EERelRank: Entity-Enhanced Relation Ranking using LLMs

Anonymous ACL submission

Abstract

Large Language Models (LLMs), such as GPT-3.5 and GPT-4, have demonstrated remarkable performance in various natural language processing tasks. However, their efficacy in specialized domains, particularly in challenging tasks like relation extraction in finance, remains limited. To tackle the complexity of understanding domain-specific relations and the semantic ambiguity among relation types, we propose incorporating semantics into relation labels, referred to as "definitions", enabling a more nuanced understanding of closely related concepts. Consequently, we recast the problem as a ranking task, wherein relations are ranked based on their definitions, and the top-ranked relation is selected as the prediction. Through extensive experiments, we demonstrate that our model excels in handling domain-specific and closely related semantic relations, and achieves state-of-the-art performance on the largest financial relation extraction dataset, REFinD.[1]

1 Background

Recent approaches in natural language processing (NLP) have been significantly shaped by the emergence of large language models (LLMs) like ChatGPT (GPT-3.5) (OpenAI, 2023a) and GPT4 (OpenAI, 2023b). These LLMs have revolutionized the approach to tackling NLP tasks by leveraging prompts or instructions in-context (Brown et al., 2020). Although they have demonstrated remarkable performance across various NLP tasks (Radford et al., 2019), their efficacy in information retrieval tasks within specialized domains such as finance is limited (Li et al., 2023). For instance, the pre-trained language encoder LUKE-base (Yamada et al., 2020) surpasses both GPT4 and ChatGPT by at least 26.77 in Macro F1 for the relation extraction task (Kaur et al., 2023; Li et al., 2023). As a result, there is a growing interest in developing task-specific architectures based on pre-trained language encoders (Vaswani et al., 2017), and fine-tuning these models on specialized domain data (Tang et al., 2023). Hence this paper focuses on utilizing pre-trained language encoders to address the relation extraction task in the financial domain.

The objective of relation extraction is to discern the semantic relationship between entities in text, such as identifying the "acquired_by" relation between "Microsoft" and "Github" from the text "Microsoft is the largest stakeholder in Github, having acquired 100% of its shares.[2] This area has significant applications in finance, as the extracted information can provide valuable insights into an entity (e.g., determining if Github's revenue will increase due to new leadership and available resources). Although some datasets have been developed for this task (Sharma et al., 2022; Kaur et al., 2023; Wu et al., 2020; Jabbari et al., 2020), the REFinD dataset by Kaur et al. (2023) poses a unique challenge. Specifically, the dataset encompasses a wide range of relation types, including closely related semantic relations like "shares_of" and "acquired_by", as found between "Microsoft" and "Github" in our running example. This presents a challenge, as the definitions used to label instances of these relations may have overlapping properties or be subjectively interpreted, resulting in semantic ambiguity (Kaur et al., 2023). Furthermore, understanding the meaning of domain-specific relations can be difficult, which adds to the complexity. Due to semantic ambiguity of some relation types and the challenge of grasping domain-specific relations, the current state-of-the-art still remains LUKE-base or LUKE-large (Yamada et al., 2020), as reported by Kaur et al. (2023).

To address this challenge, we consider the unique discriminative properties of each relation by

[1] Our code will be made publicly available upon acceptance.

[2] We use this text as our running example.

incorporating label semantics (Zhang et al., 2022; Aly et al., 2021). Incorporating label semantics represents a significant paradigm shift from the traditional approach of treating labels as mere symbols. This has particularly been transformative for tasks such as semantic role labeling (Zhang et al., 2022), named entity recognition (Aly et al., 2021) and entity linking (Logeswaran et al., 2019), by providing more nuanced understanding of labels. However, in our case, manually defining these relations may necessitate expert human annotators in the finance domain, which could be challenging to acquire. As a solution, we employ LLMs and utilize a natural language description of the task (zero-shot prompt) to generate relation definitions. Consequently, we recast relation extraction as a ranking problem. This approach is more fitting as it assesses the likelihood of the text conveying a specific relation in a relative sense rather than an absolute one, as demonstrated in the running example. To achieve this, we perform relation ranking in a dense space defined by a Bi-encoder (Humeau et al., 2019), which independently embeds the input text and definitions of the relation using pre-trained language encoders (Yamada et al., 2020). In summary, the contributions of this paper include:

- Incorporating label semantics to address the challenge of relation extraction within financial domain
- Introducing EERelRank, an Entity Enhanced Relation Ranking model that leverage LLMs to generate relation definitions and solves the problem as a ranking task.

2 The EERelRank Approach

Problem Statement Given a sentence $X = \{w_i\}_{i=1}^{|X|}$, the entity e is a word phrase present within the sequence of words in X. Given two entities e_1 and e_2 of sentence X, we denote their start and end positions using special entity tags $[e_1], [/e_1]$ and $[e_2], [/e_2]$, respectively. The objective of relation extraction is, given sentence X, to predict the relationship between entities, represented as $(t_1, t_2, rel) \in \mathcal{R}$, where t_1 and t_2 correspond to the respective entity types of e_1 and e_2 respectively, rel signifies the relation between these entities, and $\mathcal{R} = \{r_i\}_{i=1}^{|\mathcal{R}|}$ is a predefined set of possible entity relation types. For instance, the entity relation (*person, organization, employee_of*) $\in \mathcal{R}$ implies that a person is employed by an organization, with both entities being referenced within the sentence.

Although this problem is usually solved as a classification task, we now cast the problem as a ranking task. That is, given sentence X, rank all the relation types $\mathcal{R}$ according to the similarities between X and the definition of each relation type, denoted as $D_r = \{w_i\}_{i=1}^{|D_r|}$. The relation type whose definition aligns most closely with X is then selected as the final prediction for that sentence. To improve the contextual understanding of X and to better align it with its corresponding relation definition, we incorporate entity types t_1 and t_2 in X. Then the sentence X is formatted as follows:

$$X = \{w_1, \ldots, t_1, [e_1], e_1, [/e_1], \ldots, t_2, [e_2], e_2, [/e_2], \ldots, w_{|X|}\} \quad (1)$$

2.1 Generating Relation Definitions

We consider a generation problem, where we have access to the entity relation space $\mathcal{R}$ and a generator g_r. The function g_r is designed to transform each relation $r \in \mathcal{R}$ into a descriptive definition $D_r = g_r(p_r)$, where p_r is a prompt corresponding to r. We utilize an LLM, specifically GPT-4, to model g_r. We feed this model with a zero-shot prompt, relying entirely on the pre-existing knowledge of the LLM about the entity relations. This approach does not require any additional training, making use of the inherent capabilities of the LLM to understand and generate descriptions of entity relations.

Zero-shot Prompt *Given the relation type {t_1,t_2,rel}, where {t_1} and {t_2} are subject and object entity types, define the relation {rel} that exist between these entity types in the financial domain.*

While other types of prompts such as few-shot (Brown et al., 2020) and chain-of-thought (Wei et al., 2022) have demonstrated improvements in generation, the selection of the most appropriate examples that align with the dataset's distribution remains a significant research challenge (Zhao et al., 2021). Furthermore, the chain-of-thought approach necessitates some pre-existing knowledge of the relations as they are interpreted in specific domains, such as finance. For instance, the acquisition of more than 50% of a company's shares typically signifies a company takeover. Therefore, to effectively predict these relationships without the need for explicit prompts for each entity relation, we opt for a zero-shot approach.

2.2 Bi-Encoder Architecture

We employ a Bi-encoder (Humeau et al., 2019), which consists of two separate encoders ENC_{θ_X} and $\text{ENC}_{\theta_{D_r}}$. Both ENC_{θ_X} and $\text{ENC}_{\theta_{D_r}}$ are responsible for encoding the text X and each relation definition D_r, respectively, into the same representation space $\mathbb{R}^d$, where d is the dimension of the space. While both encoders share the same architecture, they have independent parameters θ_X and θ_{D_r}. Similar to previous work (Kaur et al., 2023), we use LUKE-base or LUKE-large to model the encoders.[3] Accordingly, we use the special tokens [CLS] and [SEP] at the beginning and end of both X and D_r.[4] We then apply the respective encoders on these text pieces as follows:

$$\mathbf{H}^X = \text{ENC}_{\theta_X}(X), \quad \mathbf{H}^{D_r} = \text{ENC}_{\theta_{D_r}}(D_r) \quad (2)$$

where $\mathbf{H}^X \in \mathbb{R}^{|X| \times d}$ and $\mathbf{H}^{D_r} \in \mathbb{R}^{|D_r| \times d}$ are the token-level embeddings extracted at the last layer of the respective encoders. Unlike other relation extraction models (Soares et al., 2019), we do not use the concatenation of the heads of the special entity tags $[e_1]$ and $[e_2]$ to represent the text X. This is because our goal is to compute the similarity between the text and relation definition, which requires that both text pieces have the same embedding dimension. The concatenation operator however doubles the dimension of the text embedding, requiring dimensionality reduction solutions to facilitate the similarity computation. From our preliminary investigations and experiments, using different dimensionality reduction solutions such as linear transformation, mean-pool or max-pool may compromise the availability of the pretrained knowledge or increase the training difficulty of the encoders. Therefore, we use the [CLS] representations of $\mathbf{H}^X$ and $\mathbf{H}^{D_r}$, denoted as $\mathbf{h}^X_{[\text{CLS}]} \in \mathbb{R}^d$ and $\mathbf{h}^{D_r}_{[\text{CLS}]} \in \mathbb{R}^d$, to represent the text and relation definition embeddings, respectively.

2.3 Relation Scoring and Training Objective

Given the [CLS] embeddings for the text $\mathbf{h}^X_{[\text{CLS}]}$ and each relation definition $\mathbf{h}^{D_r}_{[\text{CLS}]}$, the task of ranking relations is reduced to computing a similarity score between these embeddings. The similarity score for the entity relation $r \in \mathcal{R}$ with respect to X is then given by the dot product:

$$f(X, D_r) = (\mathbf{h}^X_{[\text{CLS}]})^T \mathbf{h}^{D_r}_{[\text{CLS}]} \quad (3)$$

where $f(\cdot;\cdot)$ denote a dot product function. Let r^+ be the positive entity relation in $\mathcal{R}$ that aligns with X. We then define the other $|\mathcal{R}| - 1$ relations in $\mathcal{R}$ as negative entity relations, denoted as r^-. Our goal is to achieve high scores for $f(X, D_{r^+})$ and low scores for all $f(X, D_{r^-})$. Then for each positive pair (X, D_{r^+}) and associated $|\mathcal{R}| - 1$ negative pairs (X, D_{r^-}), the optimization objective is formulated as minimizing the NCE loss function $\mathcal{L}$ (Oord et al., 2018):

$$\mathcal{L} = -\log \frac{\exp(f(X, D_{r^+}))}{\exp(f(X, D_{r^+})) + \sum_{r^- \in \mathcal{R}} \exp(f(X, D_{r^-}))} \quad (4)$$

Note that if $\mathcal{R}$ is a large set, then negative sampling becomes necessary to reduce the computational cost.

3 Experimental Results

We now present evaluation results of our proposed models, LB-EERelRank and LL-EERelRank,[5] along with its ablated models LB-RelRank that excludes entity types in text. As a baseline we ran the LUKE models and reported the results.

Experimental Setup We use the REFinD dataset (Kaur et al., 2023), which consists of a set of 22 entity relation types, including a "no_relation" to indicate that no relation exists between two marked entities. In total, there are 28,676 instances across these entity relation types, making it the largest financial relation extraction dataset to date. We present the detailed data statistics, along with the relation definitions generated by GPT-4 in Appendix A. We use the Macro-F1 as the evaluation metric and evaluate on the test set. All hyperparameters are set empirically and manually tuned on the development set to select the best model.[6]

[3] We use LUKE-base/large since it achieves SOTA on relation extraction.

[4] For brevity, we use the same notation X and D_r to denote the text and relation definition with special tokens

[5] We use the notations LB and LL to represent LUKE-base and LUKE-large pre-trained encoders, respectively.

[6] Following previous work (Kaur et al., 2023), we use the LUKE-base/large pretrained model as a backbone for our models. The max length for text and relation definitions are 256 and 70 respectively, with a batch size of 8. We use an Adam optimizer and search for learning rate (lr) in {5e-6, ... 1e-5}. Best lr for LB-RelRank and LL-RelRank is 1e-5, while LB-EERelRank and LL-EERelRank is 9e-6. All models are built with Pytorch are trained in parallel on a g5.24xlarge GPU AWS instance.

Model	Macro-F1
LB	56.3 ± 2.1
LB-RelRank	57.3 ± 1.1
LB-EERelRank	$\mathbf{61.2 \pm 2.1}$
LL	56.5 ± 1.1
LL-RelRank	58.6 ± 2.2
LL-EERelRank	$\mathbf{61.4 \pm 0.6}$

Table 1: Performance comparison among models

Main Results Our primary findings are consolidated in Table 1. The results presented are the average of five iterations, each using a unique random seed. The highest performing results are highlighted in bold. The prediction performance is evaluated in both LUKE-base (LB) and LUKE-large (LL) settings. Our preliminary observations indicate that LB-RelRank/LL-RelRank, which integrates relation definitions, achieves an improvement of at least 1% over the baseline models LB and LL. Furthermore, the inclusion of entity type information in the text assists the model in aligning more accurately with the relation definitions, resulting in an additional improvement of at least 2.8%. In subsequent experiments (refer to Section 3), we will demonstrate how this strategy enables the model to achieve such significant gains by distinguishing between closely related semantic relations. In summary, our model outperforms the current state-of-the-art on the REFinD dataset, surpassing the benchmark set in (Kaur et al., 2023), thereby establishing a new standard in performance.

	LB	LB-RelRank	LB-EERelRank
ORG/ ORG	43.5	54.0	**62.6**
ORG/ GPE	83.6	84.0	**86.5**
ORG/ MONEY	41.9	**45.7**	42.8
ORG/ DATE	**39.1**	34.6	27.2
PER/ ORG	81.9	**85.9**	85.7
PER/ UNIV	57.0	48.2	**68.6**
PER/ GOV	75.0	**76.2**	**75.0**
PER/ TITLE	28.3	**43.1**	41.4

Table 2: Macro F1 performance for different entity groups.

Performance on Entity Groups Instead of validating our model across all entity relations, we aim to assess its performance on various entity groups. Accordingly, we applied our best models for LB, LB-RelRank, and LB-EERelRank to each entity group to observe their performance. Table 2 presents the results. Overall, we observe that LB performs best on 1, LB-RelRank on 4 and LB-EERelRank on 5 entity groups. While these results indicate the importance of incorporating relation definintions and entity types, we discovered that the high gains achieved by our models were on entity groups that are more specialized on the financial domain, e.g, ORG/ORG entity group.

Case Study: We conducted a case study to illustrate the significance of relation definitions and integrating entity types. Our findings revealed instances where LB failed to infer certain relations, such as (org, money, profit_of) and (org, gpe, operations_in), while LB-RelRank successfully predicted these relations by integrating relation definitions. For instance, specific phrases like 'Profit' and 'lending activity,' contextually referring to subject and object entities, are semantically close to the phrases 'net earnings' and 'business activity' found within the definitions of (org, money, profit_of) and (org, gpe, operations_in) respectively. This allowed LB-RelRank to accurately make these inferences. However, in some cases, LB-RelRank failed to identify correct relations, whereas LB-EERelRank accurately predicted them by incorporating entity type information in the text. We believe that incorporating entity-type information guided the model in identifying the correct relation within the entity group. Detailed case examples illustrating these observations are available in Table 3 in the Appendix.

Confusion Matrix We analyze the prediction errors of LB-RelRank and LB-RelRank (types). For LB-RelRank, the majority of errors are observed in the (org, gpe), (per, org), (per, univ), and (per, gov) groups, aligning with the findings presented in Table 2. Upon incorporating entity types, the LB-RelRank (types) model effectively distinguishes between relations in (per, gov) versus (per, org) and (per, org) versus (per, univ), thereby reducing errors. Comprehensive insights into these findings are detailed in Figures 1 and 2 in the Appendix.

4 Conclusion and Future Work

We proposed a bi-encoder model for relation extraction in the finance domain, named EERelRank, which incorporates both entity types and relation definitions from LLMs. Unlike previous SOTA (Kaur et al., 2023) that models the problem as a classification task, we treat the problem as a ranking task. We apply EERelRank on the ReFinD dataset using both LUKE-base/large encoders, and our results achieve state-of-the-art. In the future, we intend to explore other advanced prompt settings to obtain more expressive relation definitions.

5 Limitations

The potential limitation of our work is that the relation descriptions generated by these LLMs are highly sensitive to the choice of verbalizers in the prompt (Webson and Pavlick, 2022). In the future, we will consider more robust approaches, for example, few-shot and chain-of-thought to interact with LLMs.

6 Ethics Statement

This work relies on the use of already published dataset REFinD, containing text obtained from the SEC website. SEC filings are freely and publicly available, but do contain names and other identifying details such as current and former job titles, schools attended, and board or professional associations for typically high-ranking officers in publicly traded companies. This information is similar to what could be found about public figures mentioned in websites such as Wikipedia and thus we do not anticipate any harm to persons mentioned in our data beyond what could be learned from reading the public financial statements themselves. Additionally, we have made no attempt to aggregate or include non-public information about any individual or entity in this research. As REFinD employed in this research is a public dataset derived from these reports, it ensures no harm to individuals or disclosure of proprietary information.

This paper also utilizes Large Language Models (LLM) for content generation, raising concerns about the possibility of generating text with misinformation and hallucinations. Apart from concerns associated with LLM usage, we do not foresee any additional ethical issues arising from this study.

References

Rami Aly, Andreas Vlachos, and Ryan McDonald. 2021. Leveraging type descriptions for zero-shot named entity recognition and classification. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1516–1528, Online. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. Poly-encoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring. *International Conference on Learning Representations*.

Ali Jabbari, Olivier Sauvage, Hamada Zeine, and Hamza Chergui. 2020. A french corpus and annotation schema for named entity recognition and relation extraction of financial news. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2293–2299.

Simerjot Kaur, Charese Smiley, Akshat Gupta, Joy Sain, Dongsheng Wang, Suchetha Siddagangappa, Toyin Aguda, and Sameena Shah. 2023. Refind: Relation extraction financial dataset. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, page 3054–3063, New York, NY, USA. Association for Computing Machinery.

Xianzhi Li, Samuel Chan, Xiaodan Zhu, Yulong Pei, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah. 2023. Are chatgpt and gpt-4 general-purpose solvers for financial text analytics? a study on several typical tasks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 408–422.

Lajanugen Logeswaran, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, Jacob Devlin, and Honglak Lee. 2019. Zero-shot entity linking by reading entity descriptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3449–3460, Florence, Italy. Association for Computational Linguistics.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.

OpenAI. 2023a. Chatgpt. `https://openai.com/blog/chatgpt/`. Accessed: [2023].

OpenAI. 2023b. Gpt-4 technical report. `https://cdn.openai.com/papers/gpt-4.pdf`. Accessed: [2023].

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Soumya Sharma, Tapas Nayak, Arusarka Bose, Ajay Kumar Meena, Koustuv Dasgupta, Niloy Ganguly, and Pawan Goyal. 2022. Finred: A dataset for relation extraction in financial domain. In *Companion Proceedings of the Web Conference 2022*, pages 595–597.

Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. 2023. Does synthetic data generation of llms help clinical text mining? *arXiv e-prints*, pages arXiv–2303.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Albert Webson and Ellie Pavlick. 2022. Do prompt-based models really understand the meaning of their prompts? In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Haoyu Wu, Qing Lei, Xinyue Zhang, and Zhengqian Luo. 2020. Creating a large-scale financial news corpus for relation extraction. In *2020 3rd International Conference on Artificial Intelligence and Big Data (ICAIBD)*, pages 259–263. IEEE.

Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. Luke: Deep contextualized entity representations with entity-aware self-attention. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454.

Li Zhang, Ishan Jindal, and Yunyao Li. 2022. Label definitions improve semantic role labeling. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5613–5620.

Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.

A Appendix

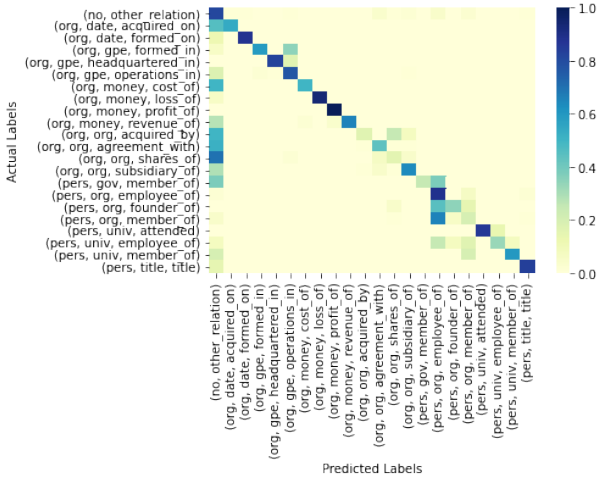

Figure 1: Confusion matrix for each relation in REFinD based on LB-RelRank predictions (normalized)

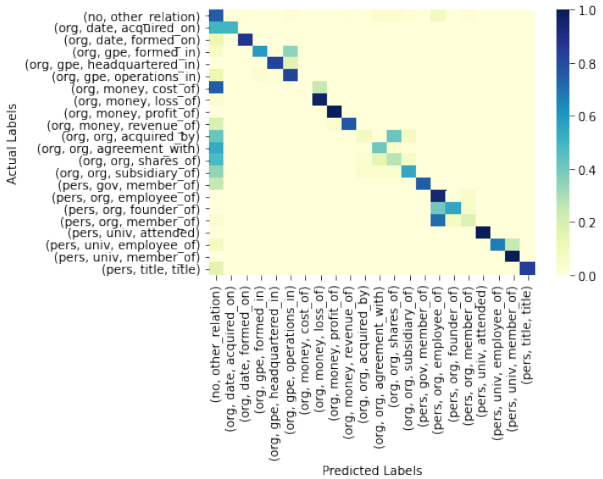

Figure 2: Confusion matrix for each relation in REFinD based on LB-EERelRank predictions (normalized)

Feature Importance	Input Sentence	Actual Label	Predicted Label		
			LB	LB-RelRank	LB-EERelRank
Effect of Relation Definition	**Makkantotti Group Corp.** revenues for the nine months period ended December 31, 2015 were $8,405; Cost of Sales was $912 for the nine months period ended December 31, 2015, and Gross Profit was **$7,493**.	(org, money, profit_of)	(no, other_relation)	(org, money, profit_of)	-
	Sunnyside Bancorp, Inc.'s lending activity is primarily concentrated in loans collateralized by real estate in **the State of New York.**	(org, gpe, operations_in)	(no, other_relation)	(org, gpe, operations_in)	-
Effect of Entity Type Marker	**Capstone Systems Inc** are entirely dependent on the services of Capstone Systems Inc sole officer and director, **Jure Perko**, to build Capstone Systems Inc customer base.	(pers, org, employee_of)	-	(no, other_relation)	(pers, org, employee_of)
	Mr. **Fischer** also was active in residential real estate from 1989 to 1997, and since 1980, has had an active **real estate broker**'s license.	(pers, title, title)	-	(no, other_relation)	(pers, title, title)

Table 3: Qualitative analysis of our proposed model with the baseline model. The **subject** and **object** entities are highlighted in the input sentences.

Entity Relation	#Train	#Dev	#Test	#Total
(no, other_relation)	9128	1965	1953	13046
(pers, title, title)	3126	671	671	4468
(org, gpe, operations_in)	2832	606	605	4043
(pers, org, employee_of)	1733	372	374	2479
(org, org, agreement_with)	653	141	141	935
(org, date, formed_on)	448	96	96	640
(pers, org, member_of)	441	94	95	630
(org, org, subsidiary_of)	386	82	83	551
(org, org, shares_of)	286	61	61	408
(org, money, revenue_of)	217	47	47	311
(org, money, loss_of)	141	30	31	202
(org, gpe, headquartered_in)	135	29	29	193
(org, date, acquired_on)	134	28	24	186
(pers, org, founder_of)	92	19	20	131
(org, gpe, formed_in)	81	17	17	115
(org, org, acquired_by)	55	11	12	78
(pers, univ, employee_of)	53	11	12	76
(pers, gov, agy, member_of)	40	8	8	56
(pers, univ, attended)	30	6	7	43
(pers, univ, member_of)	23	5	5	33
(org, money, profit_of)	20	4	5	29
(org, money, cost_of)	16	3	4	23
Total	**20070**	**4306**	**4300**	**28676**

Table 4: REFinD: dataset statistics for the individual entity relation types.

Entity Relation	Definition
(no, other_relation)	This relation means no relation between two entities.
(pers, title, title)	In the financial domain, this relation signifies that a PERSON holds a specific TITLE or position within a financial organization or institution. It indicates the role or rank the person occupies, such as CEO, CFO, or Financial Analyst.
(org, gpe, operations_in)	This relation indicates the geographical areas where a financial ORGANIZATION conducts its business activities or has a significant presence. It could refer to cities, countries, or regions where the organization's services and products are available or where it has business facilities.
(pers, org, employee_of)	This relation illustrates that a PERSON is an employee of a particular ORGANIZATION within the financial sector. The organization could be banks, investment firms, or any financial institutions.
(org, org, agreement_with)	Refers to contractual or mutual agreements made between two organizations for various purposes like partnerships or collaborations within the financial domain.
(org, date, formed_on)	This relation specifies the date on which an ORGANIZATION in the financial domain was established or created. It marks the inception point of the organization, signifying the beginning of its operational existence in the financial sector.
(pers, org, member_of)	In the financial context, this relation signifies that a PERSON is a member of an ORGANIZATION or GOV. AGENCY related to finance, such as financial regulatory bodies or financial committees.
(org, org, subsidiary_of)	This relation denotes that a financial ORGANIZATION is a subsidiary of another, meaning it is controlled or owned to a significant extent by the other organization. It establishes a hierarchical relationship between two entities.
(org, org, shares_of)	This relation indicates that a financial ORGANIZATION holds shares or partial ownership in another organization. It signifies an investment or financial stake in another entity.
(org, money, revenue_of)	This relation represents the total income or earnings generated by a financial ORGANIZATION within a specific period. It reflects the organization's financial inflow from its business activities.
(org, money, loss_of)	This relation indicates the financial deficit experienced by an ORGANIZATION when its total costs exceed total revenues. It represents a negative financial result from business activities.
(org, gpe, headquartered_in)	This relation specifies the location where the main office or central administration of a financial ORGANIZATION is situated. It is the primary place where key managerial decisions of the organization are made, and often serves as the official address of the entity.
(org, date, acquired_on)	This relation indicates the specific date on which an ORGANIZATION within the financial sector was acquired by another entity. It marks a significant event where the organization underwent a change in ownership or control, altering its organizational structure or strategy in the financial industry.
(pers, org, founder_of)	This relation denotes that a PERSON has established or initiated an ORGANIZATION in the financial sector, like a startup in fintech, a bank, or an investment firm.
(org, gpe, formed_in)	This relation refers to the specific location where a financial ORGANIZATION was established or originated. It signifies the geographical starting point of the organization's business journey.
(org, org, acquired_by)	This relation signifies that a financial ORGANIZATION has been purchased or taken over by another organization. It marks a change in ownership and possibly in the operational direction of the acquired organization.
(pers, univ, employee_of)	In this relation, a PERSON is employed by a UNIVERSITY, likely in a role such as a lecturer, researcher, or administrative staff, having formal employment ties.
(pers, gov, member_of)	In the financial context, this relation signifies that a PERSON is a member of an ORGANIZATION or GOV. AGENCY related to finance, such as financial regulatory bodies or financial committees.
(pers, univ, attended)	It indicates that a PERSON has attended a UNIVERSITY where they may have studied subjects related to finance, such as economics, business administration, or financial engineering.
(pers, univ, member_of)	Indicates that a PERSON is affiliated or associated with a UNIVERSITY, perhaps as a part of a university's community, club, or a special committee, without necessarily being a formal employee.
(org, money, profit_of)	This relation denotes the net earnings of a financial ORGANIZATION after deducting all costs and expenses. It signifies the financial gain resulting from the organization's business operations and investments
(org, money, cost_of)	This relation signifies the total expenses or expenditures incurred by a financial ORGANIZATION in conducting its business activities. It includes all the money spent on operations, investments, and other business-related expenses.

Table 5: Entity-relation types and their Descriptions as generated by GPT-4 using zero-shot prompting