

How Can We Know What Language Models Know?

Zhengbao Jiang^{1*} Frank F. Xu^{1*} Jun Araki² Graham Neubig¹

Language Technologies Institute, Carnegie Mellon University¹

Bosch Research North America²

{zhengbaj, fangzhex, gneubig}@cs.cmu.edu jun.araki@us.bosch.com

Abstract

Recent work has presented intriguing results examining the knowledge contained in language models (LM) by having the LM fill in the blanks of prompts such as “*Obama is a _ by profession*”. These prompts are usually manually created, and quite possibly sub-optimal; another prompt such as “*Obama worked as a _*” may result in more accurately predicting the correct profession. Because of this, given an inappropriate prompt, we might fail to retrieve facts that the LM *does* know, and thus any given prompt only provides a lower bound estimate of the knowledge contained in an LM. In this paper, we attempt to more accurately estimate the knowledge contained in LMs by automatically discovering better prompts to use in this querying process. Specifically, we propose mining-based and paraphrasing-based methods to automatically generate high-quality and diverse prompts and ensemble methods to combine answers from different prompts. Extensive experiments on the LAMA benchmark for extracting relational knowledge from LMs demonstrate that our methods can improve accuracy from 31.1% to 38.1%, providing a tighter lower bound on what LMs know. We have released the code and the resulting LM Prompt And Query Archive (LPAQA) at <https://github.com/jzbjyb/LPAQA>.

1 Introduction

Recent years have seen the primary role of language models (LM) transition from generating or evaluating the fluency of natural text (Mikolov and Zweig, 2012; Merity et al., 2018; Melis et al., 2018; Gamon et al., 2005) to being a powerful tool for text understanding. This understanding has mainly been achieved through the use of language modeling as a pre-training task for *feature extractors*,

*The first two authors contributed equally.

		Prompts		
	manual	DirectX is developed by	<u>Y_{man}</u>	
	mined	<u>Y_{mine}</u> released the DirectX		
	paraphrased	DirectX is created by	<u>Y_{para}</u>	
Top 5 predictions and log probabilities				
	Y _{man}	Y _{mine}	Y _{para}	
1	<u>Intel</u> -1.06	<u>Microsoft</u> -1.77	<u>Microsoft</u> -2.23	
2	<u>Microsoft</u> -2.21	They -2.43	Intel -2.30	
3	IBM -2.76	It -2.80	default -2.96	
4	Google -3.40	Sega -3.01	Apple -3.44	
5	Nokia -3.58	Sony -3.19	Google -3.45	

Figure 1: Top-5 predictions and their log probabilities using different prompts (manual, mined, and paraphrased) to query BERT. Correct answer is underlined.

where the hidden vectors learned through a language modeling objective are then used in downstream language understanding systems (Dai and Le, 2015; Melamud et al., 2016; Peters et al., 2018; Devlin et al., 2019).

Interestingly, it is also becoming apparent that LMs¹ *themselves* can be used as a tool for text understanding by formulating queries in natural language and either generating textual answers directly (McCann et al., 2018; Radford et al., 2019), or assessing multiple choices and picking the most likely one (Zweig and Burges, 2011; Rajani et al., 2019). For example, LMs have been used to answer factoid questions (Radford et al., 2019), answer common sense queries (Trinh and Le, 2018; Sap et al., 2019), or extract factual knowledge about relations between entities (Petroni et al., 2019; Baldini Soares et al., 2019). Regardless of the end task, the knowledge contained in LMs is probed by providing a prompt, and letting the LM either generate the continuation of a prefix (e.g. “*Barack Obama was born in _*”), or predict missing words

¹Technically bidirectional models like BERT and ELMo do not directly define a probability distribution over text, which is the underlying definition of an LM. Nonetheless, we call them LMs for simplicity.

in a cloze-style template (e.g., “*Barack Obama is a _ by profession*”).

However, while this paradigm has been used to achieve a number of intriguing results regarding the knowledge expressed by LMs, they all rely on prompts that were manually created based on the intuition of the experimenter. Thus it is quite possible that a fact that the LM *does* know cannot be retrieved due to the prompts not being effective queries for the fact. Thus, existing results are simply a *lower bound* on the extent of knowledge contained in LMs, and in fact, LMs may be even more knowledgeable than these initial results indicate. In this paper we ask the question: “How can we tighten this lower bound and get a more accurate estimate of the knowledge contained in state-of-the-art LMs?” This is interesting both scientifically, as a probe of the knowledge that LMs contain, and from an engineering perspective, as it will result in higher recall when using LMs as part of a knowledge extraction system.

In particular, we focus on the setting of [Petroni et al. \(2019\)](#) who examine extracting knowledge regarding the relations between entities (definitions in § 2). We propose two automatic methods to systematically improve the breadth and quality of the prompts used to query the existence of a relation (§ 3). Specifically, as shown in [Figure 1](#), these are *mining-based* methods inspired by previous relation extraction methods ([Ravichandran and Hovy, 2002](#)), and *paraphrasing-based* methods that take a seed prompt (either manually created or automatically mined), and paraphrase it into several other semantically similar expressions. Further, we note that each prompt generated above can be viewed as an “expert” to retrieve knowledge from LMs, and different experts may work better when querying for particular subject-object pairs. We thus investigate lightweight ensemble methods to combine the answers from different prompts together (§ 4).

We experiment on the LAMA benchmark ([Petroni et al., 2019](#)), which is an English-language benchmark devised to test the ability of LMs to retrieve relations between entities (§ 5). We first demonstrate that improved prompts significantly improve accuracy on this task, with the one-best prompt extracted by our method raising accuracy from 31.1% to 34.1% on BERT-base ([Devlin et al., 2019](#)), with similar gains being obtained with BERT-large as well. We further demonstrate that using a diversity of prompts through ensembling

further improves accuracy to 38.1%. We perform extensive analysis and glean insights about how to best query the knowledge stored in LMs, which both serves as useful ablation studies for our proposed methods, and elucidates potential future directions for incorporating knowledge into LMs themselves. Finally, we have released the resulting LM Prompt And Query Archive (LPAQA) to facilitate future experiments on probing knowledge contained in LMs.

2 Knowledge Retrieval from LMs

Retrieving factual knowledge from LMs is quite different from querying standard declarative knowledge bases (KB). In standard KBs, users formulate their information needs as a structured query defined by the KB schema and query language. For example, `SELECT ?y WHERE {wd:Q76 wdt:P19 ?y}` is a SPARQL query to search the birth place of `Barack.Obama`. In contrast, LMs must be queried by natural language prompts, such as “*Barack Obama was born in _*”, and the word assigned the highest probability in the blank will be returned as the answer. Unlike deterministic queries on KBs, this provides no guarantees of correctness or success.

While the idea of prompts is common to methods for extracting many varieties of knowledge from LMs, in this paper we specifically follow the formulation of [Petroni et al. \(2019\)](#), where factual knowledge is in the form of triples $\langle x, r, y \rangle$. Here x indicates the subject, y indicates the object, and r is their corresponding relation. To query the LM, r is associated with a cloze-style prompt t_r consisting of a sequence of tokens, two of which are placeholders for subjects and objects (e.g., “ *x plays at y position*”). The existence of the fact in the LM is assessed by replacing x with the surface form of the subject, and letting the model predict the missing object (e.g., “*Jordan plays at _ position*”):²

$$\hat{y} = \arg \max_{y' \in \mathcal{V}} P_{\text{LM}}(y'|x, t_r),$$

where \mathcal{V} is the vocabulary, and $P_{\text{LM}}(y'|x, t_r)$ is the LM probability of predicting y' in the blank conditioned on the other tokens (i.e., the subject and

²We can also go the other way around by filling in the objects and predicting the missing subjects. Since our focus is on improving prompts, we choose to be consistent with [Petroni et al. \(2019\)](#) to make a fair comparison, and leave exploring other settings to future work.

the prompt).³ We say that an LM has knowledge of a fact if \hat{y} is the same as the ground-truth y . Because we would like our prompts to most effectively elicit any knowledge that may be contained in the LM itself, a “good” prompt should trigger the LM to predict the ground-truth objects as often as possible.

In previous work (McCann et al., 2018; Radford et al., 2019; Petroni et al., 2019), t_r has been a single manually defined prompt based on the intuition of the experimenter. As noted in the introduction, this method has no guarantee of being optimal, and thus in the following sections we propose methods that *learn* effective prompts from a small set of training data consisting of gold subject-object pairs for each relation.

3 Prompt Generation

First, we tackle prompt generation: the task of generating a set of prompts $\{t_{r,i}\}_{i=1}^T$ for each relation r , where at least some of the prompts effectively trigger LMs to predict ground-truth objects. We employ two practical methods to either mine prompt candidates from a large corpus (§ 3.1) or diversify a seed prompt through paraphrasing (§ 3.2).

3.1 Mining-based Generation

Our first method is inspired by template-based relation extraction methods (Ravichandran and Hovy, 2002), which are based on the observation that words in the middle of the subject x and object y in a large corpus often describe the relation r . Based on this intuition, we first identify all the Wikipedia sentences that contain both subjects and objects of a specific relation r , then use those words in the middle as prompts. For example, “*Barack Obama was born in Hawaii*” is converted into a prompt “*x was born in y*” by replacing the subject and the object with placeholders. To remove noise, we rank all the unique prompts based on their frequencies and use only the top T most frequent ones.⁴

Notably, this variety of mining-based method does not rely on any manually-created prompts,

³We restrict to masked LMs in this paper because the missing slot might not be the last token in the sentence and computing this probability in traditional left-to-right LMs using Bayes’ theorem is not tractable.

⁴Words along the dependency path between the entities might be even more indicative of the relation, as noted by Toutanova et al. (2015). It is quite possible that using these techniques may further improve results, but we did not test these at this time due to the increased complexity and computational load resulting from parsing the whole corpus.

and can thus be flexibly applied to any relation where we can obtain a set of subject-object pairs. It will also result in diverse prompts, covering a wide variety of ways that the relation may be expressed in actual text. However, it may also be prone to noise, as many prompts acquired in this way may not be very indicative of the relation (e.g. “ x, y ”), even if they are frequent.

3.2 Paraphrasing-based Generation

Our second method for generating prompts is more targeted – it aims to improve lexical diversity while remaining relatively faithful to the original prompt. Specifically, we do so by performing paraphrasing over the original prompt into other semantically similar or identical expressions. For example, if our original prompt is “*x shares border with y*”, it may be paraphrased into “*x has a common border with y*” and “*x adjoins y*”. This is conceptually similar to query expansion techniques used in information retrieval that reformulate a given query to improve retrieval performance (Carpineto and Romano, 2012).

While many methods could be used for paraphrasing, we follow the simple method of using back-translation (Prabhumoye et al., 2018) to first translate the initial prompt into B candidates in another language, each of which is then back-translated into B candidates in the original language. We then rank B^2 candidates based on their round-trip probability (i.e., $P_{\text{forward}}(\bar{t}|\hat{t}) \cdot P_{\text{backward}}(t|\bar{t})$, where \hat{t} is the initial prompt, \bar{t} is the translated prompt in the other language, and t is the final prompt), and keep the top T prompts.

4 Prompt Selection and Ensembling

In the previous section, we described methods to generate a set of candidate prompts $\{t_{r,i}\}_{i=1}^T$ for a particular relation r . Each of these prompts may be more or less effective at eliciting knowledge from the LM, and thus it is necessary to decide how to use these generated prompts at test time. In this section, we describe three methods to do so.

4.1 Top-1 Prompt Selection

For each prompt, we can measure its accuracy of predicting the ground-truth objects (on a training dataset) using:

$$A(t_{r,i}) = \frac{\sum_{\langle x,y \rangle \in r} \delta(y = \arg \max_{y'} P_{\text{LM}}(y'|x, t_{r,i}))}{|r|},$$

where $\delta(\cdot)$ is Kronecker’s delta function, returning 1 if the internal condition is true, and 0 otherwise. In the simplest method for querying the LM, we choose the prompt with the highest accuracy and query using only this prompt.

4.2 Rank-based Ensemble

Next we examine methods that use not only the top-1 prompt, but combine together multiple prompts. The advantage to this is that the LM may have observed different entity pairs in different contexts within its training data, and having a variety of prompts may allow for elicitation of knowledge that appeared in these different contexts.

Our first method for ensembling is a parameter-free method that averages the predictions of the top-ranked prompts. We rank all the prompts based on their accuracy of predicting the objects, and use the average log probabilities from the top K prompts to calculate the probability of the object:

$$s(y|x, r) = \sum_{i=1}^K \frac{1}{K} \log P(y|x, t_{r,i}), \quad (1)$$

$$P(y|x, r) = \text{softmax}(s(\cdot|x, r))_y, \quad (2)$$

where $t_{r,i}$ is the prompt ranked at the i -th position. Intuitively, due to the fact that we are combining together scores in the log space, this has the effect of penalizing objects that are very unlikely given any certain prompt in the collection. We also compare with linear combination in ablations in § 6.2.

4.3 Optimized Ensemble

The above method treats the top K prompts equally, which might be sub-optimal if some prompts are more reliable than others. Thus, we also propose a method that directly optimizes prompt weights. Formally, we re-define the score in Equation 1 as:

$$s(y|x, r) = \sum_{i=1}^T P_{\theta_r}(t_{r,i}|r) \log P(y|x, t_{r,i}), \quad (3)$$

where $P_{\theta_r}(t_{r,i}|r) = \text{softmax}(\theta_r)$ is a distribution over prompts parameterized by θ_r , a T -sized real-value vector. θ_r is optimized to maximize the probability of the gold-standard object $P(y|x, r)$ over training data.

5 Main Experiments

5.1 Experimental Settings

In this section, we assess the extent to which our prompts can improve fact prediction performance,

Properties	T-REx	T-REx-train
#subject-object pairs	830.2	948.7
#unique subjects	767.8	880.1
#unique objects	150.9	354.6
object entropy	3.6	4.4

Table 1: Dataset statistics. All the values are average across 41 relations.

raising the lower bound on the knowledge we can discern is contained in LMs.

Dataset As data, we use the T-REx subset (EISa-har et al., 2018) of the LAMA benchmark (Petroni et al., 2019), which has a broader set of 41 relations (compared to the Google-RE subset which only covers 3). Each relation is associated with at most 1000 subject-object pairs from Wikidata, and a single manually designed prompt. To learn to mine prompts (§ 3.1), rank prompts (§ 4.2), or learn ensemble weights (§ 4.3), we create a separate training set of subject-object pairs also from Wikidata for each relation that has no overlap with the T-REx dataset. We denote the training set as T-REx-train. For consistency with the T-REx dataset in LAMA, T-REx-train also is chosen to contain only single-token objects. The statistics of these datasets are summarized in Table 1.

Models As the models to probe, we use BERT-base and BERT-large (Devlin et al., 2019).

Evaluation Metrics We use two metrics to evaluate the success of prompts in probing LMs. The first evaluation metric, *micro-averaged accuracy*, follows the LAMA benchmark⁵ in calculating the accuracy of all subject-object pairs for each relation, then averages these relation-level accuracies. However, we found that the object distributions of some relations are extremely skewed, e.g. more than half of the objects in relation `native_language` are French. This can lead to deceptively high scores, even for a majority-class baseline that picks the most common object for each relation, which achieves a score of 22.0%. To mitigate this problem, we also report *macro-averaged accuracy*, which computes accuracy for each unique object separately, then averages them together to get the relation-level accuracy. This is a much stricter metric, with the majority-class

⁵In LAMA, it is called “P@1.”

baseline only achieving a score of 2.2%.

Methods We attempted different methods for prompt generation and selection/ensembling, and compare them with the manually designed prompts used in [Petroni et al. \(2019\)](#). **Majority** refers to predicting the majority object for each relation, as mentioned above. **Man** is the baseline from [Petroni et al. \(2019\)](#) that only uses the manually designed prompts for retrieval. **Mine** (§ 3.1) uses the prompts mined from Wikipedia, and **Mine+Man** combines them with the manual prompts. **Mine+Para** (§ 3.2) paraphrases the highest-ranked mined prompt for each relation, while **Man+Para** uses the manual one instead. Those prompts are combined either by averaging the log probabilities from the **TopK** highest-ranked prompts (§ 4.2) or the weights after optimization (§ 4.3; **Opti.**). **Oracle** represents the upper bound of the performance of the generated prompts, where a fact is judged as correct if *any* one of the prompts allows the LM to successfully predict the object.

Implementation Details We keep $T = 30$ prompts either generated through mining or paraphrasing in all experiments, and the number of candidates in back-translation is set to $B = 7$. We use the round-trip English-German neural machine translation models pre-trained on WMT’19 ([Ng et al., 2019](#)) for back-translation.⁶

5.2 Evaluation Results

Micro- and macro-averaged accuracy of different methods are reported in [Tables 2 and 3](#) respectively.

Single Prompt Experiments When only one prompt is used (in the first **Top1** column in both tables), the best paraphrase of the manual prompt improves the micro-averaged accuracy from 31.1% to 34.1% on BERT-base, and from 32.3% to 35.9% on BERT-large. This demonstrates that the manually created prompts are a somewhat weak lower bound; there are other prompts that further improve the ability to query knowledge from LMs. However, the manual prompts are indeed a strong baseline, often superior to the best mined prompts or their best paraphrases.

[Table 4](#) shows some of the mined prompts and that resulted in a large performance gain compared to the manual ones. For the relation `religion`, “*x who converted to y*” improved 60.0% over the

⁶<https://github.com/pytorch/fairseq/tree/master/examples/wmt19>

Prompts	Top1	Top3	Top5	Opti.	Oracle
<i>BERT-base (Man=31.1)</i>					
Mine	30.7	32.7	31.2	36.9	45.1
Mine+Man	31.9	34.5	33.8	38.1	47.9
Mine+Para	30.7	33.0	33.7	33.6	45.0
Man+Para	<i>34.1</i>	35.8	36.6	37.3	47.9
<i>BERT-large (Man=32.3)</i>					
Mine	34.4	33.8	33.1	40.4	47.9
Mine+Man	<i>36.0</i>	38.6	37.1	41.9	50.8
Mine+Para	32.1	35.0	36.1	37.0	47.3
Man+Para	35.9	37.3	38.0	38.8	50.0

Table 2: Micro-averaged accuracy of different methods (%). **Majority** gives us 22.0%. Italic indicates best single-prompt accuracy, and bold indicates the best non-oracle accuracy overall.

manually defined prompt of “*x is affiliated with the y religion*”, and for the relation `subclass_of`, “*x is a type of y*” raised the accuracy by 22.7% over “*x is a subclass of y*”. It can be seen that the largest gains from using mined prompts seem to occur in cases where the manually defined prompt is more complicated syntactically (e.g. the former), or when it uses less common wording (e.g. the latter) than the mined prompt.

Prompt Ensembling Next we turn to experiments that use multiple prompts to query the LM. Comparing the single-prompt results in Column 1 to the ensembled results in the following three columns, we can see that ensembling multiple prompts almost always leads to better performance. The simple average used in **Top3** and **Top5** outperforms **Top1** across different prompt generation methods. The optimized ensemble further raises micro-averaged accuracy to 36.9% and 40.4% on BERT-base and BERT-large respectively, outperforming the rank-based ensemble by a large margin. These two sets of results demonstrate that diverse prompts can indeed query the LM in different ways, and that the optimization-based method is able to find weights that effectively combine different prompts together.

We list the learned weights of top-3 mined prompts and micro-averaged accuracy gain over only using the top-1 prompt in [Table 5](#). Weights tend to concentrate on one particular prompt, and the other prompts serve as complements. We also depict the performance of the rank-based ensem-

Prompts	Top1	Top3	Top5	Opti.	Oracle
<i>BERT-base (Man=22.8)</i>					
Mine	21.2	22.1	21.4	24.0	32.2
Mine+Man	22.0	24.0	23.4	25.2	34.6
Mine+Para	20.2	22.1	22.6	22.6	32.2
Man+Para	22.8	23.8	24.6	25.0	34.9
<i>BERT-large (Man=25.7)</i>					
Mine	24.8	25.0	24.1	27.7	36.4
Mine+Man	27.0	27.6	26.8	29.5	38.9
Mine+Para	23.4	24.8	25.7	25.8	36.2
Man+Para	25.9	27.8	28.3	28.0	39.3

Table 3: Macro-averaged accuracy of different methods (%). **Majority** gives us 2.2%. *Italic* indicates best single-prompt accuracy, and **bold** indicates the best non-oracle accuracy overall.




Figure 2: Performance of different types of prompts.

ble method with respect to the number of prompts in Figure 2. For mined prompts, top-2 or top-3 usually gives us the best results, while for paraphrased prompts, top-5 is the best. Incorporating more prompts does not always improve accuracy, a finding consistent with the rapidly decreasing weights learned by the optimization-based method.

Mining vs. Paraphrasing For the rank-based ensembles (**Top1**, **3**, **5**), prompts generated by paraphrasing usually perform better than mined prompts, while for the optimization-based ensemble (**Opti.**), mined prompts perform better. We conjecture this is because mined prompts exhibit more variation compared to paraphrases, and proper weighting is of central importance. This difference in the variation can be observed in the average edit distance between the different prompts, which is 3.27 and 2.73 for mined and paraphrased prompts respectively. However, the improvement led by ensembling paraphrases is still significant over just using one prompt (**Top1** vs. **Opti.**), raising micro-accuracy from 30.7% to 33.6% on BERT-base, and




Figure 3: Correlation of edit distance between prompts and their prediction divergence.

from 32.1% to 37.0% on BERT-large. This indicates that even small modifications to prompts can result in relatively large changes in predictions. Table 6 demonstrates cases where modification of one word (either function or content word) leads to significant accuracy improvements, indicating that large-scale LMs are still brittle to small changes in the ways they are queried.

Micro vs. Macro Comparing Table 2 and Table 3, we can see that macro-averaged accuracy is much lower than micro-averaged accuracy, indicating that macro-averaged accuracy is a more challenging metric that evaluates how many unique objects LMs know. Our optimization-based method improves macro-averaged accuracy from 22.8% to 24.0% on BERT-base, and from 25.7% to 27.7% on BERT-base. This again confirms the effectiveness of ensembling multiple prompts, but the gains are somewhat smaller. Notably, in our optimization-based methods, the ensemble weights are optimized on each example in the training set, which is more conducive to optimizing micro-averaged accuracy. Optimization to improve macro-averaged accuracy is potentially an interesting direction for future work that may result in prompts more generally applicable to different types of objects.

5.3 Analysis

Next, we perform further analysis to better understand what type of prompts proved most suitable for facilitating retrieval of knowledge from LMs.

Prediction Consistency by Prompt We first analyze the conditions under which prompts will yield different predictions. We define the divergence between predictions of two prompts $t_{r,i}$ and $t_{r,j}$ using

ID	Relations	Manual Prompts	Mined Prompts	Acc. Gain
P140	religion	x is affiliated with the y religion	x who converted to y	+60.0
P159	headquarters location	The headquarter of x is in y	x is based in y	+4.9
P20	place of death	x died in y	x died at his home in y	+4.6
P264	record label	x is represented by music label y	x recorded for y	+17.2
P279	subclass of	x is a subclass of y	x is a type of y	+22.7
P39	position held	x has the position of y	x is elected y	+7.9

Table 4: Micro-accuracy gain (%) of the mined prompts over the manual prompts.

ID	Relations	Prompts and Weights	Acc. Gain
P127	owned by	x is owned by y .485 x was acquired by y .151 x division of y .151	+7.0
P140	religion	x who converted to y .615 y tirthankara x .190 y dedicated to x .110	+12.2
P176	manufacturer	y introduced the x .594 y announced the x .286 x attributed to the y .111	+7.0

Table 5: Weights of top-3 mined prompts, and the micro-accuracy gain (%) over only using the top-1 prompt.

ID	Modifications	Acc. Gain
P413	x plays in → at y position	+23.2
P495	x was created → made in y	+10.8
P495	x was → is created in y	+10.0
P361	x is a part of y	+2.7
P413	x plays in y position	+2.2

Table 6: Small modifications (**update**, **insert**, and **delete**) in paraphrase lead to large accuracy gain (%).

the following equation:

$$\text{Div}(t_{r,i}, t_{r,j}) = \frac{\sum_{(x,y) \in r} \delta(R(x,y,t_{r,i}) \neq R(x,y,t_{r,j}))}{|r|},$$

where $R(x, y, t_{r,i}) = 1$ if prompt $t_{r,i}$ can successfully predict y and 0 otherwise, and $\delta(\cdot)$ is Kronecker’s delta. For each relation, we normalize the edit distance of two prompts into $[0, 1]$ and bucket the normalized distance into 5 bins with intervals of 0.2. We plot a box chart for each bin to visualize the distribution of prediction divergence in Figure 3, with the green triangles representing mean values and the green bars in the box representing median values. As the edit distance becomes larger, the divergence increases, which confirms our intuition that very different prompts tend to cause different prediction results. The Pearson correlation coefficient is 0.25, which shows that there is a weak correlation between these two quantities.

POS-based Analysis Next, we try to examine which types of prompts tend to be effective in the

x/y V y/x	x/y V P y/x	x/y V W* P y/x
V = verb particle? adv?		
W = (noun adj adv pron det)		
P = (prep particle inf. marker)		

Table 7: Three part-of-speech-based regular expressions used in ReVerb to identify relational phrases.

abstract by examining the part-of-speech (POS) patterns of prompts that successfully extract knowledge from LMs. In open information extraction systems (Banko et al., 2007), manually defined patterns are often leveraged to filter out noisy relational phrases. For example, ReVerb (Fader et al., 2011) incorporates three syntactic constraints listed in Table 7 to improve the coherence and informativeness of the mined relational phrases. To test whether these patterns are also indicative of the ability of a prompt to retrieve knowledge from LMs, we use these three patterns to group prompts generated by our methods into four clusters, where one cluster is “other” containing prompts that do not match any pattern. We then calculate the rank of each prompt within the top 30 extracted prompts, and plot the distribution of rank using box plots in Figure 4.⁷ From this, we can see that the average rank of prompts matching these patterns is

⁷We use the ranking position of a prompt to represent its quality instead of its accuracy because accuracy distributions of different relations might span different ranges, making accuracy not directly comparable across relations.




Figure 4: Ranking position distribution of prompts with different patterns. The lower is the better.

better than those in the “other” group, confirming our intuitions that good prompts should conform with those patterns. Some of the best performing prompts’ POS signatures are “ x VBD VBN IN y ” (e.g., “ x was born in y ”) and “ x VBZ DT NN IN y ” (e.g., “ x is the capital of y ”).

Cross-model Consistency Finally, it is of interest to know whether the prompts that we are extracting are highly tailored to a specific model, or whether they can generalize across models. To do so, we use both the BERT-base and BERT-large models, and compare when the optimization-based ensembles are trained on the same model, or when they are trained on one model and tested on the other model. As shown in Table 8, we found that in general that there is usually some drop in performance in the cross-model scenario (third and fifth columns), but the losses tend to be small, and the highest performance when querying BERT-base is actually achieved by the weights optimized on BERT-large. Notably, the best accuracies of 38.5% and 40.4% with the weights optimized on the other model are still much higher than those obtained by the manual prompts (31.1% and 32.3% respectively), indicating that optimized prompts still afford large gains in performance over the previous method.

6 Omitted Design Elements

Finally, in addition to the elements of our main proposed methodology in § 3 and § 4, we experimented with a few additional methods that did not prove highly effective, and thus were omitted from our final design. We briefly describe these below, along with cursory experimental results.

Test Train	BERT-base		BERT-large	
	base	large	large	base
mine	36.9	36.6	40.4	39.0
mine+man	38.1	38.5	41.9	40.0
mine+para	33.6	33.7	37.0	35.3
man+para	37.3	35.6	38.8	37.5

Table 8: Micro-accuracy (%) of cross-model optimization. The first row is the model to test, and the second row is the model on which prompt weights are learned.

Prompts	Top1	Top3	Top5	Opti.	Oracle
before	31.9	34.5	33.8	38.1	47.9
after	30.2	32.5	34.7	37.5	50.8

Table 9: Micro-accuracy (%) before and after LM-aware prompt fine-tuning.

6.1 LM-aware Prompt Generation

We examined methods to generate prompts by solving an optimization problem that maximizes the probability of producing the ground-truth objects with respect to the prompts:

$$t_r^* = \arg \max_{t_r} P_{\text{LM}}(y|x, t_r),$$

where $P_{\text{LM}}(y|x, t_r)$ is parameterized with a pre-trained LM. In other words, this method directly searches for a prompt that causes the LM to assign ground-truth objects the highest probability.

Solving this problem of finding text sequences that optimize some continuous objective has been studied both in the context of end-to-end sequence generation (Hoang et al., 2017), and in the context of making small changes to an existing input in the context of adversarial attacks (Ebrahimi et al., 2018; Wallace et al., 2019). However, we found that directly optimizing prompts guided by gradients was unstable and usually yielded unreadable snippets in our preliminary experiments. Thus, we instead resorted to a more straightforward hill-climbing method that starts with an initial prompt, then masks out one token at a time and replaces it with the most probable token conditioned on the other tokens, inspired by the mask-predict decoding algorithm used in non-autoregressive machine




Figure 5: Performance of two interpolation methods.

translation (Ghazvininejad et al., 2019):⁸

$$P_{\text{LM}}(w_i | t_r \setminus i) = \frac{\sum_{\langle x, y \rangle \in r} P_{\text{LM}}(w_i | x, t_r \setminus i, y)}{|r|},$$

where w_i is the i -th token in the prompt and $t_r \setminus i$ is the prompt with the i -th token masked out. We followed a simple rule that modifies a prompt from left to right, and this is repeated until convergence.

We used this method to refine all the mined and manual prompts on the T-REx-train dataset, and display their performance on the T-REx dataset in Table 9. After fine-tuning, the oracle performance increased significantly, while the ensemble performances (both rank-based and optimization-based) dropped slightly. This indicates that LM-aware fine-tuning has the potential to discover better prompts, but some portion of the refined prompts may have over-fit to the training set upon which they were optimized.

6.2 Linear vs. Log-linear Combination

As mentioned in § 4.2, we use log-linear combination of probabilities in the ensemble methods in our main experiments. However, it is also possible to calculate probabilities through regular linear interpolation:

$$P(y|x, r) = \sum_{i=1}^K \frac{1}{K} P_{\text{LM}}(y|x, t_{r,i}) \quad (4)$$

We compare these two ways to combine predictions from multiple mined prompts in Figure 5 (§ 4.2). We assume that log-linear combination outperforms linear combination because log probabilities make it possible to penalize objects that are very unlikely given any certain prompt.

⁸In theory, this algorithm can be applied to both masked LMs like BERT and traditional left-to-right LMs, since the masked probability can be computed using Bayes’ theorem for traditional LMs. However, in practice, due to the large size of vocabulary, it can only be approximated with beam search, or computed with more complicated continuous optimization algorithms (Hoang et al., 2017).

Features	Mine		Paraphrase	
	macro	micro	macro	micro
forward	38.1	25.2	37.3	25.0
+backward	38.2	25.5	37.4	25.2

Table 10: Performance (%) of using forward and backward features with BERT-base.

6.3 Forward and Backward Probabilities

Finally, given class imbalance and the propensity of the model to over-predict the majority object, we examine a method to encourage the model to predict subject-object pairs that are more strongly aligned. Inspired by the maximum mutual information objective used in Li et al. (2016a), we add the backward log probability $\log P_{\text{LM}}(x|y, t_{r,i})$ of each prompt to our optimization-based scoring function in Equation 3. Due to the large search space for objects, we turn to an approximation approach that only computes backward probability for the most probable B objects given by the forward probability at both training and test time. As shown in Table 10, the improvement resulting from backward probability is small, indicating that a diversity-promoting scoring function might not be necessary for knowledge retrieval from LMs.

7 Related Work

Much work has focused on understanding the internal representations in neural NLP models (Belinkov and Glass, 2019), either by using extrinsic probing tasks to examine whether certain linguistic properties can be predicted from those representations (Shi et al., 2016; Linzen et al., 2016; Belinkov et al., 2017), or by ablations to the models to investigate how behavior varies (Li et al., 2016b; Smith et al., 2017). For contextualized representations in particular, a broad suite of NLP tasks are used to analyze both syntactic and semantic properties, providing evidence that contextualized representations encode linguistic knowledge in different layers (Hewitt and Manning, 2019; Tenney et al., 2019a,b; Jawahar et al., 2019; Goldberg, 2019).

Different from analyses probing the representations themselves, our work follows Petroni et al. (2019) in probing for factual knowledge with manually defined prompts. Somewhat unsurprisingly, some have found that the knowledge contained by pre-trained LMs is still limited, with performance becoming lower when tested on hard-to-guess facts

(Poerner et al., 2019), leading to a conclusion that using LMs as reliable knowledge sources may still be far from reality. However, as our work points out, sub-optimal manually created prompts may be significantly under-estimating the true performance obtainable by LMs in these settings.

Orthogonally, some previous works integrate external knowledge bases so that the language generation process is explicitly conditioned on symbolic knowledge (Ahn et al., 2016; Yang et al., 2017; IV et al., 2019; Hayashi et al., 2020). Similar extensions have been applied to large-scale pre-trained LMs like BERT, where contextualized word representations are enhanced with external entity embeddings either at training time or solely at test time (Peters et al., 2019; Poerner et al., 2019). In contrast, we focus on better knowledge retrieval methods through prompts from pre-trained LMs as-is, without modifying them.

8 Conclusion

In this paper, we examined the importance of the prompts used in retrieving factual knowledge from language models. We propose mining-based and paraphrasing-based methods to systematically generate diverse prompts to query specific pieces of relational knowledge. Those prompts, when combined together, improve factual knowledge retrieval accuracy by 7%, outperforming manually designed prompts by a large margin. Our analysis indicates that LMs are indeed more knowledgeable than initially indicated by previous results, but they are also quite sensitive to how we query them. This indicates potential future directions such as (1) more robust LMs that can be queried in different ways but still return similar results, (2) methods to incorporate factual knowledge in LMs, and (3) further improvements in optimizing methods to query LMs for knowledge. Finally, we have released all our learned prompts to the community as the LM Prompt and Query Archive (LPAQA), available at: <https://github.com/jzbyjyb/LPAQA>.

Acknowledgments

This work was supported by a gift from Bosch Research and NSF Award No. 1815287. We would like to thank Paul Michel and Hiroaki Hayashi for their insightful comments and suggestions.

References

- Sungjin Ahn, Heeyoul Choi, Tanel Pärnamaa, and Yoshua Bengio. 2016. [A neural knowledge language model](#). *CoRR*, abs/1608.00318.
- 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*, pages 2895–2905, Florence, Italy. Association for Computational Linguistics.
- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. 2007. [Open information extraction from the web](#). In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 2670–2676.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov and James R. Glass. 2019. [Analysis methods in neural language processing: A survey](#). *TACL*, 7:49–72.
- Claudio Carpineto and Giovanni Romano. 2012. [A survey of automatic query expansion in information retrieval](#). *ACM Comput. Surv.*, 44(1):1:1–1:50.
- Andrew M. Dai and Quoc V. Le. 2015. [Semi-supervised sequence learning](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3079–3087.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Hady ElSahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon S. Hare, Frédérique 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, May 7-12, 2018*.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pages 1535–1545.
- Michael Gamon, Anthony Aue, and Martine Smets. 2005. Sentence-level MT evaluation without reference translations: Beyond language modeling. In *Proceedings of EAMT*, pages 103–111.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-predict: Parallel decoding of conditional masked language models. 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 6114–6123, Hong Kong, China. Association for Computational Linguistics.
- Yoav Goldberg. 2019. Assessing bert’s syntactic abilities. *CoRR*, abs/1901.05287.
- Hiroaki Hayashi, Zecong Hu, Chenyan Xiong, and Graham Neubig. 2020. Latent relation language models. In *Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, New York, USA.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4129–4138.
- Cong Duy Vu Hoang, Gholamreza Haffari, and Trevor Cohn. 2017. Towards decoding as continuous optimisation in neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 146–156, Copenhagen, Denmark. Association for Computational Linguistics.
- Robert L. Logan IV, Nelson F. Liu, Matthew E. Peters, Matt Gardner, and Sameer Singh. 2019. Barack’s wife hillary: Using knowledge graphs for fact-aware language modeling. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5962–5971.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does BERT learn about the structure of language? In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3651–3657.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. Understanding neural networks through representation erasure. *CoRR*, abs/1612.08220.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of lstms to learn syntax-sensitive dependencies. *TACL*, 4:521–535.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. The natural language decathlon: Multitask learning as question answering. *arXiv preprint arXiv:1806.08730*.
- Oren Melamud, Jacob Goldberger, and Ido Dagan. 2016. context2vec: Learning generic context embedding with bidirectional LSTM. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, CoNLL 2016, Berlin, Germany, August 11-12, 2016*, pages 51–61.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. On the state of the art of evaluation in neural language models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. Regularizing and optimizing LSTM language models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Tomas Mikolov and Geoffrey Zweig. 2012. Context dependent recurrent neural network language model. In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 234–239. IEEE.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. Facebook fair’s WMT19 news translation task submission. In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 314–319.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Lan-*

- guage Technologies, *NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237.
- Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. 2019. [Knowledge enhanced contextual word representations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 43–54, Hong Kong, China. 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.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2019. [Bert is not a knowledge base \(yet\): Factual knowledge vs. name-based reasoning in unsupervised qa](#). *arXiv preprint arXiv:1911.03681*.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. [Style transfer through back-translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876, Melbourne, Australia. 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).
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain yourself! leveraging language models for commonsense reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Deepak Ravichandran and Eduard Hovy. 2002. [Learning surface text patterns for a question answering system](#). In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 41–47. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. [Atomic: An atlas of machine commonsense for if-then reasoning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3027–3035.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. [Does string-based neural MT learn source syntax?](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1526–1534, Austin, Texas. Association for Computational Linguistics.
- Noah A. Smith, Chris Dyer, Miguel Ballesteros, Graham Neubig, Lingpeng Kong, and Adhiguna Kuncoro. 2017. [What do recurrent neural network grammars learn about syntax?](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pages 1249–1258.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4593–4601.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. [Representing text for joint embedding of text and knowledge bases](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1499–1509.
- Trieu H Trinh and Quoc V Le. 2018. [A simple method for commonsense reasoning](#). *arXiv preprint arXiv:1806.02847*.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). 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 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Zichao Yang, Phil Blunsom, Chris Dyer, and Wang Ling. 2017. [Reference-aware language models](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1850–1859.
- Geoffrey Zweig and Christopher JC Burges. 2011. [The microsoft research sentence completion challenge](#). *Microsoft Research, Redmond, WA, USA, Tech. Rep. MSR-TR-2011-129*.