

Comparing Template-based and Template-free Language Model Probing

Sagi Shaier,[▽] Kevin Bennett,[◇] Lawrence E Hunter,[†] Katharina von der Wense^{▽◇}

[▽]University of Colorado Boulder

[†]Independent Scholar

[◇]Memorial Healthcare System

[◇]Johannes Gutenberg University Mainz

[▽]E-mail: {sagi.shaier, katharina.kann}@colorado.edu

[†]E-mail: Prof.Larry.Hunter@gmail.com

[◇]E-mail: kevbennett@mhs.net

Abstract

The differences between cloze-task language model (LM) probing with 1) expert-made templates and 2) naturally-occurring text have often been overlooked. Here, we evaluate 16 different LMs on 10 probing English datasets – 4 template-based and 6 template-free – in general and biomedical domains to answer the following research questions: (RQ1) Do model rankings differ between the two approaches? (RQ2) Do models’ absolute scores differ between the two approaches? (RQ3) Do the answers to RQ1 and RQ2 differ between general and domain-specific models? Our findings are: 1) Template-free and template-based approaches often rank models differently, except for the top domain-specific models. 2) Scores decrease by up to 42% Acc@1 when comparing parallel template-free and template-based prompts. 3) Perplexity is negatively correlated with accuracy in the template-free approach, but, counter-intuitively, they are positively correlated for template-based probing. 4) Models tend to predict the same answers frequently across prompts for template-based probing, which is less common when employing template-free techniques. Code and data can be found here: https://github.com/Shaier/probing_template_based_template_free.git.

1 Introduction

In the past few years there has been a growing interest in understanding what parametric knowledge language models (LMs) contain (Jiang et al., 2020). One standard approach of probing LMs for knowledge consists of using “fill-in-the-blank” cloze statements (Shin et al., 2020; Kassner and Schütze, 2020; Sung et al., 2021; Petroni et al., 2019; Meng et al., 2021), where models are tasked with predicting a masked entity given a prompt, e.g., “Dante was born in [MASK]”.

Template-based Probing	Template-free Probing
Template: “[X] (born [MASK])”	N/A
Peter F. Martin (born [MASK])	Peter F Martin (born [MASK]) is an American politician [...]
Dennis B. Sullivan (born [MASK])	Sullivan was born in Chippewa Falls Wisconsin in [MASK]
Tan Jiexi (born [MASK])	Tan Jiexi (born December 2, [MASK] in Shenzhen, China), is a Chinese singer-songwriter
Tasos Neroutsos (born [MASK])	Neroutsos was born in Athens in [MASK] to a wealthy family

Table 1: Template-based and template-free probing examples from the LAMA: Google-RE dataset. In the template-based approach, each template is used to create many prompts which are identical except for the subject entity. In comparison, the template-free approach does not use templates and prompts LMs with an often unique prompt per entity.

While this has been studied in various settings, including multilingual (Kassner et al., 2021), single token predictions (Zhong et al., 2021; Petroni et al., 2019; Sung et al., 2021; Bouraoui et al., 2019), multi-token predictions (Kassner et al., 2021), and prompt optimization (Zhong et al., 2021; Shin et al., 2020), the differences between the two types of prompts – template-based and template-free; see Table 1 – have been overlooked so far.

Although in both methods whether a LM knows a fact is defined by its ability to successfully predict the masked object in a prompt (Petroni et al., 2019), the template-based approach uses templates (which are often manually created) to create the prompts, where *each template is used to create many prompts which are identical except for the subject entity*. In comparison, the template-free approach does not use templates and prompts LMs with an often *unique prompt per entity*.

Each of the two methods has its pros and cons. For example, while the template-based approach generally guarantees that the prompt is evaluating the required knowledge, it requires expensive domain experts. And while the prompts in the

template-free approach are more similar to the training data LMs are trained on, as they come from real-world text and not from artificial templates, they may contain additional irrelevant information. **We hypothesize that this may result in different rankings for the same models when being prompted via the two approaches.**

Here, we 1) evaluate 16 different LMs on 10 probing datasets (4 template-based and 6 template-free) in multiple domains; 2) we propose a method to create template-free domain-specific datasets and use it to develop the first template-free biomedical probing dataset, which allows us to compare the effect of the two probing approaches in two different domains; 3) ask the following research questions: (RQ1) Do model rankings differ between template-based and template-free probing? (RQ2) Do models’ absolute scores differ between the two approaches? (RQ3) Do the answers to the two previous questions differ between general and domain-specific models?

Our study’s results can be summarized as follows: 1) There is a discrepancy in ranking models between template-free and template-based methods, except for the top domain-specific models. 2) Scores decrease by up to 42% Acc@1 when comparing parallel template-free and template-based prompts (i.e., similar subject entities and masked objects). 3) Perplexity is negatively correlated with accuracy in the template-free approach, but, counter-intuitively, they are positively correlated for template-based probing. 4) Models have a tendency to predict similar objects to various prompts, even when the subjects change, when utilizing template-based probing, which is less common when employing template-free techniques.

2 Related Work

Knowledge Probing The premise of knowledge probing is that during training (or pretraining), LMs learn and memorize facts. A standard approach to uncover such information is to use cloze-style “fill-in-the-blank” statements, where models are tasked with predicting the masked entity (Shin et al., 2020; Kassner and Schütze, 2020; Sung et al., 2021; Petroni et al., 2019; Meng et al., 2021). However, previous research suggest that prompt-based knowledge probing methods are often inaccurate, unreliable, and inconsistent (Yao et al., 2022; Jiang et al., 2019; Cao et al., 2022), in addition to requiring a significant amount of manual

effort from experts (Shin et al., 2020). To that end, many work on improving different aspects of probing, e.g., Zhong et al. (2021); Petroni et al. (2019); Sung et al. (2021); Bouraoui et al. (2019) solely focus on predicting single-token entities, and Kassner et al. (2021) expand the approach to multi-token predictions. Furthermore, while many use experts to create templates, Zhong et al. (2021); Shin et al. (2020) use automatic methods to find optimal prompts. There is also work on multilingual probing (Kassner et al., 2021), which expands the probing methods to multiple languages, in addition to expanding the general probing to domain-specific settings, such as biomedicine (Meng et al., 2021; Sung et al., 2021).

Template-based Knowledge Probing The general approach of using template-based probing is to use expert-made templates. The idea is that LMs store relational knowledge, which can be used to populate knowledge bases (KBs; Petroni et al., 2019), which are normally stored as triples in the format of (subject, relation, object). Some of the first ones that have used such template-based probing method also develop the LAMA dataset (Petroni et al., 2019). While showing that LMs do in fact store relational knowledge, they also state that the template design has an impact on the results. This is also shown in domain-specific probing, e.g., biomedicine, where most predictions are highly correlated with prompt templates (Sung et al., 2021). Many others use template-based probing methods (Meng et al., 2021; Sung et al., 2021; Bouraoui et al., 2019; Kassner et al., 2021; Heinzlerling and Inui, 2021).

Template-free Knowledge Probing The difference between template-based and template-free probing is the prompt that is used to query the model. While template-based approaches use expert-made templates, the template-free method uses naturally-occurring text. While the LAMA dataset (Petroni et al., 2019) aims at template-based probing, it also contains **parallel** template-free prompts for two datasets: Google-RE and SQuAD; see Table 1. While most prior probing work features template-based probing, the idea of masking tokens and tasking models with predicting them is what the masked language modeling task (Devlin et al., 2019) is based on. Template-free probing works similarly, but the masks are placed strategically within the sentence.

3 Experiments

The LMs’ input for both template-free and template-based probing is a prompt with one masked entity; see Table 1. The models are then tasked with predicting the masked entity.

While most previous work focuses solely on single-token mask prediction, many entities are composed of more than one token. Hence, we follow Kassner et al. (2021) that expand the probing technique to multi-token prediction and show that it is a better method to investigate knowledge captured by LMs. In particular, we use the same alternative to the “fill-in-the-blank” querying by framing the task as entity ranking. However, while Kassner et al. (2021) limit the prediction to entities of the type required by the prompt, we relax this limitation because 1) few existing datasets have entity type information, and using external resource to classify entity types may be inaccurate and skew results; 2) this simplifies the problem for the models which, again, may skew results. Instead, we allow our models to predict any entity from the dataset’s entity list.

Lastly, we experiment in both general and biomedical domains to analyze whether our experiments generalize. Hence, a key portion of the experiments pertains to biomedical models, as in addition to generic English models, we use both biomedical fine-tuned models and biomedical models.

Evaluation Metric Following prior work by Sung et al. (2021), we use top- k accuracy (Acc@ k), wherein a score of 1 is given if the correct entity appears among the top k predicted entities, and 0 otherwise. Since entities are often related to numerous other entities (N -to- M connections), we use Acc@1, Acc@5, and Acc@10.

3.1 Models

We evaluate 16 different LMs belonging to 3 categories: 1) trained exclusively on generic English text; 2) pretrained on generic English text and fine-tuned on biomedical text; 3) trained only on biomedical text; see Table 2 for an overview.

Generic English Models We experiment with 7 generic English models: DistilBERT (Sanh et al., 2020), BERT-base/large (Devlin et al., 2019), RoBERTa-base/large (Liu et al., 2019), ALBERT-base/large (Lan et al., 2020).

Biomedical Fine-tuned Models We probe 6 models which have been pretrained on generic text, followed by finetuning: PMC RoBERTa¹, COVID Bert,² BlueBert (Peng et al., 2019), Bio Discharge Summary BERT (Alsentzer et al., 2019), Bio ClinicalBERT (Alsentzer et al., 2019), and BioMed-RoBERTa (Gururangan et al., 2020).

Biomedical Models We further experiment with PubMedBERT (Gu et al., 2020), Bioformer (Fang et al., 2023), and BioM-ELECTRA (Alrowili and Shanker, 2021).

3.2 Datasets

3.2.1 Template-based Probing

Comparative Toxicogenomics Database The Comparative Toxicogenomics Database (CTD) is a biomedical database with relations and interactions between biomedical entities. We use the same subset as Sung et al. (2021), which contains template-based prompts that were manually curated.

Biomedical Wikidata The Wikidata dataset from Sung et al. (2021) contains template-based prompts and is based on a general knowledge base. We use the same subset of it as Sung et al. (2021) which only contains biomedical entities and relations that were manually curated.

Google-RE (Templates) Google-RE (Petroni et al., 2019) contains 6.11K template-based prompts from Wikipedia and 3 relations.

T-REx (Templates) The T-REx dataset from Petroni et al. (2019) is based on a subset of Wikidata triples and contain 41 relations. The authors manually define a template for each relation which result in 1.3M template-based prompts.

3.2.2 Template-free Probing

Google-RE (Template-free) While the Google-RE dataset from Petroni et al. (2019) contains 6.11K template-prompts from Wikipedia, **each prompt is manually aligned by the creators of the dataset to text** from Wikipedia that supports it. We use the latter as template-free prompts; see Table 1 for examples.

¹<https://huggingface.co/raynardj/pmc-med-bio-mlm-roberta-large>

²<https://huggingface.co/mrm8488/bioclinicalBERT-fine-tuned-covid-papers>

Model	Parameters	Data
*PubMedBERT	109M	PubMed abstracts+PMC full-text articles (3.2B words/21GB)
*Bioformer	42M	33M PubMed abstracts+1M PMC full-text articles
*BioM-ELECTRA-Generator	49M	PubMed Abstracts
†BioMed-RoBERTa	124M	RoBERTa (160GB)+Semantic Scholar corpus (2.68M papers/47GB)
†COVID Bert	108M	N/A
†BlueBert	109M	Bert+PubMed abstracts+MIMIC-III clinical notes (4500M words/27GB)
†Bio Discharge Summary BERT	108M	Biobert (18B words)+MIMIC III discharge summaries (880M words)
†PMC RoBERTa	355M	RoBERTa (160GB)+ PMC and PubMd abstracts
†Bio ClinicalBERT	108M	Biobert (18B words)+MIMIC notes (880M words)
◊RoBERTa-base	124M	BookCorpus, English Wikipedia, CC-News, OpenWebText, Stories (160GB)
◊RoBERTa-large	355M	BookCorpus, English Wikipedia, CC-News, OpenWebText, Stories (160GB)
◊BERT-base	109M	BookCorpus, English Wikipedia (16GB)
◊BERT-large	334M	BookCorpus, English Wikipedia (16GB)
◊ALBERT-base	12M	BookCorpus, English Wikipedia (16GB)
◊ALBERT-base	18M	BookCorpus, English Wikipedia (16GB)
◊DistilBERT	12M	BookCorpus, English Wikipedia (16GB)

Table 2: Models, number of parameters, and their training data. # Parameters were taken directly from the Huggingface implementation. GB/# words are taken from the authors’ reports; N/A=no information regarding the training data has been provided by the authors. In blue (*) we have models that were only trained on biomedical text. In green (†) we have models that were trained on generic English text and fine-tuned on biomedical text. In red (◊) we have models that were only trained on generic English text.

T-REx (Template-free) While the T-REx dataset from [Petroni et al. \(2019\)](#) contains 1.3M templates from Wikidata, **each prompt is automatically aligned by the creators of the dataset to natural text** from Wikipedia that supports it and which we use for template-free probing.

ConceptNet The ConceptNet dataset ([Petroni et al., 2019](#)) contains 29.8K natural prompts from Open Mind Common Sense, covering 16 relations.

SQuAD The SQuAD dataset from [Petroni et al. \(2019\)](#) contains 305 template-free prompts from the SQuAD dataset ([Rajpurkar et al., 2016](#)): the authors select a subset of 305 context-insensitive questions from the SQuAD validation set and manually modify them to be a cloze-style question.

LIPID We further experiment with our novel template-free biomedical Probing Dataset (LIPID), composed of 88,666 template-free prompts from PubMed abstracts which we split into two datasets: chemicals and genes. The chemical portion contains 46,827 chemical-related prompts and 1870 unique chemical entities, where the gene portion contains 41,839 gene-related prompts and 2591 unique gene entities. While entity-centric cloze-style QA datasets have previously been proposed for biomedicine, such as BioRead ([Pappas et al., 2018](#)) and BioMRC ([Pappas et al., 2020](#)), *we create the first template-free dataset for biomedical probing, which allows us to compare the effect of the two probing approaches in 2 different domains.* Furthermore, to encourage more research on template-free

probing, we propose an approach to develop such domain-specific datasets composed of four steps.

3.2.3 The Creation of LIPID

We create LIPID, a template-free dataset for probing models with prompts from the biomedical domain, via four steps, which we describe below: 1) retrieving a collection of biomedical text, 2) using a list of biomedical entities to select sentences, 3) filtering the resulting sentences using a list of keywords, and 4) entity masking. The creation of our dataset *takes about a day*, which consists of automatically downloading six months worth of PubMed publications, parsing, filtering, and masking.

Biomedical Text Retrieval It is important to ensure that the LMs were not trained on the test data used for probing. For that, we choose to use PubMed abstracts³ which were submitted after December 2021, which is the publication date of the most recent Biomedical LM we will probe. Such separation between the dates ensures that our questions and contexts which are used to prompt the LMs for knowledge are **entirely unseen** to all our models during training.

Biomedical Entities We use a list of 1870 unique chemicals and 2591 unique genes taken from the ChemDNER dataset ([Krallinger et al., 2015](#)) and retrieve sentences that include exactly one of those entities.

³<https://ftp.ncbi.nlm.nih.gov/pubmed/updatefiles/>

Quality Control Since we care about sentences that are facts (e.g., “Penicillin is used to treat certain infections”) rather than hypotheses, suppositions, or various other sentence forms (e.g., “We examine the effect penicillin has on infections”), we filter the resulting sentences from the previous step using a simple list of keywords we create. For example, we remove all sentences that contain parentheses, as often the entity will precede its short notation (e.g., “penicillin (PCN)”) which will most likely reveal to the model the identity of the masked entity. The simple list of keywords is: “here”, “we”, “investigate”, “study”, “propose”, “outline”, “(”, “our”, “performed”, “suggest”, and “However.” Finally, two annotators – one of which is a medical expert and the second is a CS PhD student – review 200 random prompts and evaluate the number of non-factual statements. For the chemical portion of the data the average is 92.5, and for the gene portion the average is 96.0.

Masking Lastly, we mask the entity in each sentence with a masking token. For example, the sentence “Penicillin is used to treat infections.” becomes “[MASK] is used to treat infections.”

LIPID Statistics Example prompts from our datasets are as follows. From the chemical portion: “A key to longevity assurance is the nutrient-sensing [MASK] pathway”. From the gene portion: “Amyloid- β is a product of the processing of the amyloid precursor protein, encoded by the [MASK] gene on chromosome 21”. On average, each chemical and gene entity appears 25.04 and 16.14 times, respectively, with standard deviations of 91.85 and 57.58. The maximum number of times any chemical or gene entity appears is 1669 and 1541 times, respectively. The minimum number of times any entity appears is 1 for both chemicals and genes.

4 Results

Tables 3 and 4 show our main results on the template-based and template-free datasets.

Model Rankings In both the template-based and template-free biomedical datasets – CTD, Biomed-Wikidata, and our novel LIPID datasets – PubMedBERT performs best, followed by Bioformer and BioM-ELECTRA: both techniques clearly separate models that were trained solely on biomedical data (in blue) as opposed to those who were

fine-tuned on biomedicine (in green) or general domain (in red). In comparison, on general-domain datasets, BERT-large performs best, followed by either BERT-base or DistilBERT. Furthermore, the top-5 general-domain models are roughly the same across all general-domain datasets. However, as the rank increases, the pattern is less obvious and there is no clear separation between models in blue and those in green, e.g., while PubMedBERT is the 6th best, followed by Bioformer as 7th or 8th, its rank changes to 15 (i.e., second to last) on Google-RE template-based. This is especially surprising as on the same Google-RE dataset, but in the template-free setting, Bioformer ranks 7th. Similarly, large changes in ranking can be seen for RoBERTa-base, which moves from rank 12 for template-free to 7 in the template-based Google-RE. This is also visible for the T-REx dataset, where RoBERTa-large moves from rank 7 in the template-free to 12 in the template-based setting. Similar ranking differences between datasets also appear in general models that are fine-tuned on biomedical data (in green). Notably, model rankings change between two general-domain datasets (e.g., Google-RE and T-REx), between domain-specific datasets (e.g., CTD and Biomed-Wikidata), and between both the template-free and template-based approaches (e.g., both Google-RE and T-REx settings).

Model Scores Since both Google-RE and T-REx are composed of parallel template-free and template-based datasets in which each template has a corresponding template-free text, see Table 1, we can directly compare models’ scores across them.

We find substantial different scores between the template-free and template-based datasets. For example, the average Acc@1 on the template-free datasets Google-RE and T-REx are 0.094 and 0.21, respectively. These scores change to 0.025 and 0.11 when the dataset is converted to template-based.

We see the largest performance difference in BERT-large, which obtains a score of 0.72 Acc@1 on template-free T-REx, but a score of 0.3 Acc@1 on the corresponding template-based data.

While T-REx and Google-RE are the only parallel datasets we have, allowing us to directly compare between the datasets, we can also see that the scores are different in general between template-based and template-free datasets: e.g., the average Acc@1 on the CTD and Biomed-Wikidata are 0.002 and 0.011, while on our LIPID datasets of the same biomedical domain, the average Acc@1

are 0.12 on genes and 0.18 on chemicals.

Another strange model behavior we see for the template-based datasets is the effect of model size: larger models generally perform better than their smaller counterparts. This can be seen across all base and large models on all template-free datasets. However, this is not the case in the template-based datasets: e.g., on CTD and Biomed-Wikidata, ALBERT-base outperforms its larger counterpart, and, on Google-RE, RoBERTa-base outperforms RoBERTa-large. We can further see this pattern with BERT on the Biomed-Wikidata dataset, where on T-REx and Google-RE the BERT-base version performs as well as BERT-large.

5 Discussion and Analysis

We now discuss and investigate why the aforementioned differences in scores and rankings occur between datasets and the two probing approaches.

5.1 Vocabulary

One obvious reason for the difference in scores between the datasets is the vocabulary of the models. For example, models trained on PubMed full text articles (e.g., PubMedBert and Bioformer) might have very uncommon chemical and gene names in their vocabularies in comparison to models trained on generic English text. This may result in biomedical models scoring lower on general domain datasets, and vice versa. However, framing both probing methods as the entity ranking method described in Section 3 by averaging log probabilities of the individual tokens of each entity should mitigate this effect. And while the vocabulary may have some effect on models' scores, it is important to note that on the two parallel datasets – Google-RE and T-REx, all models stay the same and the only different variable is the probing method. However, the ranking still changes.

5.2 Reused Templates

Another possible reason for the difference between the probing techniques is the variability of the templates. As experts are often required to create templates, the number of different templates is low. For example, the Google-RE dataset is composed of five templates, but Petroni et al. (2019) only use three of these. This may skew results, as, e.g., models may score lower on templates they are not familiar with.

To further analyze this, we evaluate the average Acc@10 over all models on each of the three templates Petroni et al. (2019) use from the Google-RE template-based dataset. We find that the score on the template "ENTITY (born [MASK])" is 8.9, the score on the template "ENTITY was born in [MASK]" 0.0, and the score on the template "ENTITY died in [MASK]" is 0.03. This highlights that models do in fact struggle with some templates more than with others.

5.3 Different Data Distribution

The data distribution of the templates differs from real-world text which LMs are often trained on. For example, the template-based prompts in Table 1 are not full sentences, which standard LMs are trained on. This by itself may result in skewed results, but, in general, models trained on the masked language modeling task may be getting the answer wrong because they are unfamiliar with such text structure, rather than lacking knowledge of the domain itself (See Section 5.2).

5.4 Overconfident Models

We note from Table 4 in Sung et al. (2021)'s paper that models often make the same predictions for a given template, even when the entities change (e.g., on the CTD dataset ESR1 is always first, followed by NR1I2). Manually analyzing such behavior, we find similar patterns in our various datasets. Additionally, we find that models that do not do that (e.g., Bioformer) often perform much better than those models that do. We also find that models are far more likely to generate similar answers to template-based prompts than to template-free ones. This is in line with work that show that models often rely on simple heuristics and keywords from the data for prediction (McCoy et al., 2019; Kassner and Schütze, 2020; Gururangan et al., 2018). This is a significant issue for template-based probing, as the keywords are the same for each template (e.g., "born-in"). In comparison, template-free probes often provide a unique prompt per entity. This seems neither a result of using Kassner et al. (2021)'s probing method nor a domain-specific issue, as this can be seen in both the generic domain and also the biomedical domain using the original probing technique (i.e., not entity ranking) in Sung et al. (2021)'s paper.

To quantify this behavior, we calculate the number of times each entity appears in each model's top-10 predictions. Figure 1 shows our result for

Model	Google-RE				T-REx				Biomed-Wikidata				CTD			
	A@1	A@5	A@10	R	A@1	A@5	A@10	R	A@1	A@5	A@10	R	A@1	A@5	A@10	R
*PubMedBERT	5.4e ⁻³	1.7e ⁻²	3.2e ⁻²	8	1.1e ⁻¹	2.1e ⁻¹	2.7e ⁻¹	6	4.4e ⁻²	1.3e ⁻¹	1.9e ⁻¹	1	7.8e ⁻³	2.9e ⁻²	4.5e ⁻²	1
*Bioformer	9.8e ⁻⁴	6.2e ⁻³	1.3e ⁻²	15	9.3e ⁻²	1.7e ⁻¹	2.1e ⁻¹	7	3.7e ⁻²	1.0e ⁻¹	1.6e ⁻¹	2	6.0e ⁻³	2.3e ⁻²	3.6e ⁻²	2
*BioM-ELECTRA	1.3e ⁻³	7.8e ⁻²	1.5e ⁻²	13	5.4e ⁻²	1.3e ⁻¹	1.8e ⁻¹	9	3.0e ⁻²	1.0e ⁻¹	1.5e ⁻¹	3	3.4e ⁻³	1.4e ⁻²	2.6e ⁻²	3
†BioMed-RoBERTa	1.2e ⁻²	3.0e ⁻²	4.3e ⁻²	6	5.1e ⁻²	1.2e ⁻¹	1.7e ⁻¹	11	2.3e ⁻³	1.4e ⁻²	3.4e ⁻²	12	3.6e ⁻⁴	3.7e ⁻³	8.4e ⁻³	6
†COVID Bert	9.8e ⁻⁴	4.5e ⁻³	1.2e ⁻²	16	2.8e ⁻²	6.8e ⁻²	1.1e ⁻¹	13	8.1e ⁻³	4.3e ⁻²	6.4e ⁻²	6	4.0e ⁻³	1.0e ⁻²	2.0e ⁻²	4
†BlueBert	1.6e ⁻³	9.9e ⁻³	1.6e ⁻²	12	3.5e ⁻²	1.0e ⁻¹	1.5e ⁻¹	14	1.3e ⁻²	5.1e ⁻²	8.7e ⁻²	4	5.4e ⁻⁴	5.4e ⁻³	9.8e ⁻³	12
†Discharge BERT	9.8e ⁻⁴	7.2e ⁻³	1.3e ⁻²	14	2.1e ⁻²	6.2e ⁻²	9.7e ⁻²	15	6.8e ⁻³	3.7e ⁻²	6.0e ⁻²	10	3.9e ⁻³	1.2e ⁻²	2.0e ⁻²	5
†PMC RoBERTa	3.2e ⁻³	1.7e ⁻²	3.5e ⁻²	10	4.0e ⁻²	9.4e ⁻²	1.3e ⁻¹	12	2.0e ⁻³	1.7e ⁻²	3.0e ⁻²	14	8.1e ⁻⁴	3.7e ⁻³	8.5e ⁻³	11
†Bio ClinicalBERT	1.9e ⁻³	5.0e ⁻³	1.0e ⁻²	11	1.2e ⁻²	4.1e ⁻²	6.6e ⁻²	16	7.8e ⁻³	3.6e ⁻²	6.2e ⁻²	7	1.9e ⁻³	1.0e ⁻²	1.5e ⁻²	7
◊RoBERTa-base	6.3e ⁻³	2.5e ⁻²	4.9e ⁻²	7	5.3e ⁻²	9.4e ⁻²	1.2e ⁻¹	9	1.3e ⁻³	2.2e ⁻²	3.6e ⁻²	15	3.6e ⁻⁴	4.4e ⁻³	7.9e ⁻³	16
◊RoBERTa-large	4.2e ⁻³	2.3e ⁻²	4.2e ⁻²	9	5.6e ⁻²	1.1e ⁻¹	1.5e ⁻¹	8	2.0e ⁻³	2.0e ⁻²	3.7e ⁻²	13	5.4e ⁻⁴	3.7e ⁻³	7.4e ⁻³	13
◊BERT-base	1.1e ⁻¹	2.3e ⁻¹	3.2e ⁻¹	2	3.0e ⁻¹	5.3e ⁻¹	6.4e ⁻¹	2	7.8e ⁻³	3.4e ⁻²	6.0e ⁻²	8	8.1e ⁻⁴	3.9e ⁻³	9.6e ⁻³	10
◊BERT-large	1.1e ⁻¹	2.4e ⁻¹	3.3e ⁻¹	1	3.0e ⁻¹	5.3e ⁻¹	6.5e ⁻¹	1	7.5e ⁻³	3.9e ⁻²	5.2e ⁻²	9	1.3e ⁻³	5.3e ⁻³	1.1e ⁻²	8
◊ALBERT-base	2.2e ⁻²	7.4e ⁻²	1.2e ⁻¹	5	1.4e ⁻¹	3.0e ⁻¹	4.2e ⁻¹	5	3.0e ⁻³	3.3e ⁻²	5.6e ⁻²	11	4.5e ⁻⁴	3.8e ⁻³	1.0e ⁻²	15
◊ALBERT-large	2.6e ⁻²	8.4e ⁻²	1.2e ⁻¹	4	2.1e ⁻¹	4.0e ⁻¹	5.2e ⁻¹	4	1.0e ⁻³	1.9e ⁻²	3.5e ⁻²	16	4.5e ⁻⁴	2.2e ⁻³	6.6e ⁻³	14
◊DistilBERT	1.0e ⁻¹	2.1e ⁻¹	2.9e ⁻¹	3	2.7e ⁻¹	5.2e ⁻¹	6.4e ⁻¹	3	1.1e ⁻²	4.2e ⁻²	6.8e ⁻²	5	9.0e ⁻⁴	4.7e ⁻³	1.1e ⁻²	9
Average	2.5e ⁻²	6.6e ⁻²	9.1e ⁻²	1.1e ⁻¹	2.1e ⁻¹	2.8e ⁻¹	1.1e ⁻¹	1.1e ⁻²	4.6e ⁻²	7.3e ⁻²	2.0e ⁻³	8.6e ⁻³	1.5e ⁻²			

Table 3: **Template-based** results. We report Acc@1/Acc@5/Acc@10 of each model and the macro average, the ranking of it based on its Acc@1 score (or Acc@5/10 if there is a tie) for each dataset column “(R)”. A=Acc.

Model	Google-RE				ConceptNet				SQuAD				T-REx				LIPID-Gene			LIPID-Chem				
	A@1	A@5	A@10	R	A@1	A@5	A@10	R	A@1	A@5	A@10	R	A@1	A@5	A@10	R	A@1	A@5	A@10	R				
*PubMedBERT	1.0e ⁻¹	2.3e ⁻¹	2.9e ⁻¹	6	8.9e ⁻²	1.7e ⁻¹	2.1e ⁻¹	6	1.8e ⁻¹	3.7e ⁻¹	4.5e ⁻¹	6	1.9e ⁻¹	3.4e ⁻¹	4.0e ⁻¹	6	4.3e ⁻¹	5.7e ⁻¹	6.1e ⁻¹	1	4.8e ⁻¹	6.4e ⁻¹	6.9e ⁻¹	1
*Bioformer	6.3e ⁻²	1.5e ⁻¹	1.9e ⁻¹	7	5.2e ⁻²	1.2e ⁻¹	1.5e ⁻¹	8	1.2e ⁻¹	2.7e ⁻¹	3.4e ⁻¹	8	1.4e ⁻¹	2.7e ⁻¹	3.1e ⁻¹	8	3.8e ⁻¹	5.2e ⁻¹	5.7e ⁻¹	2	4.3e ⁻¹	6.1e ⁻¹	6.6e ⁻¹	2
*BioM-ELECTRA	5.1e ⁻²	1.0e ⁻¹	1.4e ⁻¹	8	5.1e ⁻²	1.1e ⁻¹	1.4e ⁻¹	9	1.0e ⁻¹	2.6e ⁻¹	3.2e ⁻¹	11	1.2e ⁻¹	2.3e ⁻¹	2.8e ⁻¹	11	1.3e ⁻¹	4.8e ⁻¹	5.2e ⁻¹	3	4.3e ⁻¹	6.0e ⁻¹	6.5e ⁻¹	3
†BioMed-RoBERTa	4.6e ⁻²	1.1e ⁻¹	1.6e ⁻¹	10	3.4e ⁻²	6.4e ⁻²	8.2e ⁻²	14	1.2e ⁻¹	2.6e ⁻¹	3.3e ⁻¹	9	1.2e ⁻¹	2.4e ⁻¹	3.0e ⁻¹	10	1.0e ⁻¹	8.8e ⁻¹	1.1e ⁻¹	13	1.3e ⁻¹	2.7e ⁻¹	7.8e ⁻¹	13
†COVID Bert	3.2e ⁻³	1.7e ⁻²	3.4e ⁻²	13	1.3e ⁻²	1.2e ⁻¹	1.7e ⁻¹	7	8.2e ⁻²	2.2e ⁻¹	2.7e ⁻¹	13	1.3e ⁻¹	1.9e ⁻¹	2.0e ⁻¹	13	2.0e ⁻¹	2.3e ⁻¹	2.3e ⁻¹	4	2.2e ⁻¹	3.2e ⁻¹	3.6e ⁻¹	4
†BlueBert	3.2e ⁻³	1.6e ⁻²	3.4e ⁻²	14	1.3e ⁻²	1.1e ⁻¹	1.2e ⁻¹	11	1.2e ⁻¹	1.2e ⁻¹	1.5e ⁻¹	11	1.3e ⁻¹	1.3e ⁻¹	1.3e ⁻¹	11	1.6e ⁻¹	1.6e ⁻¹	1.6e ⁻¹	1	1.2e ⁻¹	1.2e ⁻¹	2.1e ⁻¹	10
†Discharge BERT	9.4e ⁻⁴	7.7e ⁻³	1.6e ⁻²	16	1.6e ⁻²	1.1e ⁻¹	1.5e ⁻¹	10	1.5e ⁻¹	1.6e ⁻¹	2.2e ⁻¹	14	1.4e ⁻¹	1.4e ⁻¹	1.2e ⁻¹	14	1.4e ⁻¹	1.5e ⁻¹	1.7e ⁻¹	5	1.8e ⁻¹	2.6e ⁻¹	3.0e ⁻¹	7
†PMC RoBERTa	4.9e ⁻²	1.1e ⁻¹	1.8e ⁻¹	9	3.0e ⁻²	6.0e ⁻²	7.8e ⁻²	15	1.1e ⁻¹	2.7e ⁻¹	3.5e ⁻¹	10	1.0e ⁻¹	2.5e ⁻¹	3.2e ⁻¹	9	2.9e ⁻¹	5.7e ⁻¹	7.2e ⁻¹	14	1.2e ⁻¹	4.3e ⁻¹	5.7e ⁻¹	15
†Bio ClinicalBERT	1.8e ⁻²	7.9e ⁻²	1.3e ⁻¹	13	1.3e ⁻¹	8.3e ⁻²	1.1e ⁻¹	11	1.1e ⁻¹	1.1e ⁻¹	1.8e ⁻¹	16	1.3e ⁻¹	8.1e ⁻¹	1.1e ⁻¹	11	1.6e ⁻¹	1.0e ⁻¹	1.2e ⁻¹	10	1.0e ⁻¹	2.0e ⁻¹	2.3e ⁻¹	11
◊RoBERTa-base	3.2e ⁻²	9.1e ⁻²	1.4e ⁻¹	12	1.9e ⁻²	4.4e ⁻²	5.8e ⁻²	16	1.0e ⁻¹	2.0e ⁻¹	3.0e ⁻¹	12	1.9e ⁻¹	2.5e ⁻¹	2.9e ⁻¹	12	1.2e ⁻¹	3.9e ⁻¹	5.0e ⁻¹	16	1.6e ⁻¹	2.7e ⁻¹	3.6e ⁻¹	16
◊RoBERTa-large	4.5e ⁻²	1.2e ⁻¹	1.9e ⁻¹	11	1.3e ⁻¹	7.2e ⁻²	9.3e ⁻²	13	1.4e ⁻¹	3.0e ⁻¹	4.0e ⁻¹	7	1.4e ⁻¹	2.7e ⁻¹	3.3e ⁻¹	7	2.9e ⁻¹	5.6e ⁻¹	7.1e ⁻¹	15	1.5e ⁻¹	4.4e ⁻¹	5.9e ⁻¹	14
◊BERT-base	2.5e ⁻¹	4.6e ⁻¹	5.7e ⁻¹	2	1.4e ⁻¹	2.8e ⁻¹	3.4e ⁻¹	2	3.6e ⁻¹	6.9e ⁻¹	8.1e ⁻¹	3	5.2e ⁻¹	7.8e ⁻¹	8.4e ⁻¹	2	6.7e ⁻¹	9.5e ⁻¹	1.0e ⁻¹	11	1.1e ⁻¹	2.5e ⁻¹	2.8e ⁻¹	8
◊BERT-large	2.7e ⁻¹	4.8e ⁻¹	5.9e ⁻¹	1	1.7e ⁻¹	3.0e ⁻¹	3.7e ⁻¹	1	4.4e ⁻¹	7.7e ⁻¹	8.6e ⁻¹	1	5.6e ⁻¹	8.0e ⁻¹	8.6e ⁻¹	1	7.5e ⁻¹	1.0e ⁻¹	1.2e ⁻¹	9	2.0e ⁻¹	2.7e ⁻¹	3.0e ⁻¹	5
◊ALBERT-base	1.4e ⁻¹	3.2e ⁻¹	4.1e ⁻¹	5	1.1e ⁻¹	2.2e ⁻¹	2.8e ⁻¹	5	2.5e ⁻¹	5.3e ⁻¹	6.2e ⁻¹	5	3.3e ⁻¹	5.8e ⁻¹	6.8e ⁻¹	5	7.6e ⁻¹	1.3e ⁻¹	1.6e ⁻¹	8	1.2e ⁻¹	2.0e ⁻¹	2.5e ⁻¹	12
◊ALBERT-large	2.0e ⁻¹	3.9e ⁻¹	4.8e ⁻¹	4	1.4e ⁻¹	2.7e ⁻¹	3.3e ⁻¹	3	3.0e ⁻¹	5.8e ⁻¹	7.1e ⁻¹	4	4.1e ⁻¹	6.6e ⁻¹	7.5e ⁻¹	4	9.6e ⁻¹	1.6e ⁻¹	1.9e ⁻¹	6	1.4e ⁻¹	2.3e ⁻¹	2.8e ⁻¹	9
◊DistilBERT	2.5e ⁻¹	4.6e ⁻¹	5.6e ⁻¹	3	1.3e ⁻¹	2.7e ⁻¹	3.4e ⁻¹	4	3.8e ⁻¹	7.3e ⁻¹	8.2e ⁻¹	2	5.0e ⁻¹	7.6e ⁻¹	8.3e ⁻¹	3	8.8e ⁻¹	1.3e ⁻¹	1.5e ⁻¹	7	1.9e ⁻¹	2.8e ⁻¹	3.2e ⁻¹	6
Average	9.4e ⁻²	1.9e ⁻¹	2.4e ⁻¹	7.3e ⁻²	1.4e ⁻¹	1.8e ⁻¹	1.7e ⁻¹	3.6e ⁻¹	4.5e ⁻¹	2.1e ⁻¹	3.6e ⁻¹	4.2e ⁻¹	1.2e ⁻¹	1.8e ⁻¹	2.1e ⁻¹	1.8e ⁻¹	1.8e ⁻¹	2.6e ⁻¹	2.9e ⁻¹					

Table 4: **Template-free** results. We report Acc@1/Acc@5/Acc@10 of each model and the macro average, the ranking of it based on its Acc@1 score (or Acc@5/10 if there is a tie) for each dataset column “(R)”. A=Acc.

the Google-RE datasets. Two obvious things appear: 1) models that score highest on the datasets, such as ALBERT base and large, BERT base and large, and DistilBERT, predict the same entities the least on both the template-free and template-based datasets; 2) converting the template-free prompts to template-based result in an increase of the amount of times each model predicts similar entities. Notably, Bio ClinicalBERT has an increase of roughly 40% in the amount of times it predicts the top-1 entity, where Discharge BERT has an increase of roughly 30%. We can also see that the amount of unique entities each model predicts is also significantly reduced when converting to templates. For example, Albert-large goes from 0.6% to 0.24% on top-1 and Bio ClinicalBERT goes from 0.77% to 0.15% on top-10.

5.5 Pseudo-perplexity

We further measure the models’ average certainty for both probing techniques. As perplexity is undefined for masked LMs like the ones we evaluate, we follow Salazar et al. (2020)’s approach to compute a model’s pseudo-perplexity. We create t copies of a sentence, with t being the number of tokens in the sentence, and mask one token at a time. Then, we pass the token IDs per sentence to the models, and get the average negative log-likelihood for each token. Summing the above and taking an exponentiated average results in:

$$\text{PPL}(X) = \exp\left\{-\frac{1}{t} \sum_{i=1}^t \log(p_{\theta}(x_i | x \neq i))\right\}$$

where $\log(p_{\theta}(x_i | x \neq i))$ is the log-likelihood of the i th token, where $[0 \leq i \leq t]$, conditioned on the remaining tokens $x \neq i$.

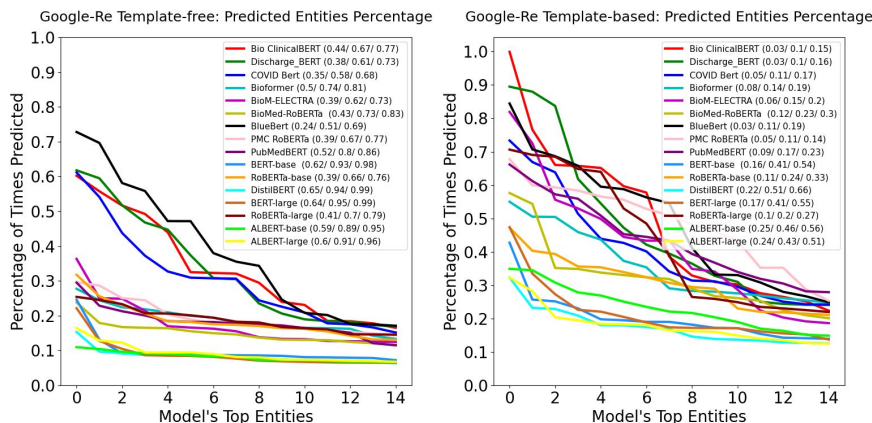


Figure 1: Template-free vs. template-based: We evaluate the percentage of times each entity appears in the top 10 predictions for each prompt. We show the results for the top 15 most frequent entities. Next to each model’s name we also add the percentage of unique entities it predicts over all prompts for top 1, 5, and 10.

Results can be seen in Figure 2 and Table 5. For template-free datasets, model perplexity decreases as accuracy increases. Surprisingly, however, we see the opposite for template-based datasets. This is unexpected, as it suggests that, as models get less certain about their answers, they perform better. A potential explanation is that models that are less certain about their answers are less likely to predict similar entities for the same template.

We also find a strange behavior for the template-based datasets regarding model size. Larger models generally perform better than smaller ones. However, we find that many times the smaller models have a lower perplexity. We only find one such occurrence – the BERT models — in the template-free ConceptNet.

5.6 Amount of Entities

While the entity ranking technique we use (see Section 3) allows us to circumvent the fact that real-world entities are often composed of more than one token, it is also susceptible to the amount of entities: as models are tasked with ranking entities, datasets with a larger number of unique entities are harder as there are more options. To quantify the effect the number of entities has on model accuracy we select three datasets – SQuAD, Biomed-Wikidata, and CTD — and increase the amount of entities as follows: for SQuAD, we add all entities from the development set (8827 unique entities in comparison to 303), for Biomed-Wikidata, we add all entities from the CTD dataset (4514 in comparison to 1267), and for the CTD dataset we add all en-




Figure 2: Template-free vs. Template-based: Average Acc@1 vs average Perplexity per model, over datasets. Template-based Pearson’s correlation coefficient: 0.83, p-value=0.16. Template-free Pearson’s correlation coefficient: 0.60, p-value=0.20.

tities from the Biomed-Wikidata dataset (4514 in comparison to 3251).

Table 6 shows our results. There is a performance drop across all models and datasets. While for the CTD and Wikidata datasets that drop is relatively low, on the SQuAD dataset the effect is more significant. This difference may be a result of the number of entities added to each dataset or the difficulty of the task. While the number of entities may have some effect on models scores, it is important to note that on the two parallel datasets – Google-RE and T-REx, the number of entities stays the same and the only different variable is the probing method. However, the ranking still changes.

Model	Google-RE (TB)	Google-RE (TF)	T-REX (TB)	T-REX (TF)	ConceptNet (TF)	SQuAD (TF)	LIPID-Gene (TF)	LIPID-Chem (TF)	Biomed-Wikidata (TB)	CTD (TB)
*PubMedBERT	76.80	23.75	255.02	46.86	129.59	90.35	10.39	13.65	10.80	50.16
*Bioformer	121.00	36.19	401.90	67.22	201.22	125.88	9.53	11.62	15.94	76.44
*BioM-ELECTRA	46.25	18.79	152.26	35.41	130.00	51.28	12.24	27.12	8.46	21.65
†BioMed-RoBERTa	172.12	25.55	335.73	20.93	343.86	73.18	8.56	10.46	18.17	26.32
†COVID Bert	84.62	25.65	212.26	43.26	247.78	88.75	5.16	5.72	5.81	11.19
†BlueBert	2345.89	719.22	2060.30	6360.65	971.47	964.89	40.15	45.83	18.67	65.21
†Discharge BERT	497.91	124.49	634.86	118.29	462.68	306.32	13.44	14.67	13.72	33.60
†PMC RoBERTa	53.79	7.68	66.05	4.87	81.76	16.05	6.32	6.62	12.04	22.13
†Bio ClinicalBERT	451.46	159.26	662.10	155.75	452.21	345.26	19.02	20.71	21.33	57.18
◊RoBERTa-base	63.99	9.82	105.80	8.11	150.24	24.34	11.26	11.62	19.05	42.09
◊RoBERTa-large	51.81	7.63	63.82	5.00	75.84	16.22	7.21	7.48	13.32	27.02
◊BERT-base	24.51	7.36	35.57	7.44	126.63	25.23	14.43	15.90	7.09	17.61
◊BERT-large	25.14	6.83	36.23	6.06	187.94	23.44	13.24	13.95	6.02	14.99
◊ALBERT-base	423.59	71.01	878.23	90.15	3720.06	163.88	197.73	193.15	13.32	365.75
◊ALBERT-large	260.91	27.51	406.03	33.37	235.64	120.33	85.15	86.55	95.03	380.45
◊DistilBERT	25.96	8.77	33.50	10.40	115.47	31.69	15.22	15.64	9.75	26.99
Average	295.36	79.97	396.23	43.54 ⁴	477.02	154.19	29.31	31.29	24.66	77.42

Table 5: **Perplexity** results. We report average perplexity for each model in addition to the average on each dataset. TF=Template-free. TB=Template-based.

Model	SQuAD Original			SQuAD Modified			Biomed-Wikidata Original			Biomed-Wikidata Modified			CTD Original			CTD Modified		
	A@1	A@5	A@10	A@1	A@5	A@10	A@1	A@5	A@10	A@1	A@5	A@10	A@1	A@5	A@10	A@1	A@5	A@10
*PubMedBERT	1.8e ⁻¹	3.7e ⁻¹	4.5e ⁻¹	9.9e ⁻²	2.2e ⁻¹	2.7e ⁻¹	4.4e ⁻²	1.3e ⁻¹	1.9e ⁻¹	4.2e ⁻²	1.0e ⁻¹	1.6e ⁻¹	7.8e ⁻³	2.9e ⁻²	4.5e ⁻²	7.3e ⁻³	2.2e ⁻²	3.7e ⁻²
*Bioformer	1.2e ⁻¹	2.7e ⁻¹	3.4e ⁻¹	6.2e ⁻²	1.4e ⁻¹	1.7e ⁻¹	3.7e ⁻²	1.0e ⁻¹	1.6e ⁻¹	3.7e ⁻²	1.0e ⁻¹	1.6e ⁻¹	6.0e ⁻³	2.3e ⁻²	3.6e ⁻²	3.4e ⁻³	1.6e ⁻²	2.7e ⁻²
*BioM-ELECTRA	1.0e ⁻¹	2.6e ⁻¹	3.2e ⁻¹	3.9e ⁻²	8.2e ⁻²	1.2e ⁻¹	3.0e ⁻²	1.0e ⁻¹	1.5e ⁻¹	3.0e ⁻²	8.3e ⁻²	1.2e ⁻¹	3.4e ⁻³	1.4e ⁻²	2.6e ⁻²	2.9e ⁻³	1.1e ⁻²	1.9e ⁻²
†BioMed-RoBERTa	1.2e ⁻¹	2.6e ⁻¹	3.3e ⁻¹	6.6e ⁻³	9.9e ⁻³	9.9e ⁻³	2.3e ⁻³	1.4e ⁻²	3.4e ⁻²	2.3e ⁻³	1.4e ⁻²	3.4e ⁻²	3.6e ⁻⁴	3.7e ⁻³	8.4e ⁻³	3.6e ⁻⁴	1.7e ⁻³	3.9e ⁻³
†COVID Bert	8.2e ⁻²	2.2e ⁻¹	2.7e ⁻¹	3.6e ⁻²	7.6e ⁻²	1.0e ⁻¹	8.1e ⁻³	4.3e ⁻²	6.4e ⁻²	7.1e ⁻³	3.2e ⁻²	4.4e ⁻²	4.0e ⁻³	1.0e ⁻²	2.0e ⁻²	2.5e ⁻³	8.8e ⁻³	1.4e ⁻²
†BlueBert	4.9e ⁻²	1.5e ⁻¹	2.2e ⁻¹	2.3e ⁻²	3.3e ⁻²	4.3e ⁻²	1.3e ⁻²	5.1e ⁻²	8.7e ⁻²	1.2e ⁻²	4.1e ⁻²	6.6e ⁻²	5.4e ⁻⁴	5.4e ⁻³	9.8e ⁻³	3.6e ⁻⁴	2.8e ⁻³	5.5e ⁻³
†Discharge BERT	5.9e ⁻²	1.6e ⁻¹	2.2e ⁻¹	2.3e ⁻²	4.6e ⁻²	4.6e ⁻²	6.8e ⁻³	3.7e ⁻²	6.0e ⁻²	6.8e ⁻³	2.9e ⁻²	4.8e ⁻²	3.9e ⁻³	1.2e ⁻²	2.0e ⁻²	2.4e ⁻³	4.9e ⁻³	1.0e ⁻²
†PMC RoBERTa	1.1e ⁻¹	2.7e ⁻¹	3.5e ⁻¹	0.0	6.6e ⁻³	9.9e ⁻³	2.0e ⁻³	1.7e ⁻²	3.0e ⁻²	2.0e ⁻³	1.6e ⁻²	2.9e ⁻²	8.1e ⁻⁴	3.7e ⁻³	8.5e ⁻³	4.5e ⁻⁴	2.5e ⁻³	4.5e ⁻³
†Bio ClinicalBERT	3.6e ⁻²	1.1e ⁻¹	1.8e ⁻¹	9.9e ⁻³	2.9e ⁻²	4.6e ⁻²	7.8e ⁻³	3.6e ⁻²	6.2e ⁻²	7.8e ⁻³	2.6e ⁻²	4.5e ⁻²	1.9e ⁻³	1.0e ⁻²	1.5e ⁻²	1.8e ⁻³	7.7e ⁻³	1.0e ⁻²
◊RoBERTa-base	1.0e ⁻¹	2.0e ⁻¹	3.0e ⁻¹	0.0	0.0	0.0	1.3e ⁻³	2.2e ⁻²	3.6e ⁻²	1.3e ⁻³	1.9e ⁻²	3.1e ⁻²	3.6e ⁻⁴	4.4e ⁻³	7.9e ⁻³	3.6e ⁻⁴	3.3e ⁻³	5.8e ⁻³
◊RoBERTa-large	1.4e ⁻¹	3.0e ⁻¹	4.0e ⁻¹	0.0	6.6e ⁻³	6.6e ⁻³	2.0e ⁻³	2.0e ⁻²	3.7e ⁻²	2.0e ⁻³	1.9e ⁻²	3.5e ⁻²	5.4e ⁻⁴	3.7e ⁻³	7.4e ⁻³	0.0	1.6e ⁻³	3.7e ⁻³
◊BERT-base	3.6e ⁻¹	6.9e ⁻¹	8.1e ⁻¹	2.2e ⁻¹	4.2e ⁻¹	5.0e ⁻¹	7.8e ⁻³	3.4e ⁻²	6.0e ⁻²	7.5e ⁻³	2.3e ⁻²	4.0e ⁻²	8.1e ⁻⁴	3.9e ⁻³	9.6e ⁻³	0.0	1.5e ⁻³	5.0e ⁻³
◊BERT-large	4.4e ⁻¹	7.7e ⁻¹	8.6e ⁻¹	2.7e ⁻¹	5.3e ⁻¹	6.1e ⁻¹	7.5e ⁻³	3.9e ⁻²	5.2e ⁻²	6.8e ⁻³	2.6e ⁻²	4.0e ⁻²	1.3e ⁻³	5.3e ⁻³	1.0e ⁻²	1.8e ⁻⁴	2.3e ⁻³	6.0e ⁻³
◊ALBERT-base	2.5e ⁻¹	5.3e ⁻¹	6.2e ⁻¹	1.6e ⁻¹	3.6e ⁻¹	4.5e ⁻¹	3.0e ⁻³	3.3e ⁻²	5.6e ⁻²	2.7e ⁻³	1.8e ⁻²	3.2e ⁻²	4.5e ⁻⁴	3.8e ⁻³	1.0e ⁻²	3.6e ⁻⁴	1.0e ⁻³	5.0e ⁻³
◊ALBERT-large	3.0e ⁻¹	5.8e ⁻¹	7.1e ⁻¹	2.1e ⁻¹	4.1e ⁻¹	5.0e ⁻¹	1.0e ⁻³	1.9e ⁻²	3.5e ⁻²	1.0e ⁻³	1.0e ⁻³	2.1e ⁻²	4.5e ⁻⁴	2.2e ⁻³	6.6e ⁻³	1.8e ⁻⁴	1.6e ⁻³	2.9e ⁻³
◊DistilBERT	3.8e ⁻¹	7.3e ⁻¹	8.2e ⁻¹	2.4e ⁻¹	4.9e ⁻¹	5.6e ⁻¹	1.1e ⁻²	4.2e ⁻²	6.8e ⁻²	1.0e ⁻²	2.7e ⁻²	5.4e ⁻²	9.0e ⁻⁴	4.7e ⁻³	1.1e ⁻²	8.1e ⁻⁴	2.2e ⁻³	4.7e ⁻³
Average	1.7e ⁻¹	3.6e ⁻¹	4.5e ⁻¹	8.7e ⁻²	1.7e ⁻¹	2.1e ⁻¹	1.1e ⁻²	4.6e ⁻²	7.3e ⁻²	1.1e ⁻²	3.5e ⁻²	5.9e ⁻²	2.0e ⁻³	8.6e ⁻³	1.5e ⁻²	1.4e ⁻³	5.6e ⁻³	1.0e ⁻²

Table 6: Accuracy (A@k) of our models for different numbers of entities.

6 Which Method Should We Use?

The obvious question is: *Which probing technique should we use: template-based or template-free?*

As described in Petroni et al. (2019), the cloze-task measures the lower bound for what LMs know. From that regard, we find that the template-free approach results in a higher lower bound of knowledge, and hence, we conclude that a better method to evaluate the *amount of model knowledge* is the template-free approach. From a cost perspective, it is also much cheaper to develop template-free datasets, as they do not require domain experts.

Lastly, our analyses suggest that the two techniques may evaluate *different kinds of knowledge*. It is, e.g., unclear why smaller models often have better performance (e.g., on perplexity or Acc@K) than their larger counterparts in the template-based approach, but almost always lower performance

using the template-free approach. This suggests that it is best to use multiple probing methods to assess the factual information these models contain.

7 Conclusion

Our study demonstrates that the choice of probing technique – template-based vs. template-free – significantly affects the assessment of LMs knowledge. Using "fill-in-the-blank" cloze statements, we compare 16 LMs across 10 probing datasets, and find substantial disparities between the two approaches. We further propose a method to create template-free domain-specific datasets and use it to develop LIPID: the first template-free biomedical probing dataset, making it possible to compare the effect of the two probing techniques in different domains. Our findings emphasize the necessity of employing multiple evaluation methods to obtain a comprehensive understanding of LMs' knowledge.

Limitations While we try to be extensive in our analysis of the potential reasons for the differences between template-base and template-free probing, other reasons may exist. That being said, we find many patterns that highlight that existing approaches are significantly different in the amount of knowledge they are capable of extracting which opens avenues for future work. Additionally, a lot of research these days is focusing on extremely large models with billions of parameters. While we expect our results to generalize as we evaluate a significant number of systems, more work can be done on evaluating template-free and template-based probing of such extremely large models.

Ethics Statement The motivation for this paper is to develop a better evaluation strategy of the amount of LMs’ parametric knowledge. We believe that it is crucial that future work continues to evaluate and improve models’ factual knowledge in biomedicine as well as other domains.

Acknowledgments

We thank the reviewers for their comments and great suggestions. The authors acknowledge financial support from NIH grants OT2TR003422 and R01LM013400.

References

- Sultan Alrowili and Vijay Shanker. 2021. [BioM-transformers: Building large biomedical language models with BERT, ALBERT and ELECTRA](#). In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 221–227, Online. Association for Computational Linguistics.
- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Zied Bouraoui, Jose Camacho-Collados, and Steven Schockaert. 2019. [Inducing relational knowledge from bert](#).
- Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu, and Le Sun. 2022. [Can prompt probe pretrained language models? understanding the invisible risks from a causal view](#).
- 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.

- Li Fang, Qingyu Chen, Chih-Hsuan Wei, Zhiyong Lu, and Kai Wang. 2023. [Bioformer: an efficient transformer language model for biomedical text mining](#).
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. [Domain-specific language model pretraining for biomedical natural language processing](#).
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of ACL*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Benjamin Heinzerling and Kentaro Inui. 2021. [Language models as knowledge bases: On entity representations, storage capacity, and paraphrased queries](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1772–1791, Online. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2019. [How can we know what language models know?](#)
- 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.
- Nora Kassner, Philipp Dufter, and Hinrich Schütze. 2021. [Multilingual LAMA: Investigating knowledge in multilingual pretrained language models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3250–3258, Online. 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.

- Martin Krallinger, Obdulia Rabal, Florian Leitner, Miguel Vazquez, David Salgado, Zhiyong Lu, Robert Leaman, Yanan Lu, Donghong Ji, Daniel M Lowe, Roger A Sayle, Riza Theresa Batista-Navarro, Rafal Rak, Torsten Huber, Tim Rocktäschel, Sérgio Matos, David Campos, Buzhou Tang, Hua Xu, Tsendsuren Munkhdalai, Keun Ho Ryu, SV Ramanan, Senthil Nathan, Slavko Žitnik, Marko Bajec, Lutz Weber, Matthias Irmer, Saber A Akhondi, Jan A Kors, Shuo Xu, Xin An, Utpal Kumar Sikdar, Asif Ekbal, Masaharu Yoshioka, Thaer M Dieb, Miji Choi, Karin Verspoor, Madian Khabsa, C Lee Giles, Hongfang Liu, Komandur Elayavilli Ravikumar, Andre Lamurias, Francisco M Couto, Hong-Jie Dai, Richard Tzong-Han Tsai, Caglar Ata, Tolga Can, Anabel Usié, Rui Alves, Isabel Segura-Bedmar, Paloma Martínez, Julen Oyarzabal, and Alfonso Valencia. 2015. **The CHEMDNER corpus of chemicals and drugs and its annotation principles.** *Journal of Cheminformatics*, 7(S1).
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. **Albert: A lite bert for self-supervised learning of language representations.**
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach.** *ArXiv*, abs/1907.11692.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. **Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Zaiqiao Meng, Fangyu Liu, Ehsan Shareghi, Yixuan Su, Charlotte Collins, and Nigel Collier. 2021. **Rewire-then-probe: A contrastive recipe for probing biomedical knowledge of pre-trained language models.**
- Dimitris Pappas, Ion Androutsopoulos, and Haris Papatheodorou. 2018. **BioRead: A new dataset for biomedical reading comprehension.** In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Dimitris Pappas, Petros Stavropoulos, Ion Androutsopoulos, and Ryan McDonald. 2020. **BioMRC: A dataset for biomedical machine reading comprehension.** In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 140–149, Online. Association for Computational Linguistics.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. **Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets.** *ArXiv*, abs/1906.05474.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2019. **Language models as knowledge bases?**
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. **SQuAD: 100,000+ questions for machine comprehension of text.** In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. **Masked language model scoring.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. **Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.**
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. **AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.
- Mujeen Sung, Jinhyuk Lee, Sean Yi, Minji Jeon, Sungdong Kim, and Jaewoo Kang. 2021. **Can language models be biomedical knowledge bases?** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4723–4734, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zonghai Yao, Yi Cao, Zhichao Yang, and Hong Yu. 2022. **Context variance evaluation of pretrained language models for prompt-based biomedical knowledge probing.**
- Zexuan Zhong, Dan Friedman, and Danqi Chen. 2021. **Factual probing is [MASK]: Learning vs. learning to recall.** In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5017–5033, Online. Association for Computational Linguistics.