
How Do LLMs Ask Questions? A Pragmatic Comparison with Human Question-Asking

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Question asking is a key linguistic and cognitive skill that supports collaboration
2 and diverse social actions. However, large language models (LLMs) often underuse
3 questions in their outputs, leading to misunderstandings or unproductive outputs.
4 Toward bridging this gap, we analyze human questions from a real-world social
5 setting—Reddit—and compare them to LLM-generated questions. We use a
6 pragmatics-based taxonomy of social actions to examine six open- and closed-
7 source model families. Our analysis shows that LLMs fall short of capturing the
8 diversity and balance of human question-asking, with significant differences in
9 question type distributions. Prompting often introduces prompt-specific biases that
10 diverge from human patterns, while the effects of instruction-tuning are model-
11 dependent and inconsistent across social functions. These findings highlight the
12 need for more fine-grained approaches to align LLMs with human-like questioning
13 behavior.

14 1 Introduction

15 Question asking is a hallmark of human intelligence, reflecting critical thinking [Albergaria-Almeida,
16 2011], curiosity [Jirout and Klahr, 2012], and creativity [Acar et al., 2023]. It also underpins
17 collaboration [Sperber et al., 2010] and serves diverse social functions [Stivers et al., 2010]. Yet large
18 language models (LLMs) rarely ask questions in their outputs [Shaikh et al., 2024], even as they
19 become more agentic and integrated into collaborative workflows [Wang et al., 2020]. Prior work has
20 linked this limitation to preference-tuning data, focusing mainly on clarification questions [Shaikh
21 et al., 2024, Zhang et al., 2024, Andukuri et al., 2024] or on a question’s utility for producing answers
22 [Zhang et al., 2024]. While clarification is valuable for problem-solving, human question asking
23 is far more dynamic—seeking information, suggesting alternatives, and challenging assumptions
24 [Stivers et al., 2010]. This highlights a gap in understanding the broader communicative, creative,
25 and curiosity-driven functions of questions.

26 To address this gap, we analyze and interpret question-asking behavior in both humans and language
27 models (LMs). Specifically, we investigate: (1) *How do humans and LMs differ in their use of*
28 *question types?* and (2) *How do current alignment methods affect the question-asking behavior of*
29 *LMs, and do they improve alignment with human patterns?* To answer these questions, we construct a
30 dataset from a Reddit community, `r/NoStupidQuestions`. Using a question type taxonomy based
31 on social actions [Stivers and Enfield, 2010], we categorize the questions posed by both humans and
32 LMs in response to these posts. We then investigate whether aligning to enhance question asking
33 leads to human-like question asking behavior.

34 Our findings reveal that LMs fail to replicate the diversity and balance of human question-asking
35 behavior. Their question distributions differ significantly from those of humans, often focusing
36 disproportionately on rhetorical questions and information requests. Prompting techniques help

Table 1: Question types along with their definitions and corresponding examples

Social Action	Definition	Example
Request for information	Genuine questions seeking information.	"Are you busy tonight?"
Other initiation of repair	Includes open-class repair initiators and partial repeats.	"Huh?" / "He went where?"
Request for confirmation	Asserts a proposition for confirmation.	"So you're coming tomorrow night."
Assessment	Evaluations formatted to seek agreement.	"Isn't it beautiful out today?"
Suggestion/Offer/Request	Suggests, proposes, offers, or requests something.	"Did you want some?"
Rhetorical question	May seek a response but not an answer; often expresses opinion.	"Everything comes out in the wash, doesn't it?"
Outloud	Directed to no one in particular, often quieter, not designed to secure a response.	"Now where are my keys."
Other	Does not fit other categories; action specified separately.	Case-specific action

elicit more questions but frequently introduce prompt-specific biases—such as overproducing certain question types or relying on repetitive lexical patterns—ultimately failing to align with the nuanced characteristics of human questioning. We also observe that the effect of instruction-tuning is highly model-dependent and varies across social functions of questions. Instruction-tuning does not consistently improve human alignment, underscoring the need for a more fine-grained analysis of its impact on different question types.

2 Related Works

Previous studies have identified question-asking as a key deficiency in LLMs [Bai et al., 2024, Shaikh et al., 2024], with most prior work focusing on clarification questions by constructing benchmarks [Li et al., 2025a, Aliannejadi et al., 2021, Guo et al., 2021] and proposing methods for improvement [Kobalczyk et al., 2025, Zhang et al., 2025, Testoni and Fernández, 2024, Zhang and Choi, 2025]. Beyond clarification, other studies have explored question-asking as a tool for downstream applications such as preference elicitation [Li et al., 2025b, Andukuri et al., 2024], information-seeking [Meng et al., 2023], and medical reasoning [Li et al., 2025c]. However, Shaikh et al. [2024] show that prompting interventions, while increasing grounding acts like clarification and follow-up questions, often have minimal or negative effects on alignment with human behavior; replicating Zephyr’s [Tunstall et al., 2023] training further revealed that contemporary preference datasets reduce question frequency due to both the scarcity of questions in training data and annotator dispreference. In this work, we extend the focus beyond clarification and task-specific settings by analyzing three model families at different training stages and investigating LMs’ question-asking abilities through a comprehensive taxonomy of human question-asking [Stivers and Enfield, 2010], highlighting the diverse social functions questions serve in natural communication.

3 Methods

Question types taxonomy By adopting a coding scheme developed for human questions in natural conversation [Stivers and Enfield, 2010], we aim to investigate whether LMs can generate the diverse range of questions humans ask for dynamic purposes. This extends the scope beyond previously researched LLM questioning abilities, such as clarification or follow-up questions. In our analysis, we focus on the social action dimension, as it encompasses diverse social goals of questions, especially fitting in general question asking setting reflected in our current dataset. The taxonomy classifies the social actions performed by a question into seven categories as specified in Table 1. If a question does not fit any of the predefined social action types, it is labeled as "Other/" followed by a description of the specific social action being performed.

Automatic question type annotation and aggregation We automatically label questions from both humans and LLMs using three-shot prompting with OpenAI’s o3 model, which extracts questions from text and classifies them according to Stivers and Enfield [2010]. The full prompt is provided in Appendix B. To evaluate annotation quality, we compare model labels with human annotations on 136 questions¹, obtaining a Cohen’s Kappa of 0.78 and a macro-averaged F1 of 0.79, indicating substantial agreement. Additionally, we define *question proportion* as the average query-level proportion of each question type, computed by assigning each question in a post a weight of $1/n$ when n questions

¹4 from each of 29 model-prompt pairs plus 20 human-generated

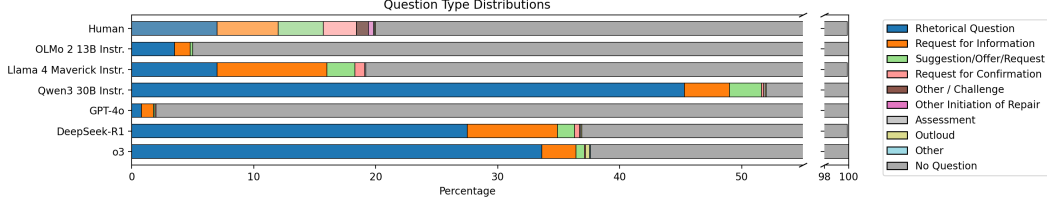


Figure 1: Question type distributions of humans and LLMs. LLMs show reduced diversity and skewed distributions compared to humans. Chi-square tests reveal significant differences in five of six models ($p_{\text{adj}} < 0.05$), except OLMo-2-13B-Instruct ($p_{\text{adj}} = 0.18$).

are present. This metric is used to compare humans and LLMs, as well as among LLM variants at different training stages.

Dataset We collect 3,564 posts and 38,205 comments from the r/NoStupidQuestions subreddit² in 2024. We include all 1,782 posts with multiple question comments and sample an equal number without question comments to balance the dataset. Comments with question marks (excluding those with URLs) are identified as questions. All comments on the selected posts are included, except those removed by administrators. This setup ensures a sufficient and unbiased sample of human-generated questions for comparison with LLMs.

LLM question generation We analyze questions generated by LLMs in response to Reddit posts in our test set. The models include OLMo-2-13B (Base, SFT, and Instruct), LLaMA-4-Maverick-17B-128E (Pretrained and Instruct), Qwen-3-30B (Base, A3B, and Instruct), GPT-4o-2024-08-06, DeepSeek-R1-0528, and o3-2025-04-16.³ We use four prompt variations: (1) Naive, which provides only the Reddit post’s title and body; (2) Question, which encourages asking questions when needed; (3) Taxonomy, which includes our question type taxonomy for guidance; and (4) Imitate, which instructs the models to emulate human Reddit users. Full prompt details are in Appendix A.

4 Results

LLMs fail to match human question diversity and distribution. Figure 1 compares question type distributions of humans and instruction-tuned LMs using the naive prompt (Appendix 4 shows all models). LMs either ask far fewer questions than humans (e.g., OLMo-2-13B-Instruct, GPT-4o) or disproportionately favor rhetorical questions and requests for information (e.g., LLaMA-4-Maverick, Qwen-3-30B, Deepseek-R1, O3). Across all models, Suggestions/Offer/Requests, Requests for Confirmation, Other/Initiation of Repair, and especially Other/Challenge questions (<0.04%) are markedly underproduced. Chi-square tests show significant differences between humans and five of six LMs ($p_{\text{adj}} < 0.05$), with OLMo-2-13B-Instruct as the exception ($p_{\text{adj}} = 0.18$) (Appendix D). Pairwise comparisons among the six LMs reveal no significant differences for 9 of 15 pairs⁴. Overall, LMs fail to align with human question-asking behavior in both diversity and distribution, while exhibiting broadly similar patterns across models.

Prompting is not sufficient for aligning LLM question asking behavior. We examine the effects of three prompt variations—Question, Imitate, and Taxonomy—on aligning LLM question types toward human-like patterns. **Quantitative analysis** reveals statistically significant differences between human and model question type distributions across all prompts ($p_{\text{adj}} < 0.01$, Chi-square tests). As expected, all three prompts increase Requests for Information, Suggestions/Offer/Requests, and Requests for Confirmation—categories typically underused by LMs—across most models (Figures 2a), 5a, 5b). However, the Question prompt disproportionately boosts information requests,

²<https://www.reddit.com/r/NoStupidQuestions/>

³We exclude other baseline methods designed to improve LLM question-asking capabilities, as they are typically trained for specific tasks (e.g., clarification or preference elicitation) and do not generate natural responses for our test set.

⁴Significant differences appear for OLMo–Qwen, LLaMA–Qwen, LLaMA–O3, Qwen–R1, Qwen–O3, and R1–O3 ($p_{\text{adj}} < 0.01$).

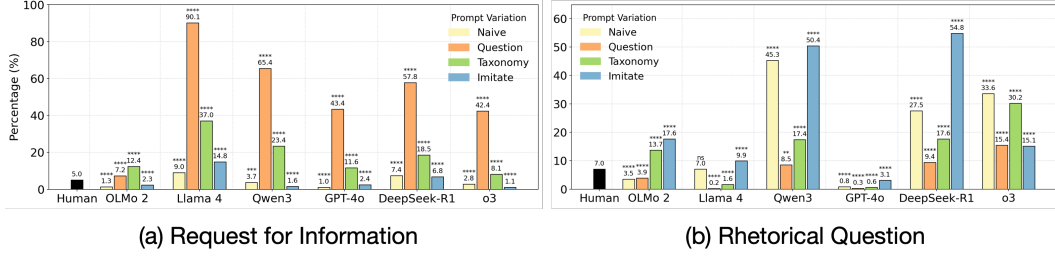


Figure 2: Question type distributions across prompting conditions. LLMs exhibit prompt-induced biases and fail to match human question patterns. Significance levels from two-sided paired t-tests with Benjamini–Hochberg correction are shown above each bar.

exceeding human levels by up to 85.1 percentage points, while the Imitate prompt inflates rhetorical and assessment questions beyond human levels (Figure 2b). None of the prompts consistently increase Other/Challenge questions, which remain below human proportions (Figure 5d).

In our **qualitative analysis** of sampled questions, we find that human Request for Information questions are typically brief and standalone, whereas LMs often embed such questions within or at the end of longer responses. These questions frequently follow templated patterns—such as offering to elaborate on related topics (e.g., “Would you like to know more about . . . ?”), soliciting the user’s opinion (e.g., “What do you think?”), or seeking validation (e.g., “Do any of these perspectives resonate with you?”). Model-generated rhetorical questions differ in purpose from human ones. Both humans and LMs use rhetorical questions for hypophora (e.g., “Am I the audience for his films? Definitely not.”), but humans more often assert opinions (e.g., “Does it even matter?”), whereas LMs rarely do so under the Naive, Question, or Taxonomy prompts. Instead, LMs tend to produce reflection-oriented prompts (e.g., “Ask yourself: What is the evidence for these thoughts?”). Under the Imitate prompt, however, models become highly biased toward assertions (e.g., “Who wouldn’t want to be able to do that?”). Assessment questions also diverge: LMs frequently rely on fixed forms (e.g., “. . . , right?”, “. . . , huh?”), while humans use more varied structures (e.g., “. . . , isn’t it?”, “Aren’t they . . . ?”, “. . . , ay?”). Overall, prompting techniques induce biases that hinder LMs from aligning with the balanced distribution and nuanced characteristics of human questions.

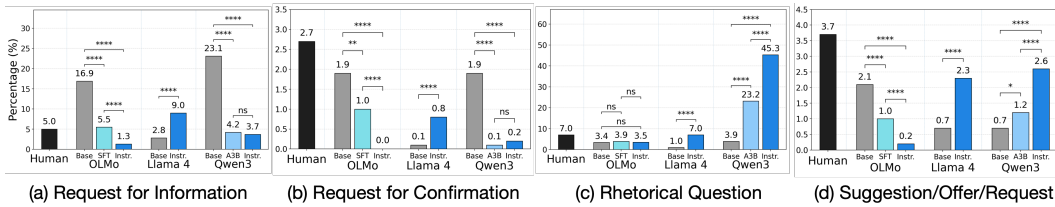


Figure 3: Question type proportions across LLM training stages. The effects of instruction-tuning are family-specific and vary by social action type. Statistical significance from two-sided paired t-tests (Benjamini–Hochberg corrected) is shown above each bar.

The impact of instruction tuning is specific to model and communicative purpose. We investigate the effect of instruction tuning across three LLM families—OLMo-2-13B, LLaMA-4-Maverick, and Qwen-3-30B—with different training stages available for generation. Patterns vary by question type: for epistemic questions (Information Requests, Confirmations), OLMo-2-13B and Qwen-3-30B base models ask more than their SFT or instruction-tuned counterparts ($p < 0.01$), while LLaMA-4-Maverick asks more in the instruct variant ($p < 10^{-4}$) (Figure 3a,b). For Rhetorical and Suggestion/Offer/Request questions, instruction-tuned LLaMA and Qwen models produce more than their base versions ($p < 10^{-4}$), whereas OLMo-2-13B-Base exceeds its tuned variants only for Suggestion/Offer/Requests ($p < 10^{-4}$) and shows no significant differences in rhetorical questions (Figure 3c,d). These results indicate that instruction tuning shapes question-asking behavior in family- and function-specific ways.

References

- Selcuk Acar, Kelly Berthiaume, and Rebecca Johnson. What kind of questions do creative people ask? *Journal of Creativity*, 33(3):100062, 2023.
- Patrícia Albergaria-Almeida. Critical thinking, questioning and creativity as components of intelligence. *Procedia-Social and Behavioral Sciences*, 30:357–362, 2011.
- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail Burtsev. Building and evaluating open-domain dialogue corpora with clarifying questions. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4473–4484, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.367. URL <https://aclanthology.org/2021.emnlp-main.367/>.
- Chinmaya Andukuri, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah D Goodman. Star-gate: Teaching language models to ask clarifying questions. *arXiv preprint arXiv:2403.19154*, 2024.
- Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, and Wanli Ouyang. MT-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7421–7454, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.401. URL <https://aclanthology.org/2024.acl-long.401/>.
- Meiqi Guo, Mingda Zhang, Siva Reddy, and Malihe Alikhani. Abg-coQA: Clarifying ambiguity in conversational question answering. In *3rd Conference on Automated Knowledge Base Construction*, 2021. URL <https://openreview.net/forum?id=S1DZ1o8FsJU>.
- Jamie Jirout and David Klahr. Children’s scientific curiosity: In search of an operational definition of an elusive concept. *Developmental review*, 32(2):125–160, 2012.
- Kasia Kobalczuk, Nicolás Astorga, Tennison Liu, and Mihaela van der Schaar. Active task disambiguation with LLMs. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=JAMxRSXLfz>.
- Belinda Z. Li, Been Kim, and Zi Wang. Questbench: Can llms ask the right question to acquire information in reasoning tasks?, 2025a. URL <https://arxiv.org/abs/2503.22674>.
- Belinda Z. Li, Alex Tamkin, Noah Goodman, and Jacob Andreas. Eliciting human preferences with language models. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL <https://openreview.net/forum?id=LvDwwAgMEW>.
- Shuyue Stella Li, Jimin Mun, Faeze Brahman, Jonathan Ilgen, Yulia Tsvetkov, and Maarten Sap. Aligning llms to ask good questions a case study in clinical reasoning. *CoRR*, abs/2502.14860, February 2025c. URL <https://doi.org/10.48550/arXiv.2502.14860>.
- Yan Meng, Liangming Pan, Yixin Cao, and Min-Yen Kan. FollowupQG: Towards information-seeking follow-up question generation. In Jong C. Park, Yuki Arase, Baotian Hu, Wei Lu, Derry Wijaya, Ayu Purwarianti, and Adila Alfa Krisnadhi, editors, *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 252–271, Nusa Dua, Bali, November 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.ijcnlp-main.17. URL <https://aclanthology.org/2023.ijcnlp-main.17/>.
- Omar Shaikh, Kristina Gligoric, Ashna Khetan, Matthias Gerstgrasser, Diyi Yang, and Dan Jurafsky. Grounding gaps in language model generations. In Kevin Duh, Helena Gomez, and Steven Bethard, editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6279–6296, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.348. URL <https://aclanthology.org/2024.naacl-long.348/>.

- 189 Dan Sperber, Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgi,
190 and Deirdre Wilson. Epistemic vigilance. *Mind & language*, 25(4):359–393, 2010.
- 191 Tanya Stivers and N.J. Enfield. A coding scheme for question–response sequences in conversation.
192 *Journal of Pragmatics*, 42(10):2620–2626, 2010. ISSN 0378-2166. doi: [https://doi.org/10.](https://doi.org/10.1016/j.pragma.2010.04.002)
193 [1016/j.pragma.2010.04.002](https://doi.org/10.1016/j.pragma.2010.04.002). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S0378216610000986)
194 [pii/S0378216610000986](https://www.sciencedirect.com/science/article/pii/S0378216610000986). Question-Response Sequences in Conversation across Ten Languages.
- 195 Tanya Stivers, Nick J Enfield, and Stephen C Levinson. Question-response sequences in conversation
196 across ten languages: An introduction. *Journal of Pragmatics*, 42:2615–2619, 2010.
- 197 Alberto Testoni and Raquel Fernández. Asking the right question at the right time: Human and
198 model uncertainty guidance to ask clarification questions. In Yvette Graham and Matthew Purver,
199 editors, *Proceedings of the 18th Conference of the European Chapter of the Association for*
200 *Computational Linguistics (Volume 1: Long Papers)*, pages 258–275, St. Julian’s, Malta, March
201 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.eacl-long.16. URL
202 <https://aclanthology.org/2024.eacl-long.16/>.
- 203 Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada,
204 Shengyi Huang, Leandro Von Werra, Clémentine Fourier, Nathan Habib, et al. Zephyr: Direct
205 distillation of lm alignment. *arXiv preprint arXiv:2310.16944*, 2023.
- 206 Dakuo Wang, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi,
207 and Qianying Wang. From human-human collaboration to human-ai collaboration: Designing ai
208 systems that can work together with people. In *Extended Abstracts of the 2020 CHI Conference*
209 *on Human Factors in Computing Systems*, CHI EA ’20, page 1–6, New York, NY, USA, 2020.
210 Association for Computing Machinery. ISBN 9781450368193. doi: 10.1145/3334480.3381069.
211 URL <https://doi.org/10.1145/3334480.3381069>.
- 212 Michael JQ Zhang and Eunsol Choi. Clarify when necessary: Resolving ambiguity through interaction
213 with LMs. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *Findings of the Association for*
214 *Computational Linguistics: NAACL 2025*, pages 5526–5543, Albuquerque, New Mexico, April
215 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/
216 2025.findings-naacl.306. URL <https://aclanthology.org/2025.findings-naacl.306/>.
- 217 Michael JQ Zhang, W Bradley Knox, and Eunsol Choi. Modeling future conversation turns to teach
218 llms to ask clarifying questions. *arXiv preprint arXiv:2410.13788*, 2024.
- 219 Michael JQ Zhang, W. Bradley Knox, and Eunsol Choi. Modeling future conversation turns to
220 teach LLMs to ask clarifying questions. In *The Thirteenth International Conference on Learning*
221 *Representations*, 2025. URL <https://openreview.net/forum?id=cwuSAR7EKd>.

222 A Prompts for LLM generations

223 A.1 Question Prompt

```
User Query: {title}
{selftext}
```

```
Converse with the user to deliver an effective response to their
query. Feel free to ask any questions if necessary.
```

225 A.2 Taxonomy Prompt

```
User Query: {title}
{selftext}
```

Respond to the user's query. If needed, consider using an appropriate question from the taxonomy below to guide the conversation:

- Request for Information: Questions genuinely seeking information, with no other primary action apparent.
- Other Initiation of Repair: Questions including open class repair initiators as well as partial repeats.
- Request for Confirmation: Questions asserting a proposition and seeking confirmation.
- Assessment: Evaluations that were formatted to seek agreement.
- Suggestion/Offer/Request: Questions proposing, suggesting, offering, or explicitly requesting something.
- Rhetorical Question: Questions that assert opinions or points, not seeking genuine answers.
- Outloud: Questions not directed toward anyone, without intent to receive a response
- Other: Questions not fitting clearly into any category above. (e.g., Challenge, Pre-invitation, etc.).

227

228 A.3 Imitate Prompt

You are a Reddit user browsing r/NoStupidQuestions.

Post Title: {title}
Post Body: {selftext}

Write a natural and authentic Reddit comment responding to this post.

- Use casual and informal language, like a normal Redditor.
- Avoid sounding like a chatbot or being overly formal.

229

230 B Automatic annotation prompt

231 The full prompt used for automatic annotation is provided below.

You are provided with multiple Reddit post and comment pairs. Each comment
→ contains at least one question.

Your task is to process each pair by performing the following steps:

1. Extract all questions asked to the poster from the comments that meet
→ these criteria:
 - Include questions if they meet any of the following criteria:
 - Are formal questions (grammatically marked as interrogative).
 - Are functional questions (aiming to elicit information, confirmation,
→ or agreement), even if not formally interrogative.
 - Are newsmarks (e.g., "Really?", "Is it?", "Yeah?") as these typically
→ seek confirmation.
 - Exclude questions if they meet any of the following criteria:
 - Solely seek acknowledgment during storytelling (e.g., "And it was a
→ Weight Watchers recipe right?").
 - Appear only as reported speech (e.g., "Then he said, 'Aren't you gonna
→ come over?')").
2. Using the full context (both the original post and the comment), assign
→ each extracted question according to three criteria below. If a single
→ comment includes multiple questions, categorize each one individually
→ into a single category.

232

Social Action:

- Request for Information - Questions genuinely seeking information, with no
 - other primary action apparent. (Ambiguous cases like "Are you busy tonight?"-a potential pre-invitation-should be coded as "Other / Pre-invitation".)
- Other Initiation of Repair - Questions including open-class repair
 - initiators ("Huh?" or "What?") as well as partial repeats ("He went where?"). If it seemed that the repair was more a challenge than an initiation of repair, "Other / Challenge" should be coded.
- Request for Confirmation - Questions asserting a proposition and seeking
 - confirmation (e.g., "So you're coming tomorrow night.").
- Assessment - Evaluations that are formatted to seek agreement such as
 - "Isn't it beautiful out today?" or "She's such a pretty girl, isn't she?"
- Suggestion/Offer/Request - Questions proposing, suggesting, offering, or
 - explicitly requesting something (e.g., "Did you want some cereal?").
- Rhetorical Question - Questions that assert opinions or points, not seeking
 - genuine answers (e.g., "Everything comes out in the wash, doesn't it?").
- Outloud - Questions not directed toward anyone, without intent to receive a
 - response (e.g., "Now where are my keys?" while looking in a bag).
- Other / [Specific Social Action] - Questions not fitting clearly into any
 - category above. Specify explicitly the social action performed after a slash ("/") (e.g., "Other / Challenge," "Other / Pre-invitation," etc.).

Example output:

```
[
  {
    "post_title": "Why do some people eat food directly from the container
    ↪ instead of using a plate?",
    "post_body": "I've noticed that certain individuals, when eating foods
    ↪ like ice cream, chips, or even leftovers, will consume it straight
    ↪ from the container rather than transferring it to a plate or bowl. Is
    ↪ there a specific reason for this behavior, or is it simply a matter
    ↪ of convenience or personal preference?",
    "comment": "Why dirty a plate when you can just eat straight from the
    ↪ container? Less dishes to wash, and sometimes it just feels more
    ↪ satisfying to dig into a tub of ice cream with a spoon. Plus, if
    ↪ you're just snacking or eating alone, who cares, right?",
    "extracted_questions": ["Why dirty a plate when you can just eat straight
    ↪ from the container?", "Plus, if you're just snacking or eating alone,
    ↪ who cares, right?"],
    "logical_semantic_structure": ["Polar Question", "Polar Question"],
    "through_produced_multi_question": "Yes",
    "social_action": ["Rhetorical Question", "Rhetorical Question"]
  },
  {
    "post_title": "Should I get pizza or wingstop today? ",
    "post_body": "I can't decide. I don't feel like leaving so I could
    ↪ doordash wingstop or I could get my ass up and go pick up a pizza
    ↪ since there is a domino's close by ",
    "comment": "Doesn't domino's have wings too? You can get both there.",
    "extracted_questions": ["Doesn't domino's have wings too?"],
    "logical_semantic_structure": ["Polar Question"],
    "through_produced_multi_question": "No",
    "social_action": ["Request for Confirmation"]
  },
  {
    "post_title": "What does cilantro even taste like to normal people?",
    "post_body": "I'm fairly certain I have that Cilantro-Soap gene you'll
    ↪ see people talk about sometimes. It's soapy, overly aromatic, and
    ↪ vile. But not ONCE have I seen a normal person describe the taste of
    ↪ cilantro. Why are you all obsessed with putting it in everything?
    ↪ What does it taste similar to?",
```



```

    "comment": "I hate when it overpowers the dish.",
    "extracted_questions": [],
    "logical_semantic_structure": [],
    "through_produced_multi_question": "N/A",
    "social_action": []
  }
]

Return the following JSON list, with the empty fields filled in for each JSON
↪ object.

{
  "id": "{id}",
  "post_title": "{post_title}",
  "post_body": "{post_body}",
  "comment": "{comment}",
  "extracted_questions":
  "logical_semantic_structure":
  "through_produced_multi_question":
  "social_action":
},
...
{
  "id": "{id}",
  "post_title": "{post_title}",
  "post_body": "{post_body}",
  "comment": "{comment}",
  "extracted_questions":
  "logical_semantic_structure":
  "through_produced_multi_question":
  "social_action":
}]

```

234

235 C Model and Human Generations Table

Field	Content
Post Title	Why do some people eat food directly from the container instead of using a plate?
Post Body	I've noticed that certain individuals, when eating foods like ice cream, chips, or even leftovers, will consume it straight from the container rather than transferring it to a plate or bowl. Is there a specific reason for this behavior, or is it simply a matter of convenience or personal preference?

GPT-4o

There are several reasons why people might choose to eat directly from the container instead of using a plate or bowl. Here are a few possibilities:

1. **Convenience:** Eating directly from the container can save time and effort. It eliminates the need to wash extra dishes, which can be appealing, especially if someone is in a hurry or doesn't want to deal with cleanup.
 2. **Portion Control:** Some people might find it easier to control their portions by eating directly from the container. For example, with ice cream, they might feel they can stop eating when they feel satisfied rather than finishing a pre-scooped bowl.
 3. **Casual Setting:** In informal or relaxed settings, people might feel more comfortable eating directly from the container. It can feel more laid-back and less formal than using a plate.
 4. **Habit or Preference:** For some, it might simply be a habit or personal preference. They might enjoy the experience of eating directly from the container or find it more satisfying.
 5. **Social Influence:** If someone sees others around them eating directly from containers, they might be more inclined to do the same. It can be a social norm in certain groups or situations.
- Do any of these reasons resonate with what you've observed, or is there a specific context you're curious about?

OLMO-instruct

Eating directly from the container seems to be a blend of convenience, habit, and individual preference. Here are a few reasons why some might opt for this method:

1. **Convenience:** Sometimes, the thought of washing extra dishes or taking the time to transfer food into another container just isn't appealing. Especially if someone is eating quickly or on-the-go, opting for the container is a time-saving choice.
2. **Portion Control:** For some, it might be a way to monitor how much they are eating, especially if the container has clear serving sizes marked on it.
3. **Taste Preference:** There's a belief that eating directly from the container can change the taste experience slightly due to the container's material or the way it maintains temperature (e.g., ice cream from a tub).
4. **Easier Access:** Containers, especially those designed for takeout or single servings, are often more portable and provide direct access to the food, which can be appealing.
5. **Mindless Eating:** This behavior can also be a part of mindless eating, where the individual is not fully conscious of how much they are eating because they are not seeing the food portioned out.
6. **Personal Habit:** For some, it could simply be a long-established habit that they grew up with or formed over time.

It's important to note that occasionally eating from containers is generally harmless. However, making it a frequent habit might lead to overeating without recognizing it, as there's less visual cue to how much has been consumed. Furthermore, it's critical to wash containers properly to prevent the spread of germs and bacteria.

So, while eating from the container is often a matter of convenience and personal preference, it's also worth considering how it aligns with one's eating habits and health goals. Would you like to delve further into any aspect of this topic?

OLMO

Assistant: This is an interesting observation! While the answer may vary from person to person, here are a few possible reasons why some people eat food directly from the container:

1. **Convenience:** It might be quicker and more convenient to eat straight from the container, especially if they are in a hurry or multitasking.
2. **Portion Control:** Eating directly from a container can help with portion control, particularly for foods that are often consumed in small servings, like nuts or candies.
3. **Sanitation:** Some people might believe that eating directly from the container is cleaner or less prone to contamination than using a plate or bowl.
4. **Personal Preference:** Ultimately, it might just be a personal preference or habit for some individuals. They might find it enjoyable or comfortable to eat straight from the container.

Do any of these reasons resonate with you, or are there other factors that you think might contribute to this behavior?

DeepSeek-R1

That's a great observation! Eating straight from the container is surprisingly common, and there are several reasons why people do it, often a mix of practicality and psychology:

1. **Convenience & Laziness (The Big One):** **Less Cleanup:** This is probably the top reason. Using a plate or bowl means washing an extra dish. Grabbing the tub, bag, or box and a spoon/fork means only cleaning that one utensil (or sometimes none, if it's finger food like chips). **Speed:** It's simply faster. No need to transfer food, especially when you just want a quick snack.
2. **Portion Control (or Lack Thereof):** **"Just a Taste" Mentality:** People often intend to have just a little bit. Pouring chips into a bowl feels like committing to a larger portion. Eating a few straight from the bag feels more contained (even if it often leads to eating more!). **Avoiding Over-Serving:** Conversely, for some *trying* to control portions, eating directly *can* be a (sometimes flawed) strategy. "I'll just have two spoonfuls of ice cream from the tub" avoids the temptation of scooping a huge bowl. (Though it often backfires!).
3. **Sensory & Experiential Factors:** **Texture/Temperature:** Ice cream straight from the tub can feel denser, colder, and more satisfying to some than when slightly melted in a bowl. Chips might stay crisper in the bag. **Mindless Eating:** Eating from a large container can facilitate zoning out while watching TV, reading, or working. The food is just *there*, requiring minimal attention. **Comfort/Indulgence:** There's a certain rebellious or comforting feeling about eating directly from the source, especially with treats like ice cream. It feels less formal, more indulgent.
4. **Context & Situation:** **Alone vs. With Others:** This behavior is *far* more common when eating alone. Doing it in front of others can be seen as rude or unhygienic (especially with shared containers). **Fatigue/Late Night:** When tired (e.g., late-night snack), the effort of getting a plate feels like too much. **"My House, My Rules":** When people are in their own private space, social norms relax. They feel free to do what's easiest/most comfortable for them.
5. **Habit & Upbringing:** For some, it's simply how they've always done it or how they saw family members do it. It becomes the default.

To make my answer even more relevant to you:

Is there a specific food or situation where you've noticed this most? (e.g., Is it always ice cream? Only when people are alone? With shared snacks?) **Are you asking out of curiosity, or is it something you do yourself and wonder about?** (Understanding your perspective helps tailor the answer!)

In essence: While convenience and laziness are major drivers, it's often layered with psychological factors like perceived portion control, sensory preference, and the comfort of informality, especially in private settings. It's usually not *just* one reason, but a combination that makes the path of least resistance (and fewest dishes) appealing!

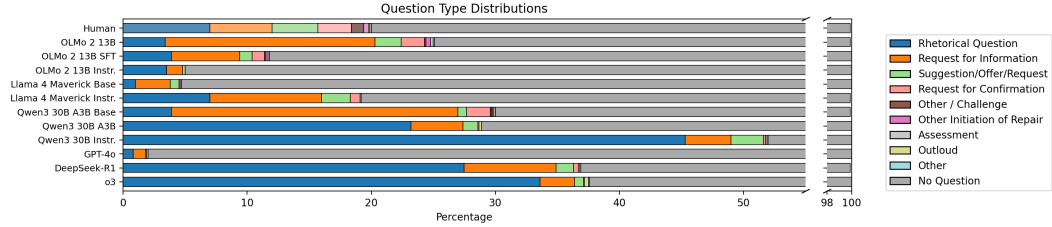


Figure 4: Distributions of question types produced by all LLMs under the naive prompt and the humans. Compared to humans, LLMs exhibit limited diversity and skewed distributions, often overproducing Rhetorical and Request for Information questions while underproducing suggestions, confirmations, and challenges.

Human Perhaps an equally valid question is the inverse. Why do certain individuals transfer their food to a plate or a bowl before eating it? Is there a specific reason for this behavior, or is it simply a matter of convenience or personal preference?

D Chi-squared test details

We compare each model’s question-type distribution with the human’s using a two-stage cluster-bootstrap Chi-squared test. We retain the six question types that reach $\geq 1\%$ frequency in any system—Request for Information, Request for Confirmation, Rhetorical Question, Suggestion/Offer/Request, Assessment, and Other / Challenge. Each of 10000 bootstrap replicates resamples the 3564 posts in the test-set with replacement; within every sampled post, we draw one response per system and tally the six labels. The resulting Monte-Carlo p-values are FDR-corrected (Benjamini–Hochberg). Cluster resampling controls within-post dependence among the responses.

E Question Type Distributions Across All Models

In Figure 4, we present the full distributions of question types generated by all evaluated LLMs under the naive prompt and the humans. This expanded view reinforces our earlier findings that LLMs exhibit less diverse and imbalanced question-asking behavior compared to humans.

F Additional question type comparisons across prompting conditions

Figure 5 presents additional comparisons of question type distributions between humans and the LLMs across the various prompting techniques. These comparisons focus on Suggestion/Offer/Request, Request for Confirmation, Assessment, and Other/Challenge questions, which were not covered in the main figure.

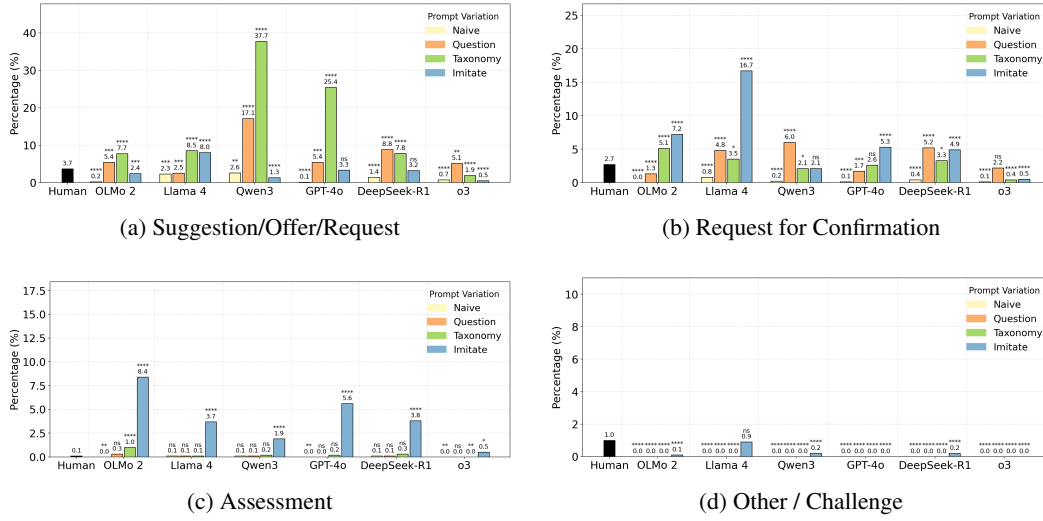


Figure 5: Additional comparisons of question type distributions between humans and LLMs across different prompting conditions. Statistical significance of differences between human and model outputs is indicated above each bar, based on two-sided paired t-tests with Benjamini–Hochberg correction.