

# Don't Hallucinate, Abstain: Identifying LLM Knowledge Gaps via Multi-LLM Collaboration

Shangbin Feng<sup>1</sup> Weijia Shi<sup>1</sup> Yike Wang<sup>2</sup>  
Wenxuan Ding<sup>3</sup> Vidhisha Balachandran<sup>4</sup> Yulia Tsvetkov<sup>1</sup>

<sup>1</sup>University of Washington <sup>2</sup>University of California, Berkeley

<sup>3</sup>The Hong Kong University of Science and Technology <sup>4</sup>Carnegie Mellon University

shangbin@cs.washington.edu

## Abstract

Despite efforts to expand the knowledge of large language models (LLMs), knowledge gaps—missing or outdated information in LLMs—might always persist given the evolving nature of knowledge. In this work, we study approaches to identify LLM knowledge gaps and *abstain* from answering questions when knowledge gaps are present. We first adapt existing approaches to model calibration or adaptation through fine-tuning/prompting and analyze their ability to abstain from generating low-confidence outputs. Motivated by their failures in self-reflection and over-reliance on held-out sets, we propose two novel approaches that are based on model collaboration, i.e., LLMs probing other LLMs for knowledge gaps, either cooperatively or competitively. Extensive experiments with three LLMs on four QA tasks featuring diverse knowledge domains demonstrate that both cooperative and competitive approaches to unveiling LLM knowledge gaps achieve up to 19.3% improvements on abstain accuracy against the strongest baseline. Further analysis reveals that our abstention methods pinpoint failure cases in retrieval augmentation and knowledge gaps in multi-hop reasoning.<sup>1</sup>

## 1 Introduction

LLMs demonstrate impressive capabilities of encoding real-world knowledge in model parameters and leveraging it to aid knowledge-intensive tasks (Petroni et al., 2019; Brown et al., 2020; Yu et al., 2023a). But when such knowledge is missing or unreliable, they resort to hallucinations (Ji et al., 2023) and biases (Feng et al., 2023a), while still “speaking with confidence.” A growing body of work seeks to expand LLM knowledge through retrieval augmentation (Guu et al., 2020; Borgeaud et al., 2022; Khattab et al., 2022; Shi et al., 2023; Chen et al., 2023), search engine integration (Nakano

et al., 2021; Press et al., 2023), and multi-LM collaboration (Luo et al., 2023; Feng et al., 2023b). However, LLM knowledge gaps might always persist due to the ever-evolving nature of knowledge (Kandpal et al., 2023; Mallen et al., 2023; De Cao et al., 2021; Hernandez et al., 2023; Kasai et al., 2024). Consequently, we posit that *abstaining from generating low-confidence outputs should be a part of LLMs’ functionality*, and ask a crucial research question: *how to identify knowledge gaps in LLMs?* Developing and evaluating robust mechanisms to address the abstain problem improves LLM reliability, reduces hallucinations, and mitigates biases due to model uncertainty.

We hypothesize that there are three ways to operationalize the abstain problem: (1) from the *data* perspective, LLMs should abstain when knowledge is missing in training data or curated from unreliable sources; (2) from the *modeling* perspective, LLMs should abstain when knowledge is not encoded in model parameters and cannot be inferred; (3) from the *usage* perspective, LLMs should abstain when generated outputs would provide an incorrect answer. Since users are directly interacting with and impacted by LLM-generated texts, not model parameters or training data, we formulate and evaluate LLM abstention focusing on the factual correctness of LLM generations.

Specifically, we focus on **abstaining in question-answering** (AbstainQA): given an LLM and a question  $q$ , we aim to develop an abstain function  $f(\text{LLM}, q) \rightarrow \{true, false\}$  to indicate whether the LLMs should abstain based on the limitations of their internal knowledge (Kamath et al., 2020; Jiang et al., 2021; Whitehead et al., 2022). To facilitate a comprehensive understanding of existing approaches that can be employed for LLM abstention, we first identify 11 methods from previous works and adapt them to incorporate the abstain functionality. We categorize these baselines into (1) *calibration-based*, such as employing tempera-

<sup>1</sup>Code and data will be publicly available at <https://github.com/BunsenFeng/AbstainQA>.

ture scaling (Guo et al., 2017; Jiang et al., 2021) to calibrate models and set a confidence score threshold for abstention; (2) *training-based*, such as employing linear probing on LLMs’ hidden layers to evaluate the veracity of generated texts (Slobodkin et al., 2023; Azaria and Mitchell, 2023); (3) *prompting-based*, such as prompting LLMs to request more information before answering questions (Feng et al., 2023b); and (4) *self-consistency based*, such as generating multiple chains-of-thought (Wei et al., 2022).

Among the 11 baselines, a majority and often the stronger approaches would require a *held-out set* for training and hyperparameter tuning, potentially harming generalization across knowledge domains. In addition, these approaches often rely on a “self-reflection” assumption that *a single LLM* could be employed or adapted to evaluate its own generated texts; however, challenges such as hallucination and confirmation biases cast doubt on this assumption (Ji et al., 2023; Xie et al., 2023).

We propose **multi-LLM collaboration-based approaches**, enabling robust LLM abstention through multi-LLM collaboration to reflect on generated text in cooperative or competitive settings while removing the need for held-out sets. For multiple LLMs to work in *cooperation* (COOPERATE), the LLM employs other models to provide feedback on the proposed answer and reasoning, and synthesizes the outputs into an overall abstain decision. For a *competitive* setting (COMPETE), the LLM is challenged by other LLMs with conflicting knowledge, producing an abstain decision based on whether it continues to prefer the proposed answer in the presence of conflicting evidence.

We evaluate the baselines and proposed collaboration approaches for abstention with three LLMs, on four knowledge-intensive QA tasks spanning diverse knowledge domains and reasoning. Extensive experiments demonstrate that COOPERATE and COMPETE are strong abstain mechanisms that outperform all baselines in 9 of the 12 settings across tasks and models, achieving an improvement of up to 19.3% in abstain accuracy. Further analysis reveals that the proposed collaboration-based approaches could help identify failure cases of retrieval augmentation, pinpoint knowledge gaps in multi-hop reasoning, and more. Our contributions are (1) a critical evaluation and typology of diverse existing methods, repurposed to identify knowledge gaps in LLMs, as well as (2) two novel, robust multi-LLM collaboration methods to detect LLM

knowledge gaps, COOPERATE and COMPETE.

## 2 Identifying Knowledge Gaps in LLMs

We propose to repurpose existing approaches to enable abstain functionality in the face of knowledge gaps in LLMs. We conceptually categorize these approaches into *calibration* (§2.1), *training* (§2.2), *prompting* (§2.3), and *self-consistency* (§2.4) methods. Empirically observing that these approaches are dependent on external knowledge/data sets and often suffer from instability, we propose a new category based on *multi-LLM collaboration* (§2.5).

### 2.1 Calibration-Based

**Token Probability** LLMs are often inherently calibrated to different extents (Radford et al., 2019; Liang et al., 2023) and token probabilities might be employed off-the-shelf. Let  $(q, \bar{a}) \in \mathcal{H}$  be a held-out set of questions and gold answers for hyperparameter tuning, while LLM produces an answer  $\mathbf{a} = \text{LLM}(q)$  with a confidence score based on token probabilities  $p(\mathbf{a}) \in [0, 1]$ . We set a threshold  $p^*$  by minimizing abstain errors on  $\mathcal{H}$ , formally:

$$p^* = \arg \min_{p' \in [0, 1]} \sum_{\mathbf{a} = \bar{\mathbf{a}}} \mathbb{1}(p(\mathbf{a}) < p') + \sum_{\mathbf{a} \neq \bar{\mathbf{a}}} \mathbb{1}(p(\mathbf{a}) \geq p')$$

where  $\mathbb{1}$  denotes the indicator function. At inference time, the LLM abstains when  $p(\mathbf{a}) < p^*$ .

**Temperature Scaling** Before setting a threshold  $p^*$  on confidence scores, temperature scaling (Guo et al., 2017; Jiang et al., 2021) first reshapes the token probability distributions by introducing a scalar temperature  $\tau > 0$ . Given an answer probability distribution  $\mathbf{a}$ , we transform it into  $\text{softmax}(\mathbf{a}/\tau)$ , then optimize  $\tau$  on the held-out set  $\mathcal{H}$  with cross-entropy loss to obtain  $\tau^*$ . A probability threshold  $p^*$  is then set over  $\{\text{softmax}(\text{LLM}(q)/\tau^*)\}_{q \in \mathcal{H}}$  and applied at inference time.

**Ask for Calibration** In addition to token probabilities, LLMs could also verbalize confidence scores: we follow Tian et al. (2023) to obtain confidence scores  $p(\mathbf{a})$  by employing a two-step process: 1) generating an answer with the LLM and 2) eliciting a confidence score by prompting with “Provide the probability that your guess is correct.”<sup>2</sup> A threshold  $p^*$  is similarly set over verbalized scores and applied at inference time.

<sup>2</sup>Full prompt templates are in Appendix B.

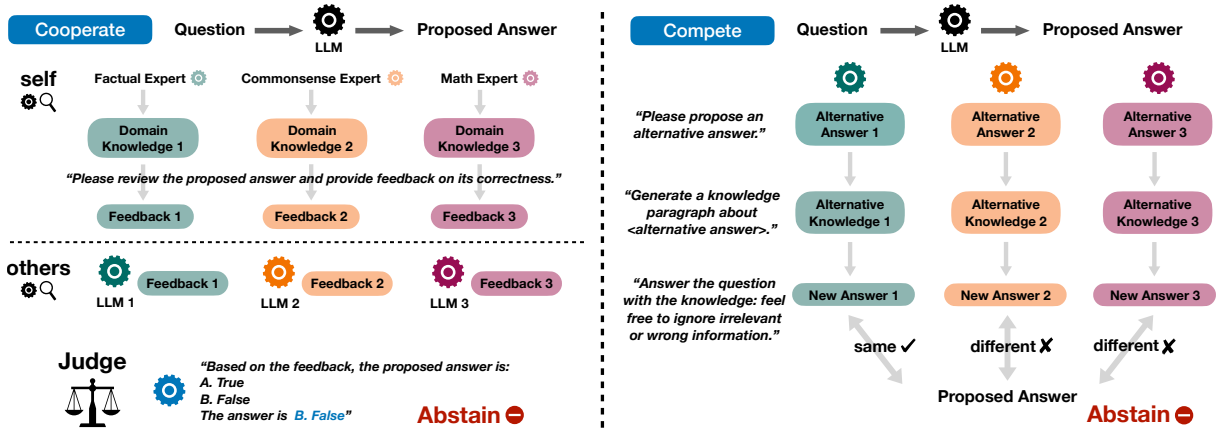


Figure 1: Overview of collaboration-based approaches for LLM abstention: COOPERATE and COMPETE.

## 2.2 Training-Based

**Hidden Layers** Previous works demonstrate that probing LLMs’ hidden layer representations might yield insights into their knowledge and factuality (Slobodkin et al., 2023; Azaria and Mitchell, 2023; CH-Wang et al., 2023). We extract the hidden embeddings  $\mathbf{e}_q = \text{LLM-hidden}(q)$ , then conduct linear probing to predict  $\ell \in \{\text{accept}, \text{reject}\}$  through  $p(\ell | q) \propto \exp(\text{FFN}(\mathbf{e}_q))$ . We optimize the cross-entropy loss on the held-out set  $\mathcal{H}$ . At inference time, the LLM should abstain if  $\arg \max_{\ell} p(\ell | q)$  indicates that the answer should be rejected.

**External Verifier** Training an external verifier to judge the veracity of LLM-generated text has shown great potential in math reasoning (Cobbe et al., 2021) and general knowledge domains (Kadavath et al., 2022; Cohen et al., 2023). We train an encoder-based language model  $\text{LM-enc}$  and employ the [CLS] token for classification. Formally,  $p(\ell | q) \propto \exp(\text{LM-enc}(q; \text{LLM}(q)))$  is optimized on  $\mathcal{H}$  with the cross-entropy loss and an abstain decision is achieved through  $\arg \max_{\ell} p(\ell | q)$  at inference time. Here both the question and the generated answer are employed to train the external verifier, different from the *hidden layers* approach where only the intermediate encodings of the question are used for linear probing.

**Instruction Tuning** A promising way to adapt LLMs is through tuning LLMs on instruction triplets in the format of  $\{\text{instruction}|\text{input}|\text{output}\}$  (Ouyang et al., 2022). Here we investigate whether abstaining could be baked into existing LLMs through instruction tuning. We first employ an instruction to encourage LLMs to abstain when necessary. Given the held-out set  $\mathcal{H}$ , we first evaluate the

correctness of the generated answer  $a = \text{LLM}(q)$ : if it is correct, we use  $\{\text{instruction}|q|a\}$  as an instruction tuning instance; otherwise, we replace the incorrect answer  $a$  with an abstain message and train on  $\{\text{instruction}|q|\text{abstain message}\}$ . In this way, the LLM abstains if the abstain message is generated at inference time for a given  $q$ .

## 2.3 Prompting-Based

**Self-Reflect** Previous studies show that LLMs might have preliminary capabilities of judging and evaluating their own answer (Kadavath et al., 2022). To this end, we prompt the LLM to self-reflect directly after its generated answer with “The above answer is: A. True B. False”. LLMs should abstain when they deem the generated answer as false.

**More Information** Since the internal knowledge of LLMs is often incomplete and noisy, existing works explore how to selectively request external information to aid LLM generation (Xu et al., 2023; Asai et al., 2023; Feng et al., 2023b). Following Feng et al. (2023b), we append a prompt about whether more information is needed to answer a given question: “Do you need more information to answer this question? (Yes or No)”. The LLM should abstain if more information is requested.

**Generate and Match** There might be gaps between LLM generation and “understanding” (West et al., 2023), which leads to hallucinations when questions are presented in a multiple-choice format (Fu et al., 2023). Instead of directly choosing from one of the options, we instruct the LLM to first generate an answer *without* access to potential options. We then prompt the LLM to evaluate whether the generated answer matches any of the options: the LLM should abstain if there’s no match between

the generated answer and options.

## 2.4 Consistency-Based

**None-of-the-Above** Previous works have investigated LLMs’ robustness to possibilities of *non-of-the-above* (*nota*) in QA tasks (Kadavath et al., 2022; Ding et al., 2023). We propose an abstain mechanism by adding an extra *nota* option to every question: the LLM should abstain if *nota* is selected as the answer, indicating its low consistency in answering the question.

**Self-Consistency Threshold** Self-consistency (Wang et al., 2022) aims to generate multiple chain-of-thought (Wei et al., 2022) reasoning paths and conduct majority voting to determine the answer. We develop an abstain mechanism by proposing *plurality index*: given a question  $q$  along with the  $k$  generated reasoning paths and answers  $\{(\mathbf{a}_i, \mathbf{p}_i)\}_{i=1,\dots,k}$ , the plurality index is defined as:

$$\text{plu}(\text{LLM}, \mathbf{q}, k) = \max_{\mathbf{a}_i} \sum_{j=1,\dots,k} \mathbb{1}(\mathbf{a}_j = \mathbf{a}_i)$$

where it indicates how large the plurality is among all  $k$  reasoning paths. A threshold  $\tau^* \in [0, 1]$  is then optimized on  $\mathcal{H}$  and applied at inference time to abstain when  $\text{plu}(\text{LLM}, \mathbf{q}, k) < \tau^* \cdot k$ .

## 2.5 Proposed: Multi-LLM Collaboration

Existing approaches often adapt an LLM to evaluate/reflect on its own generated texts, while challenges such as confirmation biases and hallucination would result in unreliable self-evaluation (Kadavath et al., 2022; Ji et al., 2023; Xie et al., 2023). In addition, the stronger approaches rely on held-out sets and result in weaker generalization across knowledge domains. We propose to enable robust LLM abstention through **multi-LLM collaboration**, i.e., employing multiple LLMs that interact with each other to evaluate the veracity of generated texts and make abstain decisions (Figure 1).

**Cooperate** LLMs often come with varying knowledge coverage with the potential of complementing each other (Yu et al., 2023a; Du et al., 2023; Bansal et al., 2023), while the internal knowledge of one LLM could also be diversified and specialized (Kang et al., 2023; Si et al., 2023). We tap into this knowledge variation in the COOPERATE approach by having different expert LLMs generate feedback on LLM-proposed answers and employing an LLM as the final judge to decide

whether to abstain. Formally, given question  $q$  and answer  $\mathbf{a} = \text{LLM}(q)$ , we obtain a set of natural language feedback  $\{\mathbf{f}_1, \dots, \mathbf{f}_k\}$  from expert LLMs and employ a judge LLM to summarize and synthesize the feedbacks into an abstain decision  $\text{LLM}(q, \mathbf{a}, \{\mathbf{f}_1, \dots, \mathbf{f}_k\}) \rightarrow \{\text{accept}, \text{reject}\}$ .

We employ two modes to generate feedback: for *self*, the same LLM is specialized into experts on domains  $\{d_1, \dots, d_k\}$  through prompting-based self-specialization: we prompt the LLM to generate a knowledge passage  $\text{knowledge}_i$  about  $q$  with a focus on domain  $d_i$ . A domain-specific feedback is then generated by prepending the knowledge passage:  $\mathbf{f}_i = \text{LLM}(\text{knowledge}_i, q, \mathbf{a})$ , and prompting the model to respond as a reviewer. For *others*, we use separate LLMs  $\{\text{LLM}_1, \dots, \text{LLM}_k\}$  to provide feedback  $\mathbf{f}_i = \text{LLM}_i(q, \mathbf{a})$ , aiming to identify the knowledge gaps in each other and complement the abstain process.

**Compete** LLMs have varying preferences for knowledge conflicts, i.e. when there is a conflict between LLM’s internal parametric knowledge and external information provided in the prompt (Xie et al., 2023; Wang et al., 2023b). We hypothesize that an LLM should abstain when it is impacted by conflicting information generated by other LLMs and does not stick to the original answer. Concretely, for question  $q$  and LLM-generated answer  $\mathbf{a} = \text{LLM}(q)$ , we prompt other LLMs to generate  $k$  alternative answers  $\{\mathbf{a}'_1, \dots, \mathbf{a}'_k\}$  along with a knowledge passage  $\text{knowledge}_i$  about  $\mathbf{a}'_i$ . We then instruct the LLM to answer  $q$  again with conflicting information prepended:  $\tilde{\mathbf{a}}_i = \text{LLM}(\text{knowledge}_i, q)$ . If  $\mathbf{a} = \tilde{\mathbf{a}}_i$ , the LLM sticks to the original answer and should be thus accepted; otherwise, the LLM is swayed by conflicting information generated by other LLMs, betraying its low confidence. This process is repeated for  $k$  alternative answers and the LLM should abstain if the answer changes in a majority of cases.

## 3 Experiment Settings

**Models** We evaluate LLM abstention baselines and our proposed collaboration-based approaches with three LLMs featuring varying sizes and openness: *Mistral-7B* (Jiang et al., 2023), *LLaMA2-70B* (Touvron et al., 2023), and *ChatGPT*. We posit that a robust abstain mechanism should ideally work for all LLMs, weak and strong. We set the default sampling temperature to 0.1, and employ 0.7 where multiple runs are required.



Method	MMLU				K-Crosswords				Hellaswag				Propaganda			
	R-Acc	ER	A-Acc	A-F1	R-Acc	ER	A-Acc	A-F1	R-Acc	ER	A-Acc	A-F1	R-Acc	ER	A-Acc	A-F1
<b>MISTRAL-7B</b>																
PROBS	.570	.109	.608	.456	.251	-.351	.397	.422	.456	-.041	.599	.659	.337	-.150	.590	.680
TEMP.	.565	.104	.601	.426	.250	-.452	.303	.182	.451	-.047	.601	.659	.340	-.170	.585	.661
ASK CALI.	.648	.141	.639	.647	.233	-.023	.705	.825	.455	-.017	.616	.734	.231	-.035	.705	.824
HIDDEN	.449	-.085	.420	.137	.104	-.788	.107	.009	.369	-.198	.424	.336	.223	-.510	.240	.084
VERIFIER	.568	.083	.586	.534	.208	-.083	<b>.805</b>	<b>.889</b>	.405	-.080	.550	.628	<b>.448</b>	<b>-.015</b>	<b>.760</b>	<b>.853</b>
INSTRUCT	.709	.198	.688	.693	-	-	-	-	.616	.075	.701	.771	-	-	-	-
REFLECT	.498	-.002	.495	.477	.111	-.686	.203	.207	.371	-.150	.477	.500	.216	-.275	.490	.602
MOREINFO	.504	.006	.500	.216	.272	-.455	.273	.003	.379	-.227	.396	.120	.219	-.540	.240	.073
GEN+MATCH	.511	.021	.515	.129	.111	-.694	.193	.188	.377	-.174	.458	.415	.226	-.545	.230	.013
NOTA	.516	.029	.528	.163	.098	-.797	.102	.011	.371	-.244	.387	.105	.259	-.410	.340	.267
SC THRES.	.604	.140	.641	.551	.210	-.090	.793	.880	.493	-.004	.614	.713	.273	-.100	.685	.799
COOP-SELF	.571	.059	.564	.601	.260	-.437	.313	.311	.406	-.042	.601	.719	.297	-.095	.680	.797
COOP-OTHERS	.688	<b>.213</b>	<b>.712</b>	.692	.266	<b>-.022</b>	.761	.852	<b>.626</b>	<b>.092</b>	<b>.725</b>	<b>.783</b>	.182	-.140	.625	.757
COMPETE	<b>.735</b>	.140	.640	<b>.700</b>	<b>.289</b>	-.129	.597	.722	.573	.032	.658	.766	.302	-.055	.700	.805
<b>LLAMA2-70B</b>																
ASK CALI.	.624	.025	.435	.568	-	-	-	-	<b>.944</b>	.032	.498	.649	-	-	-	-
HIDDEN	.473	-.018	.400	.446	.282	-.265	.423	.466	.507	.007	.497	.479	-	-	-	-
VERIFIER	.665	.201	.609	.511	.443	-.056	.634	.694	.522	.009	.504	.617	.259	<b>-.065</b>	<b>.755</b>	<b>.855</b>
INSTRUCT	.745	.216	.628	<b>.640</b>	.288	<b>.024</b>	.606	<b>.772</b>	.475	-.008	.487	.615	-	-	-	-
REFLECT	.616	.121	.529	.469	.383	-.155	.528	.537	.569	.076	.560	.530	.198	-.520	.275	.225
MOREINFO	.590	.177	.590	.042	.306	-.382	.308	.020	.516	.032	.518	.044	-	-	-	-
GEN+MATCH	.667	.050	.450	.560	.248	-.111	.573	.708	.484	-.004	.477	.614	.082	-.205	.620	.759
NOTA	.592	.167	.583	.181	.323	-.295	.388	.280	.516	.027	.522	.236	.185	-.580	.225	.124
SC THRES.	.684	.247	<b>.656</b>	.534	.426	-.090	.590	.617	.667	.100	.590	<b>.655</b>	<b>.412</b>	<b>-.030</b>	<b>.760</b>	<b>.852</b>
COOP-SELF	.615	.150	.550	.400	<b>.463</b>	-.030	.640	.714	.649	.110	.600	.643	.222	-.250	.500	.615
COOP-OTHERS	.694	<b>.262</b>	<b>.676</b>	.562	.402	-.063	.636	.757	<b>.700</b>	<b>.238</b>	<b>.704</b>	<b>.677</b>	<b>.329</b>	-.125	.675	.774
COMPETE	<b>.782</b>	.148	.552	<b>.608</b>	.323	-.080	<b>.642</b>	<b>.760</b>	.611	.047	.525	.625	.161	-.210	.595	.729
<b>CHATGPT</b>																
PROBS	.774	.421	.715	.457	.600	.187	.587	.122	.750	.278	.599	.476	.333	-.015	.625	.765
TEMP.	.769	.419	.716	.452	.616	.214	.619	.216	.750	.278	.595	.468	.250	-.010	.630	.772
ASK CALI.	.694	.385	.690	.006	.601	.202	.601	.010	.672	.344	.672	.006	.444	-.015	.580	.712
VERIFIER	.788	.301	.599	.483	-	-	-	-	.667	.305	.634	.120	.441	-.035	.570	.672
INSTRUCT	.840	.422	<b>.748</b>	.579	.752	<b>.310</b>	<b>.709</b>	.627	.817	<b>.429</b>	<b>.751</b>	.614	.605	.045	.645	.744
REFLECT	.752	.336	.630	.411	.784	.239	.641	.633	.754	.377	.701	.487	.571	.015	.615	.742
MOREINFO	.721	.246	.546	.390	.605	.145	.553	.380	.675	.224	.548	.339	.416	-.145	.470	.293
GEN+MATCH	.737	.350	.652	.383	.660	.083	.486	.550	.712	.182	.506	.447	.365	-.115	.490	.568
NOTA	.719	.389	.692	.260	.644	.163	.565	.480	.689	.307	.628	.268	.400	-.120	.485	.488
SC THRES.	.766	.424	.743	.447	.637	.216	.622	.382	.749	.366	.688	.468	.399	-.160	.440	.309
COOP-SELF	.841	<b>.436</b>	.726	.578	.794	.175	.646	<b>.646</b>	.878	.344	.670	<b>.628</b>	.684	.070	<b>.710</b>	<b>.802</b>
COOP-OTHERS	.780	.362	.660	.479	.659	.109	.509	.536	.790	.350	.676	.565	<b>.790</b>	<b>.321</b>	.647	.543
COMPETE	<b>.947</b>	.306	.602	<b>.583</b>	<b>.875</b>	.034	.441	.589	<b>.939</b>	.172	.490	.545	.611	.040	<b>.670</b>	<b>.795</b>

Table 1: Performance of abstain strategies on four datasets and three LLMs. Best results in **bold** and second best in underline. Approaches are color-coded per category: **calibration**, **training**, **prompting**, **consistency**, and **collaboration**. Certain incompatible cases, e.g. EMBEDDING with the black-box CHATGPT, are omitted. “-” indicates that this approach fails to produce meaningful abstain decisions: almost always abstain, didn’t follow instructions, etc. COOPERATE and COMPETE achieve the best performance in 9 of the 12 settings in terms of reliable accuracy.

**Tasks and Datasets** We evaluate LLM abstention with four datasets spanning diverse knowledge domains and reasoning scenarios: 1) *MMLU* (Hendrycks et al., 2020) is a multiple-choice dataset for general knowledge QA; 2) *Knowledge Crosswords* (Ding et al., 2023) is a structured QA dataset that focuses on multi-hop and compositional knowledge reasoning; 3) *Hellaswag* (Zellers et al., 2019) is a natural language inference dataset that tests commonsense knowledge and reasoning; 4) *Propaganda* (Piskorski et al., 2023) tasks LLMs with identifying the 23 propaganda tactics in a long news article based on their internal knowledge. The datasets are all in English. We create

splits of held-out sets and test sets for the four datasets with details in Appendix B.1. We employ LLMs to answer these questions with zero-shot prompting and greedy decoding to obtain the ground truth of whether the LLM “knows” the answer and whether it should abstain. We posit that a robust abstain mechanism should work for knowledge tasks across domains and reasoning contexts.

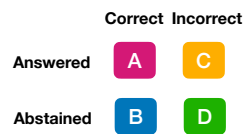


Figure 2: Four outcomes of AbstainQA.

**Evaluation Metrics** We illustrate the four scenarios under AbstainQA in Figure 2 and employ the following evaluation metrics. (1) *Reliable Accu-*

racy (R-Acc):  $\frac{A}{A+C}$  indicates to what extent could LLM-generated answers (not abstained) be trusted. i.e., *out of all questions answered, how many are correct?* (2) *Effective Reliability* (ER) (Whitehead et al., 2022; Si et al., 2023):  $\frac{A-C}{A+B+C+D}$  strikes a balance between reliability and coverage. i.e., *out of all questions, what proportion more are answered correctly than incorrectly?* (3) *Abstain Accuracy* (A-Acc):  $\frac{A+D}{A+B+C+D}$  evaluates whether the abstain decisions are correct: LLMs should abstain when it would provide an incorrect answer and vice versa. (4) *Abstain F1* (A-F1): *harmonic-mean*(precision, recall), where precision =  $\frac{D}{B+D}$  and recall =  $\frac{D}{C+D}$ , a balanced metric between reliability and answer coverage.

## 4 Results

We present the performance of abstain mechanisms with three LLMs on four tasks in Table 1.

**COOPERATE and COMPETE are state-of-the-art abstain mechanisms.** Our proposed collaboration-based approaches outperform the strongest baseline in 9 out of the 12 settings (3 models and 4 datasets), achieving an average improvement of 10.1% on reliable accuracy scores. We find that COOPERATE works better with stronger models such as ChatGPT: we hypothesize that the complexities of the multi-step feedback generation warrant a stronger base LLM. COMPETE emphasizes reliability and greatly avoids wrong answers evident in the high reliable accuracy scores: we observe that LLMs stick to the original answers only in the most confident scenarios, making COMPETE ideal for abstention when reliability is paramount.

**Calibration > Training > Consistency > Prompting** On average, the four types of approaches achieved 0.595, 0.576, 0.553, and 0.475 A-Acc scores. While simple prompting-based approaches struggle to elicit self-reflection, our proposed COOPERATE and COMPETE promote collaborative abstention and greatly improve performance. Among all baselines, *instruction tuning* and *self-consistency threshold* stand out as two strong approaches: however, they both rely on a held-out set for training and hyperparameter tuning: we further investigate their generalizability in Section 5.

**Abstain capabilities potentially correlate with base LLM utility.** Abstain accuracy (A-Acc) is a metric independent of LLMs’ underlying performance on a QA dataset and solely evaluates ab-

stention quality. For A-Acc, *Mistral-7B*, *LLaMA2-70B*, and *ChatGPT* achieve 0.524, 0.537, and 0.616 A-Acc scores on average across approaches and datasets, which aligns well with their relevant superiority on standard benchmarks. In addition, our proposed collaboration-based approaches also work better with the strongest *ChatGPT*, while also improving the abstention performance of weaker LLMs over baselines.

**Going beyond generic QA sheds light on the limitations of existing approaches.** While most baselines could function properly on MMLU, some collapse (“-” in Table 1) in other knowledge domains or reasoning contexts: ASK CALIBRATE, among others, struggle to follow instructions when the prompt context is long in detecting propaganda tactics; INSTRUCTION TUNING collapses and produces 100% abstention due to imbalances in the held-out set in K-Crosswords where problems are much harder and the LLM should mostly abstain; calibration approaches struggle to set a good threshold when the model’s base performance is too low. However, our proposed collaboration-based approaches work well with all four tasks, especially with the strongest *ChatGPT* language model, indicating its robustness and broad compatibility for knowledge-intensive tasks across domains and reasoning contexts.

## 5 Analysis

**Abstain Absolute** Aside from QA datasets where abstain decisions should be made based on the potential correctness of answers, we investigate scenarios where LLMs should 100% abstain. Specifically, we employ AmbigQA (Min et al., 2020), where LLMs should always abstain from ambiguous and underspecified questions. We also curate ElectionQA23, a QA dataset focusing on elections in 2023 around the globe (details in Appendix B.1) and LLMs should abstain due to the temporal mismatch of training data cutoff, *i.e.* if the LLM has only seen training data before 2023. Other potential *abstain absolute* scenarios include known unknowns (Amayuelas et al., 2023) and questions with false/underspecified premises (Patidar et al., 2023; Hu et al., 2023). Figure 3 demonstrates that our proposed COMPETE consistently abstains the most in *abstain absolute* scenarios across three LLMs and two datasets: we observe that the proposed alternative answers in COMPETE often provide contradictory contexts, thus the LLMs suc-

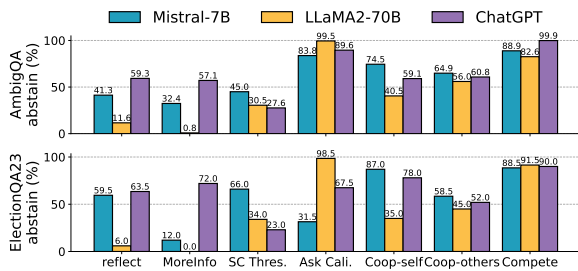


Figure 3: Performance of abstain mechanisms in the *abstain absolute* scenarios where the LLM should abstain for 100% of questions. COMPETE achieves the highest abstention rate on average across LLMs and datasets.

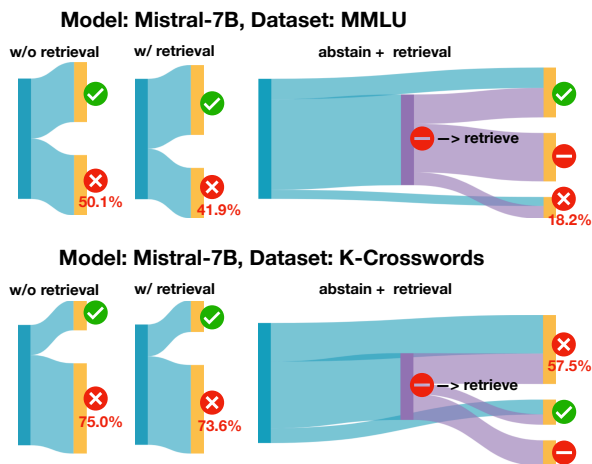


Figure 4: Performance of COMPETE with MISTRAL-7B in the two-step abstain and retrieval setting. The proposed abstain-retrieve-abstain pipeline successfully reduces the incorrect rate by at least 21.2%.

cessfully abstain due to the knowledge conflicts. We further divide the ElectionQA23 dataset into subsets of which continent is the election held in and present model performance in Figure 5. It is illustrated that LLMs tend to abstain less and assume the outcomes of future elections for Africa and Asia, indicating fairness concerns of abstain mechanisms that might underserve marginalized communities and underdeveloped countries.

**Abstain and Retrieval** Retrieval-augmented language models (Khattab et al., 2022; Shi et al., 2023; Baek et al., 2023; Yu et al., 2023b; Thakur et al., 2023) have become the go-to approach for augmenting LLMs’ knowledge access, while retrieving from an external corpus could also introduce noisy documents and lead to wrong answers (Wang et al., 2023c; Asai et al., 2023). We investigate whether abstain mechanisms could work with retrieval, specifically when questions in the MMLU dataset are presented together with a re-

Method	Ret.	R-Acc	ER	A-Acc	A-F1
SC THRES.	✗	0.604	0.140	0.641	0.551
	✓	0.634	0.160	0.652	0.611
ASK CALI.	✗	0.648	0.141	0.639	0.647
	✓	0.515	0.027	0.512	0.156
COOP-SELF	✗	0.571	0.059	0.564	0.601
	✓	0.567	0.048	0.536	0.589
COMPETE	✗	0.735	0.140	0.640	0.700
	✓	0.720	0.128	0.618	0.670

Table 2: Performance of abstain strategies on MMLU with *Mistral-7B* with or without retrieval (Ret.).

Method	MMLU	K-Cross.	Hellaswag	Propaganda
MOREINFO	0.4433	0.7281	0.5830	0.6143
PROBS	0.4253	0.6643	0.4992	0.5302
SC TRES.	0.2854	0.3428	0.2146	0.2750
ASK CALI.	<b>0.1747</b>	0.2664	<b>0.2088</b>	0.3755
VERIFIER	0.3343	0.2000	0.3624	0.2112
COOP-SELF	0.4740	<b>0.1131</b>	0.3434	0.2263
COOP-OTHERS	0.3560	0.5330	0.2164	0.2380
COMPETE	0.4484	0.3929	0.4014	<b>0.1935</b>

Table 3: Abstain estimated calibration error (ECE) of approaches with *Mistral-7B*, the lower the better. The calibration of abstention scores for approaches could vary significantly across tasks and datasets.

trieved document from the WikiSearch API. Table 2 demonstrates that while baselines such as ASK CALIBRATE are negatively impacted by retrieval, our proposed COOPERATE and COMPETE are in general robust towards additional prepended context through retrieval.

Building upon the findings that collaboration-based approaches are robust with retrieved contexts, we propose a 2-step abstain-based framework to selectively activate retrieval and identify retrieval failures. Specifically, abstain decisions are first made based on QA without retrieval. If the LLM abstains, retrieval is then requested and the abstain mechanism is applied to QA with the retrieved context: if the LLM still decides to abstain, it indicates that the retrieval has failed to find relevant/helpful documents. We evaluate this framework on MMLU and Knowledge Crosswords datasets while comparing it with no retrieval and full retrieval settings. Results in Figure 4 demonstrate that this abstain-retrieve-abstain pipeline successfully cuts back the error rate by up to 56.6%, highlighting the potential of abstain mechanisms in identifying retrieval failures and enhancing the reliability of retrieval-augmented LLMs.

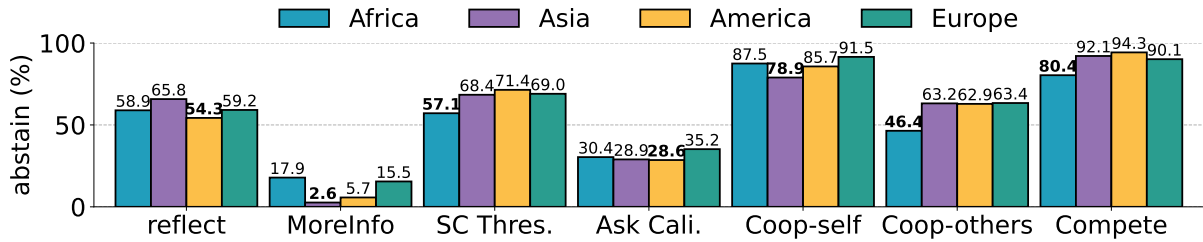


Figure 5: Abstain rate on Election23QA with *Mistral-7B* divided into where the election takes place. The lowest abstain rate for each approach among continents is highlighted in **bold**.

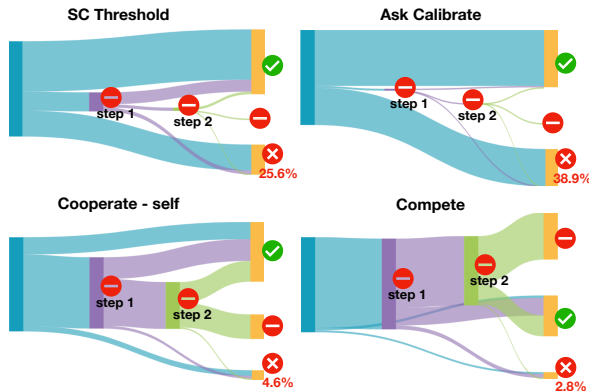


Figure 6: Performance of abstain mechanisms with ChatGPT solving 3-hop questions in K-Crosswords. COOPERATE and COMPETE are better at localizing knowledge gaps in multi-hop reasoning steps and reducing the incorrect rate by at least 82%.

**Abstain in Multi-Hop** In multi-hop knowledge reasoning, LLMs’ knowledge is often limited in some, but not all, of the reasoning steps. We investigate whether our proposed abstain strategies could pinpoint the specific steps where LLMs should abstain due to knowledge limitations: we employ the three-hop subset of the K-Crosswords dataset, conduct abstain mechanisms on each of the hops separately, and present results in Figure 6. It is demonstrated that our proposed COOPERATE and COMPETE are better in localizing knowledge limitations, with a final error rate of 67.2% to 81.2% lower than the strongest baselines.

**Abstain ECE** Aside from a binary decision, we posit that abstain mechanisms could also provide a continuous abstain likelihood based on token probabilities, heuristics, and more<sup>3</sup>. We calculate the *abstain ECE*, estimated calibration error (Guo et al., 2017) but with abstain probabilities and abstain accuracy, on the MMLU dataset and present in Table 3. It is demonstrated that the calibra-

tion of abstain likelihood generally correlates with the abstain performance, with ASK CALIBRATE and our proposed COOPERATE and COMPETE being the top-performing approaches. This indicates that our proposed collaboration-based approaches could also provide well-calibrated abstain likelihood scores where fine-grained decision-making is required.

**Precision and Recall** We observe that among the 14 abstain mechanisms, some emphasize answer coverage while others emphasize answer reliability. We illustrate the abstain precision and recall of approaches in Figure 7: It is illustrated that while most approaches have higher precision than recall and thus lean towards answer coverage, COMPETE has a much higher abstain recall and stresses answer reliability. We argue that these strategies could be ensembled to achieve balanced abstain decisions or employed based on whether the application scenario should optimize coverage or reliability.

## 6 Related Work

Previous works on knowing what language models know have focused on calibration, prompting, or training. *Calibration-based* approaches attempt to extract confidence scores from models to gauge their uncertainty (Sun et al., 2022; Kuhn et al., 2022; Zhou et al., 2023a; Liu et al., 2023a). Previous works have evaluated (Radford et al., 2019; Liang et al., 2023; Tao et al., 2023; He et al., 2023) and improved the calibration of language models (Desai and Durrett, 2020; Kong et al., 2020; Jagannatha and Yu, 2020; Kamath et al., 2020; Jiang et al., 2021; Mielke et al., 2022; Lin et al., 2022; Tian et al., 2023; Zhou et al., 2023b), while the calibration quality could vary for different downstream tasks (Desai and Durrett, 2020; Wang et al., 2020; Stengel-Eskin and Van Durme, 2023; Kalai and Vempala, 2023). In the setting of AbstainQA where

<sup>3</sup>Details in Appendix B.



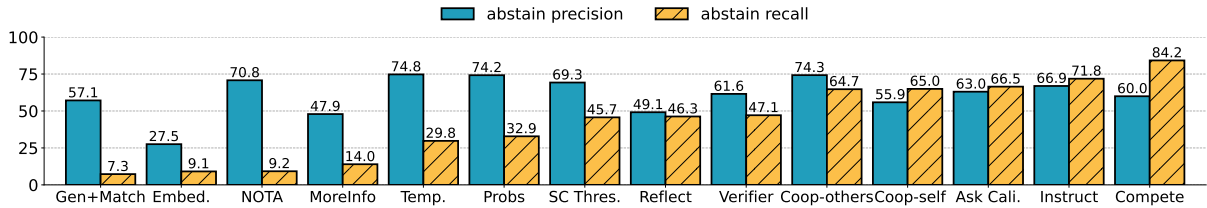


Figure 7: Abstain precision and recall of approaches on MMLU with *Mistral-7B*. While most approaches feature higher precision, COMPETE achieves the best abstention recall and could be employed where reliability is paramount.

LLMs decide whether to abstain or not given a question, calibration-based approaches would rely on a held-out set to set a threshold over calibrated confidence scores: such a threshold and reliance on a held-out set could jeopardize the generalization of calibration-based approaches across knowledge domains and reasoning contexts.

*Prompting-based* approaches leverage the instruction-following abilities of LLMs and employ instructions to induce self-reflection and gauge whether the generated answer should be trusted. Kadavath et al. (2022) investigates whether a “none-of-the-above” option or self-evaluation prompting would induce good estimations of LLMs’ internal factuality. Huang et al. (2023a) follows existing works (Kim et al., 2024; Shinn et al., 2023) to evaluate whether LLMs could self-correct their own reasoning with a three-step prompting strategy. Feng et al. (2023b) prompts LLM to reflect on whether more information is needed for a given question to selectively solicit external knowledge. These approaches, among other prompting-based methods for acknowledging knowledge limitations (Wang et al., 2023a; Si et al., 2023), mostly rely on the hope that LLMs could indeed carry out self-reflection through simple prompting, while their ability on this front might be subject to confounders and spurious correlations (Kadavath et al., 2022).

*Training-based* approaches focus on empowering LLMs with the ability to acknowledge uncertainty or abstain through training with special data or objectives. Cobbe et al. (2021) proposes to train an external verifier to evaluate the correctness of LLM-generated answers in math reasoning. A series of works also attempt to probe the hidden representations of LLMs for signals of factuality and hallucination (Slobodkin et al., 2023; Azaria and Mitchell, 2023). Contemporary works have gained interest in improving LLMs’ factuality and self-knowledge through instruction tuning and alignment (Zhang et al., 2023; Yang et al., 2023; Sun et al., 2023; Bashlovkina et al., 2023), as well

as promoting abstention from a safety perspective (Huang et al., 2023b; Liu et al., 2023b). While training-based approaches have shown great potential, they would often require substantial computational resources to fine-tune LLMs and might struggle to generalize across knowledge domains and contexts due to their reliance on the patterns and examples in the training subset.

In this work, we propose COOPERATE and COMPETE, two collaboration-based approaches to enable robust LLM abstention through multi-LLM collaboration. These approaches have shown great potential in improving LLM abstention in not only direct QA, but also retrieval-augmented QA and multi-hop reasoning scenarios.

## 7 Conclusion

We investigate AbstainQA, a setting where LLMs should abstain from answering questions incorrectly. We curate a taxonomy of 11 abstain baselines across four categories and propose COOPERATE and COMPETE, two novel abstain mechanisms that promote mechanistic reflection through multi-LLM collaboration, in cooperation or competition. Extensive experiments on four datasets demonstrate that COOPERATE and COMPETE advances the state-of-the-art in AbstainQA, with the potential to improve retrieval-augmented LLMs, multi-hop knowledge reasoning, and more.

## Limitations

While we primarily investigated and evaluated LLM abstention from a knowledge perspective, the problem of abstaining is also important from a safety perspective, preventing LLMs from generating harmful content or biased responses (Huang et al., 2023b; Liu et al., 2023b). There are also intersections between the knowledge and safety aspects of abstention, such as abstaining from underspecified questions based on stereotypes of demographic attributes. We plan to expand on this work and pro-

pose unified abstain strategies for knowledge and safety scenarios in the future.

The abstain mechanisms and experiments in this work are exhaustive to the extent that our computational budget permits, while we could not cover every single approach and proposal from existing works. Nevertheless, we will make the evaluation data, code, and framework publicly available so that new abstain approaches, as well as datasets, can be seamlessly added to advance LLM abstention research.

Our proposed collaboration-based approaches require prompting multiple LLMs for feedback and could have more computational overhead, while the exact cost would depend on the inference costs of other LLMs. We provide a summary of inference costs in Table 6 and Appendix A. We posit that the value of multi-LLM collaboration for abstention lies in using smaller models to supervise larger models in judging veracity and preventing hallucinations. When we serve a user-facing LLM with hundreds of billions of parameters, our approach enables the incorporation of a few 7B LLMs for abstention guidance and mitigate hallucinations: having a few extra 7B models does not add greatly to the cost of a hundreds-of-billions LLM, while significantly boosting its factuality and reliability. By sampling multiple feedbacks to make abstain decisions, randomness is also introduced and there might be variation across model inference runs.

In this work, we focus on a conceptual “abstain” functionality and develop robust approaches to identify the knowledge gaps in large language models. However, “abstain” doesn’t necessarily have to be part of the user interface of LLMs: it could be an internal component of an LLM-based system, followed by approaches to find better knowledge sources to fill the knowledge gaps or methods to incorporate the uncertainty to model responses to enhance trustworthiness (Zhou et al., 2024). Future work could further explore LLM pipelines where these abstain approaches are one component towards reliable question answering.

We acknowledge the potential impact of test set pollution in the training of the adopted LLMs. While we do not know the exact training composition of ChatGPT and the other two LLMs (Mistral-7B, LLaMA2-70B) are not explicitly trained on test sets, they might still be unintentionally included in the training data. Our results and findings should be thus interpreted with this limitation in mind.

## Ethics Statement

Figure 5 highlights that LLM abstain decisions could also have fairness implications, as they fail to abstain and assume the outcomes of future elections more frequently for African and Asian countries. We similarly expect varying levels of abstain performance on questions and prompts regarding different demographics, communities, and perspectives, potentially due to LLMs’ internal social biases (Blodgett et al., 2020; Jin et al., 2021; Bender et al., 2021; Shaikh et al., 2023; Feng et al., 2023a). We plan to examine the intersections of LLM social biases and their abstention abilities in future work, especially in critical domains such as hate speech and misinformation where failing to abstain from generating problematic content could have grave consequences for marginalized communities.

## Acknowledgements

This material is based upon work supported by the National Science Foundation under CAREER Grant No. IIS2142739, NSF Grant No. IIS2203097, and the DARPA CMO under Contract No. HR001120C0124. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies. We also gratefully acknowledge support from Alfred P. Sloan Foundation Fellowship.

## References

- Alfonso Amayuelas, Liangming Pan, Wenhui Chen, and William Wang. 2023. Knowledge of knowledge: Exploring known-unknowns uncertainty with large language models. *arXiv preprint arXiv:2305.13712*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*.
- Amos Azaria and Tom Mitchell. 2023. The internal state of an llm knows when it’s lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976.
- Jinheon Baek, Soyeong Jeong, Minki Kang, Jong C Park, and Sung Hwang. 2023. Knowledge-augmented language model verification. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1720–1736.
- Rachit Bansal, Bidisha Samanta, Siddharth Dalmia, Nitish Gupta, Sriram Ganapathy, Abhishek Bapna, Praateek Jain, and Partha Talukdar. 2023. Llm augmented

- llms: Expanding capabilities through composition. In *The Twelfth International Conference on Learning Representations*.
- Vasilisa Bashlovkina, Zhaobin Kuang, Riley Matthews, Edward Clifford, Yennie Jun, William W Cohen, and Simon Baumgartner. 2023. Trusted source alignment in large language models. *arXiv preprint arXiv:2311.06697*.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack Rae, Erich Elsen, and Laurent Sifre. 2022. Improving language models by retrieving from trillions of tokens. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2206–2240. PMLR.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sky CH-Wang, Benjamin Van Durme, Jason Eisner, and Chris Kedzie. 2023. Do androids know they’re only dreaming of electric sheep? *arXiv preprint arXiv:2312.17249*.
- Tong Chen, Hongwei Wang, Sihao Chen, Wenhao Yu, Kaixin Ma, Xinran Zhao, Dong Yu, and Hongming Zhang. 2023. Dense x retrieval: What retrieval granularity should we use? *arXiv preprint arXiv:2312.06648*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. Crawling the internal knowledge-base of language models. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1811–1824.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506.
- Shrey Desai and Greg Durrett. 2020. Calibration of pre-trained transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.
- Wenxuan Ding, Shangbin Feng, Yuhan Liu, Zhaoxuan Tan, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2023. Knowledge crosswords: Geometric reasoning over structured knowledge with large language models. *arXiv preprint arXiv:2310.01290*.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multi-agent debate. *arXiv preprint arXiv:2305.14325*.
- Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023a. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, Toronto, Canada. Association for Computational Linguistics.
- Shangbin Feng, Weijia Shi, Yuyang Bai, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2023b. Knowledge card: Filling llms’ knowledge gaps with plug-in specialized language models. In *The Twelfth International Conference on Learning Representations*.
- Xingyu Fu, Sheng Zhang, Gukyeong Kwon, Pramuditha Perera, Henghui Zhu, Yuhao Zhang, Alexander Hanbo Li, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Patrick Ng, Dan Roth, and Bing Xiang. 2023. Generate then select: Open-ended visual question answering guided by world knowledge. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2333–2346, Toronto, Canada. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.



- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Guande He, Peng Cui, Jianfei Chen, Wenbo Hu, and Jun Zhu. 2023. Investigating uncertainty calibration of aligned language models under the multiple-choice setting. *arXiv preprint arXiv:2310.11732*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Evan Hernandez, Belinda Z Li, and Jacob Andreas. 2023. Measuring and manipulating knowledge representations in language models. *arXiv preprint arXiv:2304.00740*.
- Shengding Hu, Yifan Luo, Huadong Wang, Xingyi Cheng, Zhiyuan Liu, and Maosong Sun. 2023. Won’t get fooled again: Answering questions with false premises. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5626–5643, Toronto, Canada. Association for Computational Linguistics.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023a. Large language models cannot self-correct reasoning yet. In *The Twelfth International Conference on Learning Representations*.
- Xiaowei Huang, Wenjie Ruan, Wei Huang, Gao Jin, Yizhen Dong, Changshun Wu, Saddek Bensalem, Ronghui Mu, Yi Qi, Xingyu Zhao, Kaiwen Cai, Yanghao Zhang, Sihao Wu, Peipei Xu, Dengyu Wu, André Freitas, and Mustafa A. Mustafa. 2023b. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *ArXiv*, abs/2305.11391.
- Abhyuday Jagannatha and Hong Yu. 2020. [Calibrating structured output predictors for natural language processing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2078–2092, Online. Association for Computational Linguistics.
- 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. *ACM Computing Surveys*, 55(12):1–38.
- Albert Qiaoju Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L’elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *ArXiv*, abs/2310.06825.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves, and Xiang Ren. 2021. On transferability of bias mitigation effects in language model fine-tuning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3770–3783, Online. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, T. J. Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zachary Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, John Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom B. Brown, Jack Clark, Nicholas Joseph, Benjamin Mann, Sam McCandlish, Christopher Olah, and Jared Kaplan. 2022. Language models (mostly) know what they know. *ArXiv*, abs/2207.05221.
- Adam Tauman Kalai and Santosh S Vempala. 2023. Calibrated language models must hallucinate. *arXiv preprint arXiv:2311.14648*.
- Amita Kamath, Robin Jia, and Percy Liang. 2020. Selective question answering under domain shift. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5684–5696, Online. Association for Computational Linguistics.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pages 15696–15707. PMLR.
- Junmo Kang, Hongyin Luo, Yada Zhu, James Glass, David Cox, Alan Ritter, Rogerio Feris, and Leonid Karlinsky. 2023. Self-specialization: Uncovering latent expertise within large language models. *arXiv preprint arXiv:2310.00160*.
- Jungo Kasai, Keisuke Sakaguchi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir Radev, Noah A Smith, Yejin Choi, Kentaro Inui, et al. 2024. Realtime qa: What’s the answer right now? *Advances in Neural Information Processing Systems*, 36.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*.



- Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2024. Language models can solve computer tasks. *Advances in Neural Information Processing Systems*, 36.
- Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. Calibrated language model fine-tuning for in- and out-of-distribution data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1326–1340, Online. Association for Computational Linguistics.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2022. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2023. Holistic evaluation of language models. *Transactions on Machine Learning Research*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.
- Xin Liu, Muhammad Khalifa, and Lu Wang. 2023a. Litcab: Lightweight calibration of language models on outputs of varied lengths. *arXiv preprint arXiv:2310.19208*.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2023b. Trustworthy llms: a survey and guideline for evaluating large language models’ alignment. In *Socially Responsible Language Modelling Research*.
- 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 preprint arXiv:1907.11692*.
- Ziyang Luo, Can Xu, Pu Zhao, Xiubo Geng, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. Augmented large language models with parametric knowledge guiding. *arXiv preprint arXiv:2305.04757*.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822.
- Sabrina J Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. Reducing conversational agents’ overconfidence through linguistic calibration. *Transactions of the Association for Computational Linguistics*, 10:857–872.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. AmbigQA: Answering ambiguous open-domain questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Ouyang Long, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. Webgpt: Browser-assisted question-answering with human feedback. *ArXiv*, abs/2112.09332.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Mayur Patidar, Prayushi Faldu, Avinash Singh, Lovekesh Vig, Indrajit Bhattacharya, and Mausam. 2023. Do I have the knowledge to answer? investigating answerability of knowledge base questions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10341–10357, Toronto, Canada. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. Semeval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multilingual setup. In *Proceedings of the the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*.

- Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. 2023. On second thought, let’s not think step by step! bias and toxicity in zero-shot reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4454–4470, Toronto, Canada. Association for Computational Linguistics.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chenglei Si, Weijia Shi, Chen Zhao, Luke Zettlemoyer, and Jordan Boyd-Graber. 2023. Getting more out of mixture of language model reasoning experts. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8234–8249.
- Aviv Slobodkin, Omer Goldman, Avi Caciularu, Ido Dagan, and Shauli Ravfogel. 2023. The curious case of hallucinatory (un) answerability: Finding truths in the hidden states of over-confident large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3607–3625.
- Elias Stengel-Eskin and Benjamin Van Durme. 2023. Calibrated interpretation: Confidence estimation in semantic parsing. *Transactions of the Association for Computational Linguistics*, 11:1213–1231.
- Meiqi Sun, Wilson Yan, Pieter Abbeel, and Igor Mordatch. 2022. Quantifying uncertainty in foundation models via ensembles. In *NeurIPS 2022 Workshop on Robustness in Sequence Modeling*.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*.
- Linwei Tao, Younan Zhu, Haolan Guo, Mingjing Dong, and Chang Xu. 2023. A benchmark study on calibration. In *The Twelfth International Conference on Learning Representations*.
- Nandan Thakur, Luiz Bonifacio, Xinyu Zhang, Odunayo Ogundepo, Ehsan Kamaloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Boxing Chen, Mehdi Rezagholizadeh, et al. 2023. Nomiracl: Knowing when you don’t know for robust multilingual retrieval-augmented generation. *arXiv preprint arXiv:2312.11361*.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Boshi Wang, Xiang Yue, and Huan Sun. 2023a. Can ChatGPT defend its belief in truth? evaluating LLM reasoning via debate. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11865–11881, Singapore. Association for Computational Linguistics.
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020. On the inference calibration of neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Yike Wang, Shangbin Feng, Heng Wang, Weijia Shi, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2023b. Resolving knowledge conflicts in large language models. *arXiv preprint arXiv:2310.00935*.
- Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md Rizwan Parvez, and Graham Neubig. 2023c. Learning to filter context for retrieval-augmented generation. *arXiv preprint arXiv:2311.08377*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D Hwang, Liwei Jiang, Jillian Fisher, Abhilasha Ravichander, Khyathi Chandu, et al. 2023. The generative ai paradox: “what it can create, it may not understand”. In *The Twelfth International Conference on Learning Representations*.
- Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. 2022. Reliable visual question answering: Abstain rather than answer incorrectly. In *European Conference on Computer Vision*, pages 148–166. Springer.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2023. Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts. In *The Twelfth International Conference on Learning Representations*.

Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. Re-comp: Improving retrieval-augmented lms with compression and selective augmentation. *arXiv preprint arXiv:2310.04408*.

Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. 2023. Alignment for honesty. *arXiv preprint arXiv:2312.07000*.

Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, et al. 2023a. Kola: Carefully benchmarking world knowledge of large language models. In *The Twelfth International Conference on Learning Representations*.

Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2023b. Chain-of-note: Enhancing robustness in retrieval-augmented language models. *arXiv preprint arXiv:2311.09210*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800.

Hanning Zhang, Shizhe Diao, Yong Lin, Yi R Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. 2023. R-tuning: Teaching large language models to refuse unknown questions. *arXiv preprint arXiv:2311.09677*.

Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2023a. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. *arXiv preprint arXiv:2309.17249*.

Kaitlyn Zhou, Jena D Hwang, Xiang Ren, and Maarten Sap. 2024. Relying on the unreliable: The impact of language models’ reluctance to express uncertainty. *arXiv preprint arXiv:2401.06730*.

Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023b. Navigating the grey area: How expressions of uncertainty and overconfidence affect language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5506–5524, Singapore. Association for Computational Linguistics.

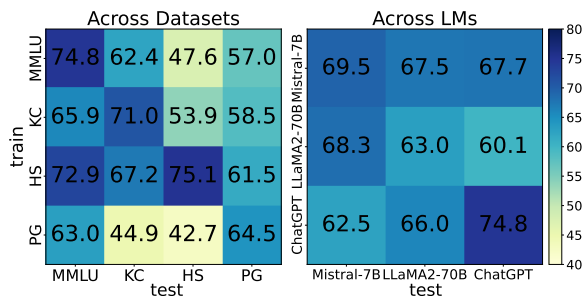


Figure 8: Performance of INSTRUCTION TUNING when trained with one dataset/LLM and tested for another. Training on a held-out set harms the generalization of abstain accuracy across tasks and LLMs.

	fail in 1/3 hops	fail in 2/3 hops	fail in 3/3 hops
abstain in 1/3 hops	<b>351</b>	153	12
abstain in 2/3 hops	91	<b>512</b>	63
abstain in 3/3 hops	30	20	<b>156</b>

Table 4: Alignment between model abstain decisions and incorrect answers regarding the number of hops in multi-hop QA. COMPETE shows 73.4% match between abstain decisions and model failures.

## A Analysis (cont.)

**Retrieval Failure (cont.)** We present more results with the two-step abstention in retrieval-augmented LLMs in Figure 10. Across three LLMs and four datasets, the abstain-retrieve-abstain pipeline successfully brings down the incorrect rate and accounts for retrieval failure.

**Abstain and Multi-Hop (cont.)** We present more results with the multi-hop abstention in Figure 11. It is demonstrated that our proposed COOPERATE and COMPETE are consistently better at pinpointing knowledge limitations in specific reasoning steps across LLMs and datasets. We further investigate the alignment between model abstain decisions and model answer incorrectness: how many hops did the LLM abstain for and how many hops did the LLM actually answer incorrectly. Table 4 indicates that model abstain decisions match model failures in 73.4% of the time, showcasing the effectiveness of multi-LLM collaboration for abstention in multi-hop problems.

**Held-Out Sets Harms Generalization** In the study we find that most of the baselines would require a held-out set of questions for model tuning, hyperparameter optimization, and more. Specifically, while INSTRUCTION TUNING is one of the strongest baselines in Section 4, its reliance on a held-out set for training might jeopardize its gen-

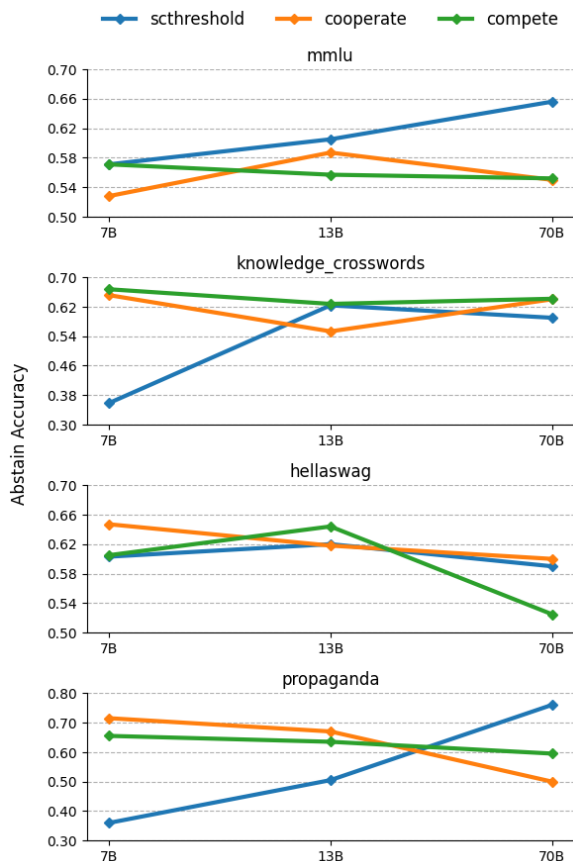


Figure 9: Performance with three abstain mechanisms with increasing size of model parameters.

eralization across knowledge domains and LLMs. To this end, we conduct a generalization study in Figure 8 where there is a transfer across datasets or LLMs between training and testing. It is illustrated that INSTRUCTION TUNING approaches struggle to generalize across knowledge domains and LLMs, resulting in performance drops sometimes as large as 33.8% in abstain accuracy. On the contrary, our proposed collaboration-based approaches do not need a held-out set and will not suffer from these generalization limitations.

**Minority Opinion** By default, we have  $k = 3$  feedback passages for each final judge LLM to consider. We find that on the MMLU dataset, the final judge follows the majority opinion in generated feedbacks 84.2% of the time. We manually examine 10 examples where the judge employed the minority opinion and find that in 8 of the 10 cases, the minority opinion indeed presents new information and should be thus taken into account, showcasing the benefit of having multiple feedbacks through multi-LLM collaboration.

Method	R-Acc	ER	A-Acc	A-F1
COOP-SELF-1	.540	.040	.544	.543
COOP-SELF-2	.564	.035	.498	.553
COOP-SELF-3	.562	.032	.504	.562
COMPETE-1	.672	.088	.590	.671
COMPETE-2	.698	.126	.623	.680
COMPETE-3	.655	.077	.577	.662

Table 5: Performance of COOPERATE and COMPETE with variations in the prompt format, on the MMLU dataset with the *Mistral-7B* language model.

**Scaling Laws** We investigate whether LLM abstention with these mechanisms could be improved through scaling, by evaluating three abstain mechanisms with varying sizes of the LLAMA2 LLM. Results in Figure 9 demonstrate that the abstain performance is not steadily increasing with model size, showcasing that empowering LLMs with abstention abilities is not solely solvable through scaling up.

**Shift in Knowledge Domains** We investigate whether LLM abstention abilities could be impacted by which knowledge domain the given question is about. Specifically, we present the abstain accuracy metric for different knowledge domains in the MMLU dataset in Figure 12. It is demonstrated that the abstention decisions are indeed of varying quality given the knowledge domain, from 25% to 87.5%, potentially informed by the LLM’s underlying knowledge abilities regarding that domain.

**Prompt Robustness** Since LLMs could be sensitive to minor changes in prompts (Sclar et al., 2023), we experiment with minor paraphrases of the prompts used in COOPERATE and COMPETE and present the results in Table 5. It is demonstrated that the proposed abstain mechanisms are relatively robust to minor changes in prompt phrasing.

**LLM Overhead** Different abstain mechanisms have different computational overhead and the most expensive part is the number of LLM inference requests for each question. We present a summary in Table 6.

**Working Examples** We present qualitative examples of COOPERATE (Tables 14 and 15 for *self*, Tables 16 and 17 for *others*) and COMPETE (Tables 18, 19, 20, and 21).



Method	#	Method	#
NOTA	2	REFLECT	2
MOREINFO	2	PROBS	3
SC THRES.	k+1	ASK CALL.	6
TEMP.	3	VERIFIER	3
INSTRUCT	2+ <i>ft</i>	EMBEDDING	4
GEN+MATCH	3	COOP-SELF	8
COOP-OTHERS	2+ <i>o</i>	COMPETE	2+ <i>o</i>

Table 6: Overhead of LLM inference requests for different abstain mechanisms. “*ft*” denotes additional fine-tuning and “*o*” denotes that other reviewer LLMs are also called once.

## B Experiment Details

### B.1 Dataset Details

We employ six datasets in the experiments, four for the main experiments, and two *abstain absolute* datasets.

- **MMLU** (Hendrycks et al., 2020): we randomly downsample the official validation and test splits into 1,000 questions each as the held-out set and test set.
- **K-Crosswords** (Ding et al., 2023): we follow the official splits to employ the validation set of 1,094 questions as the held-out set and the test set of 1,007 questions.
- **Hellaswag** (Zellers et al., 2019): we randomly sample the official validation set into 1,000 questions each as the held-out set and test set.
- **Propaganda** (Piskorski et al., 2023): we create a random split of 231 questions as the held-out set and 200 questions as the test set based on the official validation split. We randomly sample three other propaganda tactics accompanied by the correct tactic to form four-way multi-hop choice questions.
- **AmbigQA** (Min et al., 2020): we randomly downsample the official validation and test splits into 1,000 questions each as the held-out set and test set.
- **ElectionQA23**: we first collect the overview paragraphs of the Wikipedia pages of elections held in 2023<sup>4</sup> and employ ChatGPT to compose 10 questions for each election in a four-way multiple-choice format. We conduct manual

<sup>4</sup>[https://en.wikipedia.org/wiki/List\\_of\\_elections\\_in\\_2023](https://en.wikipedia.org/wiki/List_of_elections_in_2023)

evaluation to remove low-quality or ambiguous questions and create a held-out set of 67 and a test set of 200 questions.

### B.2 Model Details

We employ three large language models to evaluate abstain mechanisms: 1) *Mistral-7B*, through the MISTRALAI/MISTRAL-7B-INSTRUCT-V0.1 checkpoint on Huggingface (Wolf et al., 2019); 2) *LLaMA2-70B*, through the META-LLAMA/LLAMA-2-70B-CHAT-HF checkpoint on Huggingface; 3) *ChatGPT*, through the GPT-3.5-TURBO-INSTRUCT checkpoint on OpenAI API.

### B.3 Implementation Details

for every single baseline and approach, along with ECE score extraction, prompt for prompting-based and proposed approaches

- **Token Probability**: The token probability of the answer token(s) is employed as  $p(\mathbf{a})$ . The abstain likelihood is obtained through  $1 - p(\mathbf{a})$ .
- **Temperature Scaling**: We search for an optimal  $\tau$  value from 0.1 to 10 on the held-out set  $\mathcal{H}$ . The abstain likelihood is obtained through  $1 - p(\mathbf{a})$  after temperature scaling.
- **Ask for Calibration**: We follow Tian et al. (2023) and employ the prompt in Table 7 to elicit verbalized confidence scores. The abstain likelihood is obtained through  $1 - p(\mathbf{a})$  with verbalized confidence scores.
- **Hidden Layers**: We employ the feature-extraction pipeline with HuggingFace (Wolf et al., 2019) to extract  $\mathbf{e}_q$ . A linear layer of size (DIM, 2) is then employed for linear probing where DIM is the dimension of  $\mathbf{e}_q$ . The abstain likelihood is obtained as the  $\ell = \text{reject}$  likelihood after softmax.
- **External Verifier**: We employ RoBERTa (ROBERTA-BASE) (Liu et al., 2019) as LLM-*enc* for binary classification. The abstain likelihood is obtained as the  $\ell = \text{reject}$  likelihood with the [CLS] token.
- **Instruction Tuning**: We fine-tune ChatGPT (GPT-3.5-TURBO-INSTRUCT) with the instruction tuning triplets. The *instruction* is “Answer the following question. If you don’t have enough knowledge, abstain by saying ‘sorry, I don’t have enough knowledge to answer this question.’” The

*abstain message* is “sorry, I don’t have enough knowledge to answer this question.” The abstain likelihood is obtained as the token probability of the answer token, or that of “*sorry*” if it abstains.

- **Self-Reflect:** Following [Kadavath et al. \(2022\)](#), we employ the prompt in Table 8. The abstain likelihood is obtained as the token probability of choosing “B”.
- **More Information:** Following [Feng et al. \(2023b\)](#), we employ the prompt in Table 9. The abstain likelihood is obtained as the token probability of “yes”.
- **Generate and Match:** We employ the prompt in Table 10. The abstain likelihood is obtained as the probability of “no”, or  $1 -$  that of “yes”.
- **None-of-the-Above:** The abstain likelihood is obtained as the probability of choosing the NOTA option.
- **Self-Consistency Threshold:** We employ  $k = 5$  chain-of-thought paths by default. The abstain likelihood is obtained as  $1 - \text{plu}(\text{LLM}, \mathbf{q}, k)$ .
- **Cooperate:** We employ the prompt in Tables 11 and 12 for the *self* and *others* modes. The abstain likelihood is obtained as the token probability of choosing “B”.
- **Compete:** We employ the prompt in Table 13 and employ three alternative answers by default. The abstain likelihood is obtained as the probability of the new answer if it changes, or  $1 -$  that of the original answer if it doesn’t change, then averaged across all alternative answers.

---

Provide your best guess for the following question. Give ONLY the guess, no other words or explanation.

For example:

Guess: <most likely guess, as short as possible; not a complete sentence, just the guess!>

The question is: <question>

*[LLM-generated answer]*

Provide the probability that your guess is correct. Give ONLY the probability, no other words or explanation.

For example:

Probability: <the probability between 0.0 and 1.0 that your guess is correct, without any extra commentary whatsoever; just the probability!>

Probability: *[LLM-generated probability]*

---

Table 7: Prompt for ASK FOR CALIBRATION.

---

<question>

*[LLM-generated answer]*

The above answer is:

A. True

B. False

The answer is *[LLM-generated A/B]*

---

Table 8: Prompt for SELF-REFLECT.

---

<question>

Do you need more information to answer this question? (Yes or No)

*[LLM-generated yes/no]*

---

Table 9: Prompt for MORE INFORMATION.

---

<question without multiple-choice options>

Proposed answer: *[LLM-generated answer]*

<options>

Does the proposed answer exist in the options?

*[LLM-generated yes/no]*

---

Table 10: Prompt for GENERATE AND MATCH.

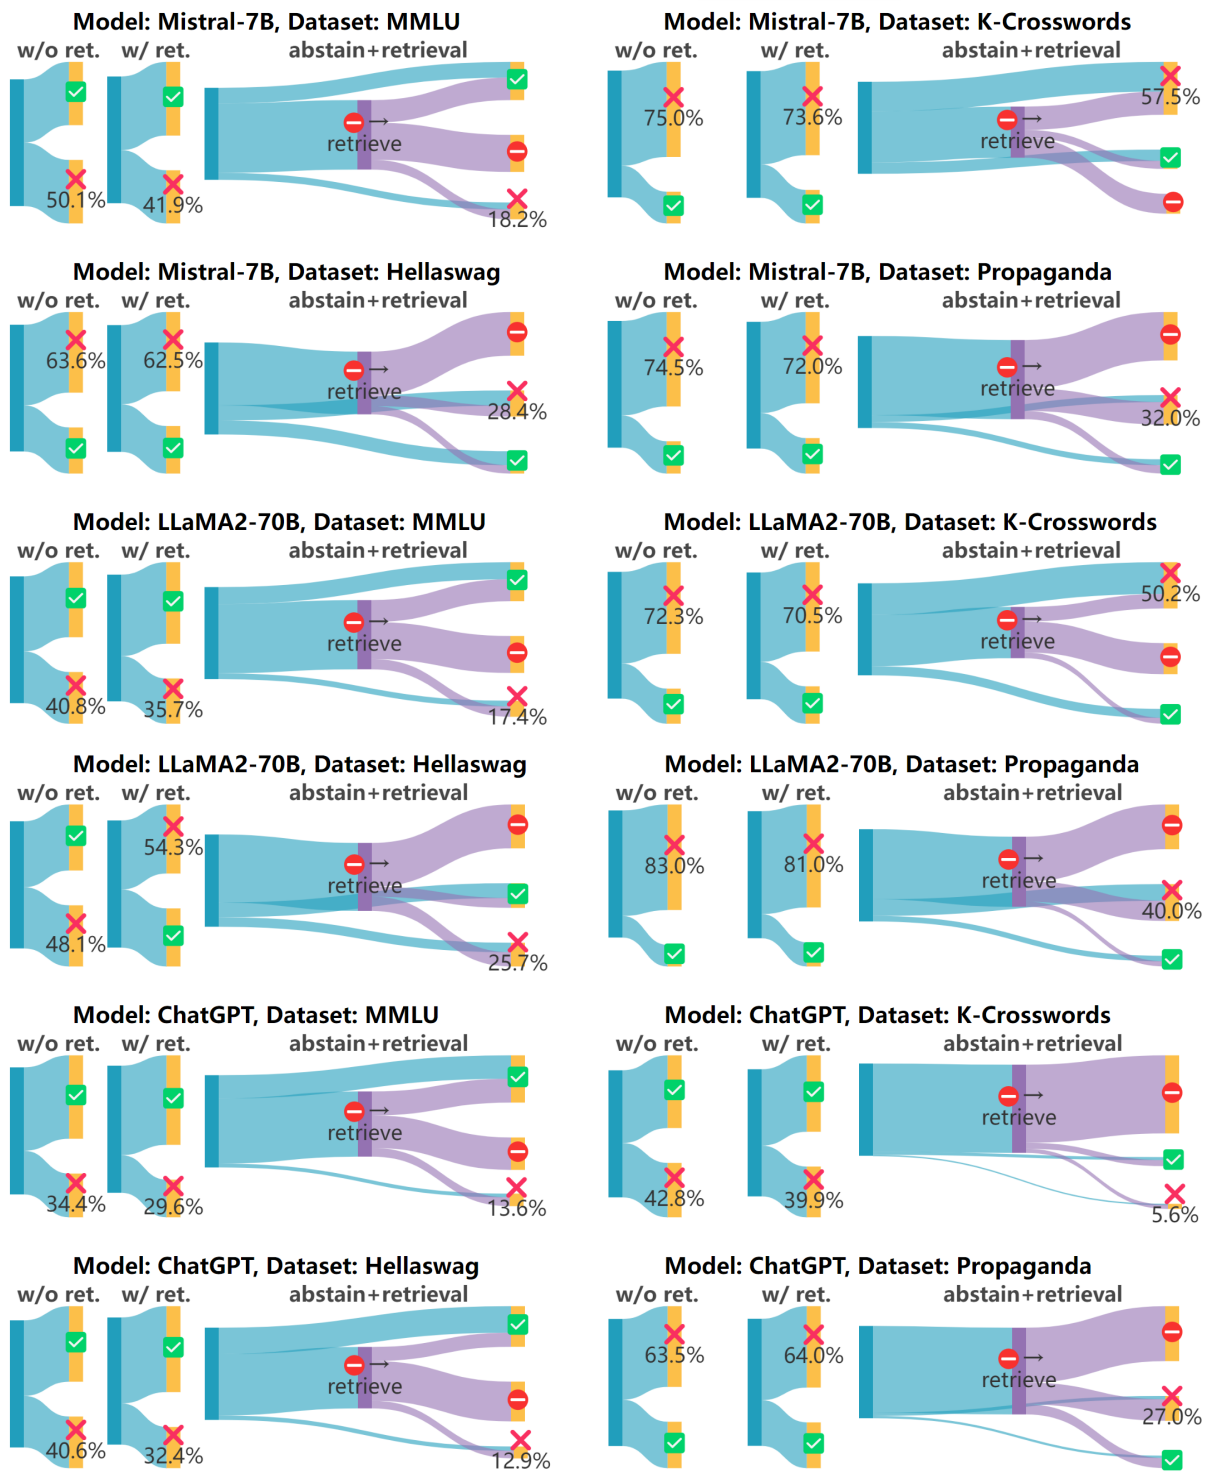


Figure 10: Model performance of *without retrieval*, *with retrieval*, and *retrieval+abstain* settings with the COMPETE approach.



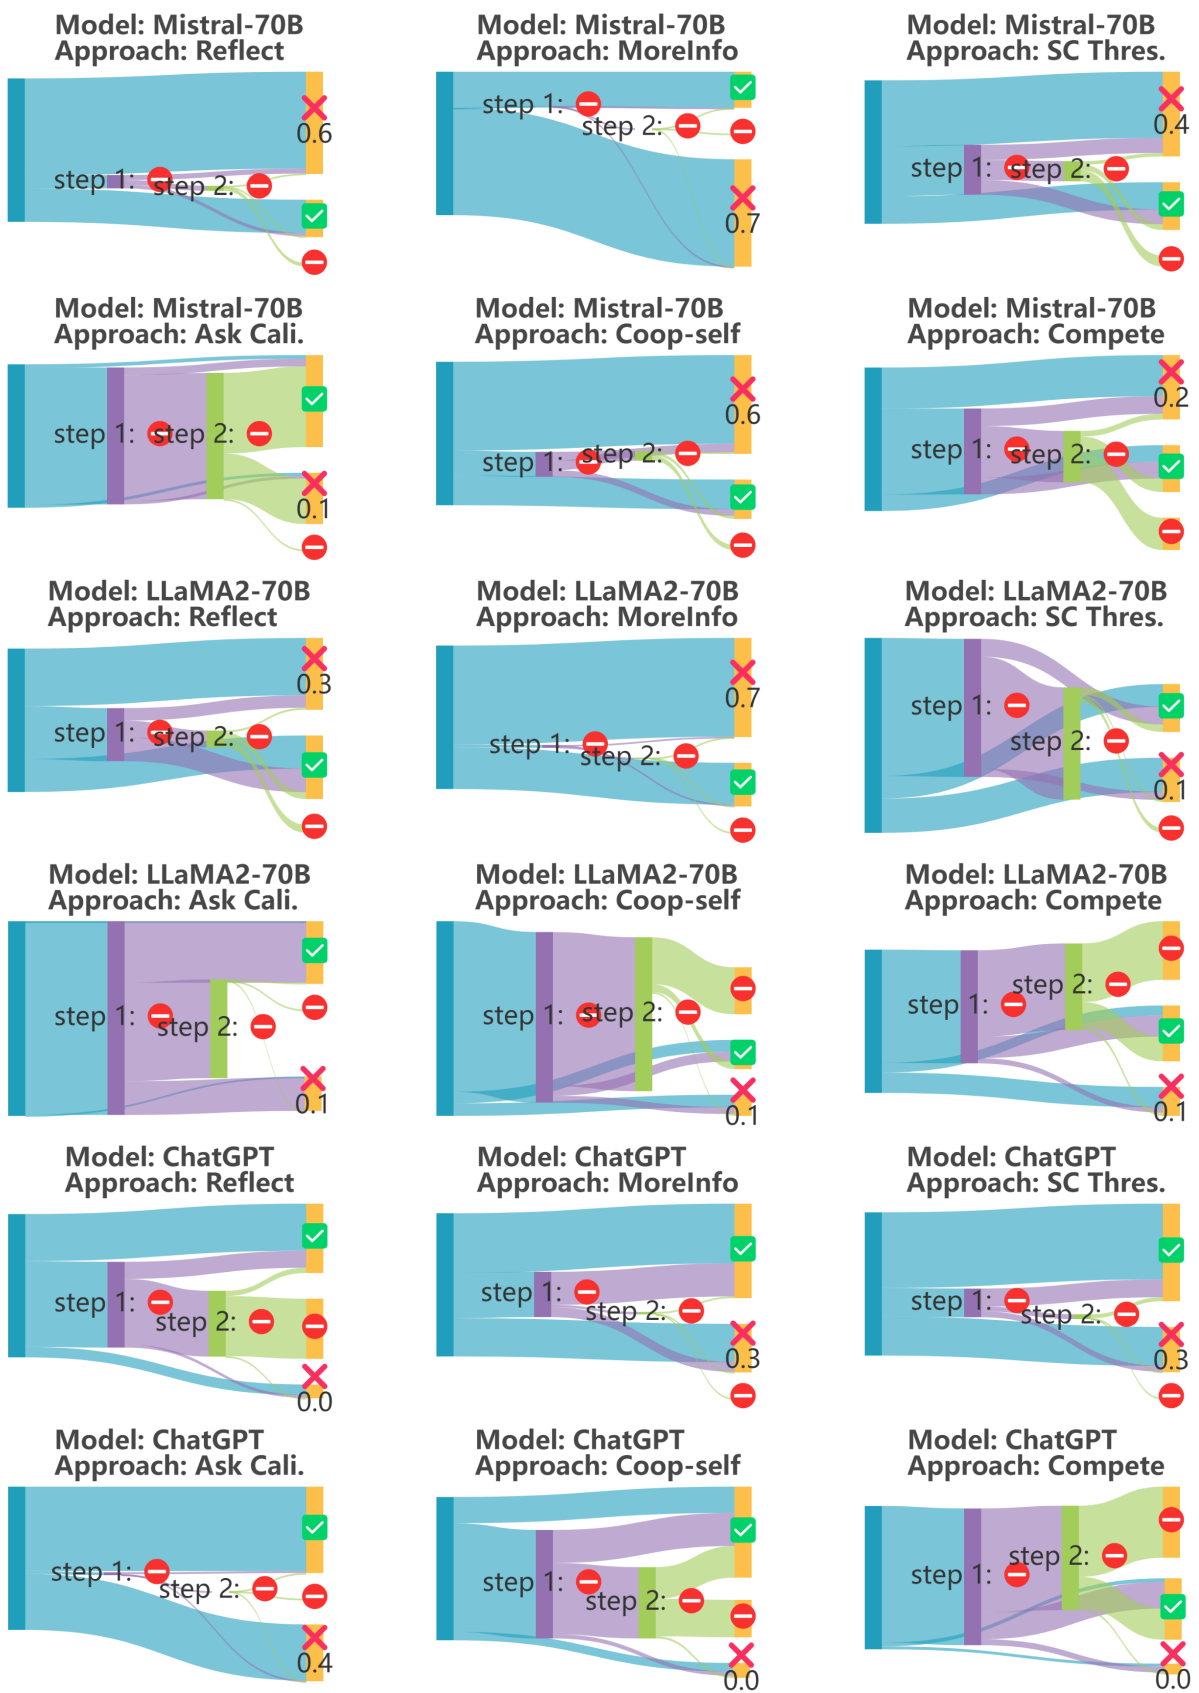


Figure 11: Model performance in the multi-hop knowledge reasoning setting.

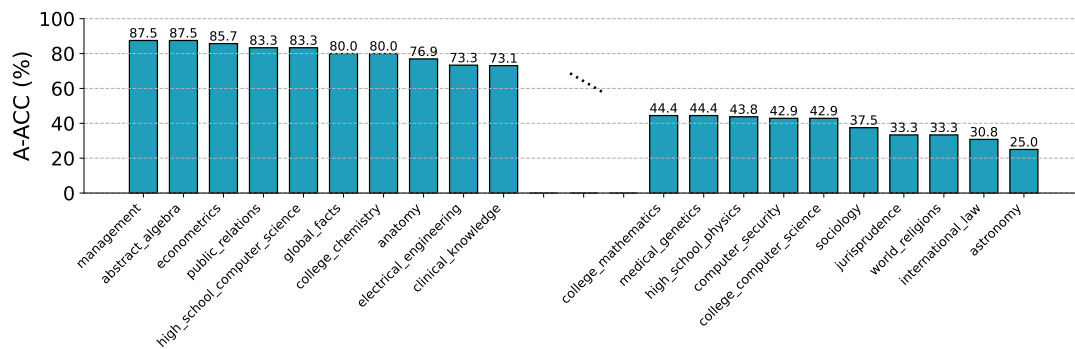


Figure 12: Performance of the COMPETE mechanism with *Mistral-7B* across the 57 subtasks in MMLU.

---

```

// obtain proposed answer

Question: <question>
Answer: [generated proposed answer]

// obtain feedback from self-specialized experts

for domain in ["factual information", "commonsense knowledge", "mathematical knowledge"]:
    Generate some knowledge about the question, focusing on <domain>: [generated domain knowledge]

    Knowledge: <generated domain knowledge>
    Question: <question>
    Answer: <generated proposed answer>
    Please review the proposed answer and provide feedback on its correctness.
    Feedback: [generated feedback]

// area-chairing for abstain decision

Question: <question>
Proposed Answer: <generated proposed answer>

Feedback 1: <generated feedback from expert 1>
...
Feedback k: <generated feedback from expert k>

Based on the feedback, the proposed answer is:
A. True
B. False
The answer is [A/B].

```

---

Table 11: Prompt for COOPERATE-*self*.

---

```

// obtain proposed answer

Question: <question>
Answer: [generated proposed answer]

// obtain feedback from other LLMs as reviewers
for llm in list_of_other_llms:

    Question: <question>
    Answer: <generated proposed answer>
    Please review the proposed answer and provide feedback on its correctness.
    Feedback: [generated feedback using llm]

// area-chairing for abstain decision

Question: <question>
Proposed Answer: <generated proposed answer>

Feedback 1: <generated feedback from llm 1>
...
Feedback k: <generated feedback from llm k>

Based on the feedback, the proposed answer is:
A. True
B. False
The answer is [A/B].

```

---

Table 12: Prompt for COOPERATE-*others*.

---

```

// obtain proposed answer

Question: <question>
Answer: [generated proposed answer]

// obtain alternative answers
if multiple-choice:
    <alternative answer> = randomly select another unchosen answer
else:
    Question: <question>
    Answer: <generated proposed answer>
    Please propose an alternative answer: [alternative answer]

// obtain a supporting paragraph for the alternative answer

Question: <question>
Generate a knowledge paragraph about <alternative answer>: [generated alternative passage]

// challenge!

Answer the question with the following knowledge: feel free to ignore irrelevant or wrong information.

Knowledge: <generated alternative passage>
Question: <question>
Answer: [new generated answer]

// abstain decision
if <new generated answer> == <generated proposed answer>:
    abstain = False
else:
    abstain = True

```

---

Table 13: Prompt for COMPETE.

---

Question: The following information pertains to Dash Co.'s utility bills: Period covered Amount Date paid April 16 - May 15 \$5000 June 1 - May 16 \$6000 June 15 - June 16 \$8000 July 1 - June 16 \$7000 July 15 - August 1 \$8000 What is the amount that Dash should report as a liability in its June 30 balance sheet?

- A: \$6,000
- B: \$7,000
- C: \$10,000
- D: \$14,000

Choose one answer from the above choices. The answer is B: \$7,000

Feedback 1: The correct answer is actually A: \$6,000. The amount that Dash should report as a liability in its June 30 balance sheet is the utility bill for the period of May 16 - June 15, which is \$6,000. The subsequent utility bills for June 16 - July 15 and July 16 - August 15 will be reported as liabilities in future balance sheets.

Feedback 2: The answer is incorrect. To determine the liability that Dash should report in its June 30 balance sheet, we need to look at the amount that is unpaid as of June 30. According to the information given, the utility bill for April 16 - May 15 was paid on June 1. Therefore, Dash should report the unpaid amount for the period May 16 - June 15 as a liability on its June 30 balance sheet. This amount is \$6,000. Thus, the correct answer is A: \$6,000.

Feedback 3: The proposed answer is incorrect.

Based on the feedback, the proposed answer is:

- A. True
- B. False

The answer is

- B. False
- 

Table 14: Example one for COOPERATE-self. The reviewers provided different feedback, while the area chair LLM merged into a final abstain decision.

---

Question: Compared with the current in an LED lamp the current in the connecting wire is

- A: less.
- B: more.
- C: the same.
- D: Need more information.

Choose one answer from the above choices. The answer is A: less.

Feedback 1: The proposed answer, A: less, is incorrect. The answer C: the same, would be more appropriate. In an LED lamp, the current flows from the power source through the connecting wire and into the LED. The current in the wire would be the same as the current in the LED if there are no resistive elements in the wire causing a drop in voltage.

Feedback 2: The proposed answer is incorrect. The statement does not provide enough information to determine the current in the connecting wire compared to the current in an LED lamp. Therefore, the correct answer is D: Need more information.

Feedback 3: The proposed answer is correct. In an LED lamp, the current in the connecting wire is typically less compared to the current in the LED itself. This is because the LED lamp is designed to limit the amount of current flowing through the LED to ensure proper operation and prevent damage. The connecting wire acts as a conductor to allow the flow of current from the power source to the LED, but it does not have as high a current as the LED itself.

Based on the feedback, the proposed answer is:

- A. True
- B. False

The answer is

- B. False.
- 

Table 15: Example two for COOPERATE-self. The reviewers provided different feedback, while the area chair LLM merged into a final abstain decision.



---

Question: The following information pertains to Dash Co.'s utility bills: Period covered Amount Date paid April 16 ,À May 15 \$5000 June 1 May 16 ,À June 15 \$6000 July 1 June 16 ,À July 15 \$8000 August 1 What is the amount that Dash should report as a liability in its June 30 balance sheet?

- A: \$6,000
- B: \$7,000
- C: \$10,000
- D: \$14,000

Choose one answer from the above choices. The answer is B: \$7,000

Feedback 1: The answer provided is not correct. The correct amount that Dash should report as a liability in its June 30 balance sheet is \$6,000. This is because the amount for the period of May 16 - June 15 has already been paid on June 1, and does not need to be reported as a liability. The only amount that is still outstanding and needs to be reported is the amount for the period of June 16 - July 15, which is \$6,000. Therefore, the correct answer is A: \$6,000.

Feedback 2: The proposed answer is incorrect. Dash should report \$7,000 as a liability in its June 30 balance sheet, but not \$6,000.

Here's the reasoning:

1. The period covered by the utility bills is May 16 ,À June 15.
2. The amount due is \$6,000, and Dash paid \$5,000 on May 16 ,À May 15 and \$1,000 on June 16 ,À June 15.
3. The date paid for the \$6,000 bill is June 16 ,À June 15 and the date paid for the \$1,000 bill is June 16 ,À June 15.
4. Since the payment of the \$1,000 bill has already

Feedback 3: The correct answer is indeed B: \$7,000.  
The explanation is as follows:

The June 30 balance sheet should report a liability for the utility bills that have been incurred but not yet

Based on the feedback, the proposed answer is:

- A. True
  - B. False
- The answer is
- B. False.

---

Table 16: Example one for COOPERATE-others. The reviewers provided different feedback, while the area chair LLM merged into a final abstain decision.

---

Question: Who kills Tony at the end of the film 'West Side Story'?

- A: Riff
- B: Chino
- C: Bernardo
- D: He kills himself

Choose one answer from the above choices. The answer is B: Chino

Feedback 1: The proposed answer is correct. Chino is the one who kills Tony at the end of the film 'West Side Story'.

Feedback 2: The answer provided is incorrect. In the film 'West Side Story', it is Bernardo who kills Tony. Thank you for providing the correct answer.

Feedback 3: The proposed answer is incorrect. Tony is killed by Bernardo, not Chino. The correct answer is C: Bernardo.

New proposed answer:

Answer: Who kills Tony at the end of the film 'West Side Story'

Based on the feedback, the proposed answer is:

- A. True
  - B. False
- The answer is
- B. False. The correct answer is C: Bernardo.

---

Table 17: Example two for COOPERATE-others. The reviewers provided different feedback, while the area chair LLM merged into a final abstain decision.

---

Question: The following information pertains to Dash Co.'s utility bills: Period covered Amount Date paid April 16 ,Ài May 15 \$5000 June 1 May 16 ,Ài June 15 \$6000 July 1 June 16 ,Ài July 15 \$8000 August 1 What is the amount that Dash should report as a liability in its June 30 balance sheet?

- A: \$6,000
- B: \$7,000
- C: \$10,000
- D: \$14,000

Generate a knowledge paragraph about A.

Dash Co. should report \$6,000 as a liability in its June 30 balance sheet. This is because the utility bills for the period from May 16 to June 15 amount to \$6,000. The balance sheet is a financial statement that provides an overview of a company's financial position at a specific point in time. Liabilities are obligations or debts owed by a company, which are classified as current liabilities if they are expected to be settled within one year. In this case, the utility bill for the period ending on June 15 falls within the one-year timeframe and should therefore be reported as a current liability. Therefore, the correct answer is A: \$6,000.

New answer: A: \$6,000.

---

Table 18: Example one for COMPETE. The LLM is impacted by the alternative answer and should abstain.

---

Question: A flashing red traffic light signifies that a driver should do what?

- A: stop
- B: speed up
- C: proceed with caution
- D: honk the horn

Generate a knowledge paragraph about D.

D: Honking the horn is the appropriate action when encountering a flashing red traffic light. This is an important traffic rule that drivers must follow to maintain safety on the roads. When encountering a flashing red light, it is crucial for all drivers to be aware of traffic signals and adhere to their instructions to ensure a smooth flow of traffic and prevent any potential hazards. Honking the horn, on the other hand, is generally used as a way to alert other drivers or pedestrians of potential dangers or to communicate with them in certain situations.

New answer: A: stop.

---

Table 19: Example two for COMPETE. The LLM is not impacted by the alternative answer.

---

Question: Kelsen's theory of law is called pure theory because Kelsen:

- A: Purely discussed jurisprudence only
- B: Defined law in according with morality and purity
- C: Separated law from religion, ethics, sociology and history
- D: Discussed law purely in terms of justice

Generate a knowledge paragraph about D.

Kelsen's theory of law is called pure theory because he discussed law purely in terms of justice. According to Kelsen, law should be analyzed and understood independently of any moral or ethical considerations. He believed that law should be studied and interpreted solely through its internal system of norms, without any reference to external moral values or social contexts.

Kelsen's pure theory of law aimed to provide a "science of law" by separating law from other disciplines such as religion, ethics, sociology, and history. He argued that law should be treated as a self-contained system, with its own hierarchy of norms and principles. This meant that the validity of laws should be determined based on their consistency with higher-level norms, rather than any inherent moral or ethical qualities.

By focusing purely on the concept of justice, Kelsen attempted to provide an objective and neutral framework for understanding and analyzing legal systems. He believed that by divorcing law from subjective moral judgments, his theory would provide a more rigorous and scientific approach to the

New answer: D: Discussed law purely in terms of justice.

---

Table 20: Example three for COMPETE. The LLM is impacted by the alternative answer and should abstain.

---

Question: Suppose you live on the Moon. How long is a day (i.e. from sunrise to sunrise)?

A: about 18 years

B: 24 hours

C: 29 Earth days

D: a year

Generate a knowledge paragraph about A.

Option A: about 18 years

If you were to live on the Moon, a day would last about 18 years. This is due to a phenomenon called tidal locking. Tidal locking occurs when the gravitational force between two objects, such as the Earth and the Moon, causes one object to always face the other with the same side. In the case of the Moon, it is tidally locked to the Earth, meaning one side is constantly facing our planet.

This tidal locking also means that the Moon's rotation period matches its orbital period around the Earth, taking approximately 27.3 Earth days to complete one full revolution. However, despite this prolonged rotation period, it is still necessary to define a day from a sunrise to a sunrise, just as we do on Earth.

Since the Moon does not have an atmosphere to scatter sunlight and bring about the distinct sunrise and sunset we are accustomed to, the definition of a lunar day can be somewhat arbitrary. One way to measure it is based on

New answer: B: 24 hours.

---

Table 21: Example four for COMPETE. The LLM is not impacted by the alternative answer.