

# Efficient Knowledge Probing of Large Language Models by Adapting Pre-trained Embeddings

Kartik Sharma<sup>1</sup>, Yiqiao Jin<sup>1</sup>, Rakshit Trivedi<sup>2</sup>, Srijan Kumar<sup>1</sup>

<sup>1</sup>Georgia Institute of Technology, <sup>2</sup>Massachusetts Institute of Technology

<sup>1</sup>{ksartik, yjin328, srijan}@gatech.edu

<sup>2</sup>triver@mit.edu

## Abstract

Large language models (LLMs) acquire knowledge across diverse domains such as science, history, and geography encountered during generative pre-training. However, due to their stochasticity, it is difficult to predict what LLMs have acquired. Prior work has developed different ways to probe this knowledge by investigating the hidden representations, crafting specific task prompts, curating representative samples, and estimating their uncertainty. However, these methods require making forward passes through the underlying model to probe the LLM’s knowledge about a specific fact, making them computationally expensive and time-consuming. To bridge this gap, we propose **PEEK** or **Proxy Embeddings to Estimate Knowledge** of LLMs, by leveraging the pre-trained embedding models that effectively encode factual knowledge as text or graphs as proxies for LLMs. First, we identify a training set of facts known by LLMs through various probing strategies and then adapt embedding models to predict the LLM outputs with a linear decoder layer. Comprehensive evaluation on 3 Wikipedia-derived datasets, 4 LLMs, and 7 embedding models shows that embeddings can predict LLM knowledge on a held-out set with up to 90% accuracy. Furthermore, we find that sentence embedding models are more suitable than graph embeddings to predict LLM knowledge, shedding light on the underlying representation of factual landscape. Thus, we believe that knowledge-adapted embeddings can be used to identify knowledge gaps in LLMs at scale and can provide deeper insights into LLMs’ internal inductive bias. The code and data is made available at <https://github.com/claws-lab/peek>.

## 1 Introduction

The internet serves as a modern-day encyclopedia and ever-expanding repository of human knowledge spanning history, science, technology, and arts [1, 2]. Large language models (LLMs) trained on such web-scale data have thus become general-purpose knowledge bases [3], but it is still unclear to map their learned knowledge landscape due to stochastic learning objectives and noisy training data. As LLMs are being adopted for knowledge-intensive applications such as healthcare [4], law [5], scientific discovery [6, 7], etc., it is essential to efficiently probe their acquired knowledge since hallucinating incorrect knowledge [8] or lack of knowledge can diminish their utility. Thus, factual gaps are often identified and fixed before deployment through extensive risk assessment manually [9, 10].

Different techniques have been proposed in the literature to probe the knowledge of LLMs. This involves crafting specific knowledge probing prompts on a subsampled set of facts [3, 11–15], investigating their hidden representations for non-factuality [16–18], and estimating their epistemic uncertainty [19–21]. However, none of these methods can uncover the knowledge landscape of LLMs without making forward passes through the LLMs or peeking into their activations. This limits their applicability to extremely large and black-box language models.

Concurrently, representation learning models have also become more powerful in learning general-purpose low-dimensional representations of real-world data such as multilingual texts [22], and knowledge graphs [23]. These encoder models are optimized via self-supervised learning on curated subsets of web data, often overlapping with the training corpora of decoder-only LLMs. More recently, pre-trained




Figure 1: Comparison of our proposed approach, Proxy Embeddings to Estimate Knowledge (PEEK) with other knowledge probing approaches.

LLMs have also been directly used to initialize these representations for various domains [24, 25]. Due to the shared pre-training data and architectural similarities, these embeddings offer an efficient and accurate means of estimating the knowledge acquired by black-box LLMs.

In this work, we thus propose **PEEK**, where we use pre-trained embedding models to estimate the knowledge of an LLM for a set of training facts and then, using these embeddings as proxies to probe the knowledge of LLMs over a database. Figure 1 illustrates our contributions in comparison to the literature. In particular,

1. We propose a **novel approach** of probing LLM knowledge through proxy embedding models that are adapted using a trainable linear head.
2. We uncover LLMs’ factual knowledge through **multiple probing strategies**, including yes/no queries and free-form generative questions.
3. We estimate knowledge **encoded** in different LLMs’ **outputs**, such as truth generation, fact generation, last layer activation, and logits.
4. **Comprehensive** experiments on 4 general-purpose LLMs, 3 general knowledge datasets, and 7 embedding models show that embeddings can effectively predict whether a fact is known by an LLM without querying it.

## 2 Related Work

**Factual knowledge benchmarks.** External resources, particularly, knowledge graphs, have been used to probe the factual knowledge learned by neural language models [26], by formulating the problem as a “fill-in-the-blank” cloze task for masked language models [3, 15, 12] and soon extended to large language models using cloze tasks, question-answering tasks, and open-ended generation tasks [14, 13, 12, 27, 28]. These benchmarks involve carefully selecting test facts from the vast sources using different techniques such as the popularity of the involved entities [13] and Thomson sampling [11]. Thus, they only probe a subset of the factual knowledge, while PEEK enables us to uncover the factual knowledge without querying by learning proxy models.




Figure 2: **Proxy Embeddings to Estimate Knowledge (PEEK)**: In this framework, pre-trained embedding models are adapted to match the LLM knowledge for a training set of facts identified using different probing mechanisms. On a held-out set, we can then predict whether an LLM knows a fact or not by using the fact’s embedding.

**Non-factuality prediction and detection.** LLMs have been shown to contain multiple factual gaps, which leads them to confabulate/hallucinate incorrect information in an unpredictable manner [8]. To understand this phenomenon, existing works have investigated their intermediate representations and found that they can be used to predict the truthfulness of a fact [16, 29, 30], how knowledgeable it is about an entity [17], and to predict whether its output will be non-factual or not [18]. In addition, other works have also explored fine-tuning additional models for the specific task of detecting hallucinations [31, 32]. Hallucinations have also been associated with the model’s epistemic uncertainty about a given fact [19]. Existing works thus aim to estimate the uncertainty by checking inconsistencies in repeated samples [20], teaching models to generate uncertainty [33], using a bigger model’s predictions as aleatoric uncertainty [19], and leveraging the logits for in-context examples [21]. Unlike these works, PEEK does not involve costly forward passes, a white-box access to the model, or extensive training of bulky modules. By adapting embedding models with a linear head, we can probe model’s knowledge with respect to large databases without querying the bulky model.

An extended discussion is provided in Appendix A.

### 3 Proxy Embeddings to Estimate Knowledge (PEEK)

As shown in Figure 2, we propose **PEEK** or proxy embeddings to estimate knowledge. More formally,

**Proposition 1 (PEEK).** *We are given a large language model  $\mathcal{M}$  and a knowledge base (a set of facts)  $\mathcal{K}$ . Let  $\mathcal{P}_{\mathcal{M}} : \mathcal{K} \rightarrow \mathcal{O}$  be a probing function that queries the LLM to determine if it knows a fact  $f \in \mathcal{K}$  and post-processes its output accordingly. Then, we propose that embeddings can be used as proxies for this function, i.e.,  $\hat{\mathcal{P}}_{\theta}(f) = \mathbf{W}^{\top} \mathcal{E}(f) \approx \mathcal{P}_{\mathcal{M}}(f)$  for all facts  $f \in \mathcal{K}$ , where  $\mathcal{E}$  is an embedding model that embeds facts  $f \in \mathcal{K}$  to a low-dimensional space  $\mathbb{R}^d$  and  $\mathbf{W} \in \mathbb{R}^d$ .*

#### 3.1 Knowledge Probing Functions

We consider four different ways of probing  $\mathcal{P}_{\mathcal{M}}$  the LLM  $\mathcal{M}$  to find whether it knows a fact or not.

**(1) Binary Generation.** The simplest way to probe the knowledge of a language model [34] is by querying it for the truthfulness of a fact. In particular, we prompt:

```

You are only supposed to respond in yes/no.
Is the following statement <bool>?
STATEMENT: <fact>
ANSWER:

```

where  $\langle \text{bool} \rangle$  is randomly sampled from  $\{\text{True}, \text{False}\}$ , and  $\langle \text{fact} \rangle$  is a fact sampled from a knowledge database. For example, “Garifuna is a population of Guatemala” is a valid  $\langle \text{fact} \rangle$ . The generated response can be then post-processed into a binary label  $y \in \mathcal{O} = \{0, 1\}$ , where  $y = 1$  corresponds to the LLM knowing the fact, *i.e.*, if the output is ‘yes’/‘no’ when  $\langle \text{bool} \rangle$  is True/False. In the case of a negative fact (a fact is not actually true),  $y = 1$  when the output is ‘yes’/‘no’ and  $\langle \text{bool} \rangle$  is False/True respectively.

**(2) Binary Logits Generation.** We consider the same setting as above but extract the logits of the expected token from the LLM, *i.e.*, ‘yes’/‘no’ when  $\langle \text{bool} \rangle$  is True/False. Thus, we obtain the probed values  $y \in \mathcal{O} = \mathbb{R}$ , representing a measure of the LLM’s “external confidence” for the fact being true/false based on  $\langle \text{bool} \rangle$ .

**(3) Binary Activation Prediction.** Hidden representations are known to possess knowledge about the truthfulness [16]. Thus, we probe this knowledge directly by prompting it with *just* the “ $\langle \text{fact} \rangle$ ” and training a linear head on top of a fixed layer of the last token of the LLM to predict 1 if  $\langle \text{fact} \rangle$  is true and 0 if it is a negative fact. We thus train the linear layer for 10 epochs on 80% of the data. Then, we predict the logits of these binary predictions from the linear layer for the complete dataset, *i.e.*,  $y \in \mathcal{O} = \mathbb{R}$ . Thus, it measures the “internal confidence” of an LLM for a fact to be true/false.

**(4) Fact Generation.** Finally, we also consider a more implicit way of probing the knowledge, where it is asked to generate a long-form biography or description of an entity/person. For example, “Tell me about Garifuna”. Here, we leverage an existing dataset, FactScore [27] that extracts these passages from LLMs and then label each atomic facts from the passage based on whether it is supported, not supported, or irrelevant corresponding to a knowledge base. By considering the supported facts as 1 and the unsupported facts as 0, we get a dataset of atomic facts  $\langle \text{fact} \rangle$  generated by LLMs that are either known by the LLM or not, labeled as  $y \in \mathcal{O} = \{0, 1\}$ .

## 3.2 Proxy Embeddings

To estimate the LLM-probed knowledge, we use two different pre-trained embedding models that encode this knowledge in a self-supervised way.

**(1) Sentence embedding models.** Facts are sentences that combine entities with specific relations. Thus, one can find meaningful representations of them by leveraging pre-trained sentence embedding models [24, 22]. These models are trained to represent “similar” sentences in a database close to each other. Two sentences can be deemed “similar” based on relevance for retrieval [35] and clustering tasks [36]. While these models were traditionally trained from scratch on contrastive loss, newer models have adapted the progress in generative LLMs by leveraging their learned parameters for encoding sentences in a low-dimensional space. More details about training paradigms and architectural similarities are provided in Appendix D.

**(2) Graph neural networks.** It has been shown that the reasoning abilities of the LLMs likely emerge due to a graphical combination of knowledge within texts [37]. Facts sampled from a knowledge graph are structured as a triple consisting of a head entity, a relation, and a tail entity. Thus, we can exploit a Graph Neural Network (GNN) encoder to embed each fact of a knowledge graph in a low-dimensional space. These are trained on a self-supervised loss to predict whether an edge exists for a head, relation, or tail triple. We employ a pre-trained foundational model for knowledge graphs, ULTRA [23], which

encodes entities using the shortest path on the graph. This ensures that triples with an edge in the graph are separated in the embedding space from those that do not form an edge.

### 3.3 Proxy Tuning

Since embeddings are not pre-trained to estimate probed knowledge from LLMs, they cannot be used directly. Thus, we adapt them by tuning their parameters to minimize a loss function that matches the LLM-probed knowledge with the embedding-based prediction for a small training dataset. Then, we consider two common approaches to efficiently tune the parameters [38]:

- **Low-rank tuning:** We update the parameters of the embedding models by adding a low-rank matrix [39], *i.e.*,  $\mathbf{W}^T \mathcal{E} + \alpha/r \Delta \Theta$ , where  $\Delta \Theta$  has rank at most  $r$ .
- **Linear-tuning:** One can also simply update the final linear layer  $\mathbf{W}$  while keeping the parameters of the embedding  $\mathcal{E}$  frozen.

**Loss function.** We tune these parameters to minimize the loss function such that  $\mathcal{P}_{\mathcal{M}}(f) \approx \hat{\mathcal{P}}_{\theta}(f)$  for all  $f \in \mathcal{K}$ . Depending on the probing function  $\mathcal{P}$ , we consider different losses. When  $\mathcal{P}_{\mathcal{M}} \rightarrow \mathcal{O} = \{0, 1\}$ , *i.e.*, for binary generation and fact generation probing, we use the binary cross-entropy loss:  $\sum_f \mathcal{P}_{\mathcal{M}}(f) \log \sigma(\hat{\mathcal{P}}_{\theta}(f)) + (1 - \mathcal{P}_{\mathcal{M}}(f)) \log(1 - \sigma(\hat{\mathcal{P}}_{\theta}(f)))$ , where  $\sigma$  is the sigmoid function. On the other hand, we use the knowledge distillation loss with temperature  $T$ :  $\sum_f \sigma(\mathcal{P}_{\mathcal{M}}(f)/T) \log(\sigma(\hat{\mathcal{P}}_{\theta}(f)/T))$ , when  $\mathcal{P}_{\mathcal{M}} \rightarrow \mathcal{O} = \mathbb{R}$ , *i.e.*, for binary logits generation and binary activation prediction.

## 4 Experimental Setup

**Datasets.** For binary prediction, logits generation, and activation prediction, we construct datasets of boolean knowledge questions by downsampling large knowledge graphs, specifically DBP100k [1] (697572 facts) and YAGO310 [2] (1089040 facts), by considering 10%, 1%, and 0.1% of the graph. Furthermore, we ensure that the proportion of facts with a particular relation type is preserved in the downsampled set to avoid sampling from a particular relation type only. Since knowledge graphs encode semantic relations differently from natural language, direct decoding into natural language is non-trivial [40, 41]. Thus, we employ a heuristic approach by manually annotating a query prompt for each relation type  $r$  using placeholders for the head ( $h$ ) and tail ( $t$ ) entities. This is an efficient strategy to generate textual prompts since common knowledge graphs often contain only hundreds of relation types compared to millions of factual relationships. Some examples of templates are provided in Table 5 (Appendix B.1).

For fact generation task, we use the ChatGPT-generated labels of FactScore for different labels [27] due to their high alignment with humans. We split all datasets into training (0.8), validation (0.1), and test (0.1) sets with respective splits and report the results on the held-out test set. Note that since we randomly split the dataset, the test splits can include new nodes testing for the inductiveness of the entity representation. For more details, refer Appendix B.2.

**Embedding models.** For LLM knowledge estimation on knowledge graphs, we consider a pre-trained graph neural network ULTRA [23] while we consider 6 state-of-the-art other sentence embedders for generic LLM knowledge estimation: MPNET [42], NVE2 [43], Linq [44], GTE [45], GIST [46], and MXBAI [47]. These are chosen for their high performance due to their high retrieval performance on knowledge-related tasks on MTEB [48] benchmark.

**Large language models.** For knowledge graphs, we estimate the knowledge of 1 open-source model: Llama-3.1-8B [49] as well as 2 proprietary models: GPT-4o and GPT-4o-mini [50]. We also consider the proprietary models to reason about our prompt with a standard chain-of-thought prompt (COT) [51]: Think step-by-step and acquire the final result using structured outputs API <sup>1</sup>. For FactScore,

<sup>1</sup> <https://platform.openai.com/docs/guides/structured-outputs#chain-of-thought>

| DBP-100k  |                |                |                |                |                |                |                 |                |                |                |                  |
|-----------|----------------|----------------|----------------|----------------|----------------|----------------|-----------------|----------------|----------------|----------------|------------------|
|           | Llama-3.1-8B   |                | GPT-4o-mini    |                | GPT-4o         |                | GPT-4o-mini COT |                | GPT-4o COT     |                | Average rank for |
|           | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$  | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC+ACC*         |
| Majority  | -              | 67.23          | -              | 66.67          | -              | 69.49          | -               | 76.27          | -              | 83.05          | 9.00             |
| Random    | -              | 42.94          | -              | 46.33          | -              | 50.85          | -               | 44.63          | -              | 49.15          | 10.00            |
| Llama-Hid | 60.70          | 73.45          | 61.67          | 66.10          | 58.59          | 73.45          | 60.58           | 67.23          | 65.86          | 80.79          | 7.40             |
| MPNET     | 66.66          | 80.94          | 73.28          | 88.53          | 68.06          | 87.57          | 73.01           | 87.79          | 69.50          | 90.00          | 5.00             |
| NVE2      | <b>80.80</b>   | <b>81.71</b>   | <b>86.47</b>   | 90.22          | 80.83          | 87.50          | <b>88.82</b>    | 90.16          | 86.89          | 91.52          | <b>1.40</b>      |
| GIST      | 51.11          | 73.45          | 71.69          | 90.42          | 70.07          | 88.65          | 71.85           | 87.28          | 70.18          | 90.81          | 5.80             |
| Linq      | 64.53          | 73.45          | 85.12          | <b>91.09</b>   | <b>82.01</b>   | <b>88.11</b>   | 86.88           | <b>90.53</b>   | <b>88.72</b>   | <b>91.62</b>   | 2.00             |
| GTE       | 62.03          | 73.45          | 81.73          | 90.42          | 76.59          | 87.57          | 85.11           | 88.77          | 85.56          | 91.49          | 3.40             |
| MXBAI     | 51.93          | 73.45          | 73.29          | 90.42          | 70.53          | 88.65          | 73.79           | 87.28          | 70.93          | 90.95          | 4.60             |
| ULTRA     | 58.58          | 80.80          | 66.32          | 88.73          | 60.08          | 71.62          | 64.82           | 87.35          | 64.78          | 90.00          | 6.40             |

| YAGO310   |                |                |                |                |                |                |                 |                |                |                |                  |
|-----------|----------------|----------------|----------------|----------------|----------------|----------------|-----------------|----------------|----------------|----------------|------------------|
|           | Llama-3.1-8B   |                | GPT-4o-mini    |                | GPT-4o         |                | GPT-4o-mini COT |                | GPT-4o COT     |                | Average rank for |
|           | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC $\uparrow$  | ACC $\uparrow$ | AUC $\uparrow$ | ACC $\uparrow$ | AUC+ACC*         |
| Majority  | -              | 68.00          | -              | 52.89          | -              | 51.24          | -               | 54.55          | -              | 60.33          | 9.00             |
| Random    | -              | 47.56          | -              | 51.24          | -              | 47.93          | -               | 52.07          | -              | 52.89          | 10.00            |
| Llama-Hid | 59.33          | 67.55          | 79.67          | <b>70.25</b>   | <b>66.51</b>   | 59.50          | 79.83           | 72.72          | 70.63          | 63.64          | 2.80             |
| MPNET     | 59.08          | 64.44          | 77.50          | 67.77          | 60.80          | 57.02          | 75.77           | 67.77          | 64.73          | 61.98          | 5.60             |
| NVE2      | 60.56          | 65.22          | <b>82.97</b>   | 69.42          | 64.65          | 57.85          | <b>81.81</b>    | <b>73.55</b>   | 65.92          | 60.06          | 3.60             |
| GIST      | 65.80          | 62.40          | 73.08          | 63.64          | 61.56          | 54.55          | 73.98           | 66.12          | 68.27          | 61.16          | 6.40             |
| Linq      | <b>69.16</b>   | 65.18          | 82.75          | 68.60          | 64.05          | <b>60.33</b>   | 78.43           | 71.07          | <b>78.13</b>   | <b>70.25</b>   | <b>1.80</b>      |
| GTE       | 59.97          | <b>71.11</b>   | 80.47          | 64.46          | 65.04          | 55.37          | 77.42           | 67.77          | 77.13          | <b>70.25</b>   | 3.60             |
| MXBAI     | 59.33          | 70.22          | 74.07          | 64.46          | 62.38          | 55.37          | 75.30           | 66.94          | 69.24          | 62.81          | 5.40             |
| ULTRA     | 56.17          | 64.33          | 75.44          | 65.29          | 62.79          | 58.68          | 72.01           | 66.12          | 66.97          | 58.68          | 6.80             |

Table 1: **Binary Generation:** Accuracy (ACC) and AUC for predicting LLM-generated binary truth for 0.1% true facts sampled from knowledge graphs. Standard deviation  $\leq 1$  for accuracy and  $\leq 5$  for AUC.

we use the LLMs used in their paper, *i.e.*, we consider the labeled facts generated by Alpaca [52], StableLM-alpha [53], ChatGPT [54], Vicuna [55], and InstructGPT [56].

**Metrics.** When the LLM probing outputs are binary (*i.e.*, for binary generation and fact generation), we consider accuracy (ACC) and area under ROC (AUC) of the trained estimation model. On the other hand, when the outputs are real-valued (*i.e.*, for binary logits generation and binary activation prediction), we use mean absolute error (MAE) between the predicted and actual values.

**Baselines.** We consider three different baselines to show the efficacy of using proxy embeddings. In particular, we employ the Majority classifier (*i.e.*, always predicting the majority class) and a Random classifier (*i.e.*, randomly predicting the class) for the binary generation task. In addition to this, we follow [16] and employ LlamaHid, that uses the hidden representations of the Llama-3.1-8B at layer 10, 20, 30 and compare the best performance.

**Hyperparameters.** For linear tuning, we find the best embedding model using performance on the validation set for the learning rate  $\in [10^{-3}, 10^{-2}]$  and the number of epochs  $\in [20, 40]$ . On the other hand, we consider number of epochs  $\in [5, 10]$  for LoRA tuning with  $\alpha = 8$  and  $r = 8$ .




Figure 3: Effect of changing the number of negative samples in GPT models for knowledge graphs.

## 5 Results

### 5.1 Binary Generation

#### 5.1.1 Base performance

First, we establish the hardness of our task of estimating the LLM-acquired knowledge by validating their accuracy for the binary generation task when compared with the true factual knowledge of the two datasets. Table 7 (App. C) shows that around 40-70 % of the true facts are accurately predicted by the model in each dataset for different sampling percentages assuming 0 negative samples. This also shows that the binary labels of knowledge are *verified to be balanced*, making the accuracy and AUC metrics suitable for this evaluation and using a majority or random classifier a suboptimal choice.

Then, we estimate the LLM’s knowledge using the binary generation probing strategy on large knowledge graphs (*i.e.*, DBP100k and YAGO310) by downsampling the graphs to 0.1% of the total size and only probing the positive facts (*i.e.*, triples with an edge). Table 1 shows the effect of linear tuning of different embedding models to match the binary generation probes from different LLMs on the test set. In almost all cases, we find that sentence embedding models are more effective at estimating the knowledge acquired by different LLMs than all the baselines including the strong and arguably more expensive Llama-Hid, achieving up to 91% accuracy and up to 88% AUC scores for GPT-4o-mini. This highlights the effectiveness of PEEK in efficiently and accurately identifying factual gaps in LLMs compared to existing methods. We find that, in most cases, embedding models are very effective in estimating LLMs’ knowledge in these datasets. Further analysis in Appendix C shows that a notable exception is Llama3.3-70B, for which all embedding models fail to estimate its knowledge. has almost random accuracy in finding the correct knowledge over these datasets, indicating that estimating it is effectively as difficult as predicting the output of a random number generator.

Nvidia’s NVE2 is the overall winner in DBP100k while Linq performs the best and second-best in YAGO310 and DBP100k for different LLMs. This shows that Linq is best at estimating the knowledge acquired by arbitrary LLMs. However, a deeper look at the results shows that the ranks are mainly due to high performance on GPT-4o-mini and the results on GPT-4o in YAGO310 and Llama models are just above random performance. This shows that there is significant room for improvement for PEEK. Another interesting observation is that pre-trained representations of a graph neural network, ULTRA, does not align with what LLMs know to be true/false, as ULTRA performs the worst among all considered embedding models. This shows that LLMs do not induce facts in the same way as a path-based graph neural network does on a knowledge graph, *i.e.*, using shortest paths in a Wikipedia-based graph.

**Negative samples** Next, we find how negative facts impact the estimation performance of the LLM knowledge in these knowledge graphs. In particular, for every positive fact/edge,  $(h, r, t)$ , we randomly sample a negative edge  $(h, r, t')$  such that no edge exists corresponding to the triple. Figure 3 shows how the AUC and accuracy changes of 3 representative embedding models: NVE2, ULTRA, and MPNET, as




Figure 4: Effect of changing the number of negative samples in Llama3.1-8B for knowledge graphs.




Figure 5: **LoRA v/s Linear tuning** on 1% DBP100k

we change the number of negative samples/edges from  $[0, 1, 10]$ . We observe that AUC scores typically increase with an increase in the number of negative samples. Assuming LLMs know the ground truth, there is a likelihood of more negative facts and thus, the increase could be due to a natural class imbalance, and thus, the models tend to favor the majority class of “false”. Furthermore, knowledge graphs tend to have missing knowledge, that is true but not modeled in the graph. To avoid these cases, we focus on 0 negative samples.

**Sampling percentage** We also verify the stability of our results by varying the sampling percentage with a Llama3.1-8B model. Figure 4 shows how AUC and accuracy vary in two representative embedding models (other models show similar results) for different sampling percentages. We can observe that the results are mostly preserved with a maximum change of 5% in the AUC and 10% in the ACC as we increase the sampling percentage. Thus, we consider a fixed sampling percentage of 0.1% for our global results since it is more efficient and scalable to proprietary models, reducing API and computational costs.

**Fine-tuning** We also analyze if LoRA tuning improves the performance of embedding models as compared to linear tuning. Figure 5 compares the AUC and accuracy of the three embedding models: NVE2, MPNET, and ULTRA for LoRA and linear tuning methods to estimate the knowledge of GPT-4o and GPT-4o-mini on 1% of DBP100k dataset. One can observe that the LoRA tuning achieves identical performance as linear tuning, showing that LoRA tuning may not be needed for this task. This further strengthens the claim that embedding models implicitly encode knowledge in a way that can easily adapt to match an LLM’s acquired knowledge space using just a linear layer.

## 5.2 Binary Logits Generation

Here, we directly adapt the embedding models to predict the logits of the correct token for a prompt (*i.e.*, “yes” token when  $\langle \text{b o o l} \rangle$  is True and “no” token when  $\langle \text{b o o l} \rangle$  is False) using the knowledge distillation loss. Table 3 reports the mean absolute error between the predicted and actual logits of

| LLM                                                   | Metric         | MPNET | NVE2         | GIST  | Linq         | GTE          | MXBAI |
|-------------------------------------------------------|----------------|-------|--------------|-------|--------------|--------------|-------|
| Alpaca                                                | AUC $\uparrow$ | 84.41 | <b>90.14</b> | 86.01 | 89.11        | 89.93        | 86.07 |
|                                                       | ACC $\uparrow$ | 76.17 | 82.38        | 78.24 | <b>83.42</b> | 82.90        | 78.76 |
| StableLM                                              | AUC $\uparrow$ | 71.17 | 86.06        | 78.93 | 85.90        | <b>86.83</b> | 79.92 |
|                                                       | ACC $\uparrow$ | 83.72 | 84.50        | 83.72 | 84.88        | <b>86.05</b> | 83.72 |
| ChatGPT                                               | AUC $\uparrow$ | 67.88 | <b>79.73</b> | 68.04 | 78.90        | 77.56        | 70.15 |
|                                                       | ACC $\uparrow$ | 76.51 | <b>80.63</b> | 74.92 | 79.37        | 78.10        | 75.24 |
| Vicuna                                                | AUC $\uparrow$ | 81.54 | 86.93        | 82.99 | <b>88.19</b> | 87.26        | 83.16 |
|                                                       | ACC $\uparrow$ | 74.18 | 79.08        | 76.09 | 79.35        | <b>79.89</b> | 76.63 |
| InstructGPT                                           | AUC $\uparrow$ | 80.79 | <b>84.80</b> | 80.94 | 84.64        | 84.58        | 80.74 |
|                                                       | ACC $\uparrow$ | 76.56 | 76.92        | 72.89 | <b>77.66</b> | 76.19        | 73.63 |
| Overall rank for $\frac{1}{2}(\text{AUC}+\text{ACC})$ |                | 6     | 2            | 5     | 1            | 3            | 4     |

Table 2: **Fact Generation:** Accuracy (ACC) and AUC for predicting (in)correct facts generated by LLMs.

| Embedding | DBP-100k    | YAGO-310    |
|-----------|-------------|-------------|
| MPNET     | 4.30        | 1.78        |
| NVE2      | 2.98        | <b>1.64</b> |
| GIST      | 2.60        | 1.81        |
| Linq      | <b>2.36</b> | 1.67        |
| GTE       | 2.88        | 1.69        |
| MXBAI     | 2.43        | 1.84        |

Table 3: **Binary logits generation:** Mean absolute error in predicting the Llama3.1-8B logits of the correct token for 1% data. Lesser is better.

Llama3.1-8B for different embeddings, as we search over temperatures  $T \in [1, 5, 10]$ . We can observe the high performance of Linq model in this case as well with NVE2 and GTE models coming close. However, the error is still quite high (all  $> 1$ ), leaving room for enough improvement in estimating the LLM’s logit-level knowledge.

### 5.3 Binary Activation Prediction

Since knowledge can be encoded in the representations and not get translated to the final logits [16], we also train an additional linear layer on the hidden activations of Llama3.1-8B to separate true and negative facts in a knowledge graph. We use 1 and 10 negative samples for training and report the best results in Table 4 compares the estimation using the last hidden layer over temperatures  $T \in [1, 5, 10]$ . We observe similar trends as above and note that error values are smaller due to more normalized logits tuned for the binary classification using the adapted linear layer. NVE2 and GTE perform the best for the two datasets. We further extend this analysis by using an intermediate 15th layer in Table 8 (Appendix D) and find the trends for 0.1% and 1% data to be similar.

### 5.4 Fact Generation

Table 2 shows the performance of embedding models to estimate the knowledge of different LLMs as found through open-ended fact generation. We report the accuracy and AUC scores on the test set by learning to label the correct facts generated by the LLM as 1 and the incorrect facts as 0. Since these facts can be formed from arbitrary English text, we omit the graph neural network, ULTRA from this

| Embedding | DBP-100k      | YAGO-310      |
|-----------|---------------|---------------|
| MPNET     | 0.5714        | 0.5722        |
| NVE2      | <b>0.4746</b> | 0.5006        |
| GIST      | 0.5724        | 0.5392        |
| Linq      | 0.4845        | 0.4895        |
| GTE       | 0.4959        | <b>0.4847</b> |
| MXBAI     | 0.5625        | 0.5353        |

Table 4: **Binary activation prediction:** Mean absolute error in predicting the adapted logits from the activation layer of Llama3.1-8B on 0.1% data. Lesser is better.

comparison. We find that the embedding models can estimate this knowledge space as well with upto 90% AUC and 86% accuracy. Results further show that Linq outperforms all the other embedding models with NVE2 coming second. This complements our findings on the binary generation task and establishes that these embeddings are quite effective in adapting to the knowledge space of different LLMs with a linear layer across different knowledge probing strategies.

## 6 Conclusion

We presented PEEK, a novel framework for estimating knowledge encoded in large language models using proxy embedding models. Our method offers a scalable and cost-effective way to assess what an LLM knows, which is essential for identifying and addressing factual gaps before deployment. Embedding models should thus be used in the pre-deployment stage of any LLM to pre-emptively find and fix factual gaps in the LLMs. Effective estimation of underlying knowledge using embeddings enables us to efficiently utilize the context in retrieval-augmented generation efficient context utilization by disregarding already-known facts. Another interesting future direction lies in analyzing how the proxy embeddings can dynamically adapt as the model continues to learn from limited examples.

## Limitations

**Confounding of other factors.** A limitation of our work is that our results can be confounded by a lot of hidden factors in the learning process and data of the two models. While these findings show significant promise, it is likely that GPT-generated knowledge or architecture itself [24] is leaked to train embedding models. Since the main contribution of our work is in identifying the potential of embedding models to estimate LLM knowledge, these factors do not undermine our results in efficient factual probing. We thus leave it for future works to analyze the mechanisms of why certain embeddings estimate knowledge better than others and provide a starting point in Section D.

**Other probing tasks.** Another limitation is that we do not consider more difficult probing tasks such as fill-in-the-blank or multiple choice question-answering since they can lead to arbitrary generation of facts which are harder to estimate through encoder-only models.

**Wikipedia.** While our benchmark focuses on Wikipedia-derived knowledge (e.g. DBP100k), we emphasize that our method is not inherently restricted to factual or closed-domain settings. We also note that harder domain-specific knowledge would potentially require training domain-specific embeddings from scratch since existing embedding models may not necessarily generalize to these. Our focus here is to show a standardized probing and estimation pipeline and we believe this would generalize well to these settings as well given such specialized embeddings exist. While Wikipedia is a widely accepted source of knowledge [57, 58], more specialized sources can be explored in the future.

## References

- [1] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, et al. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195, 2015.
- [2] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: A core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706, 2007.
- [3] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *arXiv:1909.01066*, 2019.
- [4] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.
- [5] Neel Guha, Julian Nyarko, Daniel Ho, Christopher Ré, Adam Chilton, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, et al. Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models. *Advances in Neural Information Processing Systems*, 36:44123–44279, 2023.
- [6] Parshin Shojaee, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, and Chandan K Reddy. Llm-sr: Scientific equation discovery via programming with large language models. *arXiv:2404.18400*, 2024.
- [7] Pingchuan Ma, Tsun-Hsuan Wang, Minghao Guo, Zhiqing Sun, Joshua B Tenenbaum, Daniela Rus, Chuang Gan, and Wojciech Matusik. Llm and simulation as bilevel optimizers: a new paradigm to advance physical scientific discovery. In *Proceedings of the 41st International Conference on Machine Learning*, pages 33940–33962, 2024.
- [8] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv:2311.05232*, 2023.
- [9] Toby Shevlane, Sebastian Farquhar, Ben Garfinkel, Mary Phuong, Jess Whittlestone, Jade Leung, Daniel Kokotajlo, Nahema Marchal, Markus Anderljung, Noam Kolt, et al. Model evaluation for extreme risks. *arXiv:2305.15324*, 2023.
- [10] Heidy Khlaaf, Pamela Mishkin, Joshua Achiam, Gretchen Krueger, and Miles Brundage. A hazard analysis framework for code synthesis large language models. *arXiv:2207.14157*, 2022.
- [11] Shangshang Zheng, He Bai, Yizhe Zhang, Yi Su, Xiaochuan Niu, and Navdeep Jaitly. Kglens: Towards efficient and effective knowledge probing of large language models with knowledge graphs. *arXiv:2312.11539*, 2023.
- [12] Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pages 15696–15707. PMLR, 2023.
- [13] Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu, and Xin Luna Dong. Head-to-tail: How knowledgeable are large language models (llm)? aka will llms replace knowledge graphs? *arXiv:2308.10168*, 2023.
- [14] Linhao Luo, Thuy-Trang Vu, Dinh Phung, and Gholamreza Haffari. Systematic assessment of factual knowledge in large language models. *arXiv:2310.11638*, 2023.

- [15] Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020.
- [16] Amos Azaria and Tom Mitchell. The internal state of an llm knows when it’s lying. *arXiv:2304.13734*, 2023.
- [17] Daniela Gottesman and Mor Geva. Estimating knowledge in large language models without generating a single token. *arXiv:2406.12673*, 2024.
- [18] Yanling Wang, Haoyang Li, Hao Zou, Jing Zhang, Xinlei He, Qi Li, and Ke Xu. Faclens: Transferable probe for foreseeing non-factuality in large language models. *arXiv:2406.05328*, 2024.
- [19] Gustaf Ahlritz, Tian Qin, Nikhil Vyas, Boaz Barak, and Benjamin L Edelman. Distinguishing the knowable from the unknowable with language models. *arXiv:2402.03563*, 2024.
- [20] Potsawee Manakul, Adian Liusie, and Mark JF Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv:2303.08896*, 2023.
- [21] Andrew Jesson, Nicolas Beltran Velez, Quentin Chu, Sweta Karlekar, Jannik Kossen, Yarin Gal, John P Cunningham, and David Blei. Estimating the hallucination rate of generative ai. *Advances in Neural Information Processing Systems*, 37:31154–31201, 2025.
- [22] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embedding benchmark. *arXiv:2210.07316*, 2022.
- [23] Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, and Zhaocheng Zhu. Towards foundation models for knowledge graph reasoning. *2310.04562*, 2023.
- [24] Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. Llm2vec: Large language models are secretly powerful text encoders. *arXiv:2404.05961*, 2024.
- [25] Duo Wang, Yuan Zuo, Fengzhi Li, and Junjie Wu. Llms as zero-shot graph learners: Alignment of gnn representations with llm token embeddings. *Advances in Neural Information Processing Systems*, 37:5950–5973, 2024.
- [26] Paul Youssef, Osman Alperen Koraş, Meijie Li, Jörg Schlötterer, and Christin Seifert. Give me the facts! a survey on factual knowledge probing in pre-trained language models. *arXiv:2310.16570*, 2023.
- [27] Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. *arXiv:2305.14251*, 2023.
- [28] Yuyang Bai, Shangbin Feng, Vidhisha Balachandran, Zhaoxuan Tan, Shiqi Lou, Tianxing He, and Yulia Tsvetkov. Kgquiz: Evaluating the generalization of encoded knowledge in large language models. In *Proceedings of the ACM Web Conference 2024*, pages 2226–2237, 2024.
- [29] Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36:41451–41530, 2023.
- [30] Kevin Liu, Stephen Casper, Dylan Hadfield-Menell, and Jacob Andreas. Cognitive dissonance: Why do language model outputs disagree with internal representations of truthfulness? *arXiv:2312.03729*, 2023.
- [31] Gabriel Y Arteaga, Thomas B Schön, and Nicolas Pielawski. Hallucination detection in llms: Fast and memory-efficient finetuned models. *arXiv:2409.02976*, 2024.

- [32] Yuyan Chen, Qiang Fu, Yichen Yuan, Zhihao Wen, Ge Fan, Dayiheng Liu, Dongmei Zhang, Zhixu Li, and Yanghua Xiao. Hallucination detection: Robustly discerning reliable answers in large language models. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 245–255, 2023.
- [33] Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in words. *arXiv:2205.14334*, 2022.
- [34] Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv:1905.10044*, 2019.
- [35] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv:2104.08663*, 2021.
- [36] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)*, pages 410–420, 2007.
- [37] Xinyi Wang, Alfonso Amayuelas, Kexun Zhang, Liangming Pan, Wenhua Chen, and William Yang Wang. Understanding the reasoning ability of language models from the perspective of reasoning paths aggregation. *arXiv:2402.03268*, 2024.
- [38] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv:2108.07258*, 2021.
- [39] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv:2106.09685*, 2021.
- [40] Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. Creating training corpora for NLG micro-planning. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 111–120. Association for Computational Linguistics, 2017.
- [41] Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. The webnlg challenge: Generating text from rdf data. In *Proceedings of the 10th international conference on natural language generation*, pages 124–133, 2017.
- [42] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *NeurIPS*, 33:16857–16867, 2020.
- [43] Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Nv-embed: Improved techniques for training llms as generalist embedding models. In *ICLR*, 2025.
- [44] Junseong Kim, Seolhwa Lee, Jihoon Kwon, Sangmo Gu, Yejin Kim, Minkyung Cho, Jy yong Sohn, and Chanyeol Choi. Linq-embed-mistral: Elevating text retrieval with improved gpt data through task-specific control and quality refinement. Linq AI Research Blog, 2024.
- [45] Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning. *arXiv:2308.03281*, 2023.
- [46] Aivin V. Solatorio. Gistembed: Guided in-sample selection of training negatives for text embedding fine-tuning. *arXiv:2402.16829*, 2024.

- [47] Xianming Li and Jing Li. Angle-optimized text embeddings. *arXiv preprint arXiv:2309.12871*, 2023.
- [48] Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. Mteb: Massive text embedding benchmark. In *EACL*, pages 2014–2037, 2023.
- [49] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv:2407.21783*, 2024.
- [50] OpenAI. Gpt-4 technical report. *Arxiv Preprint*, arXiv:2303.08774, 2023.
- [51] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- [52] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- [53] Marco Bellagente, Jonathan Tow, Dakota Mahan, Duy Phung, Maksym Zhuravinskyi, Reshinh Adithyan, James Baicoianu, Ben Brooks, Nathan Cooper, Ashish Datta, et al. Stable lm 2 1.6b technical report. *arXiv:2402.17834*, 2024.
- [54] OpenAI. ChatGPT: Conversational ai platform. Online, 2022.
- [55] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6, 2023.
- [56] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [57] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale dataset for fact extraction and verification. *arXiv:1803.05355*, 2018.
- [58] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv:1809.09600*, 2018.
- [59] Andrew Ng et al. Sparse autoencoder. *CS294A Lecture notes*, 72(2011):1–19, 2011.
- [60] Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2. *arXiv:2408.05147*, 2024.
- [61] Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology*, 15(2):1–38, 2024.
- [62] Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. Investigating the limitations of transformers with simple arithmetic tasks. *arXiv:2102.13019*, 2021.
- [63] Philip Quirke and Fazl Barez. Understanding addition in transformers. *arXiv:2310.13121*, 2023.

- [64] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. *arXiv:1911.00172*, 2019.
- [65] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR, 2020.
- [66] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *NeurIPS*, 33:9459–9474, 2020.
- [67] Chanwoong Yoon, Taewhoo Lee, Hyeon Hwang, Minbyul Jeong, and Jaewoo Kang. Compact: Compressing retrieved documents actively for question answering. In *EMNLP*, 2024.
- [68] Xin Cheng, Di Luo, Xiuying Chen, Lema Liu, Dongyan Zhao, and Rui Yan. Lift yourself up: Retrieval-augmented text generation with self-memory. *NeurIPS*, 36, 2024.
- [69] Fangyuan Xu, Weijia Shi, and Eunsol Choi. Recomp: Improving retrieval-augmented lms with compression and selective augmentation. In *ICLR*, 2024.
- [70] Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Halueval: A large-scale hallucination evaluation benchmark for large language models. *arXiv:2305.11747*, 2023.
- [71] Hanyu Duan, Yi Yang, and Kar Yan Tam. Do llms know about hallucination? an empirical investigation of llm’s hidden states. *arXiv:2402.09733*, 2024.
- [72] Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Akim Tsvigun, Daniil Vasilev, Rui Xing, Abdelrahman Boda Sadallah, Kirill Grishchenkov, Sergey Petrakov, et al. Benchmarking uncertainty quantification methods for large language models with lm-polygraph. *arXiv:2406.15627*, 2024.
- [73] Jonas Wallat, Jaspreet Singh, and Avishek Anand. Bertnesia: Investigating the capture and forgetting of knowledge in bert. *arXiv:2106.02902*, 2021.
- [74] Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? *arXiv:2002.08910*, 2020.
- [75] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022.
- [76] Sebastian Farquhar, Vikrant Varma, Zachary Kenton, Johannes Gasteiger, Vladimir Mikulik, and Rohin Shah. Challenges with unsupervised llm knowledge discovery. *arXiv:2312.10029*, 2023.
- [77] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model editing at scale. *arXiv:2110.11309*, 2021.
- [78] Shankar Padmanabhan, Yasumasa Onoe, Michael Zhang, Greg Durrett, and Eunsol Choi. Propagating knowledge updates to lms through distillation. *Advances in Neural Information Processing Systems*, 36, 2024.
- [79] Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv:2309.00267*, 2023.

## Appendix

### A Extended Related Work

**Interpretable LLM representations.** Prior work has explored the problem of understanding the internal mechanism of these extremely large models by learning sparse autoencoders [59, 60], prompting-based explanations [61], and understanding induction on simpler tasks [62, 63, 37]. In the same vein, the interpretability of factual knowledge induction in LLMs can also be improved by leveraging smaller proxy models with well-defined inductive biases to estimate LLM knowledge.

**Applications.** In *retrieval-augmented generation* (RAG) [64–66], where managing long contexts is a key challenge [67–69], PEEK can filter out already-known facts to the LLM, reducing context length and improving retrieval efficiency. For *hallucination detection* [70], existing works explored LLMs’ self-awareness of their own hallucinations [30, 71] and the use of fine-tuned auxiliary models [31] for uncertainty estimation [72]. Our approach complements prior work on hallucination detection by providing a scalable proxy for evaluating LLM self-knowledge. In *model editing and updating*, where factual knowledge is localized and modified within the model [73–76, 75, 77], PEEK embeddings can help identify misaligned facts, which can then be corrected via neuron-level editing [75, 77] or knowledge distillation [78]. Future works can also focus on directly aligning the LLMs to the knowledge space of facts by using embedding alignment techniques or reinforcement learning [79]. Finally, we flag that training embedding models specifically to estimate LLM knowledge can be a double-edged sword as it may lead to forgetting more useful true knowledge, resulting in bad performance on retrieval tasks. We leave it for future works to study this trade-off in more detail and how both fields can grow hand-in-hand.

### B Additional Experimental Details

#### B.1 Relation Annotation

We annotate each relation type to form a question template that encodes the semantic information of this relation. Since the number of relations are limited (at most 470), we can do this efficiently and effectively. Table 5 show some question templates used in this work while we provide the complete list in the supplementary material.

#### B.2 Dataset Details

We also note that the data splits already include new nodes that were not seen during training. Table 6 shows that our splits contain novel entities in the test set that were not seen during training.

#### B.3 Implementation details

All the experiments were conducted on Python 3.8.12 on a Ubuntu 18.04 PC with an Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz processor, 512 GB RAM, and Tesla A100 64 GB GPUs.

### C Additional Experimental Results

**Base accuracy.** Table 7 shows the base accuracy of different LLMs on the binary generation task for the two knowledge graphs at different splits. Here, we also include Llama-3.3-70B [49] for comprehensiveness. As discussed in the main paper, we include the spurious results of Binary generation on Llama-3.3-70B in Table 10.

| Dataset                                  | Relation Type       | Prompt Template                                 |
|------------------------------------------|---------------------|-------------------------------------------------|
| DBP100k<br>(470 total<br>relation types) | associatedBand      | {t} is an associated band of {h}.               |
|                                          | instrument          | {t} is the instrument of {h}.                   |
|                                          | ideology            | {t} is the ideology of {h}.                     |
|                                          | country             | {h} is in the country of {t}.                   |
|                                          | language            | {t} is a language of {h}.                       |
|                                          | occupation          | {t} has been the occupation of {h}.             |
|                                          | education           | {h} has the education in {t}.                   |
|                                          | currentMember       | {t} has been an active member of {h}.           |
|                                          | product             | {h} has the product named {t}.                  |
|                                          | birthPlace          | {t} is the birth place of {h}.                  |
|                                          | managerClub         | {h} has been the manager of the club named {t}. |
|                                          | thirdTeam           | {t} has been the third team in {h}.             |
|                                          | writer              | {t} is the writer of {h}.                       |
|                                          | president           | {t} was the president when {h} was in office.   |
|                                          | city                | {h} is in the city of {t}.                      |
| YAGO310<br>(37 total<br>relation types)  | isLocatedIn         | {h} is located in {t}.                          |
|                                          | playsFor            | {h} has played for {t}.                         |
|                                          | isAffiliatedTo      | {h} has been affiliated to {t}.                 |
|                                          | diedIn              | {h} died in {t}.                                |
|                                          | actedIn             | {h} acted in {t}.                               |
|                                          | graduatedFrom       | {h} graduated from {t}.                         |
|                                          | wasBornIn           | {h} was born in {t}.                            |
|                                          | hasGender           | {t} is the gender of {h}.                       |
|                                          | happenedIn          | {h} happened in {t}.                            |
|                                          | hasMusicalRole      | {t} is the musical role of {h}.                 |
|                                          | isConnectedTo       | {h} is connected to {t}.                        |
|                                          | isMarriedTo         | {h} is married to {t}.                          |
|                                          | participatedIn      | {h} participated in {t}.                        |
|                                          | hasOfficialLanguage | {t} is an official language of {h}.             |
|                                          | hasWonPrize         | {h} has won the {t} prize.                      |

Table 5: Exemplary templates to construct a question prompt for each relation type.

**Training.** Figure 6 shows a set of representative loss curves for linear tuning of embeddings for binary generation in 0.1% sampling of DBP100k. Similar curves are observed for other datasets and tasks.

**Binary Activation Prediction.** Table 8 shows the results on using the 15th layer of Llama-3.1-8B to probe for the knowledge. These results further complement our earlier findings of the effectiveness of the Linq embedding.

**Examples.** Table 11 lists some examples of the correct and incorrect estimation by the best performing Linq embedding for GPT-4o in DBP100k. However, it is hard to make generalizable claims from the list.

## D Discussion on Theoretical Understanding

Due to the nature of the task, it is possible that the same embeddings that estimate an LLM’s knowledge on one dataset may not necessarily estimate its knowledge on another dataset. We do not claim or recommend any one embedding model to estimate knowledge of an LLM for all datasets and our main

| Dataset  | Sample Percentage | # Train Entities | # Val Entities | # Test Entities | # Test - Train Entities |
|----------|-------------------|------------------|----------------|-----------------|-------------------------|
| DBP 100k | 0.01              | 8494             | 1218           | 1508            | 1127                    |
|          | 0.001             | 1421             | 232            | 347             | 191                     |
| YAGO310  | 0.01              | 12763            | 1973           | 2016            | 1358                    |
|          | 0.001             | 1551             | 201            | 230             | 209                     |

Table 6: Statistics of the dataset splits.

| Dataset | LLM             | Sample percentage | Accuracy |
|---------|-----------------|-------------------|----------|
| YAGO310 | GPT-4o          | 0.01              | 46.23    |
|         | GPT-4o          | 0.001             | 44.69    |
|         | GPT-4o-mini     | 0.01              | 46.43    |
|         | GPT-4o-mini     | 0.001             | 43.11    |
|         | GPT-4o COT      | 0.01              | 60.71    |
|         | GPT-4o COT      | 0.001             | 56.30    |
|         | GPT-4o-mini COT | 0.001             | 50.88    |
|         | Llama-3.1-8B    | 0.01              | 42.78    |
|         | Llama-3.1-8B    | 0.001             | 63.94    |
|         | Llama-3.3-70B   | 0.01              | 52.06    |
|         | Llama-3.3-70B   | 0.001             | 52.06    |
| DBP100k | GPT-4o          | 0.01              | 39.69    |
|         | GPT-4o          | 0.001             | 66.45    |
|         | GPT-4o-mini     | 0.01              | 39.74    |
|         | GPT-4o-mini     | 0.001             | 64.63    |
|         | GPT-4o COT      | 0.01              | 59.47    |
|         | GPT-4o COT      | 0.001             | 76.91    |
|         | GPT-4o-mini COT | 0.01              | 51.27    |
|         | GPT-4o-mini COT | 0.001             | 73.01    |
|         | Llama-3.1-8B    | 0.01              | 45.35    |
|         | Llama-3.3-70B   | 0.01              | 51.02    |
|         | Llama-3.3-70B   | 0.001             | 51.02    |

Table 7: Accuracy of LLMs on the Binary Generation task assuming 0 negative samples.

focus is in presenting a systematic evaluation framework to test it. It is thus highly difficult to theoretically understand estimation given the lack of pre-training data for both embeddings and large language models. For completeness and to explore any such correlations from the information that we know, we first categorize the embeddings in Table 9 based on what they encode and their training paradigm and data. This shows that modern embedding architectures are more powerful in estimating LLM-acquired knowledge, such as Mistral-based NVE2 and Linq and Qwen-based GTE. Furthermore, we also find that instruction-tuning and hard negative mining is generally a preferred way to train these models to estimate LLM knowledge. However, the exact performance gaps depend on what kind of negative mining is done to train these models. It can also thus, be argued that LLM-based negative mining would naturally lead to a better estimation of LLM knowledge. Furthermore, we can also note that specialized embedding models such as ULTRA for graph structures are not quite suitable to estimate how LLMs encode knowledge which leads to interesting questions about what/how they encode.




Figure 6: **Training loss curves of various embedding models to adapt with the LLM knowledge probed using binary generation on 0.1% DBP100k.**

| Embedding | DBP-100k      | YAGO-310      |
|-----------|---------------|---------------|
| NVE2      | 0.0917        | 0.0795        |
| GIST      | 0.1130        | 0.0851        |
| Linq      | <b>0.0886</b> | <b>0.0680</b> |
| MXBAI     | 0.1123        | 0.0831        |

Table 8: **Binary activation prediction:** Mean absolute error in predicting the adapted logits from the 15th layer of Llama3.1-8B on 1% data. Lesser is better.

| Embedding model | Encode          | Architecture                              | Training data                     | Training paradigm                                   |
|-----------------|-----------------|-------------------------------------------|-----------------------------------|-----------------------------------------------------|
| MPNET           | Sentence        | BERT and XLNet                            | 1B sentence pairs                 | Contrastive learning                                |
| NVE2            | Sentence        | Mistral7B-v0.1 + latent-attention pooling | Unknown                           | Instruction-tuning and hard-negative mining         |
| Linq            | Sentence        | E5-mistral7b-instruct and Mistral7B-v0.1  | LLM-generated + benchmark         | Instruction tuning + Teacher-guided negative mining |
| GTE             | Sentence        | Qwen7B                                    | Unknown                           | Unknown                                             |
| GIST            | Sentence        | BAAI/bge-large-en-v1.5                    | MEDI + mined triplets from MTEB   | Guided In-sample Selection of Training Negatives    |
| MXBAI           | Sentence        | Matryoshka + binary quantization          | Unknown                           | Instruction + contrastive-tuning                    |
| ULTRA           | Knowledge Graph | NBFNet                                    | FB15k-237, DBP100k, YAGO310, etc. | Self-supervised contrastive learning                |

Table 9: Embedding categorization for understanding knowledge estimation.

| Dataset  | Metric         | Majority | Random | Llama | MPNET        | NVE2  | GIST         | Linq  | GTE   | MXBAI | ULTRA |
|----------|----------------|----------|--------|-------|--------------|-------|--------------|-------|-------|-------|-------|
| DBP-100k | AUC $\uparrow$ | -        | -      | 67.01 | <b>51.14</b> | 51.07 | 51.23        | 47.51 | 49.79 | 49.45 | 49.45 |
|          | ACC $\uparrow$ | 52.40    | 51.29  | 61.62 | <b>51.47</b> | 50.63 | 50.81        | 50.48 | 50.76 | 50.88 | 49.35 |
| YAGO-310 | AUC $\uparrow$ | -        | -      | 58.77 | 48.50        | 50.55 | <b>50.70</b> | 47.61 | 48.94 | 50.65 | 49.75 |
|          | ACC $\uparrow$ | 50.72    | 50.45  | 52.89 | 51.14        | 51.06 | <b>51.23</b> | 47.51 | 49.79 | 49.45 | 49.45 |

Table 10: **Binary Generation on Llama-3.3-70B:** Accuracy (ACC) and AUC for predicting LLM-generated binary truth for 0.1% true facts sampled from knowledge graphs.

| Correct estimation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>University of Edinburgh is an alma mater of Dixie Carter.</p> <p>Berklee College of Music is an alma mater of Hiro Kanagawa.</p> <p>Toyota Corolla (E170/E180) has been assembled in Durban.</p> <p>Mass Mental is an associated band of Stevie Salas.</p> <p>Kerala is the birth place of Sathaar.</p> <p>3rd Marine Division fought in the battle of Calhoun.</p> <p>New Hampshire is the broadcasting area of Bally Sports Midwest.</p> <p>Ibaraki is in the country of Japan.</p> <p>Alf Kjellin is the director of Adventure Time.</p> <p>Luis Milla has been the manager of the club named England national under-21 association football team.</p> |
| Incorrect estimation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <p>Université catholique de Louvain has an affiliation with Coimbra Group.</p> <p>Sap is an associated band of The Greencards.</p> <p>Curve is an associated band of Joe Satriani.</p> <p>Canada is the birth place of Norman Jewison.</p> <p>Biddeford is a type of city.</p> <p>Rochester Americans is a former team of Bob Gracie.</p> <p>Peabody is a part of Essex County.</p> <p>Local Government Act 2000 is the name of a leader from Folkestone and Hythe.</p> <p>Sudan is the place of Mahdist War.</p> <p>George VI was the monarch when Walter Runciman, 1st Viscount Runciman of Doxford was in office.</p>                                     |

Table 11: Examples of correct and incorrect estimation by the Linq embedding for GPT-4o in DBP100k.