $\pm$ 0.0	6.1 $\pm$ 0.0	7.2 $\pm$ 0.0	70.5 $\pm$ 0.0	79.2 $\pm$ 0.0	73.2 $\pm$ 0.0	80.3 $\pm$ 0.0
	Llama-3.1-70B	27.8 $\pm$ 3.98	6.9 $\pm$ 1.08	9.1 $\pm$ 1.68	73.8 $\pm$ 0.9	80.0 $\pm$ 0.17	73.9 $\pm$ 0.26	81.2 $\pm$ 0.14
TACRED	Qwen2.5-VL-32B	34.2 $\pm$ 2.49	7.5 $\pm$ 1.2	10.1 $\pm$ 1.28	74.3 $\pm$ 0.43	79.2 $\pm$ 0.23	73.1 $\pm$ 0.29	80.0 $\pm$ 0.21
	Llama-4-Scout-17B-16E	23.4 $\pm$ 1.26	6.6 $\pm$ 0.76	8.5 $\pm$ 0.77	71.4 $\pm$ 0.58	78.2 $\pm$ 0.2	72.2 $\pm$ 0.18	78.7 $\pm$ 0.14
	Llama-3.3-70B	30.8 $\pm$ 2.13	6.5 $\pm$ 0.44	8.9 $\pm$ 0.66	74.3 $\pm$ 0.95	78.7 $\pm$ 0.2	72.8 $\pm$ 0.23	79.4 $\pm$ 0.24
	Mixtral-8x22B	16.3 $\pm$ 4.08	5.8 $\pm$ 0.83	6.8 $\pm$ 0.97	69.6 $\pm$ 1.02	77.9 $\pm$ 0.17	71.8 $\pm$ 0.14	78.5 $\pm$ 0.19
	Llama-3.1-70B	23.8 $\pm$ 3.0	6.4 $\pm$ 0.94	8.0 $\pm$ 1.07	72.0 $\pm$ 0.52	78.6 $\pm$ 0.12	72.2 $\pm$ 0.19	79.4 $\pm$ 0.09
Average Δ		-26.69	-58.5	-47.54	-20.49	-11.22	-17.52	-7.52

Table 6: (Ablation 2) Excluding entity-type aware filtering such that all relations were evaluated for each entity pair

context are ambiguous.

Impact of Entity-Type Awareness Without entity-type awareness, the model must evaluate all relations for each entity pair regardless of semantic plausibility. Naturally, precision suffers from false positives as the expanded search space introduces implausible relation candidates that dilute relevance elicitation. However, TACRED’s single-relation constraint implies that BRE always selects one final relation, regardless of how many relations are equally relevant to a sentence. Thus, BRE randomly selects a relation among the highest scores, and with diluted relevance scores, this random selection significantly reduces recall.

6 Conclusion and Future Work

We introduce Binary Relevance Elicitation (BRE), a zero-shot RC approach that addresses systematic over-prediction bias in LLM-based RC through simplified binary prompting and relevance-based filtering. BRE decomposes multi-class classification into independent binary decisions with quantitative relevance scoring, achieving state-of-the-art zero-shot performance on the TACRED family of datasets and matching or exceeding few-shot baselines. Notably, regression analysis reveals that parameter size, rather than model release date, is the primary driver of zero-shot relation classification performance, with $\log_{10}(\text{parameters})$ accounting for the majority of variance in outcomes; this finding, while counter to initial visual trends

across models, suggests that performance scales log-linearly with model capacity rather than being primarily attributable to architectural advances associated with newer model releases.

Our ablation studies further isolate the source of BRE’s effectiveness. Removing the binary True/False component while retaining relevance elicitation ranking results in negligible performance change across all models and datasets, whereas removing relevance elicitation and relying solely on binary classification leads to substantial degradation, particularly in precision and overall accuracy, as the model must select the correct relation from a much larger candidate pool (Appendix A.4). These findings confirm that relevance elicitation, rather than the binary verification step, is the primary contributor to BRE’s performance gains. Additionally, analysis of our parsing pipeline shows minimal invalid outputs and no-score rates across most models, confirming that observed performance trends reflect genuine model behavior rather than extraction artifacts (Appendix A.5). However, the comparatively higher tie rates observed across models indicate that relevance elicitation, while effective at narrowing the candidate set, does not always fully disambiguate among top-scoring relations, pointing to an opportunity for further refinement.

While BRE substantially improves zero-shot RC performance and provides a robust framework for NoTA handling, performance variation across relation types remains an open challenge. Future work should explore relation-specific solutions and hybrid approaches combining BRE’s filtering strengths with retrieval-augmented generation for ambiguous sentence-level cases, as well as methods for resolving residual ties within the relevance elicitation ranking. Additionally, given the small sample size underlying our parameter-size analysis ($N = 10$), further validation with a larger and more diverse set of models, alongside additional covariates such as training data quality and instruction tuning, would help confirm and refine this relationship. Despite these limitations, our work demonstrates that binary decomposition and systematic relevance filtering represent a significant advancement toward more reliable and interpretable zero-shot RC.

Limitations

BRE achieves state-of-the-art RC performance when dataset schemas provide well-documented

relation definitions and entity type specifications. Since relation definitions are incorporated directly into the prompt and valid relations are filtered based on entity types, incomplete or inconsistent metadata may reduce prompt quality and system performance, as demonstrated in our ablation studies.

Additionally, our sentence-level evaluation strategy introduces substantial computational costs. For each entity pair in a sentence, BRE queries each valid relation type individually, resulting in worst-case complexity of $\mathcal{O}(|C| \times |E|^2 \times |R|)$, where C represents the set of sentences, $E \times E$ is the set of entity pairs, and R is the set of candidate relations. While this granular approach enables interpretable relevance scoring, it may become prohibitive for large-scale deployment without optimization strategies such as batching or parallel processing.

Ethical Considerations

Our approach leverages LLM reasoning capabilities but inherits biases from their pre-training corpora, which predominantly reflect internet content from Western, English-speaking populations (Bender et al., 2021). These biases produce systematic performance disparities across languages, domains, and demographic groups (Navigli et al., 2023). Since RC commonly populates knowledge bases for downstream applications, unmitigated biases may propagate through entire pipelines.

References

- Toyin Aguda, Erik Wilson, Allan Anzagira, Simerjot Kaur, and Charese Smiley. 2025. Conservative bias in large language models: Measuring relation predictions. *arXiv preprint arXiv:2506.08120*.
- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. Taced revisited: A thorough evaluation of the taced relation extraction task. *arXiv preprint arXiv:2004.14855*.
- Hyeong-Ryeol Baek and Yong-Suk Choi. 2022. Enhancing targeted minority class prediction in sentence-level relation extraction. *Sensors*, 22(13):4911.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Mohan Raj Chanthran, Soon Lay Ki, Ong Huey Fang, and Bhawani Selvaratnam. 2026. Document-level zero-shot relation extraction with entity side information. *arXiv preprint arXiv:2601.07271*.

- Yew Ken Chia, Lidong Bing, Soujanya Poria, and Luo Si. 2022. Relationprompt: Leveraging prompts to generate synthetic data for zero-shot relation triplet extraction. *Findings of the Association for Computational Linguistics: ACL 2022*.
- Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and unfairness in information retrieval systems: New challenges in the llm era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6437–6447.
- Kartik Detroja, CK Bhensdadia, and Brijesh S Bhatt. 2023. A survey on relation extraction. *Intelligent Systems with Applications*, 19:200244.
- Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Joe Ellis, Jeremy Getman, Heather Simpson, Kira Griffith, Hoa Trang Dang, Ralph Grishman, Heng Ji, Catherine DePrince, Thomas Riese, and Neil Kuster. 2015. *TAC KBP 2015 Slot Descriptions*. Technical report, Linguistic Data Consortium. Version 1.0.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. T-rex: A large scale alignment of natural language with knowledge base triples. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Vladimir Geroimenko. 2025. Key challenges in prompt engineering. In *The Essential Guide to Prompt Engineering: Key Principles, Techniques, Challenges, and Security Risks*, pages 85–102. Springer.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. 2024. A survey on llm-as-a-judge. *The Innovation*.
- Bernal Jimenez Gutierrez, Nikolas McNeal, Clayton Washington, You Chen, Lang Li, Huan Sun, and Yu Su. 2022. Thinking about gpt-3 in-context learning for biomedical ie? think again. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4497–4512.
- Jiale Han, Shuai Zhao, Bo Cheng, Shengkun Ma, and Wei Lu. 2022. Generative prompt tuning for relation classification. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3170–3185.
- Tingting Hang, Shuting Liu, Jun Feng, Hamza Djigal, and Jun Huang. 2025. Few-shot relation extraction based on prompt learning: A taxonomy, survey, challenges and future directions. *ACM Computing Surveys*, 58(2):1–38.
- George Hannah, Jacopo de Berardinis, Terry R Payne, Valentina Tamma, Andrew Mitchell, Ellen Piercy, Ewan Johnson, Andrew Ng, Harry Rostron, and Boris Konev. 2025. Relrae: Llm-based relationship extraction, labelling, refinement, and evaluation. *arXiv preprint arXiv:2507.03829*.
- Pere-Lluís Hugué Cabot, Roberto Navigli, et al. 2021. Rebel: Relation extraction by end-to-end language generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2370–2381. Association for Computational Linguistics.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks. In *The Eleventh International Conference on Learning Representations*.
- Amirhossein Layegh, Amir H Payberah, and Mihail Matskin. 2024. Rea: Refine-estimate-answer prompting for zero-shot relation extraction. In *International Conference on Applications of Natural Language to Information Systems*, pages 301–316. Springer.
- Guozheng Li, Peng Wang, and Wenjun Ke. 2023. Revisiting large language models as zero-shot relation extractors. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6877–6892.
- Guozheng Li, Peng Wang, Jiajun Liu, Yikai Guo, Ke Ji, Ziyu Shang, and Zijie Xu. 2024. Meta in-context learning makes large language models better zero and few-shot relation extractors. *arXiv preprint arXiv:2404.17807*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM computing surveys*, 55(9):1–35.
- Yinhan Liu. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 364.
- Yubo Ma, Yixin Cao, Yong Ching Hong, and Aixin Sun. Large language model is not a good few-shot information extractor, but a good reranker for hard samples! In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in large language models: origins, inventory, and discussion. *ACM Journal of Data and Information Quality*, 15(2):1–21.
- Tapas Nayak, Navonil Majumder, Pawan Goyal, and Soujanya Poria. 2021. Deep neural approaches to relation triplets extraction: A comprehensive survey. *Cognitive Computation*, 13(5):1215–1232.
- Abiola Obamuyide and Andreas Vlachos. 2018. Zero-shot relation classification as textual entailment. In *Proceedings of the first workshop on fact extraction and VERification (FEVER)*, pages 72–78.

- Seongsik Park and Harksoo Kim. 2021. Improving sentence-level relation extraction through curriculum learning. *arXiv preprint arXiv:2107.09332*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. arxiv 2018. *arXiv preprint arXiv:1802.05365*, 12.
- Yongbin Qin, Weizhe Yang, Kai Wang, Ruizhang Huang, Feng Tian, Shaolin Ao, and Yanping Chen. 2021. Entity relation extraction based on entity indicators. *Symmetry*, 13(4):539.
- Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, et al. 2023. Mixture-of-experts meets instruction tuning: A winning combination for large language models. *arXiv preprint arXiv:2305.14705*.
- George Stoica, Emmanouil Antonios Platanios, and Barnabás Póczos. 2021. Re-tacred: Addressing shortcomings of the tacred dataset. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 13843–13850.
- Dianbo Sui, Xiangrong Zeng, Yubo Chen, Kang Liu, and Jun Zhao. 2023. Joint entity and relation extraction with set prediction networks. *IEEE transactions on neural networks and learning systems*, 35(9):12784–12795.
- Wei Tang, Benfeng Xu, Yuyue Zhao, Zhendong Mao, Yifeng Liu, Yong Liao, and Haiyong Xie. 2022. Unirel: Unified representation and interaction for joint relational triple extraction. In *Proceedings of the 2022 conference on empirical methods in natural language processing*, pages 7087–7099.
- Zhen Wan, Fei Cheng, Zhuoyuan Mao, Qianying Liu, Haiyue Song, Jiwei Li, and Sadao Kurohashi. 2023. Gpt-re: In-context learning for relation extraction using large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pages 3534–3547.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *Preprint*, arXiv:2409.12191.
- Benfeng Xu, Quan Wang, Yajuan Lyu, Dai Dai, Yongdong Zhang, and Zhendong Mao. 2023. S2ynre: Two-stage self-training with synthetic data for low-resource relation extraction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8186–8207.
- Liang Xu, Xiaoxuan Bu, and Xuetao Tian. 2024. Dynamic Prompt-Driven Zero-Shot Relation Extraction. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32:2900–2912.
- Kai Zhang, Bernal Jiménez Gutiérrez, and Yu Su. 2023a. Aligning instruction tasks unlocks large language models as zero-shot relation extractors. In *Findings of the association for computational linguistics: ACL 2023*, pages 794–812.
- Kai Zhang, Bernal Jimenez Gutierrez, and Yu Su. 2023b. Aligning instruction tasks unlocks large language models as zero-shot relation extractors. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45, Copenhagen, Denmark. Association for Computational Linguistics.
- Xiaoyan Zhao, Yang Deng, Min Yang, Lingzhi Wang, Rui Zhang, Hong Cheng, Wai Lam, Ying Shen, and Ruifeng Xu. 2024a. A comprehensive survey on relation extraction: Recent advances and new frontiers. *ACM Computing Surveys*, 56(11):1–39.
- Xiaoyan Zhao, Yang Deng, Min Yang, Lingzhi Wang, Rui Zhang, Hong Cheng, Wai Lam, Ying Shen, and Ruifeng Xu. 2024b. A comprehensive survey on relation extraction: Recent advances and new frontiers. *ACM Computing Surveys*, 56(11):1–39.
- Yifan Zheng, Wenjun Ke, Qi Liu, Yuting Yang, Ruizhuo Zhao, Dacheng Feng, Jianwei Zhang, and Zhi Fang. 2024. Making llms as fine-grained relation extraction data augmentor. In *IJCAI*, pages 6660–6668.
- Yijie Zhong, Yunfan Gao, Xiaolian Zhang, and Haofen Wang. 2025. Ooda: An ooda-driven diverse data augmentation framework for low-resource relation extraction. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 267–285.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, et al. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*.
- Sizhe Zhou, Yu Meng, Bowen Jin, and Jiawei Han. 2024. Grasping the Essentials: Tailoring Large Language Models for Zero-Shot Relation Extraction. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13462–13486. Association for Computational Linguistics.
- Wenxuan Zhou and Muhao Chen. 2022. An improved baseline for sentence-level relation extraction. In *Proceedings of the 2nd Conference of the Asia-Pacific*

Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pages 161–168.

Honglei Zhuang, Zhen Qin, Kai Hui, Junru Wu, Le Yan, Xuanhui Wang, and Michael Bendersky. 2024. Beyond yes and no: Improving zero-shot llm rankers via scoring fine-grained relevance labels. In *Proceedings of the 2024 conference of the North American chapter of the Association for Computational Linguistics: Human language technologies (volume 2: short papers)*, pages 358–370.

A Appendix

A.1 BRE Prompt Template

Input Prompt: Using only the context text provided, label the provided statements ‘True’ or ‘False’ followed by a numerical rating representing how relevant the statement is to the main idea in the context text.

Answer ‘True’ if the statement is explicitly stated in the context and meets the definition of the relationship or if the numerical rating is higher than 5. The rating is an integer ranging from 0 to 10 where 0 implies no relevance and 10 implies the statement is highly relevant to the context text.

Fill in the answer below:

context: c

statement: e_h and e_t are somehow connected through the relationship of ‘ r_i ’.

relationship definition: d_i

LLM Output Example: True, 10. The statement is supported in the provided context because...

Figure 5: Generation prompt template not including output formatting instructions for brevity

A.2 Parsing Scripts

All models were evaluated using standardized Python scripts that parse LLM outputs to extract binary predictions and relevance scores. While most models adhered to the specified output format placing answers at the beginning of responses, Llama-4-Scout exhibited poor formatting behavior, frequently positioning final answers at the end of its responses despite explicit instructions. To ensure fair comparison, our parsing scripts were adapted to handle this variation by searching the entire response for the required information. Figures 6 and 7 contain the script for all LLMs (excluding Llama-4-Scout) and then the one for Llama-4-Scout, respectively.

For relevance elicitation scoring included in Figure 8, the relevance elicitation measure is extracted

from the same LLM paragraph that contains the “true” or “false” classification, which increases the likelihood of extracting the correct relevance elicitation score. When a numeric score could not be extracted from the model output, we conservatively retained the relation as a potential candidate with a default score of 10 (the maximum). This default assignment strategy may contribute to both false positives through over-inclusion of uncertain relations (reducing precision) and missed correct relations when multiple candidates receive the default score and random selection chooses incorrectly (reducing recall). Despite these parsing challenges, the evaluation methodology remained consistent across all models, ensuring that performance differences reflect genuine model capabilities rather than evaluation artifacts.

A.3 Additional Relation-Specific Results

Figures 9 and 10 reveal a pattern: all evaluated models exhibit consistently low performance on approximately the same subset of relations across both Re-TACRED and TACREV datasets. This shared weakness pattern persists regardless of overall model accuracy differences or architectural variations, indicating that certain relations pose fundamental challenges beyond what model selection or prompting strategies can address. The consistency of this pattern across two independently revised versions of TACRED (which differ in their annotation corrections and label distributions) further confirms that these difficulties are robust to dataset-specific artifacts. Thus, while BRE’s relevance elicitation improves overall performance and NoTA handling, it cannot fully compensate for relations that lack clear distinguishing features or require deeper contextual reasoning beyond what current LLMs can reliably provide in zero-shot settings.

A.4 Additional Ablations

Based on the ablation results in Table 7, the performance differences between the full method and the score-only variant are minimal across all three datasets and all evaluated metrics. The average deltas are consistently near zero, ranging from -1.9 for macro-average recall to -0.0 for weighted-average F_1 , with micro- F_1 accuracy showing a negligible change of only -0.04 . This consistency holds across the models evaluated, from smaller models such as Llama-3.2-3B to larger ones such as Qwen2.5-VL-32B, and across the three relation extraction datasets (Re-TACRED,

```

def parse_binary_output(llm_output):
    llm_output = llm_output.replace('\n', ' ')
    parts = re.split(r'^\w\s', llm_output)
    for part in parts: # Assuming output starts response with 'true' or 'false'
        if 'true' in part.lower() or 'false' in part.lower():
            return 'true' in part.lower() and 'not true' not in part.lower(),
                ↳ int(parse_relevance_elicitation_score(part))
    return False # No 'true' or 'false' in the output

```

Figure 6: Python script to parse outputs for all LLM models except Llama-4-Scout-17B-16E

```

def parse_binary_output_scout(llm_output):
    parts = llm_output.split('\n\n')

    # If the output starts response as requested (with 'true' or 'false')
    if 'true' in parts[0].lower() or 'false' in parts[0].lower():
        return 'true' in parts[0].lower(), int(parse_relevance_elicitation_score(parts[0]))

    # Identify where the final answer is located
    phrases = ['correct answer', 'final answer']
    for part in parts:
        for phrase in phrases:
            if phrase.lower() in part.lower():
                sp = ' '.join(part.split('\n')[-2:])
                return 'true' in sp.lower(), int(parse_relevance_elicitation_score(sp.lower()))
    return False

```

Figure 7: Python script to parse outputs for Llama-4-Scout-17B-16E

```

def parse_relevance_elicitation_score(llm_output):
    parts = re.findall(r'[0-9]+', llm_output)
    if parts:
        result = parts[0]
    else:
        result = 10 # No confidence score found or we miss it; keep the result
    return result

```

Figure 8: Python script to parse outputs relevance elicitation scores for all LLM models

TACREV, and TACRED). Since the score-only configuration removes the binary True/False classification component while retaining the relevance elicitation ranking, the fact that performance remains largely stable indicates that the binary classification step contributes little additional discriminative value beyond what is already captured by the relevance elicitation scores. This provides strong evidence that the relevance elicitation mechanism itself, rather than the auxiliary binary verification, is the primary driver of the observed performance gains.

In contrast, the binary-only ablation shown in Table 8, which removes the relevance elicitation ranking and relies solely on True/False verification, exhibits substantial performance degradation across most metrics, with average deltas of -13.98 in weighted-average F_1 and -15.03 in micro- F_1 accuracy. Notably, this configuration shows a marked increase in macro-average recall ($+15.3$) alongside

significant decreases in precision (-5.66 macro, -7.46 weighted), a pattern consistent with the model effectively guessing among a much larger pool of candidate relations rather than benefiting from the reduced candidate space that relevance elicitation provides. Because the binary-only setup lacks the relevance elicitation step to first narrow down plausible candidates, the model is more likely to indiscriminately affirm many relations as ‘true’, inflating recall while severely compromising precision and overall accuracy. This stands in sharp contrast to the score-only ablation, where performance remained stable, and together these results demonstrate that relevance elicitation, rather than the binary classification step, is the primary contributor to the method’s performance, as it is responsible for effectively constraining the candidate set to a manageable and relevant subset prior to final classification.

Dataset	Model	macro-average			weighted-average			micro- F_1 Accuracy
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	
Re-TACRED	Llama-3.2-3B	17.55 $\pm$ 0.37	41.46 $\pm$ 1.33	20.79 $\pm$ 0.51	58.4 $\pm$ 0.43	45.75 $\pm$ 0.86	48.77 $\pm$ 0.74	41.94 $\pm$ 0.6
	Llama-3.1-8B	21.99 $\pm$ 0.24	55.4 $\pm$ 0.94	28.24 $\pm$ 0.29	64.38 $\pm$ 0.31	50.92 $\pm$ 0.79	50.41 $\pm$ 0.95	43.82 $\pm$ 0.94
	Llama-3.1-70B	61.43 $\pm$ 0.5	53.46 $\pm$ 1.3	54.18 $\pm$ 1.0	80.66 $\pm$ 0.22	75.87 $\pm$ 0.31	75.66 $\pm$ 0.23	79.02 $\pm$ 0.33
	Llama-3.3-70B	58.44 $\pm$ 2.02	62.75 $\pm$ 3.27	57.79 $\pm$ 1.6	73.23 $\pm$ 8.43	72.4 $\pm$ 2.85	71.38 $\pm$ 4.47	79.47 $\pm$ 0.34
	Llama-4-Scout-17B-16E	55.94 $\pm$ 0.53	72.77 $\pm$ 0.31	60.93 $\pm$ 0.27	86.62 $\pm$ 0.33	83.55 $\pm$ 0.29	84.23 $\pm$ 0.3	83.38 $\pm$ 0.32
	Mistral-7B	38.75 $\pm$ 1.82	57.77 $\pm$ 1.51	41.28 $\pm$ 2.82	67.29 $\pm$ 4.9	63.45 $\pm$ 2.83	63.62 $\pm$ 3.31	66.09 $\pm$ 1.79
	Mixtral-8x7B	47.2 $\pm$ 1.55	58.43 $\pm$ 1.16	46.92 $\pm$ 0.49	75.34 $\pm$ 0.38	72.46 $\pm$ 0.49	72.45 $\pm$ 0.32	71.86 $\pm$ 0.53
	Mixtral-8x22B	50.24 $\pm$ 0.49	60.83 $\pm$ 0.84	51.96 $\pm$ 0.39	81.98 $\pm$ 0.25	77.4 $\pm$ 0.26	78.03 $\pm$ 0.17	79.3 $\pm$ 0.02
	Qwen2.5-VL-32B	68.13 $\pm$ 1.22	66.11 $\pm$ 0.31	63.25 $\pm$ 0.36	86.86 $\pm$ 0.47	82.72 $\pm$ 0.07	83.44 $\pm$ 0.08	85.74 $\pm$ 0.08
	DeepSeek-R1-Distill-Qwen-32B	38.75 $\pm$ 0.182	57.77 $\pm$ 1.51	48.33 $\pm$ 0.55	67.29 $\pm$ 4.9	63.45 $\pm$ 2.83	74.15 $\pm$ 0.25	74.32 $\pm$ 0.48
TACREV	Llama-3.2-3B	21.04 $\pm$ 0.28	45.11 $\pm$ 1.56	24.26 $\pm$ 0.46	75.74 $\pm$ 0.23	55.01 $\pm$ 0.33	61.47 $\pm$ 0.17	49.92 $\pm$ 0.25
	Llama-3.1-8B	21.96 $\pm$ 0.52	63.56 $\pm$ 0.87	28.82 $\pm$ 0.62	79.27 $\pm$ 0.32	52.78 $\pm$ 1.13	59.06 $\pm$ 1.01	48.1 $\pm$ 1.01
	Llama-3.1-70B	57.0 $\pm$ 0.51	58.61 $\pm$ 1.91	54.0 $\pm$ 0.93	85.62 $\pm$ 0.32	87.4 $\pm$ 0.16	86.04 $\pm$ 0.18	85.66 $\pm$ 0.16
	Llama-3.3-70B	55.8 $\pm$ 0.39	66.66 $\pm$ 0.68	57.79 $\pm$ 0.21	87.72 $\pm$ 0.36	87.92 $\pm$ 0.28	87.52 $\pm$ 0.3	86.15 $\pm$ 0.23
	Llama-4-Scout-17B-16E	47.84 $\pm$ 0.89	72.33 $\pm$ 0.91	53.93 $\pm$ 0.7	90.37 $\pm$ 0.1	87.83 $\pm$ 0.09	88.32 $\pm$ 0.08	86.26 $\pm$ 0.11
	Mistral-7B	39.38 $\pm$ 0.38	69.72 $\pm$ 1.0	46.77 $\pm$ 0.48	84.84 $\pm$ 0.17	75.28 $\pm$ 0.33	78.38 $\pm$ 0.21	73.7 $\pm$ 0.26
	Mixtral-8x7B	46.72 $\pm$ 1.35	68.4 $\pm$ 0.62	52.35 $\pm$ 1.06	86.26 $\pm$ 0.23	80.48 $\pm$ 0.6	82.57 $\pm$ 0.43	78.64 $\pm$ 0.54
	Mixtral-8x22B	46.37 $\pm$ 0.4	68.2 $\pm$ 1.28	52.66 $\pm$ 0.64	87.94 $\pm$ 0.05	87.29 $\pm$ 0.16	87.22 $\pm$ 0.09	85.31 $\pm$ 0.15
	Qwen2.5-VL-32B	60.47 $\pm$ 0.42	70.87 $\pm$ 0.24	62.22 $\pm$ 0.26	91.29 $\pm$ 0.05	90.43 $\pm$ 0.01	90.47 $\pm$ 0.03	89.84 $\pm$ 0.02
	DeepSeek-R1-Distill-Qwen-32B	42.01 $\pm$ 0.51	69.07 $\pm$ 0.3	48.63 $\pm$ 0.38	85.57 $\pm$ 0.2	83.56 $\pm$ 0.38	83.97 $\pm$ 0.18	80.67 $\pm$ 0.44
TACRED	Llama-3.2-3B	22.82 $\pm$ 0.31	44.84 $\pm$ 0.88	25.59 $\pm$ 0.2	75.23 $\pm$ 0.29	54.55 $\pm$ 0.33	61.27 $\pm$ 0.29	50.04 $\pm$ 0.33
	Llama-3.1-8B	23.54 $\pm$ 0.24	58.42 $\pm$ 1.05	30.1 $\pm$ 0.22	78.12 $\pm$ 0.06	52.2 $\pm$ 0.9	58.69 $\pm$ 0.8	48.27 $\pm$ 0.82
	Llama-3.1-70B	55.0 $\pm$ 0.98	50.24 $\pm$ 1.38	50.09 $\pm$ 0.54	83.93 $\pm$ 0.17	84.71 $\pm$ 0.19	83.84 $\pm$ 0.14	83.29 $\pm$ 0.29
	Llama-3.3-70B	53.66 $\pm$ 0.29	60.72 $\pm$ 1.45	54.68 $\pm$ 1.01	86.75 $\pm$ 0.32	84.62 $\pm$ 0.07	85.42 $\pm$ 0.18	83.74 $\pm$ 0.1
	Llama-4-Scout-17B-16E	48.95 $\pm$ 0.39	62.52 $\pm$ 1.38	52.65 $\pm$ 0.56	88.09 $\pm$ 0.07	84.96 $\pm$ 0.14	86.07 $\pm$ 0.11	84.12 $\pm$ 0.15
	Mistral-7B	40.28 $\pm$ 0.26	59.79 $\pm$ 1.19	44.64 $\pm$ 0.22	83.29 $\pm$ 0.12	73.44 $\pm$ 0.24	77.04 $\pm$ 0.2	72.45 $\pm$ 0.27
	Mixtral-8x7B	47.36 $\pm$ 0.39	56.86 $\pm$ 0.93	48.4 $\pm$ 0.53	84.38 $\pm$ 0.17	78.72 $\pm$ 0.44	80.99 $\pm$ 0.32	77.29 $\pm$ 0.34
	Mixtral-8x22B	47.51 $\pm$ 0.31	59.53 $\pm$ 0.08	51.01 $\pm$ 0.01	85.94 $\pm$ 0.33	84.22 $\pm$ 0.02	84.86 $\pm$ 0.17	82.92 $\pm$ 0.01
	Qwen2.5-VL-32B	59.25 $\pm$ 0.38	57.62 $\pm$ 0.33	55.77 $\pm$ 0.12	88.42 $\pm$ 0.07	86.59 $\pm$ 0.04	87.09 $\pm$ 0.03	86.58 $\pm$ 0.03
	DeepSeek-R1-Distill-Qwen-32B	43.01 $\pm$ 0.65	59.61 $\pm$ 0.87	47.05 $\pm$ 0.44	83.75 $\pm$ 0.28	81.37 $\pm$ 0.3	82.16 $\pm$ 0.21	79.12 $\pm$ 0.35
Average Δ		-0.46	-1.9	-0.73	-0.32	-0.36	-0.0	-0.04

Table 7: Score-only ablation considers the relevance elicitation ranking and not the binary True/False response

A.5 Model Output Statistics

Our parsing scripts introduced minimal error into the evaluation pipeline, as reflected by the consistently low rates of invalid output and no-score across nearly all models and datasets in Table 9. Invalid output, which captures instances where the parser could not extract a valid True/False response, and no-score, which captures instances where the system defaulted to the relevance elicitation score due to a parsing failure, both remain below 1.5% for the majority of models, with several models such as Llama-3.3-70B and Qwen2.5-VL-32B showing rates of exactly 0.0%. The one notable exception is Llama-3.2-3B, which exhibits elevated rates of approximately 7-8% for both metrics across all three datasets, suggesting this smaller model has more difficulty consistently adhering to the expected output format; however, even this represents a relatively small fraction of overall outputs. This overall low incidence of parsing failures confirms that the performance trends observed in our ablations are attributable to the underlying methods rather than artifacts of our extraction process. In contrast, the tie rates are considerably higher, rang-

ing from roughly 8% to over 50% depending on the model and dataset. This indicates that while relevance elicitation is effective at narrowing the candidate relation set, it does not always resolve ambiguity down to a single top-scoring answer, leaving a non-trivial proportion of instances where multiple candidates receive equivalent scores. Taken together, these results suggest that relevance elicitation substantially improves the tractability of the classification task, but leaves room for improvement in fully disambiguating among the reduced candidate set.

Dataset	Model	macro-average			weighted-average			micro- F_1 Accuracy
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	
Re-TACRED	Llama-3.2-3B	15.63 $\pm$ 0.29	56.65 $\pm$ 1.12	19.66 $\pm$ 0.49	54.44 $\pm$ 0.75	28.04 $\pm$ 1.32	22.15 $\pm$ 2.12	19.49 $\pm$ 1.09
	Llama-3.1-8B	17.06 $\pm$ 0.16	78.5 $\pm$ 0.77	24.11 $\pm$ 0.23	53.0 $\pm$ 0.4	49.28 $\pm$ 2.15	38.92 $\pm$ 2.38	29.11 $\pm$ 1.38
	Llama-3.1-70B	40.67 $\pm$ 0.39	84.25 $\pm$ 0.49	51.34 $\pm$ 0.41	68.27 $\pm$ 0.23	90.32 $\pm$ 0.12	76.05 $\pm$ 0.15	72.71 $\pm$ 0.05
	Llama-3.3-70B	50.13 $\pm$ 1.54	75.92 $\pm$ 4.75	57.4 $\pm$ 2.02	61.11 $\pm$ 11.57	78.79 $\pm$ 3.94	66.8 $\pm$ 7.23	72.25 $\pm$ 2.81
	Llama-4-Scout-17B-16E	53.05 $\pm$ 0.54	85.16 $\pm$ 0.11	61.76 $\pm$ 0.37	80.43 $\pm$ 0.35	59.87 $\pm$ 1.55	64.25 $\pm$ 1.46	60.59 $\pm$ 1.31
	Mistral-7B	34.91 $\pm$ 0.73	75.57 $\pm$ 3.11	42.7 $\pm$ 0.7	54.68 $\pm$ 6.89	80.4 $\pm$ 5.29	63.98 $\pm$ 6.16	64.51 $\pm$ 1.56
	Mixtral-8x7B	39.7 $\pm$ 0.1	84.71 $\pm$ 0.19	48.96 $\pm$ 0.11	64.05 $\pm$ 0.32	84.18 $\pm$ 3.01	71.22 $\pm$ 1.17	66.4 $\pm$ 1.43
	Mixtral-8x22B	39.07 $\pm$ 0.36	83.18 $\pm$ 0.75	49.48 $\pm$ 0.45	69.41 $\pm$ 0.18	88.53 $\pm$ 0.36	75.61 $\pm$ 0.03	72.8 $\pm$ 0.22
	Qwen2.5-VL-32B	63.89 $\pm$ 0.12	72.5 $\pm$ 0.51	62.9 $\pm$ 0.55	83.98 $\pm$ 2.47	49.51 $\pm$ 11.46	56.93 $\pm$ 11.06	56.18 $\pm$ 10.34
	DeepSeek-R1-Distill-Qwen-32B	39.47 $\pm$ 0.35	78.54 $\pm$ 0.95	48.6 $\pm$ 0.5	66.3 $\pm$ 0.6	52.93 $\pm$ 4.12	54.79 $\pm$ 3.52	51.3 $\pm$ 3.22
TACREV	Llama-3.2-3B	19.31 $\pm$ 0.38	57.45 $\pm$ 1.17	22.02 $\pm$ 0.31	73.62 $\pm$ 1.58	17.49 $\pm$ 2.09	16.97 $\pm$ 3.38	16.18 $\pm$ 1.74
	Llama-3.1-8B	19.47 $\pm$ 0.32	80.96 $\pm$ 0.87	25.77 $\pm$ 0.46	70.65 $\pm$ 0.38	32.22 $\pm$ 4.02	33.2 $\pm$ 5.15	24.92 $\pm$ 2.69
	Llama-3.1-70B	40.38 $\pm$ 0.31	88.36 $\pm$ 0.87	51.38 $\pm$ 0.36	76.85 $\pm$ 0.11	96.98 $\pm$ 0.17	84.98 $\pm$ 0.07	79.63 $\pm$ 0.05
	Llama-3.3-70B	49.36 $\pm$ 0.31	86.53 $\pm$ 0.06	59.24 $\pm$ 0.34	79.21 $\pm$ 0.22	94.94 $\pm$ 1.27	85.88 $\pm$ 0.42	82.52 $\pm$ 0.45
	Llama-4-Scout-17B-16E	48.88 $\pm$ 0.69	87.32 $\pm$ 0.59	58.15 $\pm$ 0.56	89.07 $\pm$ 0.63	53.14 $\pm$ 2.78	62.3 $\pm$ 2.58	61.07 $\pm$ 2.14
	Mistral-7B	36.09 $\pm$ 0.13	82.59 $\pm$ 0.3	46.07 $\pm$ 0.13	75.08 $\pm$ 0.24	89.53 $\pm$ 1.62	80.65 $\pm$ 0.63	73.48 $\pm$ 0.68
	Mixtral-8x7B	41.03 $\pm$ 0.23	87.94 $\pm$ 0.2	52.17 $\pm$ 0.18	77.73 $\pm$ 0.76	85.65 $\pm$ 3.17	79.98 $\pm$ 1.29	71.73 $\pm$ 1.35
	Mixtral-8x22B	39.28 $\pm$ 0.11	88.31 $\pm$ 0.31	50.78 $\pm$ 0.11	77.25 $\pm$ 0.08	95.68 $\pm$ 0.27	84.67 $\pm$ 0.1	80.41 $\pm$ 0.1
	Qwen2.5-VL-32B	59.76 $\pm$ 0.21	80.59 $\pm$ 0.39	63.98 $\pm$ 0.28	90.68 $\pm$ 1.55	43.26 $\pm$ 10.7	53.41 $\pm$ 10.85	54.92 $\pm$ 9.25
	DeepSeek-R1-Distill-Qwen-32B	38.67 $\pm$ 0.44	83.98 $\pm$ 0.75	48.3 $\pm$ 0.46	76.51 $\pm$ 0.52	48.83 $\pm$ 3.92	55.25 $\pm$ 3.44	52.33 $\pm$ 2.86
TACRED	Llama-3.2-3B	20.76 $\pm$ 0.29	55.66 $\pm$ 1.68	23.43 $\pm$ 0.36	73.44 $\pm$ 0.28	18.93 $\pm$ 2.44	19.21 $\pm$ 4.2	18.2 $\pm$ 2.07
	Llama-3.1-8B	20.55 $\pm$ 0.07	74.94 $\pm$ 1.09	26.96 $\pm$ 0.19	71.65 $\pm$ 0.45	34.98 $\pm$ 2.75	36.99 $\pm$ 3.36	27.5 $\pm$ 1.77
	Llama-3.1-70B	40.81 $\pm$ 0.42	76.33 $\pm$ 0.82	50.13 $\pm$ 0.32	75.71 $\pm$ 0.15	94.96 $\pm$ 0.28	83.68 $\pm$ 0.04	78.57 $\pm$ 0.05
	Llama-3.3-70B	48.72 $\pm$ 0.11	74.31 $\pm$ 0.07	56.01 $\pm$ 0.14	79.16 $\pm$ 0.31	89.72 $\pm$ 0.59	83.73 $\pm$ 0.44	80.42 $\pm$ 0.41
	Llama-4-Scout-17B-16E	49.21 $\pm$ 0.34	74.8 $\pm$ 0.7	55.84 $\pm$ 0.37	87.76 $\pm$ 0.42	49.09 $\pm$ 2.03	59.16 $\pm$ 1.96	58.32 $\pm$ 1.62
	Mistral-7B	37.06 $\pm$ 0.09	74.31 $\pm$ 0.22	45.14 $\pm$ 0.11	74.15 $\pm$ 0.27	88.72 $\pm$ 2.56	79.93 $\pm$ 1.0	72.93 $\pm$ 1.06
	Mixtral-8x7B	42.25 $\pm$ 0.06	78.54 $\pm$ 0.39	50.82 $\pm$ 0.13	75.84 $\pm$ 0.33	86.55 $\pm$ 2.69	79.78 $\pm$ 1.32	71.96 $\pm$ 1.3
	Mixtral-8x22B	40.0 $\pm$ 0.13	77.67 $\pm$ 0.06	49.89 $\pm$ 0.2	76.83 $\pm$ 0.5	91.26 $\pm$ 1.87	82.78 $\pm$ 0.54	78.43 $\pm$ 0.59
	Qwen2.5-VL-32B	57.83 $\pm$ 0.09	66.26 $\pm$ 0.31	58.05 $\pm$ 0.27	86.84 $\pm$ 2.05	54.6 $\pm$ 9.05	64.84 $\pm$ 7.41	64.36 $\pm$ 6.68
	DeepSeek-R1-Distill-Qwen-32B	39.54 $\pm$ 0.33	73.66 $\pm$ 0.96	47.58 $\pm$ 0.51	76.59 $\pm$ 0.73	50.28 $\pm$ 3.26	57.27 $\pm$ 2.83	54.08 $\pm$ 2.35
Average Δ		-5.66	+15.3	-0.95	-7.46	-9.11	-13.98	-15.03

Table 8: Binary-only ablation considers the True/False binary response and not the relevance elicitation ranking

Dataset	Model	Invalid Output	No-Score	Tie Rates
Re-TACRED	Llama-3.2-3B	6.875 $\pm$ 0.092	6.7704 $\pm$ 0.0754	33.537 $\pm$ 0.4
	Llama-3.1-8B	0.9029 $\pm$ 0.0897	0.8203 $\pm$ 0.0635	53.6742 $\pm$ 1.12
	Llama-3.1-70B	1.1672 $\pm$ 0.0161	1.1638 $\pm$ 0.0175	50.8824 $\pm$ 0.2
	Llama-3.3-70B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	36.7596 $\pm$ 2.83
	Llama-4-Scout	0.5486 $\pm$ 0.1732	0.1814 $\pm$ 0.1323	20.8321 $\pm$ 0.49
	Mistral-7B	0.1472 $\pm$ 0.0116	0.0638 $\pm$ 0.0048	48.4066 $\pm$ 4.47
	Mixtral-8x7B	0.0047 $\pm$ 0.0031	0.0151 $\pm$ 0.0019	50.3324 $\pm$ 2.19
	Mixtral-8x22B	0.5088 $\pm$ 0.012	0.7329 $\pm$ 0.0053	46.5395 $\pm$ 0.6
	Qwen2.5-VL-32B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	14.1511 $\pm$ 4.73
	DeepSeek-R1-Distill-Qwen-32B	0.0033 $\pm$ 0.0014	0.048 $\pm$ 0.0102	24.0543 $\pm$ 2.23
TACREV	Llama-3.2-3B	7.8906 $\pm$ 0.1481	7.1393 $\pm$ 0.0721	25.1931 $\pm$ 0.73
	Llama-3.1-8B	1.1367 $\pm$ 0.1694	0.9979 $\pm$ 0.1001	39.2765 $\pm$ 1.15
	Llama-3.1-70B	1.0822 $\pm$ 0.0582	1.0801 $\pm$ 0.0567	40.7312 $\pm$ 0.3
	Llama-3.3-70B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	34.1769 $\pm$ 0.67
	Llama-4-Scout	1.665 $\pm$ 0.6556	1.2021 $\pm$ 0.681	12.0653 $\pm$ 0.75
	Mistral-7B	0.1417 $\pm$ 0.0051	0.0748 $\pm$ 0.0089	39.1463 $\pm$ 1.19
	Mixtral-8x7B	0.0044 $\pm$ 0.0028	0.0118 $\pm$ 0.0048	37.6169 $\pm$ 2.12
	Mixtral-8x22B	0.5256 $\pm$ 0.016	0.7697 $\pm$ 0.0339	37.4573 $\pm$ 0.3
	Qwen2.5-VL-32B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	7.9154 $\pm$ 2.08
	DeepSeek-R1-Distill-Qwen-32B	0.0047 $\pm$ 0.002	0.507 $\pm$ 0.0042	17.6066 $\pm$ 1.43
TACRED	Llama-3.2-3B	8.0146 $\pm$ 0.1861	7.3173 $\pm$ 0.07	25.9359 $\pm$ 0.82
	Llama-3.1-8B	1.118 $\pm$ 0.0914	0.977 $\pm$ 0.05	40.8356 $\pm$ 1.11
	Llama-3.1-70B	1.1151 $\pm$ 0.0371	1.1104 $\pm$ 0.04	40.8124 $\pm$ 0.41
	Llama-3.3-70B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	30.8047 $\pm$ 0.09
	Llama-4-Scout	1.4025 $\pm$ 0.5118	0.9141 $\pm$ 0.52	11.6436 $\pm$ 0.72
	Mistral-7B	0.1531 $\pm$ 0.0085	0.0806 $\pm$ 0.0097	39.4906 $\pm$ 1.67
	Mixtral-8x7B	0.0061 $\pm$ 0.0034	0.0128 $\pm$ 0.00414	39.3088 $\pm$ 0.98
	Mixtral-8x22B	0.5162 $\pm$ 0.0131	0.7156 $\pm$ 0.000725	35.1699 $\pm$ 1.5
	Qwen2.5-VL-32B	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	11.6616 $\pm$ 2.64
	DeepSeek-R1-Distill-Qwen-32B	0.0052 $\pm$ 0.0017	0.0574 $\pm$ 0.0089	18.8329 $\pm$ 1.39

Table 9: Reported average percentage % $\pm$ standard deviation




Figure 9: F_1 scores for the varying relation types in the Re-TACRED dataset. The relations are sorted based on their F_1 and the models are sorted based on their macro- F_1 scores.




Figure 10: F_1 scores for the varying relation types in the TACREV dataset. The relations are sorted based on their F_1 and the models are sorted based on their macro- F_1 scores.