

The Factuality of Large Language Models in the Legal Domain

Rajaa EL HAMDANI,¹ Thomas Bonald,¹ Fragkiskos Malliaros,^{2,3} Nils Holzenberger,¹ Fabian Suchanek¹

¹Télécom Paris, Institut Polytechnique de Paris

²Université Paris-Saclay, CentraleSupélec, ³INRIA

rajaa.elhamdani@telecom-paris.fr, thomas.bonald@telecom-paris.fr

fragkiskos.malliaros@centralesupelec.fr

nils.holzenberger@telecom-paris.fr, fabian.suchanek@telecom-paris.fr

Abstract

This paper investigates the factuality of large language models (LLMs) as knowledge bases in the legal domain, in a realistic usage scenario: we allow for acceptable variations in the answer, and let the model abstain from answering when uncertain. First, we design a dataset of diverse factual questions about case law and legislation. We then use the dataset to evaluate several LLMs under different evaluation methods, including exact, alias, and fuzzy matching. Our results show that the performance improves significantly under the alias and fuzzy matching methods. Further, we explore the impact of abstaining and in-context examples, finding that both strategies enhance precision. Finally, we demonstrate that additional pre-training on legal documents, as seen with SaulLM, further improves factual precision from 63% to 81%.

1 Introduction

Language models (LMs) store knowledge from pre-training documents, allowing them to be queried with natural language (NL) rather than a formal language, unlike structured knowledge bases (KBs) (Petroni et al., 2019; Chang and Bergen, 2024). This natural language interface democratizes access to knowledge, especially in domains where users may not master querying structured KBs (Hendrix et al., 1978). In the legal domain, this includes law experts (scholars, judges, lawyers, paralegals) and laymen seeking legal information (Paul et al., 2020). LMs can acquire and store more facts than KBs, as most knowledge is in NL documents. This reduces the cost of manually populating KBs, as LMs automatically store knowledge from pre-training documents without human supervision (AlKhamissi et al., 2022). While Wikidata is the largest structured KB, it lacks comprehensive domain-specific knowledge (Weikum et al., 2021). In contrast, Wikipedia articles offer this knowledge, which LMs can store through pre-training.

```
Your task is to answer legal questions. Provide a concise answer without an explanation.
If you don't know the answer, respond with "I don't know". Examples:
Question: Who wrote the majority opinion of the case "Trump v. Hawaii"?
Answer: John Roberts
Question: Who wrote the majority opinion of the case "McGirt v. Oklahoma"?
Answer: Neil Gorsuch
Question: Who wrote the majority opinion of the case "Limtiaco v. Camacho"?
Answer: Clarence Thomas
Question: Who wrote the majority opinion of the case "R. v. Morgentaler"?
Answer: Brian Dickson
Question: Who wrote the majority opinion of the case "CREW v. Trump"?
Answer: Pierre N. Leval
Question: Who wrote the majority opinion of the case "Lau v. Nichols"?
```

Figure 1: Example of prompt for the relation *majority opinion by*. The few shot examples are tailored for the relation *majority opinion by* and the subject's class *legal case*.

The concept of language models as knowledge bases (LM-as-KB) was introduced by (Petroni et al., 2019), demonstrating on the LAMA dataset that BERT retrieves facts better than traditional knowledge bases. Later research (Heinzerling and Inui, 2021) showed that large language models (LLMs) with more than 10B parameters can store all of Wikidata and retrieve facts using NL queries, though the accuracy can be sensitive to minor query phrasing changes. Additionally, fine-tuning LMs or using query variants improves their robustness (Jiang et al., 2020; Heinzerling and Inui, 2021; Adolphs et al., 2021). LAMA (Petroni et al., 2019) was the first dataset to benchmark LM-as-KB, limited to one-token facts for simplicity. Subsequent benchmarks include more complex queries encompassing various types of knowledge (Bisk et al., 2020; Peng et al., 2022; Lin et al., 2020; Misra et al., 2023; Elazar et al., 2021). Among these, KAMEL (Kalo and Fichtel, 2022) proved to be harder than LAMA because the dataset contains multi-token facts, facts with literal values, a broader range of knowledge, and relations of higher cardinality.

Early research on LM-as-KB primarily evaluated smaller models like BERT. However, the advent of LLMs like OpenAI's GPT, Google's PALM, and Meta's Llama (OpenAI et al., 2024; Chowdhery

et al., 2023; Touvron et al., 2023), revolutionized the field. These models exhibit emergent abilities useful for LM-as-KB, including few-shot prompting (Brown et al., 2020) and calibration (Wei et al., 2022). LLMs can learn new tasks through few-shot prompting without fine-tuning. Since LM-as-KB performance improves with prompt fine-tuning (Qin and Eisner, 2021), including query examples in the prompt can enhance the factuality of LMs. Additionally, well-calibrated models can predict the correctness of their own responses, enabling self-evaluation of answers to factual queries (Kadavath et al., 2022; Geng et al., 2024).

Despite their capabilities, LLMs are prone to *hallucination* (Ji et al., 2023; Suchanek and Luu, 2023), producing factually incorrect answers in a disturbingly self-confident tone. While many studies have examined the factual accuracy of LLMs (OpenAI et al., 2024; Wei et al., 2024; Muhlgay et al., 2024; Hu et al., 2024; Si et al., 2023; Min et al., 2023; Manakul et al., 2023; Mallen et al., 2023), few focused on domain-specific knowledge. Domain-specific knowledge poses unique challenges, such as specific named entities that are unique to the domain (e.g., lawyers, judges) and that do not appear in abundance on the Internet (and hence in training data); relations that may not even be known to the general public (e.g. the judge who issues the majority opinion of a case); facts that are specific to the domain (e.g., legal rules); and finally ways of phrasing that are typical for the jargon, but unusual on the Web as a whole (e.g., concurring or dissenting opinions). However, domain-specific applications have huge potential, aiding users in areas beyond one-size-fits-all products.

In this paper, we focus on factual answers in the legal domain, motivated by both the risks and opportunities of using LLMs in this field (Bomasani et al., 2021). While benchmarks like LawBench (Fei et al., 2023) and LegalBench (Guha et al., 2023) evaluate LLMs on various legal tasks, LawBench is specific to Chinese law, and LegalBench lacks a legal knowledge task, limiting their use for probing LLMs on legal facts. The legal domain is particularly sensitive to hallucinations, as incorrect information can lead to harmful decisions (Dahl et al., 2024). Previous work used strict evaluation criteria, accepting answers only if identical to the ground truth, ignoring valid response variations ("Samuel A. Alito, Jr" is as valid an answer as "Justice Alito"). Additionally, earlier

studies forced LLMs to provide answers, increasing hallucinations. In contrast, our approach allows models to abstain when unsure, reducing hallucinations. Accounting for acceptable variations in the phrasing of the answer and allowing the model to abstain from answering brings our evaluation closer to a realistic human use of an LLM as a KB.

In our work, we address four key research questions:

1. Are all LLMs equally affected by the shortcomings of exact matching-based evaluation, or are some more penalized?
2. Can an LLM abstain from generating incorrect answers?
3. Does few-shot prompting increase the factuality of LLMs?
4. Does training on legal documents improve the factuality of an LLM?

We present the methodology in Section 2, the results in Section 3, and our conclusions in Section 4.

2 Methodology

2.1 Dataset

Legal knowledge is primarily stored in legislation and legal cases (Raz, 2002). This study focuses on atomic information, such as the jurisdiction of a case or legislation. Even if most real-world queries may not be about atomic facts, these atomic facts are the fundamental building blocks of actionable legal information. Therefore, it is crucial to first ensure that the LLM is not hallucinating atomic facts before asking for the relevant legal rule to a legal dispute. We use Wikidata (Vrandečić and Krötzsch, 2014) to create a benchmark dataset of atomic legal facts on legislation and legal cases.

Manually navigating Wikidata’s taxonomy to identify relevant types is impractical due to its complexity (Suchanek et al., 2024). Instead, we leverage the Wikipedia Category Graph (WCG), which organizes articles by categories to facilitate navigation. We identify legal articles within two hierarchical steps of the “Law” category in the WCG and retrieve their corresponding Wikidata items. From these items, we construct a legal taxonomy and query Wikidata’s SPARQL endpoint for more relevant items. We refine the dataset by manually selecting relations specific to the legal domain and

Relation	count
applies to jurisdiction	3,155
majority opinion by	1,704
legislated by	1,345
signatory	817
language of work or name	730
amended by	518
plaintiff	342
defendant	309
Total	8,920

Table 1: Dataset statistics.

augmenting it with facts from Wikipedia infoboxes (Lehmann et al., 2015). For evaluation, we manually create question templates for each relation to query LLMs via natural language, resulting in 8,920 question-answer pairs covering various relations (see Table 1).¹

2.2 Models

Large models are costly to deploy on-premise. We argue that most, if not all, legal tech applications involve sensitive personal data that cannot be transferred to third-party APIs. Thus, we limit experiments to open-source models with fewer than 8B parameters, runnable on a local machine. We select 7 models among the highly ranked in the LMSYS Chatbot Arena Leaderboard (Chiang et al., 2024): Gemma-2B, Gemma-7B, Llama-2-7B, Llama-3-8B, Mistral-7B, Phi-3-min-4k, RecurrentGemma-2B (Team et al., 2024; Botev et al., 2024; Touvron et al., 2023). We add SaulLM (Colombo et al., 2024), which is the result of further training Mistral-7B on legal corpora. We use the instruction-tuned variant of each model.

2.3 Prompt strategy

We experiment with zero-shot and few-shot prompts to query the models. For a question q , related to a subject-relation pair (s, r) , the corresponding prompt contains 5 in-context examples² depending on the relation r and the class of the subject, say class c . In-context examples are question-answer pairs about subjects of class c and relation r . We do not query the model for questions used as in-context examples to avoid data leakage. Figure

¹Code and dataset available at <https://github.com/Rajjaa/LexFact>

²We experimented with one example. For most models, one example is not enough to learn the correct formatting of the answer.

1 shows an example prompt. We also instruct the model to answer “I don’t know” if it cannot answer the question correctly.

2.4 Evaluation methods

We are evaluating the LLM as a KB. The correctness and coverage of a KB are usually evaluated by precision P and recall R (Weikum et al., 2021), respectively. For a given KB containing a set of statements S and a set of true statements T from the ground truth, they are defined by:

$$P = \frac{S \cap T}{S} \quad \text{and} \quad R = \frac{S \cap T}{T}.$$

In our setting of LLM-as-KB, T is the set of all question-answer pairs (those of Table 1), S is the set of questions answered by the LLM,³ and $S \cap T$ is the set of correct answers. Thus, the corresponding precision and recall are given by:

$$P_{LLM} = \frac{|\text{correct answers}|}{|\text{answered questions}|}$$

$$R_{LLM} = \frac{|\text{correct answers}|}{|\text{all questions}|}$$

Most articles evaluating LLMs consider an answer correct only if it exactly matches the ground truth (Dahl et al., 2024; Petroni et al., 2019; Chang et al., 2024; Hu et al., 2024; Yu et al., 2024). However, LLMs can generate a fact in multiple forms and tend to answer verbosely, often providing explanations along with the answer. This behavior penalizes LLMs when evaluated with exact matching.

Each Wikidata item has a main surface form called a label and alternative names known as *aliases*. For example, item Q30 has the label “United States of America” and aliases like “U.S.A.” and “America”. We consider the following options for declaring an answer correct:

- (EM) **Exact matching.** The answer matches the label.
- (AM) **Alias matching.** The answer matches the label or any of its aliases. wI
- (FM) **Fuzzy matching.** The answer contains the label or any of its aliases.

³The LLM might abstain from answering a question.

Fuzzy matching accounts for correct but verbose answers. It is prone to errors, however. Take, for instance, the question “What is the legislation of the case ‘Rummel v. Estelle’?”, whose correct answer is “United States”. The answer “The case ‘Rummel v. Estelle’ applies to the state of Louisiana, United States” is false but contains the label of the true answer. We manually inspect such cases and implement post-processing rules for the answers to enhance the accuracy of fuzzy matching.

3 Results

3.1 Exact matching underestimates the performance of LLM-as-KB

The performance of the models, shown in Table 2, increases significantly when evaluated with alias matching (AM) and fuzzy matching (FM), compared to exact matching (EM). This indicates that LLMs often generate correct answers in varied surface forms that are not captured by exact matching alone. For instance, SauLLM in few-shot mode and allowed to abstain increased in precision from 36% (EM) to 81% (FM). The ranking of the models changed notably depending on the evaluation method. Under exact matching and alias matching, Mistral-7B performed poorly, ranking last with 8% precision (EM). However, under fuzzy matching, Mistral-7B’s improved significantly, ranking third with a precision of 63%.

These observations highlight the limitations of exact matching-based evaluation, which can unfairly penalize models that tend to generate more contextual, rich, and verbose answers, like SauLLM and Mistral-7B. Incorporating alias and fuzzy matching provides a better understanding of each model’s true performance as a knowledge base.

3.2 LLMs can be instructed to abstain from generating incorrect answers

Table 2 reports the abstain rate under each prompting strategy, with or without the abstain instruction. Most models abstain more after adding the abstain instruction to the prompt. Abstaining increases precision by answering fewer questions. Interestingly, despite no few-shot examples with answer “I don’t know”, the models still abstain a significant fraction of the time.

Including an abstain instruction generally improves precision across different models, as summarized in Figure 2. The instruction helps the models to refrain from providing incorrect answers

Model	Prompt	Abstain	EM		AM		FM		Abstain Rate
			P	R	P	R	P	R	
SauLLM	Zero-shot	True	3.7	0.5	4.4	0.6	<u>68.3</u>	9.9	85.6
		False	0.8	0.8	2.2	2.2	60.9	60.9	0.1
	Few-shot	True	35.6	21.4	73.2	44.1	81.2	48.9	39.8
		False	30.0	30.0	57.6	57.6	62.9	62.9	0.0
Mistral	Zero-shot	True	0.1	0.0	0.1	0.1	65.3	55.6	14.9
		False	0.0	0.0	0.0	0.0	58.5	57.7	1.4
	Few-shot	True	8.1	7.6	11.1	10.4	63.0	59.4	5.8
		False	4.0	4.0	4.9	4.9	60.0	59.5	0.8
Phi-3	Zero-shot	True	0.1	0.1	0.1	0.1	62.8	61.4	2.1
		False	0.0	0.0	0.0	0.0	62.3	62.3	0.1
	Few-shot	True	34.4	22.9	49.4	32.8	62.6	41.6	33.5
		False	27.0	26.9	36.2	36.0	57.2	56.8	0.6
Gemma-7B	Zero-shot	True	1.7	1.5	2.4	2.1	34.4	30.3	11.8
		False	0.8	0.8	2.0	2.0	35.3	35.2	0.4
	Few-shot	True	25.0	17.4	45.3	31.5	65.8	45.8	30.5
		False	2.8	2.5	3.3	3.0	32.8	29.9	8.9
Gemma-2B	Zero-shot	True	1.9	1.8	4.2	4.0	30.6	29.3	4.0
		False	3.4	3.4	6.4	6.4	36.8	36.8	0.2
	Few-shot	True	21.8	17.2	48.4	38.2	52.4	41.3	21.2
		False	17.4	16.4	37.7	35.7	41.4	39.1	5.4
RGemma-2B	Zero-shot	True	10.4	10.2	12.4	12.1	40.1	39.3	2.1
		False	9.9	9.9	10.4	10.4	40.2	40.0	0.5
	Few-shot	True	33.6	29.4	45.8	40.0	50.8	44.5	12.6
		False	33.6	<u>31.9</u>	45.9	43.6	51.0	48.5	5.0
Llama-3	Zero-shot	True	14.6	13.9	20.0	19.1	52.0	49.5	4.8
		False	13.4	13.4	19.3	19.3	52.1	52.1	0.1
	Few-shot	True	32.1	31.2	58.2	56.6	62.6	60.8	2.8
		False	33.8	33.7	57.5	<u>57.3</u>	60.3	60.0	0.4
Llama-2	Zero-shot	True	4.4	4.2	6.5	6.2	57.5	55.0	4.5
		False	0.8	0.8	1.1	1.1	59.7	59.7	0.1
	Few-shot	True	18.9	3.2	40.4	6.8	47.2	7.9	83.2
		False	19.4	19.2	36.0	35.7	42.2	41.8	1.0

Table 2: Precision and recall (in %) with exact matching (EM), alias matching (AM), and fuzzy matching (FM), for the few-shot / zero-shot models and with / without the abstain instruction. Highest scores in bold, second highest underlined.

without needing access to the probability scores of their responses. This behavior is particularly important for LLMs accessible only through APIs that do not provide probability scores. However, recall tends to decrease. The intuitive expectation is that when models are allowed to abstain, precision goes up, and recall goes down. Our results confirm this intuition, highlighting the trade-off between precision and recall when implementing an abstain option. For sensitive applications like the legal domain, prioritizing precision over recall may be more desirable to avoid the severe consequences of incorrect information.

3.3 In-context examples increase the factuality of LLMs

We evaluate if in-context examples enhance the factuality of LLMs by comparing precision scores between two prompting strategies: with (few-shot) and without (zero-shot) in-context examples, as shown in Figure 2. To isolate the impact of in-context examples, we calculated precision only for questions where the model did not abstain in either setting. Our observations indicate that in-context examples improve precision for 6 out of 8 models when instructed to abstain and for 5 out of 8 models when not prompted to abstain. The improvement is




Figure 2: Comparison of precision (FM), with and without the abstain instruction, under the two prompting strategies.

more pronounced with the abstain instruction.

These findings suggest that in-context examples can enhance the precision and, thus, the factuality of LLMs. We analyzed instances where the few-shot setting produced correct answers while the zero-shot setting did not, revealing two main benefits:

1. In-context examples help the model learn the expected type and format of answers (Min et al., 2022). For instance, when asked “applies to jurisdiction”, models in the zero-shot setting often responded with names of states. However, with in-context examples, the models correctly answered “USA”. Similarly, for “legislated by”, the model predicted presidents instead of the correct session of the U.S Congress for US legislation. So, the model learns a mapping between the type of legislation and the type of the answer.
2. In-context examples correct wrong patterns learned during pre-training. For example, case titles generally follow the format “plaintiff v. defendant”. Without in-context examples, all models incorrectly reversed these roles, which was corrected with few-shot prompts. This ability to override wrong prior patterns is explained in (Wei et al., 2023).

To summarize, in-context examples significantly enhance the precision and factuality of LLMs by guiding them on the expected type and format of answers and correcting erroneous patterns from pre-training.

3.4 Training an LLM on legal documents improves its factuality

Finally, we investigate whether training an LLM on legal documents improves its factuality by comparing the performance of SauLLM, which is Mistral-7B with additional pre-training on legal documents, against the standard Mistral-7B model. The top performance was achieved by SauLLM, outperforming Mistral-7B. Under the few-shot prompt with the instruction to abstain, SauLLM achieved a precision of 81%, significantly higher than the 63% precision of Mistral-7B. This suggests that training SauLLM on legal documents enhances its ability to provide correct answers, similar to the effect of few-shot prompting. The additional training helps SauLLM understand the context and format of legal questions more accurately, leading to better performance. Thus, training an LLM on legal documents substantially improves its factuality, particularly in terms of precision. This improvement aligns with the benefits observed from few-shot prompting, indicating that domain-specific training enables the model to generate more accurate answers.

Interestingly, while SauLLM showed a significant increase in precision, its recall did not improve as markedly. This can be explained by the much higher abstain rate of SauLLM compared to Mistral-7B.

4 Conclusions

In this paper, we explored the factuality of LLMs as KBs in the legal domain. We evaluated various models, including SauLLM, pre-trained on legal documents.

Our findings reveal that the performance of LLMs improves significantly when using alias and fuzzy matching instead of exact matching. Abstain instructions and few shot prompting increase factuality. Pre-training on legal documents, as shown by SauLLM, substantially improves the precision, highlighting the importance of domain-specific pre-training.

Our careful evaluation methods reveal that LLMs hallucinate significantly less than reported in (Dahl et al., 2024), due to our consideration of the verbose nature of LLM responses. While SauLLM achieves a high precision of 81%, this is still insufficient for high-stakes legal applications, but may suffice for research and analytics. Instances of lawyers being sanctioned for using fictitious cases

underscore the risks of factual errors.

These results highlight the potential of LLMs as supplementary tools for legal research, emphasizing the importance of accurate atomic facts. Future work should refine evaluation methods and explore domain-specific training to enhance LLM reliability.

Acknowledgements

This research was supported by funding from Labex DigiCosme, which is gratefully acknowledged.

References

- Leonard Adolphs, Shehzaad Dhuliawala, and Thomas Hofmann. 2021. How to Query Language Models? <https://doi.org/10.48550/arXiv.2108.01928> arXiv:2108.01928 [cs].
- Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona Diab, and Marjan Ghazvininejad. 2022. A Review on Language Models as Knowledge Bases. <http://arxiv.org/abs/2204.06031> arXiv:2204.06031 [cs].
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, et al. 2021. On the Opportunities and Risks of Foundation Models. *ArXiv* (2021). <https://crfm.stanford.edu/assets/report.pdf>
- Aleksandar Botev, Soham De, Samuel L Smith, Anushan Fernando, George-Cristian Muraru, Ruba Haroun, Leonard Berrada, Razvan Pascanu, Pier Giuseppe Sessa, Robert Dadashi, et al. 2024. RecurrentGemma: Moving Past Transformers for Efficient Open Language Models. *arXiv preprint arXiv:2404.07839* (2024).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, and other. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 1877–1901. <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>
- Tyler A. Chang and Benjamin K. Bergen. 2024. Language Model Behavior: A Comprehensive Survey. *Computational Linguistics* (March 2024), 1–58. https://doi.org/10.1162/coli_a_00492
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. A Survey on Evaluation of Large Language Models. *ACM Transactions on Intelligent Systems and Technology* 15, 3 (June 2024), 1–45. <https://doi.org/10.1145/3641289>
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference. arXiv:2403.04132 [cs.AI]
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* (2023).
- Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and Michael Desa. 2024. SaulLM-7B: A pioneering Large Language Model for Law. <http://arxiv.org/abs/2403.03883> arXiv:2403.03883 [cs].
- Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E Ho. 2024. Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis, Volume 16, Issue 1, 2024, Pages 64–93* (2024).
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhिलाशा Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and Improving Consistency in Pretrained Language Models. *Transactions of the Association for Computational Linguistics* 9 (Dec. 2021), 1012–1031. https://doi.org/10.1162/tacl_a_00410
- Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Songyang Zhang, Kai Chen, Zongwen Shen, and Jidong Ge. 2023. LawBench: Benchmarking Legal Knowledge of Large Language Models. *CoRR* (2023).
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024. A Survey of Confidence Estimation and Calibration in Large Language Models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 6577–6595.
- Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Ré, Adam Chilton, and other. 2023. Legalbench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. *SSRN Electronic Journal* (2023). <https://doi.org/10.2139/ssrn.4583531>

- 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*, Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty (Eds.). Association for Computational Linguistics, Online, 1772–1791. <https://doi.org/10.18653/v1/2021.eacl-main.153>
- Gary G. Hendrix, Earl D. Sacerdoti, Daniel Sagalowicz, and Jonathan Slocum. 1978. Developing a natural language interface to complex data. *ACM Transactions on Database Systems* 3, 2 (June 1978), 105–147. <https://doi.org/10.1145/320251.320253>
- Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S Yu, and Zhijiang Guo. 2024. TOWARDS UNDERSTANDING FACTUAL KNOWLEDGE OF LARGE LANGUAGE MODELS. (2024).
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *Comput. Surveys* 55, 12 (Dec. 2023), 1–38. <https://doi.org/10.1145/3571730>
- 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 (July 2020), 423–438. https://doi.org/10.1162/tacl_a_00324
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, and other. 2022. Language Models (Mostly) Know What They Know. <https://doi.org/10.48550/arXiv.2207.05221> arXiv:2207.05221 [cs].
- Jan-Christoph Kalo and Leandra Fichtel. 2022. KAMEL: Knowledge Analysis with Multitoken Entities in Language Models.. In *AKBC*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Social Work* 6, 2 (Jan. 2015), 167–195. <https://doi.org/10.3233/sw-140134> MAG ID: 1552847225.
- Bill Yuchen Lin, Seyeon Lee, Rahul Khanna, and Xiang Ren. 2020. Birds have four legs?! NumerSense: Probing Numerical Commonsense Knowledge of Pre-Trained Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 6862–6868. <https://doi.org/10.18653/v1/2020.emnlp-main.557>
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 9802–9822. <https://doi.org/10.18653/v1/2023.acl-long.546>
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 9004–9017. <https://doi.org/10.18653/v1/2023.emnlp-main.557>
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 12076–12100. <https://doi.org/10.18653/v1/2023.emnlp-main.741>
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 11048–11064. <https://doi.org/10.18653/v1/2022.emnlp-main.759>
- Kanishka Misra, Julia Rayz, and Allyson Ettinger. 2023. COMPS: Conceptual Minimal Pair Sentences for testing Robust Property Knowledge and its Inheritance in Pre-trained Language Models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, Andreas Vlachos and Isabelle Augenstein (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 2928–2949. <https://doi.org/10.18653/v1/2023.eacl-main.213>
- Dor Muhlgay, Ori Ram, Inbal Magar, Yoav Levine, Nir Ratner, Yonatan Belinkov, Omri Abend, Kevin Leyton-Brown, Amnon Shashua, and Yoav Shoham. 2024. Generating Benchmarks for Factuality Evaluation of Language Models. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, Yvette Graham and Matthew Purver (Eds.). Association for Computational Linguistics, St. Julian's, Malta, 49–66. <https://aclanthology.org/2024.eacl-long.4>

- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, and other. 2024. GPT-4 Technical Report. <https://doi.org/10.48550/arXiv.2303.08774> arXiv:2303.08774 [cs].
- Boniol Paul, Panagopoulos George, Xypolopoulos Christos, Hamdani Rajaa El, Amariles David Restrepo, and Vazirgianni Michalis. 2020. Performance in the Courtroom: Automated Processing and Visualization of Appeal Court Decisions in France. In *The 26Th Acm Sigkdd International Conference On Knowledge Discovery & Data Mining (Kdd 2020)*.
- Hao Peng, Xiaozhi Wang, Shengding Hu, Hailong Jin, Lei Hou, Juanzi Li, Zhiyuan Liu, and Qun Liu. 2022. COPEN: Probing Conceptual Knowledge in Pre-trained Language Models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 5015–5035. <https://doi.org/10.18653/v1/2022.emnlp-main.335>
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2463–2473. <https://doi.org/10.18653/v1/D19-1250>
- Guanghui Qin and Jason Eisner. 2021. Learning How to Ask: Querying LMs with Mixtures of Soft Prompts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, Online, 5203–5212. <https://doi.org/10.18653/v1/2021.naacl-main.410>
- Joseph Raz. 2002. *The authority of law: essays on law and morality* (reprinted ed.). Oxford Univ. Press, Oxford.
- Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber, and Lijuan Wang. 2023. Prompting GPT-3 To Be Reliable. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=98p5x51L5af>
- Fabian Suchanek and Anh Tuan Luu. 2023. Knowledge bases and language models: Complementing forces. In *International Joint Conference on Rules and Reasoning*. Springer, 3–15.
- Fabian M. Suchanek, Mehwish Alam, Thomas Bonald, Lihu Chen, Pierre-Henri Paris, and Jules Soria. 2024. YAGO 4.5: A Large and Clean Knowledge Base with a Rich Taxonomy. In *SIGIR*.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295* (2024).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, and other. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. <https://doi.org/10.48550/arXiv.2307.09288> arXiv:2307.09288 [cs].
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of The ACM* 57, 10 (Sept. 2014), 78–85. <https://doi.org/10.1145/2629489> MAG ID: 2080133951.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent Abilities of Large Language Models. *Transactions on Machine Learning Research* (2022).
- Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. 2023. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846* (2023).
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. Long-form factuality in large language models. <http://arxiv.org/abs/2403.18802> arXiv:2403.18802 [cs].
- Gerhard Weikum, Xin Luna Dong, Simon Razniewski, Fabian Suchanek, et al. 2021. Machine knowledge: Creation and curation of comprehensive knowledge bases. *Foundations and Trends® in Databases* 10, 2-4 (2021), 108–490.
- Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, Chunyang Li, Zheyuan Zhang, Yushi Bai, Yantao Liu, Amy Xin, Kaifeng Yun, Linlu GONG, Nianyi Lin, Jianhui Chen, Zhili Wu, Yunjia Qi, Weikai Li, Yong Guan, Kaisheng Zeng, Ji Qi, Hailong Jin, Jinxin Liu, Yu Gu, Yuan Yao, Ning Ding, Lei Hou, Zhiyuan Liu, Xu Bin, Jie Tang, and Juanzi Li. 2024. KoLA: Carefully Benchmarking World Knowledge of Large Language Models. In *The Twelfth International Conference on Learning Representations*.