

What Matters in Learning Facts in Language Models? Multifaceted Knowledge Probing with Diverse Multi-Prompt Datasets

Anonymous ACL submission

Abstract

Large language models (LLMs) face issues in handling factual knowledge, making it vital to evaluate their true ability to understand facts. In this study, we introduce knowledge probing frameworks, BELIEF(-ICL), to evaluate the knowledge understanding ability of decoder-based as well as encoder-based pre-trained LMs (PLMs) from diverse perspectives. BELIEFs utilize a multi-prompt dataset to evaluate PLM’s accuracy, consistency, and reliability in factual knowledge understanding. To enable a more reliable evaluation with BELIEFs, we semi-automatically create MyriadLAMA, which has massively diverse prompts. We validate the effectiveness of BELIEFs in correctly and comprehensively evaluating PLM’s factual understanding ability via extensive evaluations with recent LLMs. We then investigate key factors in learning facts in LLMs, and reveal the limitation of the prompt-based knowledge probing. The dataset is anonymously publicized.¹

1 Introduction

One of the strongest motivations for training a language model (LM) using a larger amount of text is to increase the LM’s ability to manage factual knowledge in knowledge-intensive tasks (Kamalloo et al., 2023). However, even if LMs are trained on massive text, they suffer from hallucinations that generate incorrect knowledge-grounded sentences (Zhang et al., 2023). Considering that large language models (LLMs) are being widely applied to real-world tasks, it is vital to understand the true knowledge capacity of LLMs and what factors influence the knowledge learning during pre-training.

However, evaluating the LLM’s knowledge capacity is still challenging. Although LAMA probing (Petroni et al., 2019) evaluates the knowledge contained in pre-trained LMs (PLMs), it provides only prediction accuracy. Some studies diversify

the LAMA probing datasets to compute prediction consistency (robustness) (Elazar et al., 2021; Jiang et al., 2020), but those datasets have either low quality or low quantity issues (§ B). Moreover, since the LAMA probe assumes encoder-based PLMs with the masked LM objective to solve fill-in-the-blank tasks, directly applying it to recent decoder-based LLMs will underestimate their ability of knowledge understanding. Although recent studies leveraged QA-based datasets to probe LLMs’ knowledge (Kalo and Fichtel, 2022; Mallen et al., 2023; Wiland et al., 2024; Maekawa et al., 2024) via prediction accuracy, they overlook other important aspects than accuracy such as robustness to diverse prompts and the reliability of predictions, which are important for real-world applications.

In this study, we introduce BELIEF (§2) and its variant BELIEF-ICL (§3), the multifaceted factual knowledge probing frameworks for encoder- and decoder-based PLMs. BELIEFs utilize diverse prompts for each fact to eliminate the specific influence of linguistic expression on the results. It enables us to evaluate PLMs’ robustness and reliability by measuring accuracy fluctuations, consistency, and overconfidence in fact prediction. Since BELIEFs require a multi-prompt probing dataset with diverse prompts for each fact, we build a new probing dataset, MyriadLAMA, to enable a more accurate and comprehensive evaluation (§4).

We applied BELIEFs to various encoder- and decoder-based PLMs, including BERT (Devlin et al., 2019) and Llama2 (Touvron et al., 2023) (§5.1). Through extensive evaluations, we verify the effectiveness of BELIEFs in accurately uncovering PLMs’ ability in factual understanding (§5). Moreover, by comparing different PLMs, we reveal insights into the factors affecting the acquisition of facts during pre-training (§6). Finally, we explore the upper limits of covered knowledge by different methods and reveal the limitations of prompt-based knowledge probing (§7).

¹<https://anonymous.4open.science/r/belief-CC8A>

2 BELIEF Framework

We first present a multifaceted factual probing framework - BELIEF, specifically for encoder-based PLMs. Using a multi-prompt knowledge probing dataset, BELIEF evaluates the facts understanding ability of PLMs from accuracy, robustness and reliability (§2.2-2.4). Here, robustness measures PLMs’ ability to maintain consistent accuracy and predictions when given different prompts in evaluation. Reliability reflects the extent to which we can trust the PLMs’ predictions.

2.1 Preliminary

To accurately evaluate the facts in PLMs, BELIEF aggregates results from multiple prompts for each fact to mitigate biases from specific linguistic expressions. This requires a dataset with varied expressions for each fact, namely multi-prompt factual probing dataset.

We follow the setting of fill-in-the-blank datasets, where each fact is represented as a **knowledge triple** ⟨subject, relation, object⟩ (e.g., ⟨Tokyo, Capital, Japan⟩). To probe PLMs for a knowledge triple, we first create a **masked prompt** (hereafter, **prompt**) (e.g., “Tokyo is the capital of [Y]”) for it and then input it into PLMs to see if they correctly predict the object token. To create such prompts, we first need a template for the relation (hereafter, **relational template**, e.g., [X] is the capital of [Y]). We then fill the template with target knowledge triples, replacing [X] with a subject expression and [Y] with a [MASK] token. A multi-prompt dataset offers diverse prompts for each fact by providing varied relational templates and entity expressions.

We denote the subject-relation pairs in dataset as T , the set of prompts for a given subject-relation pair $t \in T$ as P_t . If the output distribution corresponding to mask token of a prompt is $\mathcal{O} = \{(w_j, o_j) \mid \sum_j o_j = 1\}$, the prediction result is defined as the token $\hat{w} = \operatorname{argmax}_{w_j, (w_j, o_j) \in \mathcal{O}} o_j$.

2.2 Accuracy and its fluctuations

To correctly evaluate the accuracy of PLMs, we aggregate predictions from diverse prompts. Specifically, we randomly select one prompt for each subject-relation pair $t \in T$ to form a set of prompts for all triples $P = \{p_1, \dots, p_{|T|}\}$. By feeding these prompts P to PLMs, we can calculate one accuracy value based on their predictions. We repeat this process to collect a set of accuracies, which we then use to calculate both average and fluctuation.

Average accuracy: In BELIEF, accuracy metrics include Acc@1, which measures the rate of prompts with the correct token predicted within the top-1 output probabilities. Then we repeat this process N times to obtain a set of accuracies, denoted as $V_{\text{Acc@1}}$, where $|V_{\text{Acc@1}}| = N$. The final average accuracy is calculated as the mean value of $V_{\text{Acc@1}}$.

Fluctuation of accuracy: For $V_{\text{Acc@1}}$, we evaluate accuracy fluctuations using the range and standard deviation. The range is determined by the difference between the maximum and minimum accuracy values in $V_{\text{Acc@1}}$.

2.3 Consistency

For each subject-relation pair t , we assess the PLM’s consistency in predicting the object across different prompts in P_t . Specifically, we compute the degree of match between the prediction result \hat{w}_t^i for a given prompt p_t^i and the prediction results \hat{w}_t^j for other prompts $p_t^j \in P_t$ (where $j \neq i$), across all subject-relation pairs in T :

$$\text{Consist} = \frac{1}{|T|} \sum_{t \in T} \frac{\sum_{i,j:i \neq j, i,j \leq |P_t|} \mathbb{1}[\hat{w}_t^i = \hat{w}_t^j]}{\frac{1}{2}|P_t|(|P_t| - 1)} \quad (1)$$

2.4 Reliability

The reliability of PLMs reflects the extent to which we can trust the predictions they provide. In our study, we measure PLMs’ overconfidence level in making fact prediction, drawing from the expected error calibration metric (Desai and Durrett, 2020). Specially, we measure the difference between true prediction accuracy and models’ confidence to their predicted tokens. For each prompt, we first acquire the maximum probability (hereafter, **confidence**) from the output distribution for the mask token. Subsequently, all of the prompts are arranged in descending order based on confidence and segmented into M bins ($P^{(1)}, P^{(2)}, \dots, P^{(M)}$), with the same amount of data points in each bin. For each bin i , we compute the average accuracy $\overline{\text{Acc@1}}^{(i)}$ and average confidence $\overline{o_{max}}^{(i)}$. In our work, we use $M = 10$ for all the experiments. Finally, the PLM’s overconfidence in predicting the object is assessed by averaging differences between average confidence and accuracy across all bins:

$$\text{Overconf} = \sum_{i=1}^M \frac{|P^{(i)}|}{M} (\overline{o_{max}}^{(i)} - \overline{\text{Acc@1}}^{(i)}) \quad (2)$$

The closer the Overconf is to zero, the more aligned the model’s confidence is with its accuracy, indicating reliable confidence. A negative Overconf value means the model is underconfident.

3 BELIEF-ICL for Decoder-based LLMs

Recent LLMs are based on decoder-only Transformer architecture, and are trained to predict subsequent tokens in a sequence. This makes it challenging for them to directly predict [MASK] tokens in masked prompts, as they cannot utilize information following the [MASK] (*e.g.*, “[MASK] and Tokyo are twin cities”). To comprehensively evaluate LLMs and enable fair comparison between encoder- and decoder-based models, we extend BELIEF to LLMs by employing in-context learning (ICL), termed BELIEF-ICL.

3.1 In-context learning for fact probe

The in-context learning ability allows LLMs to perform complex tasks during inference using task-specific prompts (Brown et al., 2020). When designing ICL for evaluating factual knowledge, it is essential to consider **task instructions** and **context examples** appended to the target prompts.

1) Task instruction: We introduce the mask prediction (MP) instruction for prompting LLMs generating one word answer for the target masked prompt. The task instruction is formulated as “Predict the [MASK] in each sentence in one word.”.

2) Context settings: We propose four types of contexts to assess the impact of exemplar selection on factual knowledge probing, following the QA format outlined in InstructGPT (Ouyang et al., 2022).

zero-shot: no context;

X-random: sampling X facts from all relations as the few-shot learning examples;

X-relation: sampling X facts from the same relation but with random templates;

X-template: sampling X facts from the same relations and the same template.

In the few-shot learning settings, we ensure that the target fact is excluded in the examples. Refer to §E for examples of prompts.

3.2 Evaluation methods

Since LLMs generate responses without a token limit, matching the correct answer with the model’s output can be challenging. Variations in language expressions, such as the presence or absence of articles and singular or plural forms, complicate this

process. Additionally, the model may generate extra tokens not relevant to the [MASK] token, such as parts of the prompt. For example, for the prompt “John Lennon can play [MASK],” both “guitars” and “a guitar” should be considered correct.

To measure BELIEF metrics for LLMs, we need to compare two strings: the generated text and the correct target expression for Acc@1, and two generated texts for Consist and Overconf. Here, we first normalize the strings by tokenizing and lemmatizing them. For example, “a guitar” and “guitars” are normalized to “a, guitar” and “guitar.” If a normalized string list is included in the other (partial matching), they are considered matched.

1) Accuracy and its fluctuations: Accuracy is calculated by comparing the string generated by the model using a greedy decoding strategy to the correct answers. Notably, the matching judgment is one-directional: it only checks if the correct answer is included in the generated string. One-directional matching is adopted to avoid incorrect judgments from the model generating unrelated words. We use the same N as in §2.2 for accuracy measurement.

2) Consistency: We use bi-directional matching to evaluate the consistency (Consist) of generated sequences from two prompts.

3) Reliability: To calculate overconfidence, we need the model’s confidence (probability) in its output. However, we cannot obtain this directly from the probability over generated tokens, as LLMs can produce diverse outputs that represent the same answer. To address this, we propose an approximate measurement. For each prompt, we generate 100 samples using multinomial sampling². We then measure the matching rate between the output generated from greedy decoding and the outputs from the 100 samples. This matching rate serves as the confidence value for the prompt³. The calculation of Overconf follows the same setting in §2.4. This method can approximate the BELIEF’s Overconf calculation as BELIEF sampling answers from the output distribution. Note that, due to the high cost of generating 100 samples for each fact, we adopt a more efficient approach. We sample 10K prompts from 10K unique subject-relation pairs and only use these 10K prompts for answer sampling.

²Multinomial sampling selects a next token according to the probability, over the entire vocabulary given by the model.

³This method can approximate the overconfidence calculation in BELIEF of sampling answers from the output distribution. It makes the confidence calculated by BELIEF for encoder-based models comparable to that in BELIEF-ICL.

4 MyriadLAMA Dataset

The fairness and accuracy of BELIEF evaluation depend on the diversity and quality of multi-prompt factual probing datasets. However, existing datasets are either manually rewritten in small numbers (Elazar et al., 2021) or mined from texts (Jiang et al., 2020). The former is accurate but lacks diversity, providing an average of 7.3 prompts per fact with limited variation. For example, templates like “[X] works as [Y]” and “[X], who works as [Y]” are provided as different templates but very similar. Additionally, the number of templates is highly imbalanced, with 8 out of 46 relations having only one template, while P138⁴ has 20. The latter is diverse but includes templates that do not necessarily imply the relationship. For instance, for relation P937 (work location)⁵, the mined templates include “[X] to meet [Y].,” which significantly deviates from the original meaning. To achieve a more accurate and fair evaluation, we introduce MyriadLAMA, a new multi-prompt factual probing dataset with improved diversity while retaining quality. Refer to § B for detailed qualitative and quantitative comparisons between MyriadLAMA and prior datasets.

4.1 Dataset construction

We build MyriadLAMA by semi-automatically extending the existing single-prompt probing dataset LAMA-UHN (Petroni et al., 2020). MyriadLAMA generates multiple prompts for each fact by providing multiple, equal relational templates for each relation and varying the linguistic expressions of subjects. Additionally, MyriadLAMA offers multiple expressions for each object to cover missed facts that are correctly predicted but in different tokens. For example, for the query “John Lennon was born in [MASK]”, acceptable tokens could include “UK” and “Britain.”⁶

Specifically, we define knowledge triples that neglect the diversity of surface expressions as **unique triples** and distinguish them from **derived triples**, which embody the diverse entity expressions and relational templates in each unique triple. The triple extension methods are described below.

Extending entities: The knowledge triples in

⁴<https://www.wikidata.org/wiki/Property:P9138>

⁵<https://www.wikidata.org/wiki/Property:P937>

⁶Note that, we follow the setting in LAMA-UHN to include triples where the object is a single token according to the BERT tokenizer. During evaluation, we consider the fact to be present, if the model’s predicted token matches any of the correct tokens, regardless of which correct answer is predicted.

	LAMA-UHN	MyriadLAMA
Relational templates	41	4100
Unique triples	27,106	34,048
Derived triples	27,106	21,140,500
Subject-relation pairs	24,643	24,643
Prompts	24,643	6,492,800

Table 1: Statistics of LAMA-UHN and MyriadLAMA.

LAMA-UHN constitute a subset of the Wikipedia knowledge base T-REx (Elsahar et al., 2018). T-REx selectively includes only certain objects for subject-relation pairs. MyriadLAMA extends the unique triples in LAMA-UHN by mining T-REx using subject-relation as key to include other available objects. For example, if LAMA-UHN contains only E_{guitar} for instruments that “John Lennon” can play, we can extend the unique triple to include E_{piano}. We also extend the entity expressions using aliases obtained from Wikidata.⁷

Paraphrasing relational templates: MyriadLAMA creates a great variety of relational templates by a semi-automatic process. Firstly, we manually generate five distinct templates for each relation. They incorporate entailment expressions and diverse syntactic patterns like statements and question-answer formats to provide semantic and syntactic variations. Next, to enhance quantity and lexical diversity, we automatically paraphrase each manually created template 19 times using the GPT-4 API.⁸ Finally, all templates are filtered by human reviewers to remove low quality templates, yielding a total of 4100 templates covering 41 relations.

4.2 Dataset Statistics

We report the statistics of LAMA-UHN and MyriadLAMA, as shown in Table 1. As the result, the number of derived triples is increased from 27,106 in LAMA-UHN to 21,140,500, by combining various semi-automatically generated relational templates and the alias expressions for subject and object entities. As the prompts are generated from derived triples without considering the object expressions, the number of generated prompts are less than the number of derived triples, which is increased from 24,643 to 6,492,800.

Please refer to the appendix for details on dataset construction (§ A) and validity analysis of MyriadLAMA (§ C). Examples of extended templates are provided in § A.3.

⁷https://www.wikidata.org/wiki/Wikidata:Data_access

⁸OpenAI: gpt-4-1106-preview

LLMs (#params)	Pre-training corpora	
	size	source
BERT _{base} (110M)	3.3G	English Wikipedia and BookCorpus
BERT _{large} (336M)	3.3G	
BERT _{wwm} (336M)	3.3G	
Llama2-7B (7B)	2.0T	A collection of publicly available online data.
Llama2-13B (13B)	2.0T	
Llama2-70B (70B)	2.0T	
Llama3-8B (8B)	15T+	A collection of publicly available online data.
Llama3-70B (70B)	15T+	
Phi3-mini (3.8B)	3.3T	High-quality documents incl. textbook-like generated text, educational data and etc.
Phi3-small (7B)	4.8T	
Phi3-medium (14B)	4.8T	

Table 2: The PLMs for evaluation.

5 Effectiveness of BELIEFs

5.1 Experimental setups

We introduce the encoder- and decoder-based models used in experiments and evaluation settings. We evaluated three encoder-based models, BERT_{base}, BERT_{large}, and BERT_{wwm},⁹ and eight decoder-based models, Llama2-7B, -13B, -70B, Llama3-8B, -70B, Phi3-mini, -small, and -medium (Table 2). Refer to §F for more detailed model information.

We conduct a full-scale evaluation on LLMs with up to 8 billion parameters. Due to the computational cost of LLM inference, we use only five manually rewritten templates for the LLMs with more than 8B parameters, including Llama2-70B, Llama3-70B, and Phi3-medium.¹⁰ To calculate the average and fluctuation of accuracy (§2.2), we set a large sample number ($N = 50,000$) to provide stable and accurate accuracy measurement.

In the following sections, we analyze the evaluation results on various PLMs to deepen our understanding of how PLMs learn and represent factual knowledge. All evaluation results, including those for another family of encoder-based models, ALBERT, are presented in Section §F.3.

5.2 Do BELIEFs provide additional insights?

BELIEF can offer evaluation from diverse perspectives rather than accuracy. As shown in Table 3 (Above), the evaluation result highlights accuracy fluctuations among the BERT variants. All

⁹BERT_{wwm} masks all tokens for a single word at the same time, while BERT_{base} and BERT_{large} masks a single token.

¹⁰The partial evaluation is sufficient to compare performance across different model sizes.

¹¹X-relation cannot be applied to single-prompt dataset.

PLMs	Acc@1 \uparrow		Fluctuation \downarrow		Consist \uparrow	Overconf	
	LU	MyL	range	stdev			
BERT	BERT _{base}	.2403	.1095	.1534	.0217	.1682	.2154
	BERT _{large}	.2454	.1102	.1574	.0220	.1713	.2052
	BERT _{wwm}	.2448	.1364	.1517	.0208	.1524	.1000
Llama3-8B	zero-shot	.3708	.3427	.2864	.0350	.0240	-.1119
	4-random	.5050	.5205	.2033	.0273	.2156	-.0789
	4-relation	n/a ¹¹	.6871	.1236	.0156	.3659	-.0783
	4-template	.6490	.7268	.0220	.0026	.4015	-.0582

Table 3: Evaluation results of BERT variants (above) and Llama3-8B with BELIEF/-ICL (below). LU indicates LAMA-UHN evaluation results, and MyL represents the average Acc@1 on MyriadLAMA.




Figure 1: Calibration between confidence and Acc@1. (Left: Three BERT models types; Right: Llama3-8B with four different ICL settings.)

BERT models show low consistency and tend to be overconfident in their predictions. Figure 1 (left) depicts the relationship between confidence and Acc@1 of the BERT models, indicating low accuracy even for prompts with confident outputs. Whereas BERT_{wwm} performs better over most BELIEF metrics, BERT_{large} outperforms BERT_{wwm} on LAMA-UHN. This discrepancy arises from the limited prompts used in LAMA-UHN and the the single-faceted evaluation method. This highlights BELIEF’s effectiveness in achieving a more accurate factual probing comparison between PLMs.

5.3 Does ICL adhere to instructions?

We then explore the effectiveness of different ICL settings in extracting facts from LLMs. We evaluate the instruction adherence of these settings from two aspects: predicting facts and generating one-word answers, which are crucial since the target objects in MyriadLAMA are primarily one-word entities. We assessed the fact prediction and one-word generation abilities of Llama2-7B, Llama3-8B, and Phi3-small using average Acc@1 and the rate of one-word generation. As shown in Table 4, only under few-shot settings, LLMs demonstrate

ICL settings	Fact prediction	one-word
	(Average Acc@1)	generation
	(Llama2-7B / Llama3-8B / Phi3-small)	
zero-shot	.3385/.3427/.4258	.4802/.1572/.8883
4-random	.4816/.5205/.4889	.8058/.8147/.8913
4-relation	.6286/.6871/.6339	.9246/.9071/.9287
4-template	.6616/.7268/.6612	.9266/.9187/.9411

Table 4: Instruction following rates on Llama2-7B, Llama3-8B and Phi3- small.

a remarkable ability to understand instructions for fact prediction and generating one-word answers, indicating the effectiveness of prompting LLMs to predict mask tokens through in-context learning and the risk of evaluating LLMs under zero-shot settings. Our evaluation with a QA-style ICL settings also confirm the same insight (see § D for details). Moreover, using exemplars similar to the target prompt in the few-shot context can achieve improvement overall metrics (Table 3 Below), including the one-word generation rate (Table 4).

5.4 Can BELIEFs mitigate bias?

We explore whether BELIEF can mitigate prompt bias in evaluations. To measure prompt bias quantitatively, we use content-free prompts, where the subject is replaced by meaningless tokens (Zhao et al., 2021; Xu et al., 2024), and collect the probabilities of candidate tokens in the output distributions over the mask token.¹² We measure the bias level of the prompt using the certainty of distributions over candidate tokens. Specifically, we define bias level as follows:

$$\text{bias-level} = 1 - \frac{\mathcal{H}}{\mathcal{H}_{max}}$$

where \mathcal{H} is the entropy, and \mathcal{H}_{max} is the maximum entropy for the uniform distribution with same size.

We measure bias in both single- and multi-prompt evaluations. We use average bias-level over all relational templates to represent prompt bias existing during the single-prompt evaluation. Then, we simulate the prompt sampling for accuracy calculation in multi-prompt evaluation by averaging output distributions over templates of the same relation. We then use the bias-level of this average distribution to quantify bias in multi-prompt evaluation. In experiments with BERT models, we observe consistent improvement with multi-prompt

¹²Specifically, we adopt a similar setting to (Zhao et al., 2021), by ensembling the distribution over prompts with three content-free tokens: “N/A,” an empty string, and “?”.

PLMs	Average Acc@1 [↑]	Fluctuation [↓]		Consist [↑]	Overconf
		range	stdev		
Llama2-7B	.6699	.0257	.0034	.4174	-.0933
Llama2-13B	.7080	.0235	.0031	.4326	-.0662
Llama2-70B	.7784	.0190	.0024	.4449	-.0690
Llama3-8B	.7316	.0194	.0025	.4060	-.1119
Llama3-70B	.8211	.0139	.0017	.4636	-.0812
Phi3-mini	.6106	.0314	.0039	.3686	.0911
Phi3-small	.6668	.0306	.0039	.3667	.1221
Phi3-medium	.7100	.0207	.0025	.4009	.0317

Table 5: BELIEF-ICL evaluation on LLMs with the 4-template setting using manually-rewritten templates.

evaluation. For example, BERT_{wwm} has average bias-level of 0.304 for single- and 0.217 for multi-prompt evaluations. Taking P31:instance-of¹³ as an example, the average probability of “science” over all templates is 8.30%, but it rises to 52.79% for template: “[Y] contains [X] as one of its elements.”

6 Differentiating PLMs in Fact Probing

In this section, we compare the performance of different PLMs and explore the factors that impact factual knowledge learning during pre-training.

6.1 Factors affecting the acquisition of facts

1) Pre-training strategy From Table 3, we can observe that BERT_{wwm} outperforms BERT_{large} in terms of all metrics, while BERT_{wwm} differs from BERT_{large} only in the masking strategy during pre-training. The superiority of BERT_{wwm} likely stems from its challenging pre-training paradigm, which requires recalling whole words without sub-token information, enhancing word-level contextual understanding. This underscores the importance of pre-training strategy in knowledge acquisition.

2) Model size: Then, we compare models with difference sizes as shown in Table 5.¹⁴ We find that larger LLMs consistently achieve higher accuracy in predicting facts. Combining with the improvement from BERT_{base} to BERT_{large} from Table 3, the importance of model size in fact acquisition during pre-training is confirmed. Moreover, larger models also demonstrate better robustness in handling diverse prompts and shows better reliability.

3) Pre-training corpora: Table 5 shows that Llama3-8B outperforms the larger Llama2-13B

¹³<https://www.wikidata.org/wiki/Property:P31>

¹⁴Owing to the high computational cost of inference on large LLMs like Llama2-70B, we select only five manually rewritten templates with 4-template ICL setting for evaluation.

PLMs	Acc@1	Fluctuation		Consist	Overconf	1-word gen.
		range	stdev			
Llama2-7B	.3384	.2602	.0299	.1269	-.0731	.4752
Llama3-8B	.3427	.2864	.0350	.0240	-.1119	.1572
Phi3-mini	.3498	.2374	.0292	.1465	.1758	.8641

Table 6: BELIEF-ICL evaluation with zero-shot setting.

in fact probing. This is likely due to Llama3’s pre-training corpus being seven times larger than Llama2. Additionally, Llama3-70B surpasses Llama2-70B, confirming the importance of pre-training data volume for fact acquisition.

In the zero-shot evaluation using the entire MyriadLAMA, as shown in Table 6, Phi3-mini outperforms Llama2-7B and Llama3-8B in knowledge retrieval and fidelity to instructions (one-word generation rate). Given that Phi3-mini has about half the model size of Llama2-7B and Llama3-8B (3.8B), and model size typically enhances knowledge retrieval, this result is notable. This superior performance can be attributed to the high-quality, textbook-like material used for pre-training the Phi3 models, highlighting the significant impact of high-quality training data.

6.2 Factors affecting the reliability

Table 3 shows a significant difference in Overconf between BERT models and Llama3-8B, with BERT models being overconfident and Llama3-8B being underconfident. In this section, we explore the reasons for these differences and investigate additional factors affecting reliability beyond model size.

1) The number of output tokens: One main difference in Overconf calculation between encoder- and decoder-based PLMs is that the decoder-based PLMs will generate multiple tokens. Thus, we investigate the effect of output token count on Overconf values. We divide the MyriadLAMA prompt set into groups based on the number of tokens generated. For each group, we calculate the probability of the entire token sequence and compute Overconf for token counts from 1 to 5.¹⁵

The Overconf values for each group of 1 to 5 output tokens on Llama3-8B (4-template) are -0.1030, -0.0906, -0.0297, -0.0546, and 0.0573, showing models become more overconfident with more output tokens. This trend is consistent across models.

2) The writing style of training data: As shown

¹⁵The prompts generated within five tokens cover 98.78% of Llama3-8B’s generations with the 4-template ICL setting.




Figure 2: The knowledge sharing rate between models and ICL settings. Each cell indicates the rate of facts correctly predicted by the top-listed model relative to those captured by the left-listed model.

in Table 5, the Phi3 and BERT models, which were pre-trained on textbook-like text such as Wikipedia, exhibit significant overconfidence in their predictions. In contrast, the Llama2 and Llama3 models, which use general online data, show better calibration between accuracy and confidence (Figure 1 Right). This is probably because the Phi3 and BERT models are trained to predict textbook-like text with more expressions directly describing factual knowledge (e.g., definition sentences). This highlights the significant impact of the pre-training corpora’s writing style on PLM reliability.

6.3 How do PLMs perceive facts differently?

Finally, we measure the differences in fact coverage among models. We first collect the correctly predicted facts for each template, defining these as the model’s covered facts. Given the covered facts of two models, we measure the knowledge sharing rates using an asymmetric metric. This metric calculates the proportion of shared facts relative to each model’s total covered facts.

The results are shown in Figure 2. The average sharing rate among BERT models is 69.07%, while for Llama2-7B, Llama3-8B, and Phi3-mini in the zero-shot ICL setting, it is 67.58%. In comparison, the average sharing rate between encoder- and decoder-based PLMs reduced to 47.10%. Additionally, knowledge sharing rates in both zero-shot and 4-template settings indicate that incorporating context examples improves the knowledge elicited by the models. However, about 10% of knowledge can still only be elicited in the zero-shot setting.

PLMs	Average	Maximum	Oracle
BERT _{wwm}	.1364	.4501	.6636
Llama2-7B (zero-shot)	.3385	.6577	.8153
Llama3-8B (zero-shot)	.3427	.7099	.8756
Phi3-small (zero-shot)	.4258	.6828	.8642
Llama2-7B (4-template)	.6616	.7197	.8133
Llama3-8B (4-template)	.7268	.7731	.8628
Phi3-small (4-template)	.6612	.7181	.8346

Table 7: **Average**: mean accuracy using all templates; **Maximum**: accuracy with the best template for each relation; **Oracle**: accuracy when the best template is selected for each factual knowledge.

7 Limitation of prompt-based probing

Finally, we examine the limitation of prompt-based knowledge probing by using our massively diverse dataset. First, we gauge the average knowledge coverage rate by using the average Acc@1 (**average**). Next, for each relation, we calculate the maximum Acc@1 using the template that yields the highest accuracy,¹⁶ and use this value to estimate the upper limit of prompt-based knowledge probing (**maximum**). Finally, we approximate the upper limit of facts contained in LLMs, by considering a fact as existing if at least one prompt of this fact can produce the correct answer (**oracle**).

Table 7 shows three knowledge coverage rates on small-scale LLMs. For BERT_{wwm} and LLMs with zero-shot settings, we observe nearly a 30% increase between **average** and **maximum** accuracy, emphasizing the importance of selecting suitable templates for specific relationships and the potential gains from prompt engineering. This gap can be reduced to 5% by using few-shot settings. However, the gap between **maximum** and **oracle** accuracy mostly remains. This indicates that different facts prefer different templates, suggesting no versatile template works for all facts. Combining different templates reveals the true upper limits of PLMs in factual understanding and underscores the importance of using diverse prompts for knowledge retrieval instead of optimizing a single prompt. Refer to §F.4 for results on more models.

8 Related work

The LAMA probe was first proposed to evaluate the utility of PLMs as knowledge bases via solving the fill-in-the-blank task (Petroni et al., 2019).

¹⁶We select the prompt with the best subject expression among prompts for each fact.

Several researchers extend the LAMA probe to evaluate PLMs’ ability to understand facts from diverse linguistic aspects, such as the effect of negation/mispriming (Kassner and Schütze, 2020), distractors (Pandia and Ettinger, 2021), multilingual understanding (Keleg and Magdy, 2023; Zhao et al., 2024) and models’ consistency facing prompts with minor nuances (Fierro and Søgaard, 2022; Elazar et al., 2021). However, these studies lack the inspection of PLMs’ reliability in knowledge prediction, which is vital in deploying LLMs to real-world tasks. Moreover, solving the fill-in-the-blank task by LLMs with the causal LM objective can underestimate their knowledge understanding ability.

Recently, QA-based datasets have been developed to evaluate the knowledge recall ability of decoder-only LMs. Kalo and Fichtel (2022) created a high-quality QA prompt set, which is further extended by Wiland et al. (2024) to evaluate both causal and masked LMs. Mallen et al. (2023) and (Maekawa et al., 2024) developed QA datasets to see the impact of knowledge popularity and retrieval augmentation. Since the writing style of these datasets is limited to questions, we cannot perform reliable robustness evaluation.

9 Conclusions

This paper presents multi-faceted factual probing frameworks, BELIEF and BELIEF-ICL, for encoder-based PLMs and decoder-based LLMs, respectively. Leveraging a multi-prompt dataset, BELIEFs provides various evaluation metrics, including accuracy, consistency, and reliability, enabling a thorough assessment of PLMs’ comprehension of factual knowledge. To make BELIEF more reliable, we build a new multi-prompt dataset for knowledge probing, MyriadLAMA, featuring diverse prompts for each fact. We conducted extensive experiments of multiple encoder-based PLMs and recent LLMs. Based on the evaluation results, we verify the effectiveness of BELIEF in accurately and comprehensively probing knowledge by revealing differences between PLMs and demonstrating its ability to mitigate bias. Additionally, we identify key factors affecting the correctness and reliability of PLMs’ answers in factual probing, such as model size, pre-training strategy and corpora. Finally, by probing facts in different ways, we find that PLMs hold more knowledge than what is revealed by using the optimal template, highlighting the limitations of prompt-based factual probing.

10 Limitations

MyriadLAMA contains an extensive amount of prompts, which leads to high evaluation costs. In the future, we aim to extract a diverse yet robust subset from MyriadLAMA to enable a more efficient evaluation of factual knowledge.

References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.

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.

Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhिलाषा Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *Transactions of the Association for Computational Linguistics*, 9:1012–1031.

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).

Constanza Fierro and Anders Søgaard. 2022. [Factual consistency of multilingual pretrained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3046–3052, Dublin, Ireland. Association for Computational Linguistics.

Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.

Jan-Christoph Kalo and Leandra Fichtel. 2022. [Kamel: Knowledge analysis with multitoken entities in language models](#). In *Automated Knowledge Base Construction*.

Ehsan Kamaloo, Nouha Dziri, Charles Clarke, and Davood Rafiei. 2023. [Evaluating open-domain question answering in the era of large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5591–5606, Toronto, Canada. Association for Computational Linguistics.

Nora Kassner and Hinrich Schütze. 2020. [Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.

Amr Keleg and Walid Magdy. 2023. [DLAMA: A framework for curating culturally diverse facts for probing the knowledge of pretrained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6245–6266, Toronto, Canada. Association for Computational Linguistics.

Seiji Maekawa, Hayate Iso, Sairam Gurajada, and Nikita Bhutani. 2024. [Retrieval helps or hurts? a deeper dive into the efficacy of retrieval augmentation to language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5506–5521, Mexico City, Mexico. Association for Computational Linguistics.

Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.

Masanori Oya. 2020. [Syntactic similarity of the sentences in a multi-lingual parallel corpus based on the Euclidean distance of their dependency trees](#). In

paraphrases) for the “subjects” and objects. Specifically, we used the Wikidata identifiers of entities¹⁷ and the Wikidata API¹⁸ to retrieve the (English) alias expressions of entities. By combining the aliases of “subjects” and objects with the relation templates mentioned later, we generate numerous new derived triples. If N “subjects” and M objects are given for an unique triple, the number of derived triples according to this unique triple generated from a single relational template is $N \times M$.

A.2 Diversification of relation templates

We use a two-step procedure to create new relational templates, to enhance ensure both the quality and quantity. Initially, we manually rewrite relational templates, ensuring that every relation has five templates. Then, we employ the generative LLM (GPT4) to automatically paraphrase 19 additional templates. In total, we produce 100 templates for each relation.

Step 1: Manually rewriting relational templates.

The manual rewriting of the relational templates is performed by the first author of this paper. We create new templates by describing the relationship between subject and object from different perspectives rather than creating templates with absolutely the same meaning with original template. Utilizing the resource provided by Wikidata¹⁹, we not only paraphrase existing templates to generate new ones with diverse lexicons but also devise entailment expressions to encompass various semantic expressions that convey the same relations. These newly created templates are guaranteed to uphold relational equivalence, following the relationship between the subject and object. Taking P20 ([X] died in [Y]).²⁰ as an example, we create new templates by either changing the sentence pattern or adding type information of object (e.g, [X] resided in [Y] until death). Furthermore, we also create templates without directly using the keywords of the relation (dead/death) but in a entailment way (e.g., [X] spent the last years of life in [Y].) Moreover, we devise a question-answer style template for each relation to enhance syntactic diversity. In this template, the question incorporates the subject

¹⁷<https://www.wikidata.org/wiki/Wikidata:Identifiers>

¹⁸https://www.wikidata.org/wiki/Special:EntityData/<entity_idenfier>.json

¹⁹https://www.wikidata.org/wiki/Property:<relation_idenfier>

²⁰<https://www.wikidata.org/wiki/Property:P20>

and relation information, while the answer corresponds to the object.

Note that, during the paraphrase, we observe that some templates in LAMA-UHN only partially express the original meaning of relations defined in Wikidata. These are inappropriate for specific knowledge triples. For example, P136 describes the creative work’s genre or an artist’s field of work²¹, which the type of work includes music, film, literature, etc. However, the original templates of P136 in LAMA-UHN is “[X] plays [Y] music.,” which cannot correctly retrieve information on work other than music. For this kinds of template, we abandon the original templates and newly create five templates.

Step 2: Paraphrasing templates using GPT-4

Based on the original relation templates and the relation templates rewritten manually, we further paraphras these relation templates automatically using the GPT4-API (gpt-4-1106-preview²²) provided by OpenAPI. The instruction for paraphrasing used for GPT-4 generation is:

You are a professional tool that can paraphrase sentences into natural sentences that can correctly represent the relationship between [X] and [Y], without repetition. Make the paraphrase as diverse as possible using simple words. Please paraphrase the given sentence 19 times.

When the duplicated sentence is generated, we remove the duplication and regenerate new templates with the same instruction, until 19 different templates is generated. Furthermore, we observe that GPT-4 occasionally generates relation templates that are semantically inappropriate for specific relationships due to incorrect category information of entities. Consequently, in such instances, we refine the instructions to include the category information of the entities, ensuring accurate representation of the relationship between the subjects and the objects. For example, when paraphrasing the relational template “[X] used to work in [Y].”²³, we additionally add explicit guidance regarding the expected format and semantics of the relation templates to the above instruction, as following.

²¹<https://www.wikidata.org/wiki/Property:P136>

²²<https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>

²³<https://www.wikidata.org/wiki/Property:P937>

Be aware that [Y] is the geographic location but NOT company or organization, where persons or organizations were actively participating in employment, business or other work.

As a result, we can obtain the following paraphrased relational templates for “[X] used to work in [Y].”:

- “[X] was formerly employed in [Y].”
- “[X] once worked at [Y].”
- “[Y] was the place where [X] used to be engaged in work.”

A.3 Example of extended relational templates in MyriadLAMA

We display part of the created templates in MyriadLAMA. We randomly select two manually rewritten templates and three auto-generated templates of these two templates. We show the sampled templates for all the relations in Tables 19, 20, 21, 22, 23 and 24.

B The Advangage of MyriadLAMA

Given that our study seeks to mitigate the influence of individual prompt bias in evaluations, the availability of a wide range of prompts characterized in both quantity and diversity is crucial. The diversity ensures that different prompts can capture different aspects of the true knowledge distribution. On the other side, the quality or correctness of prompts ensure the evaluation can accurately reflect the true ability of understanding facts.

In this section, we provide a quantitative analysis of the quality and diversity of multi-prompt factual knowledge probing datasets. The comparison results demonstrate the superiority of MyriadLAMA over previous datasets, enabling more accurate and comprehensive evaluations. We conduct comparison between MyriadLAMA and other multi-prompts probing datasets, LPAQA (Jiang et al., 2020) and PARAREL (Elazar et al., 2021), from the perspective of quantity and diversity.

B.1 Diversity comparison

We measure the diversity of multi-prompt factual knowledge probing datasets from both **quantity** and **linguistic diversity**.

Specifically, we calculate the average number of prompts for each subject-relation pair as the

Dataset	Quantity \uparrow	Diversity \uparrow		
		Lexicon	Syntax	Semantic
PARAREL	7.30	.4860	.1489	11.03
LPAQA	53.27	.5449	.1713	13.55
MyriadLAMA	263.47	.6652	.2138	12.69

Table 8: Comparison between multi-prompts datasets.

quantity measure. MyriadLAMA introduces diversity into prompts by using various subject expressions and relational templates. On average, MyriadLAMA provides 2.47 expressions for each subject. In addition, we measure the **linguistic diversity** of relational templates from three aspects, as shown below:

Lexicon: We utilize the Jaccard distance of words in templates to gauge lexicon diversity.

Syntax: We adopt the syntax distance measure proposed in (Oya, 2020), which calculates the distance between dependency trees.

Semantics: We quantify semantic diversity by calculating the L2 distance of sentence embeddings given by BERT_{large}.

The results are shown in Table 8. MyriadLAMA demonstrates superior quantity and diversity compared to existing multi-prompt factual probing datasets. Although LPAQA exhibits greater semantic diversity, this is mainly due to its use of distant supervision to discover new templates. This method often results in problematic templates that inadequately describe the relationships between subjects and objects. For example, for relation P937²⁴ (“[X] used to work in [Y].”), LPAQA includes templates like “[X] to meet [Y],” which significantly deviate from the original semantic meaning. We analyze and compare template quality in the next section.

B.2 Quality comparison

In this section, we evaluate the quality of relational templates created MyriadLAMA in correctly expressing the relation between the subject and object. We manually evaluate the quality templates created in each dataset through a strict quality evaluation framework. Specifically, we evaluate each template based on its fluency and its ability to correctly express the semantic relationship between subjects

²⁴<https://www.wikidata.org/wiki/Property:P937>

and objects. Given the complex and specific constraints defined by Wikidata relations, creating perfect templates that satisfy all subjects and objects for a given relation is challenging.

B.2.1 Semantic relationship between template and relation

[Template \subseteq Relation]: If the subject and object fit the template, it is correct for the relation, but the relation’s knowledge range is broader than the template can cover. We denote such templates as **[Template \subseteq Relation]**. Using the templates in LAMA-UHN, which are often considered golden templates as examples, relation P136²⁵ uses the template “[X] plays [Y] music.” to describe creative work genres or an artist’s field. However, P136 encompasses film, literature, and other arts, not just music.

[Relation \subseteq Template]: In contrast, If the subject-object pair is true for the relation, it is also true for the template, meaning the template’s knowledge range is broader than the relation. For example, LAMA-UHN creates the template “[X] died in [Y].” for P20²⁶. While this template can be used to infer a person’s place of death, “[Y]” could also be the year “[X]” passed away.

[Relation \cap Template > 0]: Additionally, some templates do not fit neatly into either **[Relation \subseteq Template]** or **[Template \subseteq Relation]** but can still correctly describe the relationship for some subject-object pairs. For example, PARAREL, which paraphrases templates by human effort, uses “[X] is a follower of [Y].” to describe relation P140²⁷ (religion or worldview). This template is appropriate for individuals but not for organizations, and it does not fully capture the relation’s scope. Therefore, it does not fit **[Relation \subseteq Template]**. Additionally, when “[X]” is a person, “[Y]” can be either a religious figure, leader, or the religion itself, which does not satisfy **[Template \subseteq Relation]**. However, many subject-relation pairs can still be correctly captured by this template. We use **[Relation \cap Template > 0]** to denote them.

[Irrelevant]: We consider templates that do not correctly convey the relationship between subject and object as **[Irrelevant]**. For example, LPAQA mined many irrelevant templates from the corpus

²⁵<https://www.wikidata.org/wiki/Property:P1303>

²⁶<https://www.wikidata.org/wiki/Property:P20>

²⁷<https://www.wikidata.org/wiki/Property:P140>

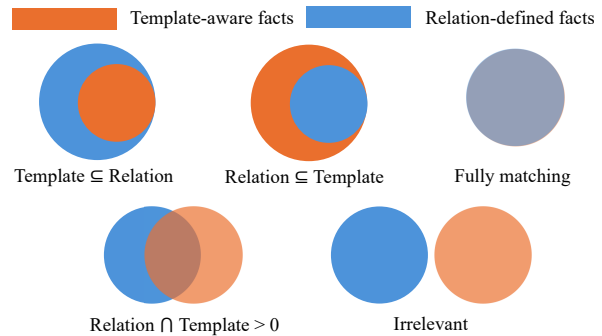


Figure 3: The relationship between the factual knowledge that can be covered by relation and template.

without careful checking, resulting in low-quality templates, such as “[X] arrived in [Y]” for P937.

[Fully matching]: Finally, we use **[Fully matching]** to denote templates that can accurately capture all the subject-object pairs fitting in the relation.

We demonstrate the five types of semantic relationships between template and relation in Figure 3 using the Venn diagram.

B.2.2 Template quality evaluation metrics

To accurately capture the fluency and the ability of created templates to correctly express the relationship between subjects and objects, we use the following metrics to score template quality. Each item is scored as either 1 or 0 based on whether the template meets the requirement.

- 1) **Fluency:** Is the template a natural sentence or noun phrase²⁸? Set 1 if the templates are natural; Otherwise 0.
- 2) **[Relation \subseteq Template]:** If the template can satisfy the definition of **[Template \subseteq Relation]**, then 1; Otherwise 0.
- 3) **[Template \subseteq Relation]:** If the template can satisfy the definition of **[Template \subseteq Relation]**, then 1; Otherwise 0.
- 4) **[Relation \cap Template > 0]:** If the template can satisfy the definition of **[Relation \cap Template > 0]**, then 1; Otherwise 0.

If either **[Relation \subseteq Template]** or **[Template \subseteq Relation]** is 1, then **[Relation \cap Template > 0]** must also be 1. If **[Relation \subseteq Template]**, **[Template \subseteq Relation]**, and **[Relation \cap Template > 0]** are all 0, the template can be classified as **[Irrelevant]**. If all three metrics are 1, the template is classified as **[Fully Matching]**.

²⁸We take the noun phrase into consideration as LPAQA created lots of noun phrase templates, such as “[X], who works as [Y].” for relation P106 (<https://www.wikidata.org/wiki/Property:P106>)

Specifically, as what PLMs see is the prompt with subject, we will consider the existence of subject when scoring the template. For example, P413²⁹ describe the position or specialism of a player on a team. While the template “[X] plays in the position of [Y].” can be too general, as it could also describe a player’s position in an orchestra, specifying “[X]” in the prompt reduces ambiguity, making it an accurate [Fully Matching] template to describe the relation.

B.2.3 Evaluation result and analysis

The comparison includes the 4 dataset: LAMA-UHN, PARAREL, LPAQA and MyriadLAMA. Considering the amount of all the templates in the three models (6654 templates in total), we randomly sample 200 templates for multi-prompt probing datasets and use all 41 templates in LAMA-UHN for evaluation. To ensure objectivity, we anonymize the source of each template and mix them together for annotator (the first author). We publicize the annotation results here³⁰. The evaluation result is shown in Table 9.

From the Table 9, We observe that our semi-automatically generated relational templates achieve quality comparable to manually created datasets like LAMA-UHN and PARAREL, while being 100 times larger than LAMA-UHN and 13.7 times larger than LPAQA. MyriadLAMA significantly outperforms LPAQA in template quality due to our two-stage template creation method.

Furthermore, Figure 4 shows the score distributions for 200 templates across the three multi-prompt datasets. It reveals that LPAQA has many low-score templates, with 0 being the most common score. Compared to PARAREL, MyriadLAMA has slightly more templates with a score of 3 but slightly fewer with a score of 4, resulting in slightly lower overall quality.

C Ablation Analysis of MyriadLAMA

In this section, we conduct ablation analysis MyriadLAMA to understand the validity of diversification on entities and templates.

C.1 Validity of extended entity expressions

We evaluate the validity of the extended entity expressions in MyriadLAMA by checking if these extensions cover facts that PLMs can capture but are

²⁹<https://www.wikidata.org/wiki/Property:P413>

³⁰<https://anonymous.4open.science/r/belief-CC8A>

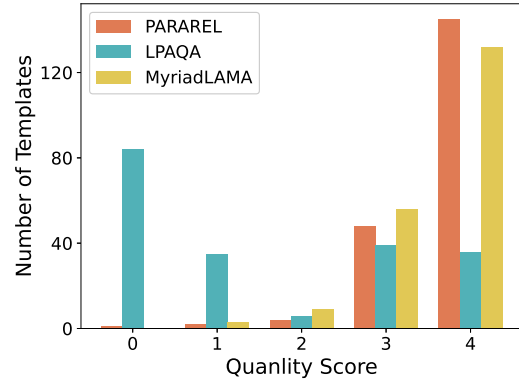


Figure 4: Quality score distributions over multi-prompt factual probing datasets.

missed in LAMA-UHN due to strict entity expression limitations. We conduct this analysis on BERT models, focusing on facts with extended subject and object expressions. MyriadLAMA contains 13,123 facts with extended subjects and 23,195 facts with extended objects. We measure the rate at which extended subjects/objects achieve higher ranks than the original expressions in the token distribution output.

The results, shown in the Table 10 below, indicate that around 50% of extended subjects and 20% of extended objects achieve higher ranks than the original entities. This suggests that many facts are missed in LAMA-UHN and other single-expression factual knowledge probing datasets.

C.2 Validity of paraphrased templates

In this section, we evaluate the validity of the relation templates in MyriadLAMA. We investigate the accuracy of each template and compare the accuracies between templates in LAMA-UHN, manually rewritten templates and auto-generate templates. Specifically, for each relation, we evaluate the accuracy (Acc@1) of all relation templates separately, and then calculate the minimum, and maximum accuracies among all templates for each relation. We then measure the dataset-level minimum/maximum accuracy by micro-averaging the template set with the minimum/maximum template accuracies (41 templates in each set). Finally, all of the template-specific accuracies are then micro-averaged to compute the average Acc@1. As indicated in Table 11, while the quality of MyriadLAMA’s prompts significantly varies, the high-quality prompts are notably superior to those of LAMA-UHN. Although the average accuracy of MyriadLAMA is lower than that of LAMA-UHN, it is considered that this

Dataset	Fluency	Template \subseteq Relation	Relation \subseteq Template	Relation \cap Template > 0	Total Average
LAMA-UHN	1	.732	.976	1	3.707
PARAREL	0.99	.790	.905	.985	3.670
LPAQA	0.57	.220	.345	.405	1.540
MyriadLAMA	1	.770	.830	.985	3.585

Table 9: The average quality metrics across templates from various factual probing datasets. LAMA-UHN has 41 templates, all of which were used for evaluation. For the other three multi-prompt probing datasets, we sampled 200 templates for comparison.

PLMs	Subject	Object
BERT _{base}	.5355	.2107
BERT _{large}	.5358	.2116
BERT _{wwm}	.5272	.1853

Table 10: Rate of prompts where the extend entities gain higher probability than the original entity.

PLMs	LAMA-UHN	MyriadLAMA		
		Min	Max	Mean
BERT _{base}	.2403	.0000	.3534	.1103
BERT _{large}	.2454	.0007	.3728	.1185
BERT _{wwm}	.2448	.0015	.3695	.1453

Table 11: Acc@1 of MyriadLAMA and LAMA-UHN

is because MyriadLAMA uses relation templates that have been semi-automatically created, whereas LAMA-UHN uses carefully selected entities and templates.

C.3 What matters to robustness? Diverse subject vs. templates

Next, we aim to investigate the factors contributing to varying performance and inconsistent predictions of prompts. MyriadLAMA creates diverse prompts for each fact by combining different subject expressions and templates. To gauge their impact on robustness, we examine both the consistency (Consist) and the accuracy range (min/max) across various expressions of subjects or relations, assessed individually. To achieve this, the complete set of prompts was partitioned into multiple subsets, with each subset containing only one expression for each unique subjects or relations. The Acc@1 of the prompts obtained in this manner is then evaluated using different variants of BERT.

The results in Table 12 indicate that while the accuracy range (min/max) and consistency (Consist) caused by aliases of subjects is less pronounced compared to diverse expressions of relational templates, its effect on factual knowledge evaluation remains significant. These findings highlight the

PLMs	Consist \uparrow		Acc@1 range (min/max)	
	Subject	Relation	Subject	Relation
BERT _{base}	.5745	.1504	.0673/.1441	.0000/.3534
BERT _{large}	.5497	.1548	.0714/.1554	.0007/.3728
BERT _{wwm}	.5005	.1057	.0831/.1884	.0015/.3695

Table 12: Diversity evaluation of subjects and relation templates

vulnerability of factual knowledge evaluation based on single prompts and underscore the significance of harnessing the diversity of prompts within MyriadLAMA for robust assessments.

C.4 Manually rewritten vs. auto-generated templates

Upon comparing relational templates generated through manual rewriting and GPT-4 auto-generation, we find that auto-generated templates exhibit comparable quality (accuracy) to manually rewritten templates; they also demonstrate less diversity in acquiring different predictions, aligning with our expectations.

To assess the validity of LLM-generated templates for knowledge probing, we rank the accuracies (Acc@1) of manually created templates against those generated by LLMs. Specifically, for each relation, we rank the 5 manual templates among all 100 templates and calculate the average rank across all manually created templates for all relations. Table 13 shows the average Acc@1 ranks of manual templates among 100 templates on BERT_{base}, BERT_{large}, BERT_{wwm}. They are 47.40, 45.64, and 44.80, respectively. These values closely approximate the average rank of 50, indicating that auto-generated templates can achieve nearly the same accuracy as manually created templates.

Furthermore, we quantify the diversity discrepancy between manually written and auto-generated templates. We categorize the auto-generated templates, including the original ones, as one group, resulting in five groups for each relation, each

PLMs	Average rank of manual prompts based on Acc@1	Consist	
		Inner-group	Inter-group
BERT _{base}	47.40	.2904	.1065
BERT _{large}	45.64	.2884	.1125
BERT _{wwm}	44.80	.2387	.0630

Table 13: Comparison between prompts generated through manual labor and LLM.

comprising 20 templates. Subsequently, we evaluate the similarity between templates within the same group and across different groups using the consistency measure (Consist), as presented in Table 13. The consistency among prompts within the same group (inner-group) is notably high, whereas prompts from different groups (inter-group) exhibit less diversity in predictions. This underscores the significance of manual phrase rewriting, which can yield more diverse prompts and facilitate a more comprehensive evaluation.

D QA-Style ICL and its Evaluation

D.1 QA-style instruction

Beside the mask-prediction-style (MP-style) ICL task, we also defined and evaluate a question-answer-style (QA-style) ICL task utilizing the QA-style relational templates available in MyriadLAMA. This is available because MyriadLAMA provides 20 QA-style templates for each relation, offering not only syntactical diversity but also suitability for the autoregressive generation process in LLMs. Each QA-style prompt adheres to a format where the subject and relation construct the question, and the object corresponds to the answer, such as “Who developed [X]? [Y].” For the QA prompt, we employ the few-shot prompt comprising X random QA pairs, following the format outlined in InstructGPT (Ouyang et al., 2022). Given that all objects in MyriadLAMA are intended to be matched with single words, we append the instruction “Answer each question in one word.” to ensure compatibility.

Given the limited number of templates (20 for each relation) in the QA-style, the evaluation of QA-style prompts represents only one-fifth of the full prompts in MyriadLAMA.

D.2 Evaluation

We measure the ability of QA-style prompts in adhering to instructions and compare it to MP-style prompts. To ensure a fair comparison between QA-

ICL settings	Fact prediction (Acc@1)		One-word generation	
	QA	MP	QA	MP
zero-shot	.4534	.5066	.5285	.4802
4-random	.5429	.5591	.7996	.8058
4-relation	.6582	.6649	.9187	.9246
4-template	.6687	.6765	.9216	.9266

Table 14: Instruction following rate on Llama2-7B.

and MP-style ICL, we conduct evaluations using shared templates in both settings on Llama2-7B, with 20 QA-style templates for each relation.

We evaluate the abilities of fact prediction and one-word generation individually on Llama2-7B using average Acc@1 and rate of the one-word generation. As demonstrated in Table 4, Llama2-7B exhibits a remarkable capability to comprehend instructions for answering questions and generating one-word answers. We observe that QA-style instructions perform better under the zero-shot setting, likely due to decoder-based PLMs’ ability to autoregressively generate text. However, this gap diminishes with the use of few-shot examples. This suggests that while MP-style prompts may slightly underestimate the knowledge in LLMs in zero-shot settings, MP-style ICL settings can demonstrate comparable or even superior performance in factual knowledge prediction compared to QA-style ICL prompts.

E Examples of BELEIF-ICL Prompts

In this section, we provide example prompts for the four patterns introduced in §3: zero-shot, X-random, X-relation, and X-template. We focus on examples where X equals 4, which is the primary setting used in our work.

E.1 zero-shot

Predict the [MASK] in each sentence in one word.

Q: [MASK] consists of LAUPT.

A:

E.2 4-random

Predict the [MASK] in each sentence in one word.

Q: [MASK] is the administrative center of Jiangsu.

A: Nanjing.

Q: Mar del Plata and [MASK] are sister cities that have been developing

1307 together.
 1308 A: Havana.
 1309 Q: Malawi has established diplomatic
 1310 ties with [MASK].
 1311 A: Australia.
 1312 Q: Which country is House of
 1313 Representatives located? [MASK].
 1314 A: Libya.
 1315 Q: [MASK] consists of LAUPT.
 1316 A:

1317 E.3 4-relation

1318 Predict the [MASK] in each sentence in
 1319 one word.
 1320 Q: What is the overarching group for
 1321 Panzer Division Kempf? [MASK].
 1322 A: Wehrmacht.
 1323 Q: To whom does Mount Bulusan relate?
 1324 [MASK].
 1325 A: Luzon.
 1326 Q: Who is responsible for Army National
 1327 Guard? [MASK].
 1328 A: National Guard.
 1329 Q: What group is pharmacy a part of?
 1330 [MASK].
 1331 A: biology.
 1332 Q: [MASK] consists of environmental
 1333 factors.
 1334 A:

1335 E.4 4-template

1336 Predict the [MASK] in each sentence in
 1337 one word.
 1338 Q: [MASK] consists of Panzer Division
 1339 Kempf.
 1340 A: Wehrmacht.
 1341 Q: [MASK] consists of Mount Bulusan.
 1342 A: Luzon.
 1343 Q: [MASK] consists of Army National Guard.
 1344 A: National Guard.
 1345 Q: [MASK] consists of pharmacy.
 1346 A: biology.
 1347 Q: [MASK] consists of environmental
 1348 factors.
 1349 A:

1350 F Experimental details

1351 In this section, we list the detailed information of
 1352 PLMs used in our study, including 3 encoder-based
 1353 models and 8 decoder-based LLMs.

F.1 Model cards

Here are the links from Hugging Face to load each model:

BERT_{base}: <https://huggingface.co/bert-base-uncased>

BERT_{large}: <https://huggingface.co/bert-large-uncased>

BERT_{wwm}: <https://huggingface.co/bert-large-uncased-whole-word-masking>

ALBERT_{base}: <https://huggingface.co/albert/albert-base-v2>

ALBERT_{large}: <https://huggingface.co/tftransformers/albert-xxlarge-v2>

Llama2-7B: <https://huggingface.co/meta-llama/Llama-2-7B>

Llama2-13B: <https://huggingface.co/meta-llama/Llama-2-13B>

Llama2-70B: <https://huggingface.co/meta-llama/Llama-2-70B>

Llama3-8B: <https://huggingface.co/meta-llama/Meta-Llama-3-8B>

Llama3-70B: <https://huggingface.co/meta-llama/Meta-Llama-3-70B>

Phi3-mini: <https://huggingface.co/microsoft/Phi-3-mini-4k-instruct>

Phi3-small: <https://huggingface.co/microsoft/Phi-3-small-8k-instruct>

Phi3-medium: <https://huggingface.co/microsoft/Phi-3-medium-4k-instruct>

F.2 Model differences

We outline the differences between PLMs in their pre-training details in Table 15, including the type of Transformer architecture, model size, and the size and resources of the pre-training corpora.

F.3 Evaluation results on all PLMs based on BELIEFs

We present all evaluation results and their computational costs in this section. In Table 16, we report the full-scale experiments using all the prompts provided by MyriadLAMA. This includes PLMs with 8B parameters or fewer, such as

LLMs	Architecture	Model size	Pre-training corpora	
			Size	Resource
BERT _{base}	Encoder-based	110M	3.3G	} Two datasets: BookCorpus, a dataset consisting of 11,038 unpublished books and English Wikipedia (excluding lists, tables and headers).
BERT _{large}	Encoder-based	336M	3.3G	
BERT _{wwm}	Encoder-based	336M	3.3G	
ALBERT _{base}	Encoder-based*	11.8M	3.3G	
ALBERT _{large}	Encoder-based*	223M	3.3G	
Llama2-7B	Decoder-based	7B	2.0T	} A collection of publicly available online data (excluding sites containing a lot of personal information. The knowledge sources for factual knowledge are upsampled).
Llama2-13B	Decoder-based	13B	2.0T	
Llama2-70B	Decoder-based	70B	2.0T	
Llama3-8B	Decoder-based	8B	15T+	} A collection of publicly available online data (details unknown, the code is 4 times that of Llama2).
Llama3-70B	Decoder-based	70B	15T+	
Phi3-mini	Decoder-based	3.8B	3.3T	} High-quality materials including educational data and code, textbook-like generated text, high-quality chats.
Phi3-small	Decoder-based	7B	4.8T	
Phi3-medium	Decoder-based	14B	4.8T	

*ALBERT use parameter compression techniques, including parameter sharing between token embeddings and transformer layers.

Table 15: The pre-training information of various PLMs used in this study

BERT_{base}, BERT_{large}, BERT_{wwm}, ALBERT_{base}, ALBERT_{large}, Llama2-7B, Llama3-8B, Phi3-mini (3.8B), and Phi3-small (7B). For decoder-based models, we conduct experiments on 4 types of ICL settings.

For PLMs with more than 8B parameters, we report the evaluation results using partial prompts from MyriadLAMA, specifically using manually-rewritten templates (5 per relation), which account for 1/20 of the prompts compared to the full-scale experiments. Meanwhile, we run these models on only two ICL settings: zero-shot and 4-template. To ensure fair comparison with models having 8B parameters or fewer, we apply the same settings to all other decoder-based LLMs. The result is shown in Table 17. We also list the approximate runtime for each experiments in these two tables. The experiments are all run on the NVIDIA RTX 6000 Ada. For experiments using model less than or equal to 8B parameters, we use single GPU to measure the consumption time. We use 2 GPUs for Llama2-13B and Phi3-small and 4 GPUs for 70B models.

Furthermore, we display the calibration level between accuracy and confidence for a straightforward inspection on Overconf metrics. We show the calibration figures of models with full-scale experiments in Figure 5 and experiments with partial prompts in Figure 6.

F.4 Knowledge coverage rate on all PLMs

We present the average, maximum, and upper limit knowledge coverage rates, as introduced in §7, for all PLMs evaluated using all templates. The results are shown in Figure 18.

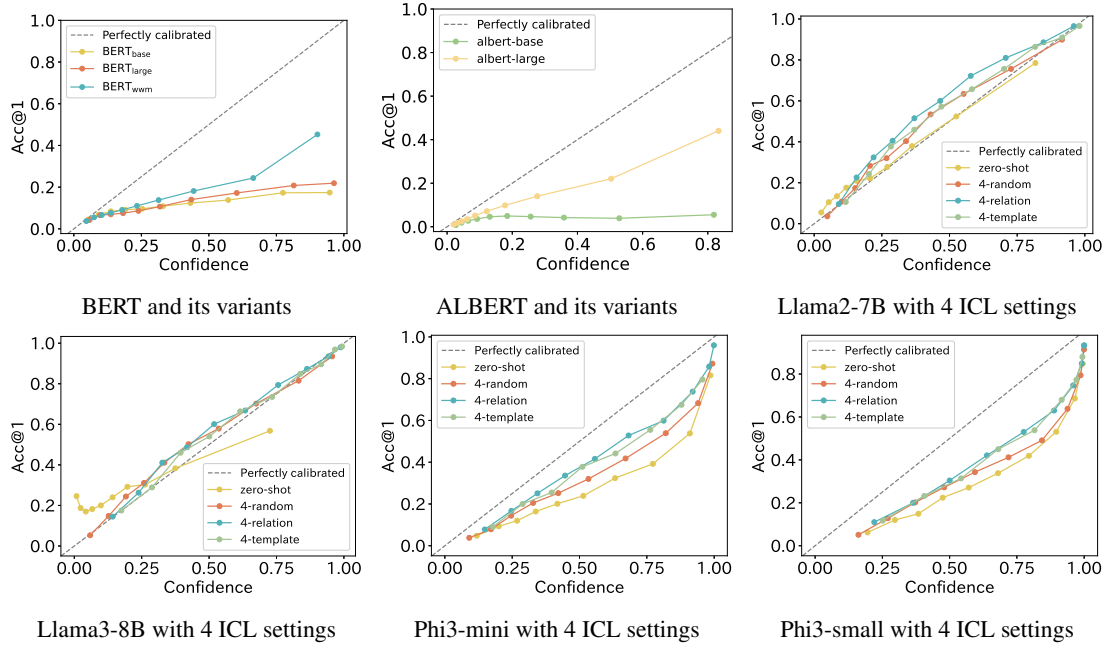


Figure 5: The calibration between accuracy and confidence of PLMs' prediction with full-scale experiments.

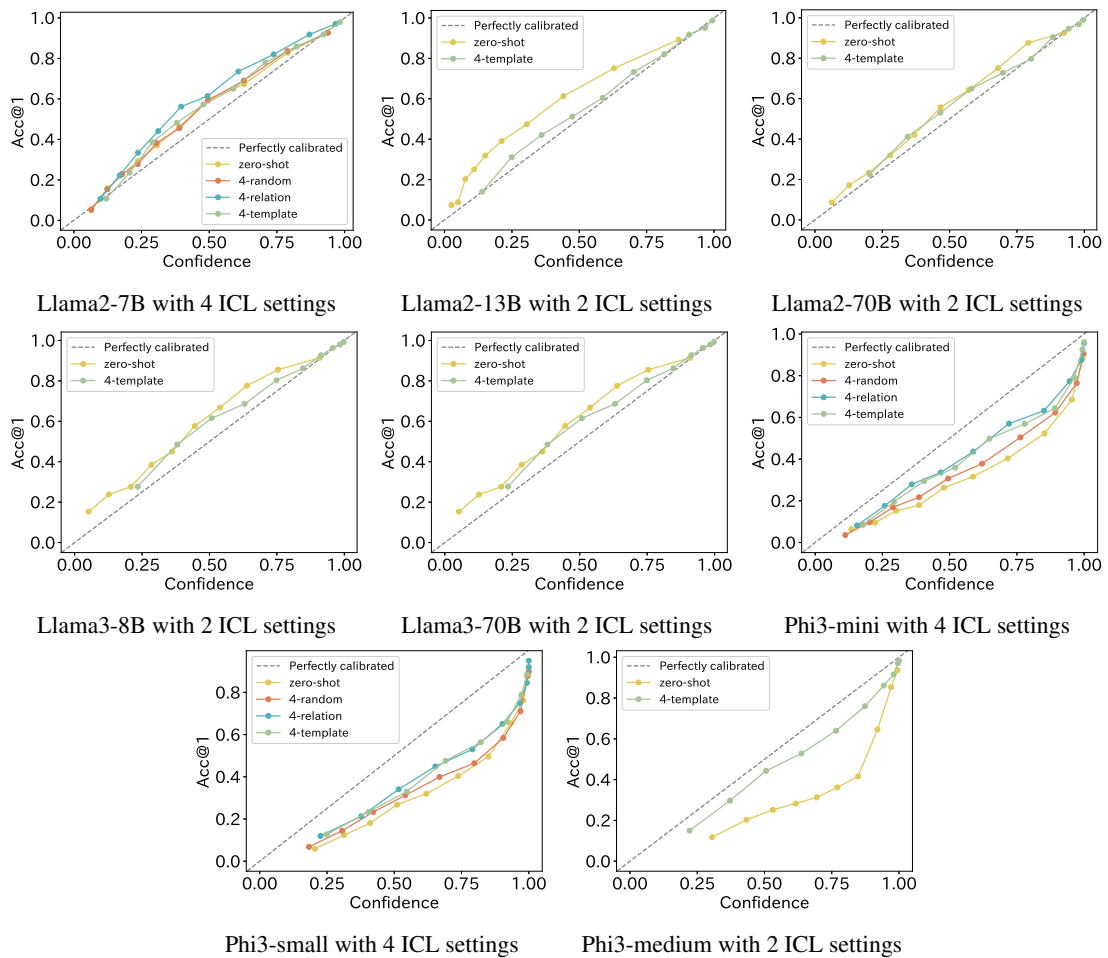


Figure 6: The calibration between accuracy and confidence of PLMs' prediction with partial-scale experiments (using manually rewritten templates only).

PLMs		Average Acc@1 [↑]	Fluctuation [↓]		Consist [↑]	Overconf	one-word Generation Rate [↑]	Runtime
			Range	Stdev				
BERT	BERT _{base}	.1095	.1534	.0217	.1682	.2154	N/A	6.3h
	BERT _{large}	.1102	.1574	.0220	.1713	.2052	N/A	} 7.4h
	BERT _{wwm}	.1364	.1517	.0208	.1524	.1000	N/A	
ALBERT	ALBERT _{base}	.0362	.0668	.0131	.1333	.1647	N/A	6.1h
	ALBERT _{large}	.0974	.1110	.0148	.0821	.0553	N/A	15.2h
Llama2-7B	zero-shot	.3385	.2602	.0299	.1269	-.1119	.4752	46.4h
	4-random	.4816	.2250	.0270	.2312	-.0894	.8247	} 47.8h
	4-relation	.6286	.1221	.0150	.3753	-.1335	.9060	
	4-template	.6616	.0294	.0036	.4163	-.0933	.9299	
Llama3-8B	zero-shot	.3427	.2864	.0350	.0240	-.1329	.1572	44.9h
	4-random	.5205	.2033	.0273	.2156	-.0796	.8147	} 82.1h
	4-relation	.6871	.1236	.0156	.3659	-.0783	.9071	
	4-template	.7268	.0220	.0026	.4015	-.0582	.9187	
Phi3-mini (3.8B)	zero-shot	.3498	.2374	.0292	.1465	.1752	.8641	30.7h
	4-random	.4193	.2324	.0269	.1649	.1189	.8184	} 32.9h
	4-relation	.5686	.1440	.0164	.2818	.0755	.8769	
	4-template	.6067	.0510	.0048	.3612	.0887	.8808	
Phi3-small (7B)	zero-shot	.4258	.2437	.0292	.1782	.2171	.8883	82.4h
	4-random	.4889	.2170	.0276	.2070	.1670	.8913	} 148h
	4-relation	.6339	.1012	.0129	.3361	.1252	.9287	
	4-template	.6612	.0360	.0043	.3626	.1279	.9411	

Table 16: Evaluation results and approximate experiment runtime on all MyriadLAMA data. Comparison of LLMs with 8B parameters or less, including BERT and its variants, across four ICL settings.

PLMs		Average Acc@1 [↑]	Fluctuation [↓]		Consist [↑]	Overconf	one-word Generation Rate [↑]	Runtime
			Range	Stdev				
zero-shot	Phi3-mini (3.8B)	.4248	.1880	.0247	.2066	.1609	.8596	1.54h
	Phi3-small (7B)	.4881	.1900	.0244	.2284	.1985	.8996	4.12h
	Llama2-7B	.4311	.2014	.0249	.1932	-.0922	.5558	2.32h
	Llama3-8B	.4224	.2820	.0353	.1269	-.1438	.1786	2.45h
	Llama2-13B	.4785	.2131	.0260	.1437	-.1673	.3185	4.84h
	Phi3-medium (14B)	.5173	.2123	.0277	.6167	.2316	.7759	4.85h
	Llama2-70B	.5675	.2126	.0280	.0359	-.0988	.6239	28.97h
	Llama3-70B	.5974	.2137	.0278	.2290	-.1438	.7790	32.55h
4-template	Phi3-mini (3.8B)	.6106	.0314	.0039	.3686	.0911	.9051	1.65h
	Phi3-small (7B)	.6668	.0306	.0039	.3666	.1222	.9413	7.40h
	Llama2-7B	.6699	.0257	.0034	.4174	-.0933	.9299	2.39h
	Llama3-8B	.7316	.0194	.0025	.4060	-.1119	.9190	4.10h
	Llama2-13B	.7080	.0235	.0031	.4326	-.0662	.9190	4.23h
	Phi3-medium (14B)	.7304	.0207	.0025	.4009	.0317	.9350	3.88h
	Llama2-70B	.7784	.0190	.0024	.4448	-.0690	.9256	21.99h
	Llama3-70B	.8211	.0139	.0017	.4636	-.0812	.9378	43.10h

Table 17: Evaluation results and approximate experiment runtime on partial MyriadLAMA data. Comparison of all LLMs (8 models) using zero-shot and 4-template ICL settings with manually rewritten templates.

PLMs		Average	Maximum	Upper Limit
BERT	BERT _{base}	.1095	.4248	.6209
	BERT _{large}	.1102	.4451	.6556
	BERT _{wwm}	.1364	.4501	.6636
ALBERT	ALBERT _{base}	.0362	.2175	.3405
	ALBERT _{large}	.0974	.3746	.5979
Llama2-7B	zero-shot	.3385	.6577	.8153
	4-random	.4816	.7026	.8587
	4-relation	.6286	.7179	.8475
	4-template	.6616	.7197	.8133
Llama3-8B	zero-shot	.3427	.7099	.8756
	4-random	.5205	.7339	.8867
	4-relation	.6871	.7733	.8934
	4-template	.7268	.7731	.8628
Phi3-mini	zero-shot	.3498	.6346	.8381
	4-random	.4193	.6506	.8423
	4-relation	.5686	.6791	.8436
	4-template	.6067	.6754	.8114
Phi3-small	zero-shot	.4258	.6828	.8642
	4-random	.4889	.7037	.8695
	4-relation	.6339	.7172	.8507
	4-template	.6612	.7181	.8346

Table 18: Knowledge coverage rate on models with full-scale experiments. **Average**: mean accuracy using all templates; **Maximum**: accuracy with the best template for each relation; **Upper Limit**: accuracy when the best template is selected for each knowledge relation.

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P19	[X] started their life in [Y].	[X] took their first steps of life in [Y]. [X] activated their life's beginning in [Y]. [X] initiated their journey of life within [Y].
	The birth of [X] occurred in [Y].	The origin of [X] took place in [Y]. The inception of [X] was within [Y]. It was in [Y] that [X] first made its appearance.
P20	[X] spent the last years of life in [Y].	In [Y], [X] spent the end of their life. [X]'s final era was in [Y]. In [Y], [X]'s life came to a close.
	[Y] is the last place where [X] lived until death.	[X] inhabited [Y] up until death. [Y] was the end-of-life dwelling for [X]. [Y] served as the last dwelling for [X] before they died.
P279	Of which class is [X] a subclass? [Y].	What is the general class that [X] is a part of as a subclass? [Y]. What larger class encompasses [X] as a subclass? [Y]. Into which class is [X] categorized as a subclass? [Y].
	[X] is also necessarily a [Y].	[X] is intrinsically a [Y]. [X] is fundamentally a [Y]. [X] is by definition a [Y].
P37	[Y] is spoken as an official language by people in [X].	[Y] is the authorized language for formal use in [X]. The official language spoken by individuals in [X] is [Y]. [X] endorses [Y] as the language for state-related communication.
	Officially, the people living in [X] use the language [Y] for communication.	In [X], the standard language for dialogue among the populace is [Y]. Residents of [X] typically converse in [Y]. The official medium of verbal exchange in [X] is the [Y] language.
P413	[X] was given the [Y] job.	[X] was selected for the [Y] position. [X] was named the new [Y]. The [Y] duties have been allocated to [X].
	[X] is a famous [Y] player.	[X] has risen to fame with their [Y] playing abilities. [X] is well-known for playing [Y]. [X] is notable for their expertise in [Y].
P449	[X] premiered on the network [Y].	[Y] was the origin of the broadcast for [X]. [X] was initially broadcasted by [Y]. The debut broadcast of [X] was on [Y].
	[Y] is the first air channel of [X].	[X] was originally brought to the public by [Y]. [X] first hit the airwaves courtesy of [Y]. [X] first reached listeners and viewers via [Y].
P47	[X] and [Y] are neighboring countries.	[X] and [Y] are countries that are in close proximity. [Y] lies in the vicinity of [X]. [Y] and [X] are countries that share a boundary.
	You can go through [X] to reach [Y].	[X] acts as a gateway to [Y]. To reach [Y], one can travel through [X]. Traveling over [X] can bring you to [Y].

Table 19: Examples of templates in MyriadLAMA - part 1

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P138	Who or what is [X] named after? [Y].	Who is the namesake behind [X]? [Y]. What is the etymology behind [X]'s name? [Y]. Who or what was [X] called after? [Y].
	[X] is named after [Y].	[X] takes its name from [Y]. [Y] is the inspiration behind the name of [X]. [X] holds the name given in tribute to [Y].
P364	[X] is created in language [Y].	[X] was composed in the [Y] language. [X] unfolds in the language known as [Y]. [X] is expressed through the language of [Y].
	[X] was written in the [Y] language.	The words of [X] are in the [Y] language. The composition of [X] is in the [Y] language. [X] was created using the [Y] language.
P463	[X] served for [Y].	[X] took part in [Y]. [X] collaborated with [Y]. [X] held a position at [Y].
	Which group or organization does [X] belong to? [Y].	[X] is part of what organization? [Y]. [X] is a member of which entity? [Y]. Can you tell me which entity [X] is a member of? [Y].
P101	Which field does [X] work in? [Y].	In what industry is [X] employed? [Y]. [X] holds a position in which field? [Y]. [X] is a professional in what sector? [Y].
	[X] is influential in the domain of [Y].	The domain of [Y] feels the considerable impact of [X]. [X] plays a pivotal role in the sphere of [Y]. [X] has a profound effect on [Y].
P106	[X] is famous for achievements as a [Y].	[X] is well-known for their accomplishments in the [Y] role. [X] is well-known for their successful career as a [Y]. [X] is a celebrated [Y] with a long list of achievements.
	[X] is a [Y] by profession.	[X] has built a career as a [Y]. [X] is employed as a [Y]. [X] carries out the role of a [Y].
P527	[Y] is a member of [X].	[X] contains [Y] as part of its composition. [Y] holds a place in [X]. [Y] is a piece of [X].
	[Y] belongs to [X].	[Y] is held by [X]. [X] has [Y] under its ownership. [Y] is one of the items owned by [X].
P530	[Y] is one of the countries [X] has diplomatic relations with.	[Y] is a member of the group of countries with which [X] conducts diplomacy. [X] has a formal diplomatic relationship with [Y], as it does with several other countries. [Y] is recognized by [X] as a diplomatic partner among other nations.
	[X] has established diplomatic ties with [Y].	[X] has initiated formal diplomatic relations with [Y]. [X] and [Y] have begun a diplomatic relationship. [X] and [Y] have set up official diplomatic links.

Table 20: Examples of templates in MyriadLAMA - part 2

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P176	[X] is a product of [Y]’s manufacturing.	The entity [Y] crafts and produces [X]. The item [X] is fabricated by [Y]. [X] is brought to life by [Y]’s manufacturing capabilities.
	Which company produced [X]? [Y].	Can you tell me who made [X]? [Y]. Which producer can be linked to [X]? [Y]. What is the producing company of [X]? [Y].
P27	[X] is a person from country [Y].	[X] is a resident of [Y]. [X] bears the nationality of [Y]. [X] is a product of [Y].
	The nationality of [X] is [Y].	[X] is a native of [Y]. [X] is identified as a national from [Y]. [Y] is the country of origin for [X].
P407	[X] is in language [Y].	The language of [X] is [Y]. The primary linguistic expression of [X] is in [Y]. [X] is articulated through the [Y] language.
	[X] is a work in the [Y] language.	The [Y] language is the linguistic fabric of [X]. [X] has been produced using the [Y] language. [X] is an example of literature in the [Y] language.
P30	On what continent is [X] located? [Y].	What’s the name of the continent that [X] calls home? [Y]. What continental landmass does [X] occupy? [Y]. [X] lies on which of the Earth’s continents? [Y].
	[X] is a part of the continent [Y].	[X] is a section of the continental land of [Y]. [X] is geographically positioned as part of continent [Y]. [X] is an integral piece of the continent [Y].
P178	[X] was originally created by [Y].	The foundation of [X] was laid by [Y]. The concept of [X] was conceived by [Y]. [X] first came into existence thanks to [Y].
	[X] is developed by [Y].	[Y] has developed [X]. [Y] is the developer behind [X]. [Y] stands as the creator of [X].
P1376	[X] is the capital of [Y].	[Y]’s governmental seat is in [X]. [X] is recognized as the official capital of [Y]. The leading city and capital of [Y] is [X].
	[X] is the administrative center of [Y].	[Y]’s administrative leadership is situated in [X]. [Y]’s administrative affairs are managed from [X]. [X] is where [Y]’s administrative management is anchored.
P131	[Y] is the place where [X] is located.	[X] resides in [Y]. [X] can be found at the location of [Y]. [X] is anchored in [Y].
	[X] is located in [Y].	[Y] is where [X] is established. [Y] contains [X]. [Y] houses [X].

Table 21: Examples of templates in MyriadLAMA - part 3

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P1412	What language does [X] use? [Y].	[X] communicates in what vernacular? [Y]. What tongue does [X] utilize? [Y]. What is the primary language for [X]? [Y].
	[Y] is the language that is used by [X].	The tongue of [X] is the language [Y]. [X] uses [Y] as its mode of speech. [Y] is the language that enables communication for [X].
P108	[X] is employed by [Y].	[Y] is the employer of [X]. [X] has a job at [Y]. [Y] is the source of employment for [X].
	Who does [X] work for? [Y].	Who does [X] report to in their job? [Y]. For whom is [X] currently working? [Y]. Who holds [X] on their team? [Y].
P136	What is the genre of [X]? [Y].	In terms of genre, how would you classify [X]? [Y]. What category of genre does [X] belong to? [Y]. In what genre category would you place [X]? [Y].
	[X] is the representative of the [Y] style.	[X] personifies the [Y] style in its purest form. [X] is the epitome of the [Y] approach. [X] is the archetype of the [Y] tradition.
P17	Which country is [X] located? [Y].	Can you identify the country where [X] is situated? [Y]. Could you specify the country of [X]'s location? [Y]. [X] can be located in what country? [Y].
	[Y] is the country in which [X] is located.	[Y] is the nation that houses [X]. [Y] encompasses the region where [X] can be found. [Y] is the setting for the location of [X].
P39	What position does [X] hold? [Y].	What position does [X] occupy? [Y]. What is the employment status of [X]? [Y]. What is the position title for [X]? [Y].
	[X] was sworn in as [Y].	[X] has been designated the official role of [Y]. [X] pledged their commitment to the role of [Y]. [X] was confirmed in the role of [Y].
P264	Which music label represents [X]? [Y].	Which label has [X] on its roster? [Y]. Who is [X]'s music label? [Y]. With whom is [X] signed for music production? [Y].
	[X] is represented by music label [Y].	The music label acting on behalf of [X] is [Y]. [Y] is the music label that has signed [X]. [X] has music label [Y] as its representative.
P276	Where is [X] located? [Y].	What's the location of [X]? [Y]. Where can [X] be found? [Y]. Where should I look for [X]? [Y].
	[X] is located in [Y].	[X] is positioned in [Y]. [X] occupies a space in [Y]. [Y] contains [X].

Table 22: Examples of templates in MyriadLAMA - part 4

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P937	[Y] is the place where [X] worked.	[X] had their employment based in [Y]. [X] found their employment setting in [Y]. [X] conducted their professional activities in [Y].
	[X] had work activity in [Y].	[X] took part in business tasks in [Y]. [X] was employed within the confines of [Y]. [X] was operational in the workforce at [Y].
P140	Which religion is [X] affiliated with? [Y].	What religious belief does [X] adhere to? [Y]. Which spiritual path is embraced by [X]? [Y]. What is the creed of [X]? [Y].
	[X] is affiliated with the [Y] religion.	[X] is part of the [Y] religious denomination. [X] is associated with the [Y] spiritual tradition. [X] adheres to the [Y] religion.
P1303	[X] is a [Y] player.	[X] specializes in the [Y]. [X] is a seasoned [Y] player. [X] is a [Y] specialist.
	[X] plays [Y].	[X] expresses their musicianship through [Y]. [X] has chosen [Y] as their musical companion. [X] is a musician who specializes in [Y].
P127	Who owns [X]? [Y].	Whose property is [X] considered to be? [Y]. Who is the legal holder of [X]? [Y]. Who has the ownership rights to [X]? [Y].
	[X] is owned by [Y].	[Y] is the proprietor of [X]. [Y] holds the title to [X]. [Y] possesses [X].
P103	[X] grew up speaking [Y] as their first language.	[X]'s formative years were shaped by speaking [Y]. [X] started their life speaking [Y]. [X]'s childhood language was [Y].
	[Y] is the mother tongue of [X].	[X] has [Y] as their original tongue. [X] was nurtured in an environment where [Y] is spoken. [X] has [Y] as the language of their upbringing.
P190	The city of [X] is twinned with [Y].	[Y] and [X] have entered into a twinning arrangement. [X] is in a twinning relationship with [Y]. A twinning link has been established between [X] and [Y].
	[X] and [Y] are sister cities that have been developing together.	[X] and [Y] have been sister cities on a shared developmental journey. The cities of [X] and [Y] have jointly progressed as sister municipalities. [X] and [Y] have been in lockstep as sister cities in their development.
P1001	[X] applies to the jurisdiction in [Y].	The jurisdiction of [Y] encompasses [X]. [X] is answerable to the legal system in [Y]. [Y] exercises legal control over [X].
	The region of [Y] uses [X] as a legal term.	[X] is a term with legal standing in [Y]. The legal system of [Y] includes [X] as an official term. [X] is employed as a juridical term in [Y].

Table 23: Examples of templates in MyriadLAMA - part 5

Relation IDs	Manual rewriting templates	GPT-4 paraphrased templates
P31	[X] is a [Y].	[X] represents a [Y]. [X] is an example of a [Y]. [X] is termed a [Y].
	Speaking of [Y], [X] is an example of it.	[X] is a particular instance that reflects [Y]. [X] is a variant that falls within the scope of [Y]. [Y] can be demonstrated through [X].
P495	[X] originates from the country of [Y].	[X] was first found in the land of [Y]. The inception of [X] is linked to the country [Y]. The origin of [X] can be traced back to [Y].
	[X] first appeared in [Y].	[X] has its roots in [Y]. [X] was first crafted in [Y]. The origin of [X] is attributed to [Y].
P159	The operation of [X] depends on the headquarters in [Y].	[X]’s functioning is reliant on the main office in [Y]. The base in [Y] is essential for [X] to function. The primary operations of [X] are contingent upon the base in [Y].
	The headquarters of [X] is in [Y].	[Y] is home to the central office of [X]. The top office of [X] is positioned in [Y]. The nerve center for [X]’s operations is based in [Y].
P36	[Y] is the administrative center of [X].	The nerve center for [X]’s administration is found in [Y]. [X]’s administrative governance is centralized in [Y]. [X] is administratively governed by [Y].
	[Y] represents the capital city for [X].	[Y] functions as [X]’s political hub. [X] uses [Y] as its head city. [X]’s administrative center is [Y].
P740	[X] started their career in [Y].	[Y] served as the starting point for [X]’s career. [X] began earning their stripes in the field of [Y]. [X] commenced their employment journey with [Y].
	The formation location of [X] is [Y].	The assembly point for [X] is [Y]. [Y] is recognized as the setting for [X]’s formation. [Y] is where [X] originates.
P361	Which entity does [X] belong to? [Y].	Who owns [X]? [Y]. What is the overarching group for [X]? [Y]. What organization encompasses [X]? [Y].
	[Y] consists of [X].	[X] is what [Y] is primarily made of. [Y] incorporates [X] within it. [Y] is structured with [X].

Table 24: Examples of templates in MyriadLAMA - part 6