

Robust Retrieval Augmented Generation for Zero-shot Slot Filling

Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Alfio Gliozzo

IBM Research AI

Thomas J. Watson Research Center, NY

Abstract

Automatically inducing high quality knowledge graphs from a given collection of documents still remains a challenging problem in AI. One way to make headway for this problem is through advancements in a related task known as slot filling. In this task, given an entity query in form of [ENTITY, SLOT, ?], a system is asked to ‘fill’ the slot by generating or extracting the missing value exploiting evidence extracted from relevant passage(s) in the given document collection. The recent works in the field try to solve this task in an end-to-end fashion using retrieval-based language models. In this paper, we present a novel approach to zero-shot slot filling that extends dense passage retrieval with hard negatives and robust training procedures for retrieval augmented generation models. Our model reports large improvements on both T-REx and zsRE slot filling datasets, improving both passage retrieval and slot value generation, and ranking at the top-1 position in the KILT leaderboard. Moreover, we demonstrate the robustness of our system showing its domain adaptation capability on a new variant of the TACRED dataset for slot filling, through a combination of zero/few-shot learning. We release the source code and pre-trained models¹.

1 Introduction

Slot filling is a sub-task of Knowledge Base Population (KBP), where the goal is to recognize a pre-determined set of relations for a given entity and use them to populate infobox like structures. This can be done by exploring the occurrences of the input entity in the corpus and gathering information about its slot fillers from the context in which it is located. A slot filling system processes and indexes a corpus of documents. Then, when prompted with an entity and a number of relations,

¹Our source code is available at: <https://github.com/IBM/kgi-slot-filling>

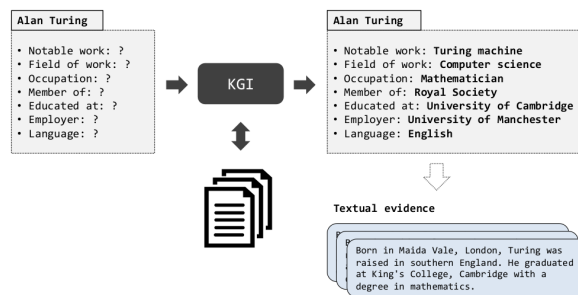


Figure 1: Slot Filling task

it fills out an infobox for the entity. Some slot filling systems provide evidence text to explain the predictions. Figure 1 illustrates the slot filling task.

Many KBP systems described in the literature commonly involve complex pipelines for named entity recognition, entity co-reference resolution and relation extraction (Ellis et al., 2015). In particular, the task of extracting relations between entities from text has been shown to be the weakest component of the chain. The community proposed different solutions to improve relation extraction performance, such as rule-based (Angeli et al., 2015), supervised (Zhang et al., 2017), or distantly supervised (Glass et al., 2018). However, all these approaches require a considerable human effort in creating hand-crafted rules, annotating training data, or building well-curated datasets for bootstrapping relation classifiers.

Recently, pre-trained language models have been used for slot filling (Petroni et al., 2020), opening a new research direction that might provide an effective solution to the aforementioned problems. In particular, the KILT benchmark (Petroni et al., 2021), standardizes two zero-shot slot filling tasks, zsRE (Levy et al., 2017) and T-REx (Elsahar et al., 2018), providing a competitive evaluation framework to drive advancements in slot filling. However, the best performance achieved by the current retrieval-based models on the two slot filling tasks in KILT are still not satisfactory. This is mainly

due to the lack of retrieval performance that affects the generation of the filler as well.

In this work, we propose KGI (Knowledge Graph Induction), a robust system for slot filling based on advanced training strategies for both Dense Passage Retrieval (DPR) and Retrieval Augmented Generation (RAG) that shows large gains on both T-REx (+38.24% KILT-F1) and zsRE (+21.25% KILT-F1) datasets if compared to previously submitted systems. We extend the training strategies of DPR with *hard negative mining* (Simo-Serra et al., 2015), demonstrating its importance in training the context encoder.

In addition, we explore the idea of adapting KGI to a new domain. The domain adaptation process consists of indexing the new corpus using our pre-trained DPR and substituting it in place of the original Wikipedia index. This enables zero-shot slot filling on the new dataset with respect to a new schema, avoiding the additional effort needed to rebuild NLP pipelines. We provide a few additional examples for each new relation, showing that zero-shot performance quickly improves with a few-shot learning setup. We explore this approach on a variant of the TACRED dataset (Alt et al., 2020) that we specifically introduce to evaluate the zero/few-shot slot filling task for domain adaptation.

The contributions of this work are as follows:

1. We describe an end-to-end solution for slot filling, called KGI, that improves the state-of-the-art in the KILT slot filling benchmarks by a large margin.
2. We demonstrate the effectiveness of hard negative mining for DPR when combined with end-to-end training for slot filling tasks.
3. We evaluate the domain adaptation of KGI using zero/few-shot slot filling, demonstrating its robustness on zero-shot TACRED, a benchmark released with this paper.
4. We publicly release the pre-trained models and source code of the KGI system.

Section 2 present an overview of the state of the art in slot filling. Section 3 describes our KGI system, providing details on the DPR and RAG models and describing our novel approach to hard negatives. Our system is evaluated in Sections 4 and 5 which include a detailed analysis. Section 6 concludes the paper and highlights some interesting direction for future work.

2 Related Work

The use of language models as sources of knowledge (Petroni et al., 2019; Roberts et al., 2020; Wang et al., 2020; Petroni et al., 2020), has opened tasks such as zero-shot slot filling to pre-trained transformers. Furthermore, the introduction of retrieval augmented language models such as RAG (Lewis et al., 2020b) and REALM (Guu et al., 2020) also permit providing textual provenance for the generated slot fillers.

KILT (Petroni et al., 2021) was introduced with a number of baseline approaches. The best performing of these is RAG (Lewis et al., 2020b). The model incorporates DPR (Karpukhin et al., 2020) to first gather evidence passages for the query, then uses a model initialized from BART (Lewis et al., 2020a) to do sequence-to-sequence generation from each evidence passage concatenated with the query in order to generate the answer. In the baseline RAG approach only the query encoder and generation component are fine-tuned on the task. The passage encoder, trained on Natural Questions (Kwiatkowski et al., 2019) is held fixed. Interestingly, while it gives the best performance of the baselines tested on the task of producing slot fillers, its performance on the retrieval metrics is worse than BM25 (Petroni et al., 2021). This suggests that fine-tuning the entire retrieval component could be beneficial. Another baseline in KILT is BART_{LARGE} fine-tuned on the slot filling tasks but without the usage of the retrieval model.

In an effort to improve the retrieval performance, Multi-task DPR (Maillard et al., 2021) used the multi-task training of the KILT suite of benchmarks to train the DPR passage and query encoder. The top-3 passages returned by the resulting passage index were then combined into a single sequence with the query and a BART model was used to produce the answer. This resulted in large gains in retrieval performance.

DensePhrases (Lee et al., 2021) is a different approach to knowledge intensive tasks with a short answer. Rather than index passages which are then consumed by a reader or generator component, it indexes the phrases in the corpus that can be potential answers to questions, or fillers for slots. Each phrase is represented by the pair of its start and end token vectors from the final layer of a transformer initialized from SpanBERT (Joshi et al., 2020).

GENRE (Cao et al., 2021) addresses the retrieval task in KILT slot filling by using a sequence-to-

sequence transformer to generate the title of the Wikipedia page where the answer can be found. This method can produce excellent scores for retrieval but it does not address the problem of producing the slot filler. It is trained on BLINK (Wu et al., 2020) and all KILT tasks jointly.

Open Retrieval Question Answering (ORQA) (Lee et al., 2019) introduced neural information retrieval for the related task of factoid question answering. Like DPR, the retrieval is based on a bi-encoder BERT (Devlin et al., 2019) model. Unlike DPR, ORQA projects the BERT [CLS] vector to a lower dimensional (128) space. It also uses the *inverse cloze* pre-training task for retrieval, while DPR does not use retrieval specific pre-training.

3 Knowledge Graph Induction

Figure 2 shows KGI, our approach to zero-shot slot filling, combining a DPR model and RAG model, both trained for slot filling. We initialize our models from the Natural Questions (Kwiatkowski et al., 2019) trained models for DPR and RAG available from Hugging Face (Wolf et al., 2020)². We then employ a two phase training procedure: first we train the DPR model, i.e. both the query and context encoder, using the KILT provenance ground truth. Then we train the sequence-to-sequence generation and further train the query encoder using only the target tail entity as the objective. It is important to note that the same query encoder component is trained in both phases.




Figure 2: KGI Architecture

3.1 DPR for Slot Filling

Our approach to DPR training for slot filling is an adaptation of the question answering training in the

²<https://github.com/huggingface/transformers>

original DPR work (Karpukhin et al., 2020). We first index the passages using a traditional keyword search engine, Anserini³. The head entity and the relation are used as a keyword query to find the top-k passages by BM25. Passages with overlapping paragraphs to the ground truth are excluded as well as passages that contain a correct answer. The remaining top ranked result is used as a hard negative for DPR training. This is the hard negative mining strategy used by DPR (Karpukhin et al., 2020) and Multi-DPR (Maillard et al., 2021).

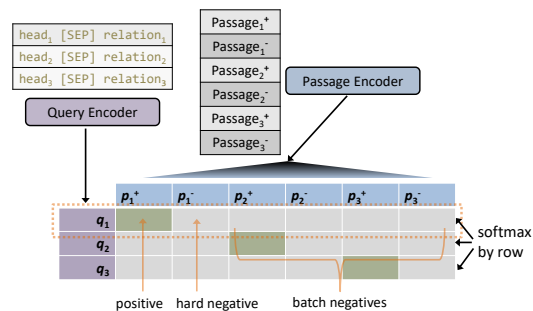


Figure 3: DPR Training

After locating a hard negative for each query, the DPR training data is a set of triples: query, positive passage (given by the KILT ground truth provenance) and the hard negative passage. Figure 3 shows the training process for DPR. For each batch of training triples, we encode the queries and passages independently. The passage and query encoders are BERT (Devlin et al., 2019) models. Then we find the inner product of all queries with all passages. The negatives for a given query are therefore the hard negative and the batch negatives, i.e. the positive and hard negative passages for other queries in the batch. After applying a softmax to the score vector for each query, the loss is the negative log-likelihood for the positive passages.

Using the trained DPR passage encoder we generate vectors for the approximately 32 million passages in our segmentation of the KILT knowledge source. Though this is a computationally expensive step, it is easily parallelized. The passage-vectors are then indexed with an ANN (Approximate Nearest Neighbors) data structure, in this case HNSW (Hierarchical Navigable Small World)(Malkov and Yashunin, 2018) using the open source FAISS library (Johnson et al., 2017)⁴. We use scalar quantization down to 8 bits to reduce the memory size.

³<https://github.com/castorini/anserini>

⁴<https://github.com/facebookresearch/faiss>

The query encoder is also trained for slot filling alongside the passage encoder. We inject the trained query encoder into the RAG model for Natural Questions. Due to the loose coupling between the query encoder and the sequence-to-sequence generation of RAG, we can update the pre-trained model’s query encoder without disrupting the quality of the generation.

Unlike previous work on zero-shot slot filling, we are training the DPR model specifically for the slot filling task. In contrast, the RAG baseline (Petroni et al., 2021) used DPR pre-trained on Natural Questions, and Multi-DPR (Maillard et al., 2021) trained on all KILT tasks jointly.

3.2 RAG for Slot Filling

Figure 4 illustrates the architecture of RAG (Lewis et al., 2020b). The RAG model is trained to predict the ground truth tail entity from the head and relation query. First the query is encoded to a vector and the top- k (we use $k = 5$) relevant passages are retrieved from the ANN index. The query is concatenated to each passage and the generator predicts a probability distribution over the possible next tokens for each sequence. These predictions are weighted according to the score between the query and passage - the inner product of the query vector and passage vector.




Figure 4: RAG Architecture

Marginalization then combines the weighted probability distributions to give a single probability distribution for the next token. This enables RAG to train the query encoder through its impact in generation, learning to give higher weight to passages that contribute to generating the correct tokens. Formally, the inputs to the BART model are sequences ($s_j = p_j$ [SEP] q) that comprise a query q plus retrieved passage p_j . The probability for each sequence is determined from the softmax over the retrieval scores (\mathbf{z}_r) for the passages. The probability for each output token t_i given the sequence s_j is a softmax over BART’s token prediction logits. Therefore the total probability for each token t_i is the log-likelihood summed over all sequences, weighted by each sequence’s probability.

$$P(s_j) = \text{softmax}(\mathbf{z}_r)_j$$

$$P(t_i|s_j) = \text{softmax}(\text{BART}(s_j)_i)_{t_i}$$

$$P(t_i) = \sum_j P(t_i|s_j) \cdot P(s_j)$$

Beam search is used at inference time to select the overall most likely tail entity. This is the standard beam search for natural language generation in deep neural networks (Sutskever et al., 2014), the only difference is in the way the next-token probabilities are obtained.

3.3 Dense Negative Sampling

As Figure 2 shows, the DPR question encoder is trained both by DPR and later by RAG. To examine the influence of this additional training from RAG on the retrieval performance, we compare retrieval metrics before and after RAG fine-tuning. Table 1 shows the large gains from training with RAG after DPR. Note that RAG training is using the weak supervision of the passage’s impact in producing the correct answer, rather than the ground truth provenance of DPR training. Since this is likely a disadvantage, we explore the other key difference with DPR and RAG training: RAG uses negatives drawn from the trained index rather than from BM25.

	T-REx		zsRE	
	R-Prec	R@5	R-Prec	R@5
DPR _{NQ}	19.50	29.80	45.49	60.77
DPR _{NQ} +RAG	53.04	65.54	68.13	79.19
DPR _{BM25}	49.02	63.34	94.55	98.17
DPR _{BM25} +RAG	65.02	75.52	96.89	98.01
DPR _{DNS}	42.62	55.09	97.53	99.30
DPR _{DNS} +RAG	74.34	82.89	98.60	99.70

Table 1: Analysis of retrieval by DPR and RAG on Dev sets

To replicate this feature of RAG in DPR, we introduce hard negatives mined from the learned index. Using the KILT trained DPR models, we index the passages. Then we gather hard negatives for DPR training, with one difference: rather than locating the hard negative passages by BM25, we find the passage by ANN search over the learned dense vector index. We train for an additional two epochs using these hard negatives. Table 1 shows the performance of the different approaches to retrieval. DPR_{NQ} is the DPR model pre-trained

Dataset	Instances			Relations		
	Train	Dev	Test	Train	Dev	Test
zsRE	148K	3724	4966	84	12	24
T-REx	2284K	5000	5000	106	104	104

Table 2: Slot filling datasets in KILT

on Natural Questions. DPR_{BM25} further trains DPR_{NQ} on the KILT data with BM25 hard negatives. Rows with +RAG further train the question encoder through RAG. The row DPR_{DNS} (**D**ense **N**egative **S**ampling) shows the performance of retrieval immediately after DNS training. Surprisingly, this results in lower performance for T-REx relative to DPR_{BM25} . However, further training the DNS model with RAG results in our best performance for both T-REx and zsRE. Since RAG does not update the context encoder, DNS training is the only training for the context encoder when negatives are drawn from the dense vector index.

After training with DNS the FAISS indexing with scalar quantization becomes prohibitively slow. We therefore remove all quantization and use four shards (the index is split into four, with the results of each query merged) for our experiments with DNS enabled KGI.

4 KILT Experiments

Table 2 gives statistics on the two zero-shot slot filling datasets in KILT. While the T-REx dataset is larger by far in the number of instances, the training sets have a similar number of distinct relations. We use only 500k training instances of T-REx in our experiments to increase the speed of experimentation.

Since the transformers for passage encoding and generation can accept a limited sequence length, we segment the documents of the KILT knowledge source (2019/08/01 Wikipedia snapshot) into passages. The ground truth provenance for the slot filling tasks is at the granularity of paragraphs, so we align our passage segmentation on paragraph boundaries when possible. If two or more paragraphs are short enough to be combined, we combine them into a single passage and if a single paragraph is too long, we truncate it.

4.1 KGI Hyperparameters

We have not done hyperparameter tuning, instead using hyperparameters similar to the original works

Hyperparameter	DPR	RAG
learn rate	5e-5	3e-5
batch size	128	128
epochs	2	1
warmup instances	0	10000
learning schedule	linear	triangular
max grad norm	1	1
weight decay	0	0
Adam epsilon	1e-8	1e-8

Table 3: KGI hyperparameters

on training DPR and RAG. Table 3 shows the hyperparameters used in our experiments. We train our models on T-REx using only the first 500k instances. For KGI_1 we use the same hyperparameters except that zsRE is trained for two epochs. In both KGI systems we use the default of five passages retrieved for each query for use in RAG.

4.2 Model Details

Number of parameters KGI is based on RAG and has the same number of parameters: $2 \times 110M$ for the $\text{BERT}_{\text{BASE}}$ query and passage encoders and $400M$ for the $\text{BART}_{\text{LARGE}}$ sequence-to-sequence generation component: $620M$ in total.

Computing infrastructure Using a single NVIDIA V100 GPU DPR training of two epochs takes approximately 24 hours for T-REx and 2 hours for zsRE. Using a single NVIDIA P100 GPU RAG training for 500k T-REx instances takes two days and 147k instances of zsRE takes 15 hours. The FAISS index on the KILT knowledge source requires a machine with large memory, we use 256GB memory - 128GB is insufficient for the indexes without scalar quantization.

4.3 Slot Filling Evaluation

As an initial experiment we tried RAG with its default index of Wikipedia, distributed through Hugging Face. We refer to this as RAG-KKS, or RAG without the KILT Knowledge Source, as reported in Table 4. Since the passages returned are not aligned to the KILT provenance ground truth, we do not report retrieval metrics for this experiment.

Motivated by the low retrieval performance reported for the RAG baseline by [Petroni et al. \(2021\)](#), we experimented with replacing the DPR retrieval with simple BM25 (RAG+BM25) over the KILT knowledge source. We provide the raw BM25 scores for the passages to the RAG model, to weight their impact in generation. We also experimented with the Natural Questions trained DPR,

	R-Prec	R@5	Acc.	F1
zsRE				
RAG-KKS			38.72	46.94
RAG+BM25	58.86	80.24	45.73	55.18
RAG+DPR _{NQ}	68.13	79.19	46.03	55.75
KGI ₀	96.24	97.53	69.58	77.24
KGI ₁	98.60	99.70	71.32	78.85
T-REx				
RAG-KKS			63.28	67.67
RAG+BM25	46.40	67.31	69.10	73.11
RAG+DPR _{NQ}	53.04	65.54	73.02	76.97
KGI ₀	61.30	71.18	76.58	80.27
KGI ₁	74.34	82.89	84.04	86.89

Table 4: Dev sets performance for different retrieval methods

with only RAG training on KILT (RAG+DPR_{NQ}).

We use the approach explained in Section 3 to train both the DPR and RAG models. KGI₀ is a version of our system using DPR with hard negative samples from BM25. The successor system, KGI₁ incorporates DPR training using DNS.

The metrics we report include accuracy and F1 on the slot filler, where F1 is based on the recall and precision of the tokens in the answer, allowing for partial credit on slot fillers. Our systems, except for RAG-KKS, also provide provenance information for the top answer. R-Precision and Recall@5 measure the quality of this provenance against the KILT ground truth provenance. Finally, KILT-Accuracy and KILT-F1 are combined metrics that measure the accuracy and F1 of the slot filler *only when the correct provenance is provided*.

Table 4 reports an evaluation on the development set, while Table 5 reports the test set performance of the top systems on the KILT leaderboard. KGI₀ and KGI₁ are our systems, while DensePhrases, GENRE, Multi-DPR, RAG for KILT and BART_{LARGE} are explained briefly in Section 2. KGI₁ gains dramatically in slot filling accuracy over the previous best systems, with gains of over 14 percentage points in zsRE and even more in T-REx. The combined metrics of KILT-AC and KILT-F1 show even larger gains, suggesting that the KGI₁ approach is effective at providing justifying evidence when generating the correct answer. We achieve gains of 21 to 41 percentage points in KILT-AC.

Relative to Multi-DPR, we see the benefit of weighting passage importance by retrieval score and marginalizing over multiple generations, com-

pared to the strategy of concatenating the top three passages and running a single sequence-to-sequence generation. GENRE is still best in retrieval for T-REx, suggesting that at least for a corpus such as Wikipedia, generating the title of the page can be very effective. A possible explanation for this behaviour is that most relations for a Wikipedia entity are mentioned in its corresponding page.

4.4 Analysis

To explore the effect of retrieval on downstream performance we consider two variants of our systems: one using random passages from the index, forcing the system to depend on implicit knowledge, and the another using passages from the ground truth provenance, to measure the upper bound performance for the ideal retrieval system. Evaluation is reported in Table 6 for 3 systems. By supplying these systems with the gold standard passages, we can see both the improvement possible through better retrieval, and the value of good retrieval during training. The best system, KGI₁ is the most effective at generating slot fillers from relevant explicit knowledge because it was trained on more cases of justifying explicit knowledge. However, given random passages it is the worst. It has sacrificed some implicit knowledge for better capabilities in using explicit knowledge.

As shown in Table 5, BART_{LARGE}, which is the best implicit-knowledge baseline system for KILT slot filling, is approximately 40 points lower in accuracy on T-REx if compared to KGI₁. To understand the impact of the explicit knowledge provided by DPR, we examine the improvement of KGI over BART_{LARGE}. We consider two main hypotheses: 1) the value of explicit knowledge depends on the relation, and 2) the value of explicit knowledge depends on the corpus frequency of the entities related.

To evaluate hypothesis 1, we consider the most frequent 20 relations in the T-REx Dev set, each occurring at least 40 times. The relations with the lowest relative performance gain are taxonomy and partonomy relations: TAXON-RANK, SUBCLASS-OF, INSTANCE-OF, PART-OF and PARENT-TAXON as well as LANGUAGES-SPOKEN,-WRITTEN-OR-SIGNED and SPORT. This suggest that essential properties of entities are well encoded in the language model itself. Inspecting the LANGUAGES-SPOKEN,-WRITTEN-OR-SIGNED we find that sur-

	R-Prec	Recall@5	Accuracy	F1	KILT-AC	KILT-F1
zsRE						
KGI ₁	98.49	99.23	72.55	77.05	72.31	76.69
KGI ₀	94.18	95.19	68.97	74.47	68.32	73.45
DensePhrases	57.43	60.47	47.42	54.75	41.34	46.79
GENRE	95.81	97.83	0.02	2.10	0.00	1.85
Multi-DPR	80.91	93.05	57.95	63.75	50.64	55.44
RAG (KILT organizers)	53.73	59.52	44.74	49.95	36.83	39.91
BART _{LARGE}	N/A	N/A	9.14	12.21	N/A	N/A
T-REx						
KGI ₁	74.36	83.14	84.36	87.24	69.14	70.58
KGI ₀	59.70	70.38	77.90	81.31	55.54	56.79
DensePhrases	37.62	40.07	53.90	61.74	27.84	32.34
GENRE	79.42	85.33	0.10	7.67	0.04	6.66
Multi-DPR	69.46	83.88	0.00	0.00	0.00	0.00
RAG (KILT organizers)	28.68	33.04	59.20	62.96	23.12	23.94
BART _{LARGE}	N/A	N/A	45.06	49.24	N/A	N/A

Table 5: KILT leaderboard top systems performance on slot filling tasks

Passages	RAG _{NQ}	KGI ₀	KGI ₁
Retrieved	70.58	76.58	84.04
Gold	88.66	89.46	90.20
Random	38.84	39.26	36.64

Table 6: T-REx Accuracy with Random and Gold Retrieval

face level information (i.e. French name vs. Russian name) is often sufficient for the correct prediction.

In contrast, the relations that gain the most from explicit knowledge are: PERFORMER, MEMBER-OF-SPORTS-TEAM, AUTHOR, PLACE-OF-BIRTH, COUNTRY-OF-ORIGIN, CAST-MEMBER, DIRECTOR. These relations are not central to the meaning of the head entity, like the taxonomy and partonomy relations, and are not typically predictable from surface-level features.

Regarding our second hypothesis, we might expect that more frequent entities have better representations in the parameters of a pre-trained language model, and that therefore the gain in performance due to use of explicit knowledge will show a strong dependence on the corpus frequency of the head or tail entity.

To test it, we group the Dev instances in T-Rex according to the decile of the head or tail entity frequency. We compute a macro-accuracy, weighting all relations equally. Figure 5 shows the macro-accuracy of BART_{LARGE} and KGI₁ for each decile of head and tail entity frequency. Although there is a general trend of higher accuracy for more fre-

quent tail entities and lower accuracy for more frequent head entities, *there is no pattern to the gain of explicit knowledge over implicit knowledge from entity frequency*. There is a similar picture when considering the decile of the *minimum* of the head or tail entity frequency. This falsifies our second hypothesis and suggests implicit knowledge is distinct in kind from explicit knowledge, rather than merely under-trained for low frequency entities.

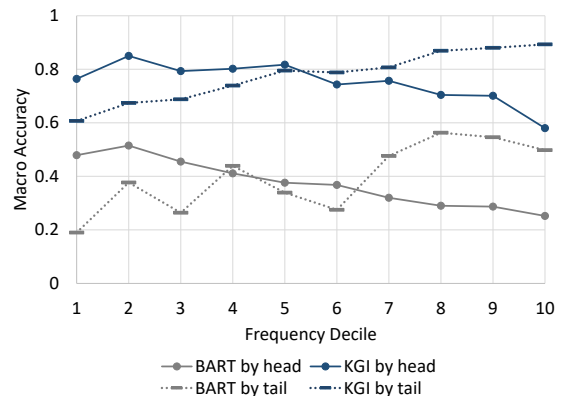


Figure 5: Performance as a function of entity frequency

5 Domain Adaptation Experiments

In this section, we evaluate the domain adaptation capability of KGI. For this purpose, we re-organize a dataset specifically designed to evaluate standard supervised relation extraction models, such as TACRED, with the aim to create a zero-shot (and few-shot) slot filling benchmark where the documents are written with a different style than Wikipedia, and the relations in the KG are different from those

in Wikidata. In order to perform an in-depth comparison and analysis, we also propose a new set of ranking baselines and use metrics which are suitable to better evaluate the slot filling task in a zero-shot setup.

5.1 Zero-shot TACRED

The TACRED dataset was originally proposed by Zhang et al. (2017) with the goal to provide a high-quality training set to supervise a relation extraction model which is shown to be competitive on TAC-KBP 2015 (Ellis et al., 2015). The target KG schema consists of two infoboxes modeling the *person* and *organization* entity types, with 41 relation types in total. For our experiments, we adopt a revisited version of TACRED (Alt et al., 2020), in which a second stage crowdsourcing is performed to further improve the quality of the annotations and resolve conflicts among relations.

In a typical supervised relation extraction setup, a model is trained to predict (i.e. classify) the right relation type given a textual passage and two entity mentions as inputs. In this paper we used the TACRED dataset as a slot filling benchmark, using the following procedure: 1) we first create the corpus by merging all the plain textual passages from the instances in the train, dev and test sets; 2) we collect the annotated triples, i.e. subject-relation-object, from the test data to come up with a ground-truth KG to be used for slot filling evaluation⁵; 3) we remove all the triples from the original test set where the subjects are pronouns. The resulting KG consists of 2673 slot filling *test* instances. Similarly, we acquire a KG from the train/dev sets to further fine-tune the KGI system as described in the next section. To enable zero-shot experiments, we also convert each relation label into a relation phrase by removing the namespaces *per:* and *org:*, and replacing the ‘_’ character with a space. Finally, for each pre-annotated entity in the corpus, we pre-compute an inverted index consisting of a list of co-occurring entities in the textual passages. We use this inverted index to compare our model with a set of ranking baselines.

An example of the obtained ground truth is illustrated in Table 7: given the query [*Dominick Dunne, employee of, ?*], a slot filling system is supposed to identify the missing slot with *Vanity Fair*, i.e. the gold standard object in the KG, by retrieving it

⁵79.5% of the overall instances are labeled as no relation. We exclude these instances from the ground truth KG, but we retain them in the textual corpus.

from the collection of passages.

5.2 Slot Filling Evaluation

Task Given a slot filling query $(e, s, ?)$ and a list of possible slot values $[v_1, \dots, v_n]$, where e is the entity as subject, s is the slot/relation and v_i are the object candidates that co-occur with e in the corpus, we can frame the zero-shot slot filling as a ranking problem: $\operatorname{argmax}_i \operatorname{score}_M(e, s, v_i)$. score_M is a function that takes as input a triple and provide a score based on the model M . Turning the slot filling into a ranking problem has two advantages: 1) we can compare the generative approach with a new set of baselines, and 2) we can limit the generation of the slot values to a pre-defined set of domain specific entities.

Models In order to adapt KGI₁, as pre-trained on T-REx, to the TACRED corpus, we indexed the textual passages using DPR, as described in Section 3. Then we replaced the original Wikipedia index with this new index. During the inference step, we restrict the generation of the slot values using the list of object candidates, i.e. the entities which co-occur with the subject from the inverted index, to facilitate comparability to a set of ranking baselines. To this aim, we adopt the technique described by Cao et al. (2021) to restrict the vocabulary of tokens during the generation.

We use three baselines to compare with our approach for this zero-shot slot filling task. **PMI** is implemented using the pointwise mutual information between e and v_i based on their co-occurrence in the corpus. Also, we train a **Word2Vec** (Mikolov et al., 2013) skip-gram model on the textual corpus, and we use it to implement the scoring function as $\cosine(e + s, v_i)$, for each candidate filler v_i . It is based on the assumption that a relation s between two (multi)word embeddings e and v can be represented as an offset vector $(v - e) = s \iff (e + s) = v$ (Rossiello et al., 2019; Vylomova et al., 2016). Finally, **GPT-2** computes the perplexity of the fragment of text by concatenating the tokens in e, s and each v_i (Radford et al., 2019).

Metrics Due to the similarity of slot filling with the knowledge base completion task, we use Mean Reciprocal Rank (MRR) and HIT@k, with $k = [1, 5, 10]$, as evaluation metrics (Bordes et al., 2013). Note that HIT@1 has the same meaning of the accuracy for the downstream task on KILT.

Subject	Relation/Slot	Object	Passage
Dominick Dunne	per:employee_of	Vanity Fair	Dominick Dunne, the author, television personality and Vanity Fair reporter who covered the trials of socialite Claus von Bulow.
Dominick Dunne	per:age	83	Dominick Dunne, 83, a crime story author, in New York City.
Dominick Dunne	per:siblings	John Gregory Dunne	Dominick Dunne, author of crime stories dies Born in 1925 in Hartford, Connecticut, was part of a famous family that included his brother, novelist and screenwriter John Gregory Dunne.
ALICO	org:country_of_headquarters	US	AIA says IPO raised 205 billion US dollars AIG said Monday it had also raised 162 billion dollars by selling unit American Life Insurance Company (ALICO) to MetLife Inc.
ALICO	org:top_members	Christopher J. Swift	Alico's chief financial officer, Christopher J. Swift, added that the bonds were issued by companies in many commercial sectors, which diversified the portfolio.
ALICO	org:parents	AIG	AIG said it had transferred ownership to the Federal Reserve Bank of parts of two subsidiaries, ALICO which is active in life assurance in the United States and AIA which provides life assurance abroad.

Table 7: Examples of annotations from TACRED dataset for both *person* and *organization* infoboxes.

	MRR	HIT@1	HIT@5	HIT@10
PMI	20.20	10.89	26.49	37.30
Word2Vec	25.24	13.83	34.92	47.60
GPT-2	17.62	8.37	23.72	35.34
KGI ₁ 0-shot	43.98	28.51	64.31	76.06
KGI ₁ 1-shot	48.86	33.89	66.63	78.75
KGI ₁ 4-shot	53.28	38.8	70.45	79.35

Table 8: Zero/few-shots results on TACRED

Hyperparameter	RAG
learn rate	3e-6
batch size	1
epochs	3
warmup instances	0
learning schedule	linear
max grad norm	1
weight decay	0
Adam epsilon	1e-8

Table 9: KGI₁ hyperparameters for TACRED few-shot

Results Table 8 reports the results of our evaluation. KGI₁ achieves substantially better performance than the aforementioned zero-shot baselines on all evaluation metrics. However, HIT@1 is $\sim 28\%$ which is significantly lower compared with the numbers reported on the datasets in KILT. This begs the question, how to further improve the transfer learning capabilities of these generative models? Interestingly, HIT@5/10 are high (i.e. $\sim 64\%/76\%$). This indicates our approach would be useful in a human-in-the-loop scenario by providing valuable candidates for the fillers that can be further validated.

For this purpose, we also conduct few-shot experiments to understand the robustness of KGI₁ by fine-tuning it with very limited amounts of training examples. We randomly pick n example(s) for each relation type from the TACRED training set,

with $n = [1, 4]$. Table 9 gives our hyperparameters for the TACRED few-shot experiments. We show that our system benefits from additional domain specific training data selected from TACRED. Just using one example and four examples per relation, HIT@1 improves ~ 5 and ~ 10 percentage points respectively.

6 Conclusion

In this paper, we presented KGI, a novel approach to zero-shot slot filling. KGI improves Dense Passage Retrieval using hard negatives from the dense index, and implements a robust training procedures for Retrieval Augmented Generation. We evaluated KGI on both T-REx and zsRE slot filling datasets, ranking at top-1 position in the KILT leaderboard with a net improvement of +38.24 and +21.25 percentage points in KILT-F1, respectively. Moreover, we proposed and release a new benchmark for zero/few-shot slot filling based on TACRED to evaluate domain adaptation where our system obtained much better zero-shot results compared with the baselines. In addition, we have observed significant improvement in results for KGI when rapidly fine-tuned in a few-shot setting. This work opens promising future research directions for slot filling and other related tasks. We plan to apply DPR with dense negative sampling to other tasks in the KILT benchmark, including dialogue and question answering. Likewise, an in-depth investigation on more effective strategies for domain adaptation, such as the combination of zero-shot and few-shot learning involving human-in-the-loop techniques, would be another interesting direction to explore.

References

- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. [TACRED revisited: A thorough evaluation of the TACRED relation extraction task](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1558–1569, Online. Association for Computational Linguistics.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. [Leveraging linguistic structure for open domain information extraction](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354, Beijing, China. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *NIPS*, pages 2787–2795.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *International Conference on Learning Representations*.
- 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joe Ellis, Jeremy Getman, Dana Fore, Neil Kuster, Zhiyi Song, Ann Bies, and Stephanie M. Strassel. 2015. [Overview of linguistic resources for the TAC KBP 2015 evaluations: Methodologies and results](#). In *TAC. NIST*.
- 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)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Michael R. Glass, Alfio Gliozzo, Oktie Hassanzadeh, Nandana Mihindukulasooriya, and Gaetano Rossiello. 2018. [Inducing implicit relations from text using distantly supervised deep nets](#). In *International Semantic Web Conference (1)*, volume 11136 of *Lecture Notes in Computer Science*, pages 38–55. Springer.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Realm: Retrieval-augmented language model pre-training](#). *arXiv preprint arXiv:2002.08909*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. [Billion-scale similarity search with gpus](#). *arXiv preprint arXiv:1702.08734*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021. [Learning dense representations of phrases at scale](#). 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 6634–6647, Online. Association for Computational Linguistics.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. [Zero-shot relation extraction via reading comprehension](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Hein-

- rich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020b. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Jean Maillard, Vladimir Karpukhin, Fabio Petroni, Wen-tau Yih, Barlas Öğuz, Veselin Stoyanov, and Gargi Ghosh. 2021. Multi-task retrieval for knowledge-intensive tasks. *arXiv preprint arXiv:2101.00117*.
- Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Fabio Petroni, Patrick S. H. Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. How context affects language models’ factual predictions. In *AKBC*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. [KILT: a benchmark for knowledge intensive language tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Gaetano Rossiello, Alfio Gliozzo, Robert Farrell, Nicolas Fauceglia, and Michael Glass. 2019. [Learning relational representations by analogy using hierarchical Siamese networks](#). 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)*, pages 3235–3245, Minneapolis, Minnesota. Association for Computational Linguistics.
- Edgar Simo-Serra, Eduard Trulls, Luis Ferraz, Iasonas Kokkinos, Pascal Fua, and Francesc Moreno-Noguer. 2015. Discriminative learning of deep convolutional feature point descriptors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 118–126.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Ekaterina Vylomova, Laura Rimell, Trevor Cohn, and Timothy Baldwin. 2016. [Take and took, gaggle and goose, book and read: Evaluating the utility of vector differences for lexical relation learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1671–1682, Berlin, Germany. Association for Computational Linguistics.
- Chenguang Wang, Xiao Liu, and Dawn Song. 2020. Language models are open knowledge graphs. *CoRR*, abs/2010.11967.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. [Scalable zero-shot entity linking with dense entity retrieval](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6397–6407, Online. 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.