

GPT-FinRE: In-context Learning for Financial Relation Extraction using Large Language Models

Pawan Kumar Rajpoot**
pawan.rajpoot2411@gmail.com
MUST Research
Bangalore, Karnataka, India

Ankur Parikh
ankur.parikh85@gmail.com
UtilizeAI Research
Bangalore, Karnataka, India

Abstract

Relation extraction (RE) is a crucial task in natural language processing (NLP) that aims to identify and classify relationships between entities mentioned in text. In the financial domain, relation extraction plays a vital role in extracting valuable information from financial documents, such as news articles, earnings reports, and company filings. This paper describes our solution to relation extraction on one such dataset REFinD. The dataset was released along with shared task as a part of the Fourth Workshop on Knowledge Discovery from Unstructured Data in Financial Services, co-located with SIGIR 2023. In this paper, we employed OpenAI models under the framework of in-context learning (ICL). We utilized two retrieval strategies to find top K relevant in-context learning demonstrations / examples from training data for a given test example. The first retrieval mechanism, we employed, is a learning-free dense retriever and the other system is a learning-based retriever. We were able to achieve 3rd rank overall (model performance and report). Our best F1-score is 0.718.

Keywords: relationship extraction, gpt, in context learning, text tagging, REFinD, KDF, SIGIR, CEIL, ICL, In Context Learning, GPT NEO, Finance, Large Language Model

ACM Reference Format:

Pawan Kumar Rajpoot and Ankur Parikh. 2023. GPT-FinRE: In-context Learning for Financial Relation Extraction using Large Language Models. In *Proceedings of The 4th Workshop on Knowledge Discovery from Unstructured Data in Financial Services (KDF @SIGIR '23)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

The emergence of large language models (LLMs) such as GPT-3 [6][12] represents a significant advancement in natural language processing (NLP). These models have expertise in variety of domains and hence they can be used

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDF @SIGIR '23, July 27, 2023, Taipei, Taiwan

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>


In the third quarter of 2016 , QCR HOLDINGS INC acquired CSB , headquartered in Ankeny , Iowa . 

Figure 1. Relation Extraction example, here both organizations are connected with "acquired by" relation.

as it is in multiple NLP tasks. Traditionally language models use separate pre-training-and fine-tuning pipelines [1] [3] [5] [4] [9] where fine-tune stage follows pre-training. Models are fine-tuned on a task-specific dataset in a fully-supervised manner. More recently a new paradigm known as in-context learning (ICL) [6][14] is being used which formulates an NLP task such that LLMs make predictions by learning from demonstrations. These demonstrations are presented to the LLMs in the context prompt itself.

Under the framework of ICL, LLMs achieve remarkable performance rivaling previous fully-supervised methods even with only a limited number of demonstrations provided in the prompt in various tasks such as solving math problems, commonsense reasoning, text classification, fact retrieval, natural language inference, and semantic parsing [6] [14] [8]. Recently, ICL based approach[16] is utilized for Relation Extraction (RE) task. RE seeks to identify a semantic relationship between a given entity pair mentioned in a sentence, which is the central task for knowledge retrieval requiring a deep understanding of natural language. The approach achieves improvements over not only existing GPT-3 baselines, but also on fully-supervised baselines. Specifically, it achieves SOTA performances on the SemEval and SciERC datasets, and competitive performances on the TACRED and ACE05 datasets.

Retrieval of examples to demonstrate is a key factor in the overall performance on these pipelines. LLMs can relate to the presented "to be predicted" data point more if the contextual examples predicted are similar to it. More relevant examples help us to leverage more out from LLMs both in terms of improvement in performance and less hallucination as examples can demonstrate model not to hallucinate in some cases.

In this paper, we employed GPT-3.5 Turbo and GPT-4 under the framework of ICL for the relation extraction task on REFinD dataset. We utilized two retrieval strategies to find top K relevant in-context learning demonstrations / examples from training data for a given test example. The first mechanism we have employed is a learning-free dense retriever. The other system we have utilized is a learning-based retriever [13].

Entity (subj)	Relation	Entity (obj)
PERSON	title	TITLE
PERSON	employee_of	ORGANIZATION
PERSON	member_of	ORGANIZATION
PERSON	founder_of	ORGANIZATION
PERSON	employee_of	UNIVERSITY
PERSON	member_of	UNIVERSITY
PERSON	attended	UNIVERSITY
PERSON	member_of	GOV. AGENCY
ORGANIZATION	formed_on	DATE
ORGANIZATION	acquired_on	DATE
ORGANIZATION	headquartered_in	GPE
ORGANIZATION	operations_in	GPE
ORGANIZATION	formed_in	GPE
ORGANIZATION	shares_of	ORGANIZATION
ORGANIZATION	subsidiary_of	ORGANIZATION
ORGANIZATION	acquired_by	ORGANIZATION
ORGANIZATION	agreement_with	ORGANIZATION
ORGANIZATION	revenue_of	MONEY
ORGANIZATION	profit_of	MONEY
ORGANIZATION	loss_of	MONEY
ORGANIZATION	cost_of	MONEY
	no/other_relation	

Figure 2. REFinD dataset relation and entity types.

2 Preliminary Background

2.1 Task Definition

As per the challenge "Relation Extraction is the task of automatically identifying and classifying the semantic relationships that exist between different entities in a given text." This shared task is a part of "Knowledge Discovery from Unstructured Data in Financial Services" (KDF) workshop which is collocated with SIGIR 2023.

Let C denote the input context and e_1 in C , e_2 in C denote the pair of entity pairs. Given a set of predefined relation classes R , relation extraction aims to predict the relation y in R between the pair of entities (e_1, e_2) within the context C , or if there is no predefined relation between them, predict y ="no relation".

2.2 Data

The dataset [18] released with this task is the largest relation extraction dataset for financial documents to date. Overall REFinD contains around 29K instances and 22 relations among 8 types of entity pairs. REFinD is created using raw text from various 10-X reports (including 10-K, 10-Q, etc. broadly known as 10-X) of publicly traded companies obtained from US Securities and Exchange Commission. Figure-2 shows different entity types and relations exist between them.

2.3 In Context Learning

In-context learning (ICL) refers to one of the core emergent abilities [17] that infers new tasks from context. We use the terms 'in-weights learning' and 'in-context learning' from prior work on sequence models [6] to distinguish between gradient-based learning with parameter updates and gradient-free learning from context, respectively. Formally, each training instance is first linearized into an input text $x = (x_1 \dots x_n)$ and an output text $y = (y_1 \dots y_n)$, where for all tokens $x_1 \dots x_n, y_1 \dots y_n$ in V and V is the vocabulary set of the

LM. Given a new test input text x -test, in-context learning defines the generation of output y as y -test \sim PLM(y -test | $x_1, y_1, \dots, x_k, y_k, x$ -test), where \sim refers to decoding strategies (e.g., greedy decoding and nuclear sampling [11]), and each in-context example $e_i = (x_i, y_i)$ is sampled from a training set D . The generation procedure is especially attractive as it eliminates the need for updating the parameters of the language model when encountering a new task, which is often expensive and impractical. Notably, the performance of ICL on downstream tasks can vary from almost random to comparable with state-of-the-art systems, depending on the quality of the retrieved in-context examples [13] [10] [19].

3 GPT-FinRE

GPT-RE is formalized under the ICL framework, using GPT models as shown in Figure-3.

3.1 Prompt Construction

We construct a prompt for each given test example, which is fed to the GPT models. Each prompt consists of the following components.

Task Description and Predefined Classes : We provide a succinct overview of the RE task description and the subset of predefined classes R , denoted by O . This subset is all possible relations exist between entity types of e_1 and e_2 . The model is explicitly asked to output the relation, which belongs to the O . Otherwise, the model will output "no relation".

K-shot Demonstrations : In the demonstration part, we reformulate each example by first showing the input prompt x -demo = Prompt(C, e_1, e_2) and the relation label y -demo.

Test Input : Similar to the demonstrations, we offer the test input prompt x -test, and GPT models are expected to generate the corresponding relation y -test.

3.2 Retrieval Systems

We have employed two retrieval strategies to find top K relevant in-context learning demonstrations / examples from training data for a given test example.

3.2.1 KNN with OpenAI Embeddings. Since ICL demonstrations closer to the test sample in the embedding space result in more consistent and robust performance [10]. We utilized the KNN to retrieve the most similar examples in the training set as the few-shot demonstrations for each test example. As this learning-free dense retriever relies on the choice of the embedding space, we used OpenAI embeddings (text-embedding-ada-002) to obtain example representations. For similarity search, we used FAISS tool [2].

3.2.2 EPR (Efficient Prompt Retrieval). This learning-based dense retriever is trained to retrieve a better singleton in-context example [13], and Top- K most similar examples are selected in the inference stage. This method for retrieving prompts for in-context learning uses annotated data and a LM. Given an input-output pair, it estimates

Retriever	LLM	F1-Score
KNN with openAI embeddings	GPT 3.5 Turbo (Examples: 5 retrieved + 5 random per possible relation)	0.643
KNN with openAI embeddings	GPT 4 (Examples: 5 retrieved + 5 random per possible relation)	0.697
EPR with GPT-Neo-2.7B	GPT 4 (Examples: 2 retrieved + 3 random per possible relation)	0.703
EPR with GPT-Neo-2.7B	GPT 4 (Examples: 5 retrieved + 4 random per possible relation)	0.718

Table 1. Our performance on test data with different combinations of retriever and LLM

the probability of the output given the input and a candidate training example as the prompt, and labels training examples as positive or negative based on this probability. It then trains an efficient dense retriever from this data, which is used to retrieve training examples as prompts at test time. Due to limited access to OpenAI, we have used the gpt-neo-2.7B model [7] as our choice of LM.

3.2.3 Random Class Examples. Along with KNN / EPR based examples, we also added K examples randomly for each possible class between two entity types to add more variety in the final prompt.

4 Experiments

Due to limited access and cost associated with OpenAI, We performed 4 primary experiments on the test dataset. We tried various rule based heuristics to improve the F1-score, but it didn't work as expected. We used retriever implementations from ¹.

5 Results

The results are shown in Table-1. Our best F1-Score is 0.718. We got 4th position in the shared-task. We find that GPT 4 performs better than GPT 3.5 Turbo. We also find that learning based retriever (EPR) outperforms learning-free retriever (KNN with OpenAI embeddings).

6 Future Work

In future we want to utilize GPT 4 for EPR. We also want to use different retrieval approaches such as Compositional Exemplars for In-context Learning (CEIL)[15].

7 Conclusion

This work explores the potential of GPT + ICL on Financial Relation Extraction (REFinD dataset). We used two retrieval mechanisms to find similar examples: (1) KNN with OpenAI Embeddings (2) EPR. We tried two different GPT models: (1) GPT 3.5 Turbo and GPT 4. The experimental results show that GPT 4 with learning based retriever EPR is giving the best F1-Score of 0.718.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the*

- North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [2] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.
- [3] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 3615–3620. <https://doi.org/10.18653/v1/D19-1371>
- [4] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. arXiv:1909.11942 [cs.CL]
- [5] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv:1910.10683 [cs.LG]
- [6] 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 (2020), 1877–1901.
- [7] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. arXiv preprint arXiv:2101.00027 (2020).
- [8] Richard Shin, Christopher H. Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. 2021. Constrained Language Models Yield Few-Shot Semantic Parsers. *CoRR* abs/2104.08768 (2021). arXiv:2104.08768 <https://arxiv.org/abs/2104.08768>
- [9] Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. A Robustly Optimized BERT Pre-training Approach with Post-training. In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*. Chinese Information Processing Society of China, Huhhot, China, 1218–1227. <https://aclanthology.org/2021.ccl-1.108>
- [10] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What Makes Good In-Context Examples for GPT-3? *CoRR* abs/2101.06804 (2021). arXiv:2101.06804 <https://arxiv.org/abs/2101.06804>
- [11] Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2022. Contrastive decoding: Open-ended text generation as optimization. arXiv preprint arXiv:2210.15097 (2022).
- [12] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh

¹<https://github.com/HKUNLP/icl-ceil>

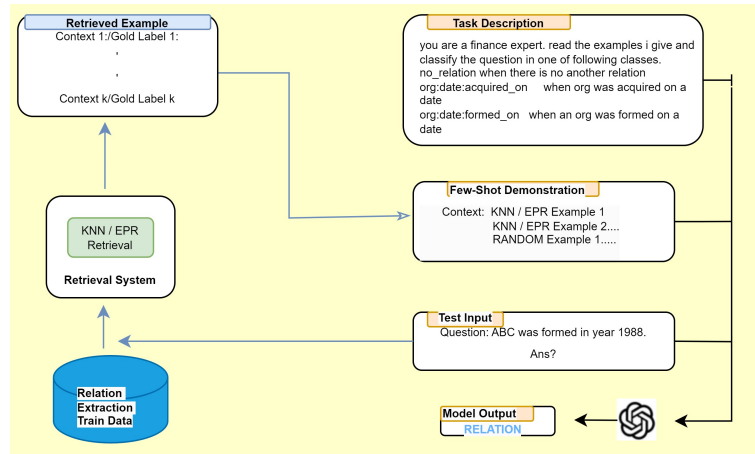


Figure 3. GPT-FinRE pipeline flow

Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. 2022. LaMDA: Language Models for Dialog Applications. *arXiv:2201.08239* [cs.CL]

- [13] Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. Learning To Retrieve Prompts for In-Context Learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 2655–2671. <https://doi.org/10.18653/v1/2022.naacl-main.191>
- [14] Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work? *arXiv preprint arXiv:2202.12837* (2022).

- [15] Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2023. Compositional Exemplars for In-context Learning. (2023). *arXiv:2302.05698* [cs.CL]
- [16] 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. *arXiv preprint arXiv:2305.02105* (2023).
- [17] Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. 2023. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846* (2023).
- [18] Simerjot Kaur, Charese Smiley, Akshat Gupta, Joy Sain, Dongsheng Wang, Suchetha Siddagangappa, Toyin Aguda, and Sameena Shah. 2023. REFinD: Relation Extraction Financial Dataset. *arXiv preprint arXiv:2305.18322* (2023).
- [19] Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye, and Lingpeng Kong. 2023. Self-Adaptive In-Context Learning: An Information Compression Perspective for In-Context Example Selection and Ordering. *arXiv:2212.10375* [cs.CL]