

---

# Zero-shot Conversational Summarization Evaluations with small Large Language Models

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Large Language Models (LLMs) exhibit powerful summarization abilities. How-  
2 ever, their capabilities on conversational summarization remains under explored.  
3 In this work we evaluate LLMs ( $\sim 10$  billion parameters) on conversational sum-  
4 marization and showcase their performance on various prompts. We show that the  
5 summaries generated by models depend on the instructions and the performance of  
6 LLMs vary with different instructions sometimes resulting steep drop in ROUGE  
7 scores if prompts are not selected carefully. We also evaluate the models with  
8 human evaluations and discuss the limitations of the models on conversational  
9 summarization.

## 10 1 Introduction

11 Large Language Models (LLMs) have been shown to exhibit impressive multitasking abilities  
12 including summarization [28]. These LLMs show great promise for automatic summarization and  
13 achieve impressive zero-shot performance [56, 48]. LLMs are ‘prompted’ to solve a task such as  
14 summarization by issuing an instruction (e.g., Generate a summary for this dialogue, Summarize:  
15 etc.) and an input (conversation, news article<sup>2</sup>, queries etc.). These instructions are rather simple  
16 natural language phrases appended as a prefix or postfix to the input to assist the LLMs in carrying  
17 out the task. Generating such instructions and prompting can be challenging for novices [53]. The  
18 brittleness of the LLMs in processing the prompts consistently is one such reason that causes a  
19 challenge [57, 51, 33, 26]. A minor variation in the prompt syntax can result in vastly different  
20 results. Prompt engineering has thus evolved to design approaches for the LLMs to produce the  
21 desired results consistently. Several resources exist for prompt engineering. Promptsource is one  
22 such popular repository consisting of 2000+ prompts for natural language prompts for various natural  
23 language tasks including summarization. [4]. In spite of these advances, the brittleness exist and  
24 more so in smaller open-sourced LLMs. One possible solution to tackle the brittleness could be to  
25 further finetune the models using the instructions for a task coupled with data points from a dataset at  
26 the developer’s disposal. This however could result in loss of generalization abilities of the LLMs  
27 even within a task (for instance, finetuning on conversational summarization from a technical design  
28 meeting dataset could result in drop in summarization capabilities on chit-chat data). In this work,  
29 we finetune the LLMs on a technical council meetings dataset and evaluate the performance on a  
30 chit-chat dataset and discuss the results.

31 This paper arises from a need for evaluating open-sourced LLMs [26] (within  $\sim 11$  billion parameters)  
32 for conversation summarization and their robustness to different prompts (The rationale for selecting  
33 models within  $\sim 11$  b parameters is discussed in detail in Sec 4.1). These specific sized LLMs provide  
34 the capability run at the edge on PCs or mobile systems preserving the privacy of the user data while  
35 benefiting the user with productivity gains. These models need to be benchmarked against the existing  
36 conversational summarization datasets by measuring their abilities on different prompts. We motivate

37 this work with the following research questions:  
38 1) How important is prompt selection for conversational summarization task? How do LLMs less  
39 than  $\sim 11$  billion parameters fare on different prompts varying slightly in syntactic structure? How do  
40 these models fare in human evaluation?  
41 2) Can finetuning the LLMs for conversational summarization a task but on a different domain  
42 improve zero shot summarization capabilities?  
43 We also discuss some limitations of the models specially focusing on hallucinations and role attribu-  
44 tions in the summaries.

## 45 **2 Related work**

46 Recent work on benchmarking the LLM abilities for news articles summarization [56] is the most  
47 related our work. The authors highlight there exists a fundamental limitation with the existing news  
48 articles summarization datasets and found that the target/gold-truth summaries were sub par for  
49 comparative evaluations. To overcome the problem the authors hired freelance writers to generate  
50 target summaries and pitch the summaries generated by the models are pitched against those written  
51 by expert freelancers. The authors found that the summaries generated by the models were on-par  
52 with the summaries generated by the experts [56] which in turn is better than the target summaries. It  
53 is important to note that the LLMs used in this work range in parameter count from 350M to 175  
54 billion [7, 16, 36, 54, 5]. In this work, we’re focused on benchmarking the LLMs on the task of  
55 conversational summarization.

56 It has been well established that training the models on instructions improves the generalization abili-  
57 ties in the LLMs (see for e.g., [52, 34, 39, 7, 36, 21, 59]. Several resources such as Promptsources[4],  
58 Natural Instructions [34], Flan [52] and, Super Natural Instructions [50] have been developed in  
59 recent times to provide researchers and engineers alike with resources to study, find and leverage  
60 instructions for a task of interest. In this work, we choose the prompts from Promptsources<sup>1</sup>.

## 61 **3 Data**

62 Several promising datasets exist for the task of conversational summarization [35, 58, 8, 25, 11,  
63 23, 20]. In this work we utilize the MeetingBank dataset [23] for instruction tuning purposes. The  
64 MeetingBank dataset is relatively new and serves as a great resource for evaluating the finetuning  
65 impacts on the models on a specialized domain. The dataset contains transcriptions from various  
66 United States cities or municipality meetings. We then evaluate the models in a zero-shot manner on  
67 Samsun corpus [20] and Dialogsum datasets [10]. We utilize the validation splits from Huggingface  
68 datasets repository for our study purposes.

69 For the purpose of finetuning the model we split the Meetingbank dataset into 95% training and 5%  
70 validation split and utilize the training set for instruction tuning. For the purposes of instruction tuning  
71 we utilize Promptsources[4] which is a crowdsourced valuable resource consisting of thousands of  
72 prompts contributed for numerous NLP tasks and datasets. We utilize the prompts for conversational  
73 summarization section (Table 4). We further instruct-tune the models using a pipeline similar to  
74 T0 [39]. For training & inferencing each transcription input is converted to a prompt for LLMs by  
75 ‘apply’ing the template with the conversations using Langchain [9].

## 76 **4 Models**

77 Conversational summarization has increasingly gained interest in a post-covid world with the nature  
78 of workplace becoming more hybrid [24]. This has amplified the need for a conversational/meeting  
79 summarization systems to run on machines with limited compute capabilities with democratization  
80 aspirations.

### 81 **4.1 Choice for parameter counts**

82 Our choice of models with  $\sim 11$  billion parameters is mainly motivated by the constraint imposed by  
83 the cost and, hardware limitations and a motivation to study less-carbon emitting approach of model

---

<sup>1</sup>we study Promptsources in this work due it’s high activity and star ratings (compared to similar resources).

84 development. These model sizes are also conducive to running on PCs and mobile systems, however  
85 there is an added impact of bigger models and more data needed for inferencing which is observed in  
86 the results noted below in the Hardware Inference section using an LLM as an example workload.  
87 **Hardware Training:** Fine-tuning a 7 billion parameter model takes approximately 66 GB of memory  
88 (RAM) at 32 bit precision and a batch size of 1 (sequence length of 2048) (see memory computation  
89 in [29]) which is slightly less than a high end GPU available in the market. While several optimization  
90 approaches exist which significantly reduce the compute requirements (e.g., Quantization [18]), these  
91 are accompanied by often deteriorated performance or comparable in the best case. Running an  
92 inference locally on client devices (without use of cloud APIs/services) even with quantization of  
93 models is memory intensive <sup>2</sup>. While bigger is often better in LLMs, it is important to study the trade-  
94 offs in models whose parameter counts are still tractable and can be deployed for inferencing (and  
95 even train) using reasonable compute budgets. **Hardware Inference:** Running inference instances  
96 is also an important consideration. For instance, we consider a GPT variant<sup>3</sup> on a state-of-the-art  
97 PC (Intel Raptor Lake 28W Core i7 system (1380P) configured with 32GB of memory and 512  
98 GB of storage). For running an inferencing pipeline that includes document ingestion, encoding,  
99 followed by autoregressive decoding with a 7B parameter model ([47]) generated a bandwidth of  
100 10 GB/s. First inference with model loading was significantly slower ~0.5 FPS. About 1/3rd of  
101 memory was consumed by text encoding and LLM operation and reduced the headroom for other  
102 applications. Optimizations to use compressed models with reducing memory interactions is the key  
103 to use these models effectively on PCs and mobile systems. Models with any greater parameter count  
104 would pose additional inferencing challenges. **Carbon emissions:** Training a 70B emits 833% more  
105 carbon and 13B parameter model emits 100% more carbon than the 7B parameter model. While, the  
106 developing organization offset the carbon, it might not be feasible for all[45]. With these experiments  
107 & observations we arrived at our model evaluations to those mentioned below. In the future we plan  
108 to extend the evaluations to models of greater parameter count.

## 109 4.2 Models selected

110 **Llama** is a foundational language model developed with publicly available datasets by Gen AI,  
111 Meta. Llama 2 [45] is a more recently released open source class of foundational large language  
112 models pretrained & finetuned with more data and higher context lengths than its predecessor  
113 Llama 1[44]. Llama 2 models achieve impressive performance on numerous public benchmarks.  
114 We choose Llama 2 as one of the models for our benchmarking purposes. The family also includes  
115 Llama-2-chat model which is developed using supervised finetuning and then iteratively refined  
116 using Reinforcement Learning with Human Feedback (RLHF) [36]. **Falcon** [43] is a family of  
117 generative LLM by Technology Innovation Institute which achieves impressive performances on  
118 many of the LLM benchmarks. The models are trained predominantly on RefinedWeb [37] along  
119 with other data from the web. The instruct version of the models are further finetuned on instructions  
120 and conversational data. **Alpaca** [41] is a instruction-tuned model based on Llama [44] which is  
121 instruction tuned using the techniques from self-instruct [49] paradigm. The alpaca model was  
122 found to exhibit behavior similar to text-davinci-003 model [1]. **Flan-T5** are a class of language  
123 models [12] where encoder-decoder transformer-based model [38] are finetuned on 1836 tasks using  
124 instruction-finetuning and chain-of-thought finetuning paradigms. The models exhibit state-of-the-art  
125 abilities in various benchmark tasks. We also benchmark the results using **GPT3.5 turbo** [2] which is  
126 a an improved GPT-3 model with instruction following ability on Natural language and code. While,  
127 the parameter count of the model is in hundreds of billion, the purpose of including the model is for  
128 comparison.

## 129 5 Experiments

130 Our goal is compare the model performance on conversational summarization (Samsun & Dialogsum)  
131 and then evaluate the model performance on different prompts. There are 6 prompts in Promptsources  
132 for conversational summarization (Table 4) that we then apply using Langchain on the each con-  
133 versation as input to the models. We then extract the output from these models as summaries. We  
134 additionally finetune existing instruction-tuned models- (Flan & Lama2-7b-chat) using task specific  
135 instructions and the data from Meetingbank dataset (Models marked as \*\*). While instruction tuning

<sup>2</sup><https://github.com/ggerganov/llama.cpp>

<sup>3</sup><https://github.com/imartinez/privateGPT>

136 has been shown to improve the overall performance in the model generalization abilities [34], we  
 137 intend to explore if the task performance (conversational summarization) utilizing the data from a  
 138 related but different dataset can further provide a boost.

139 We report Rouge scores (R1, R2, RL) [32] and BertScore [55] on the summaries generated by the  
 140 models comparing against the available ground truth summaries. We additionally report Length Ratio  
 141 (LR), Novel uni/bi (N1, N2) gram metrics which can be evaluated automatically on the Dialogsum  
 142 and the Samsun datasets which the models haven't seen during our training process <sup>4</sup>. We compute  
 143 Length Ratio = length (tokens in summaries generated)/length (tokens in the conversations), N-gram  
 144 novelty = length (N-gram tokens in summaries generated  $\cap$  tokens in conversation)/ length(tokens in  
 145 the conversations).

146 Target summaries don't always serve as a reliable ground truth for evaluation [56]. Summarization  
 147 task is subjective and also information can be presented in several different forms (sentences &  
 148 format). However, the summaries in the dataset themselves can be relied upon as one of the many  
 149 possible summaries for a given conversation. We thus experiment with humans to compare the  
 150 model generated summary with the target and report the findings here. We developed a web-interface  
 151 (Figure A.4) for collecting human preferences between target and model generated summary. The  
 152 conversation is displayed to the users along with two summaries, one model generated and the other  
 153 reference/target summaries. The summary orders are jumbled and the raters are not aware of the  
 154 source of summaries. The raters were volunteers and employees of anonymous and had education  
 155 levels of graduate degree and were aware of the task. Each user was requested to rate as many  
 156 summaries as they could. We obtain comparative ratings for Overall, Novelty, Informativeness and  
 157 Conciseness of the summaries. To further understand if the summaries contains hallucinations, correct  
 158 role attribution and if the summaries had misrepresented gender attributes we developed a similar  
 159 interface, but this time requesting the human evaluators to rate only the generated summaries if they  
 160 consist hallucinations, if the roles are properly attributed in the summaries and if the genders are  
 161 properly attributed in the summary (if applicable).

## 162 6 Results

163 **Do Instruction tuned models perform generate better summaries?** We can observe  
 164 from Table 1 that when compar-  
 165 ing the instruction tuned mod-  
 166 els (Alpaca-7b vs Llama-7b,  
 167 Falcon-7b-instruct vs Falcon-  
 168 7b) with just pre-trained model,  
 169 the instruction tuned models  
 170 generate summaries with bet-  
 171 ter Rouge-scores, BertScores,  
 172 more concise and generate  
 173 fewer novel 1-gram and 2-gram  
 174 (due to fewer hallucinations).  
 175 We also find that the sum-  
 176 maries generated by the instr-  
 177 uction tuned models are signifi-  
 178 cantly better w.r.t Rouge scores  
 179 and BERT scores ( $p < 0.001$ ,  
 180 Wilcoxon rank-sum test) than  
 181 the summaries generated by the non-instruction tuned models. To arrive at this conclusion we  
 182 compared sample-wise Rouge & BERT scores in separate experiments between the non-instruction  
 183 tuned models and their instruction tuned counterparts. This implies that instruction tuning helps  
 184 improve the conversational summarization performance which is consistent with the literature (on  
 185 a different task) [34]. Humans also rated the summaries from instruction tuned model higher than  
 186 pre-trained models (Figure 1 ).

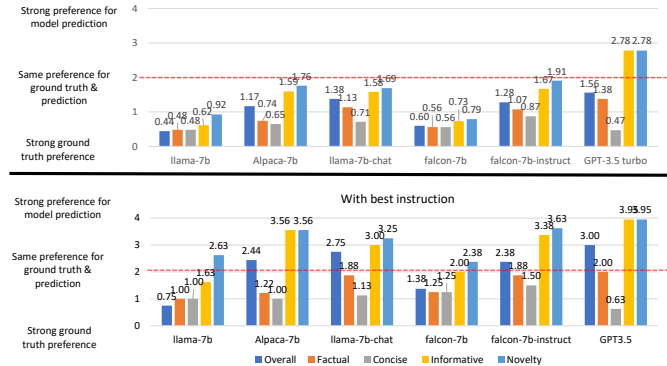


Figure 1: Shows the evaluation of the LLMs on various metrics for conversational summarization with any instruction from Promptsources vs. oracle selected best instruction.

<sup>4</sup>T0pp model however have been exposed to Samsun dataset off-the-shelf. We cannot however guarantee that the models have not been exposed to these datasets due to limited information available.

189 **Does further task-specific instruction tuning help models perform better?** We find that further  
 190 tuning the models for a task (summarization) even on a different dataset  
 191 (Meetingbank) slightly improves the performance of the models. In  
 192 Table 1 we identify task-specific instruction tuned model with ‘\*\*’  
 193 vs available off-the-shelf instruction tuned checkpoints ‘\*’ (Flan-  
 194 T5-`{MODEL_SIZE}`\* vs Flan-T5-`{MODEL_SIZE}`\*\* , Llama-7b-chat\*  
 195 vs Llama-7b-chat\*\*). We do not perform any experiments to verify if the  
 196 model’s performance on a more general abilities since our goal is to leverage  
 197 general LLMs for summarization task alone. This slight change in the summary  
 198 quality was also resulted in increase in preferences over target in human-evaluation  
 199 overall and even on their factual and conciseness (Flan-T5-base\* vs. Flan-t5-base\*\*,  
 200 Flan-T5-small\* vs. Flan-T5-small\*\* in Figure 2 ). This is important, since we’re not  
 201 pitching the models one against the other, but rather against the target summaries.  
 202  
 203  
 204  
 205  
 206  
 207  
 208  
 209  
 210  
 211

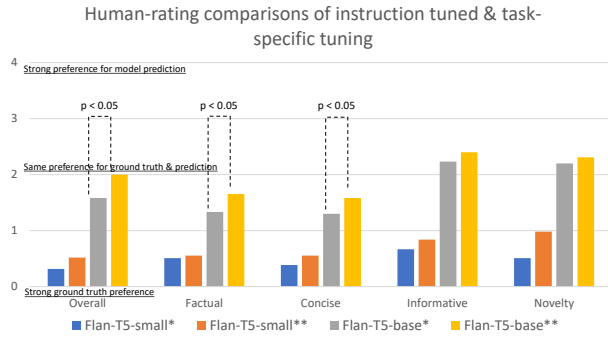


Figure 2: Shows the human evaluations on various metrics comparing instruction tuned model \* and task-specific fine-tuned model \*\* demonstrating improvements.

212 **Do summary qualities differ across instructions and models?** We also find that the model’s  
 213 summarizing abilities are highly dependent on the instruction used and the model itself. For instance,  
 214 ‘Given the above dialogue write a summary’ as an instruction gives 8% better Rouge-L scores than  
 215 ‘Generate a summary for this dialogue’ in Falcon-7b instruct tuned model compared to ‘Llama-7b’  
 216 which gives the best Rouge scores from the prompt ‘Given the above dialogue write a summary’.  
 217 We also tested this variable performance using ANOVA test (single factor,  $p < 0.0001$  ( $F=34.56$ ,  
 218  $df=6$ ,  $MS=0.029$ ) and find that the difference in Rouge scores are significantly different across  
 219 prompts. Figure 3 shows R-L box-plots for each prompt on SAMSUM corpus. We further find that  
 220 the variability in instruction tuned models is slightly lower compared to non-instruction tuned models.  
 221 This result highlights the need for future work to reduce the gap between the model performance  
 222 across prompts which could make it easier for non-experts to query the LLMs. We also find that this  
 223 variability reduces with increase in model parameters.

224 **If we were to use the best prompt, how do LLMs compare with target summaries?** We selected the  
 225 prompt yielding the highest Rouge-L scores (oracle selection) and then compare the human ratings for  
 226 the best performing prompt for each of the model. Upon selecting the best prompt we find that all the LLM  
 227 generated summaries are rated more novel than the target summaries (Figure 1 under ‘With best instruc-  
 228 tion’). All the instruction tuned  
 229 models (Alpaca, Llama-chat, Falcon-instruct) were rated more informative and were preferred  
 230 overall compared to the target summaries. However, the target summaries are rated most factually  
 231 consistent and concise compared to the generated summaries. Lower rating for the conciseness  
 232 were more obvious when we compare the LR of the generated summaries vs the target summaries  
 233 (Dailogsum target summaries LR = 14.85%, Samsun LR = 17.94%) compared to the summaries  
 234 length of the model generated summaries. Lower factual ratings can be explained by presence of  
 235 hallucinations and other errors. Table 3 shows best-possible Rouge and BERT scores on the oracle  
 236 selected best prompts used once per dataset.

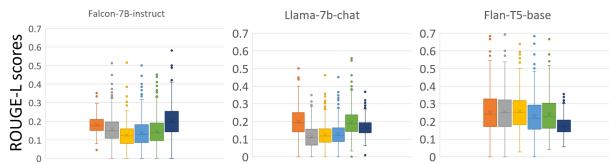


Figure 3: Box plot of Rouge scores of each sample from Samsun dataset showing variability for each prompt (Table 4). Each color represents different prompt.

237  
 238  
 239  
 240  
 241  
 242  
 243  
 244

Model	Samsun							Dialogsum						
	R-1	R-2	R-L	B-Score	LR	N1	N2	R-1	R-2	R-L	B-Score	LR	N1	N2
Llama-7b	0.14	0.02	0.10	0.701	56.58	47.06	85.70	0.16	0.02	0.11	0.736	40.24	31.22	61.87
Falcon-7b	0.14	0.02	0.10	0.703	58.80	48.74	89.68	0.16	0.02	0.11	0.737	36.32	26.61	54.37
Alpaca*	0.18	0.04	0.13	0.721	35.19	27.82	52.80	0.17	0.03	0.13	0.740	21.78	15.06	32.29
Falcon-7b-inst*	0.21	0.05	0.15	0.739	33.96	27.82	52.63	0.22	0.05	0.15	0.761	24.99	18.10	38.80
Flan-T5-small*	0.42	0.19	0.33	0.844	24.38	8.71	27.25	0.22	0.07	0.17	0.742	17.11	5.54	<b>15.62</b>
Flan-T5-small**	0.43	0.20	0.33	0.845	24.45	<b>8.70</b>	27.33	0.23	0.07	0.18	0.741	17.16	5.63	15.86
Flan-T5-base*	0.46	0.23	0.36	0.856	24.29	9.14	28.01	0.31	0.11	0.24	0.785	18.82	<b>4.80</b>	15.93
Flan-T5-base**	0.47	0.23	0.37	0.857	24.23	9.09	27.92	0.32	0.12	0.25	0.787	18.58	4.82	15.97
Flan-T5-large*	0.44	0.19	0.35	0.849	18.38	10.24	25.25	0.29	0.09	0.23	0.793	13.23	7.59	18.75
Flan-T5-large**	0.46	0.20	0.36	0.851	18.38	10.24	24.22	0.30	0.10	0.24	0.793	13.24	7.52	18.72
Flan-T5-xl*	0.48	0.26	0.39	0.862	24.37	9.87	28.21	0.37	0.16	0.29	0.824	18.10	5.71	18.13
Flan-T5-xl**	0.49	0.26	0.40	0.863	24.41	9.92	28.27	0.38	0.17	0.31	0.827	18.17	5.71	18.13
Llama-7b-chat*	0.18	0.04	0.13	0.721	32.41	25.65	48.58	0.21	0.05	0.15	0.765	31.46	22.54	47.39
Llama-7b-chat**±	0.18	0.05	0.14	0.721	31.66	25.03	47.45	0.21	0.06	0.15	0.760	30.54	22.17	46.17
GPT-3.5 turbo	0.14	0.06	0.10	0.746	189.27	80.86	195.89	0.14	0.05	0.11	0.741	140.31	62.33	150.97

Table 1: Evaluation of models on the summarization task. \* indicates the models are instruction tuned, \*\* means the models were task instructed tuned for summarization task. Llama-7b-chat was finetuned using LORA[22].

## 7 Discussions & Future work

**Do low human ratings for factual ratings correlated to hallucinations and role attributions?** Hallucinations a challenge to be addressed in LLMs. We collected human ratings to further analyse the cause for lower rating for the factual ratings in LLMs. We requested human raters to mark (Yes/Now) if the summary contained irrelevant phrases not related to the input. We sampled 286 summaries across models randomly from the dataset for this small scale study. We also requested human raters to answer yes/no question if the the summaries had incorrect role attributions and if there was assign/misrepresented gender pronouns. We find that the presence of hallucinations is strongly negatively correlated to the ratings of factual and overall ratings (PCC,  $\rho = -0.77$ ,  $\rho = -0.71$  resp.). The ratings are also weakly/moderately negatively correlated to incorrect speaker attributions ( $\rho = -0.31$ ). The raters were more forgiving to incorrect gender attributions (he instead of she) towards overall summary rating ( $\rho = -0.18$ ). This means that hallucinations and incorrect role attributions could've played a more important role in lower factual and overall ratings. Since, our sample size was low we cannot conclude with conviction. This remains an interesting area of exploration.

**Are the models with higher overlap with target summaries or input conversations rated more informative?** Rouge (p) between the input and summary generated measures what fraction of the conversation was covered in the summary. We measured the Pearson's correlation co-efficient (PCC) between these Rouge (p) scores (Table 5) and the human preference ratings to measure if any correlation exists. We find moderate correlation (PCC,  $\rho = 0.46$ ) for the informativeness between the models Rouge scores and the preference ratings. This implies that the raters tended to mark models with higher input coverage as more informative. Additionally, we also found that high rouge (p) scores correlated moderately with Overall and Factual ratings. This was also true for informativeness. This is however not all that surprising but a reinforcing indicator that summaries need to be concise while increasing the overall coverage.

Higher novelty ratings ('Select the summary which you think is a more novel way to present the conversation') was an interesting observation. The LLMs consistently performed better on novelty preferences compared to the target. We believe that novel ways of presenting the summaries (e.g., as question answering exchanges, shorter rephrased conversation versions, alternative narrations) along with novel phraseology could be contributing towards higher novelty ratings. We additionally note that the human raters being experts (have good knowledge of technology and recent advancements) rather than novices could've contributed to this. We intend to explore this in our future work.

Conversational summarization is an important application area of the LLMs and remains an interesting area of research. In this work we evaluated various LLMs on conversational summarization datasets and showed that the instruction tuning yields better results. Given the parameter sizes of the LLM models evaluated, this outcome has interesting applications to deployment of <11B parameter LLMs on PC and mobile systems. We also showed that further task-specific finetuning could yield more promising results. We further analysed the summaries via human and automatic evaluations and discussed the limitations including hallucinations and incorrect role attributions.

## 283 References

- 284 [1] Open AI. Openai chatgpt, 2023. URL <https://openai.com/blog/chatgpt>.
- 285 [2] Open AI. Openai gpt-3.5 turbo, 2023. URL <https://openai.com/blog/gpt-3-5-turbo-fine-tuning-and-api-updates>.
- 286
- 287 [3] Alex Andonian, Quentin Anthony, Stella Biderman, Sid Black, Preetham Gali, Leo Gao, Eric  
288 Hallahan, Josh Levy-Kramer, Connor Leahy, Lucas Nestler, Kip Parker, Michael Pieler, Jason  
289 Phang, Shivanshu Purohit, Hailey Schoelkopf, Dashiell Stander, Tri Songz, Curt Tigges, Ben-  
290 jamin Thérien, Phil Wang, and Samuel Weinbach. GPT-NeoX: Large Scale Autoregressive Lan-  
291 guage Modeling in PyTorch, 9 2023. URL <https://www.github.com/eleutherai/gpt-neox>.
- 292
- 293 [4] Stephen H Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V Nayak,  
294 Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, et al. Promptsources: An  
295 integrated development environment and repository for natural language prompts. *arXiv preprint*  
296 *arXiv:2202.01279*, 2022.
- 297 [5] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn  
298 Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless  
299 assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*,  
300 2022.
- 301 [6] Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric  
302 Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff,  
303 Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large  
304 language models across training and scaling, 2023.
- 305 [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
306 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
307 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 308 [8] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain,  
309 Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. The ami meeting  
310 corpus: A pre-announcement. In *International workshop on machine learning for multimodal*  
311 *interaction*, pages 28–39. Springer, 2005.
- 312 [9] Harrison Chase. Langchain, 10 2022. URL <https://github.com/hwchase17/langchain>.
- 313 [10] Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. DialogSum: A real-life scenario  
314 dialogue summarization dataset. In *Findings of the Association for Computational Linguistics:*  
315 *ACL-IJCNLP 2021*, pages 5062–5074, Online, August 2021. Association for Computational  
316 Linguistics. doi: 10.18653/v1/2021.findings-acl.449. URL <https://aclanthology.org/2021.findings-acl.449>.
- 317
- 318 [11] Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. DialogSum: A real-life scenario  
319 dialogue summarization dataset. In *Findings of the Association for Computational Linguistics:*  
320 *ACL-IJCNLP 2021*, pages 5062–5074, Online, August 2021. Association for Computational  
321 Linguistics. doi: 10.18653/v1/2021.findings-acl.449. URL <https://aclanthology.org/2021.findings-acl.449>.
- 322
- 323 [12] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li,  
324 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned  
325 language models. *arXiv preprint arXiv:2210.11416*, 2022.
- 326 [13] Together Computer. Redpajama: An open source recipe to reproduce llama training dataset,  
327 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- 328 [14] Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick  
329 Wendell, Matei Zaharia, and Reynold Xin. Free dolly: Introducing the world’s first truly open  
330 instruction-tuned llm, 2023. URL <https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm>.
- 331

- 332 [15] Databricks. Dolly-hhrlhf dataset. [https://huggingface.co/datasets/mosaicml/dolly\\_hhrlhf](https://huggingface.co/datasets/mosaicml/dolly_hhrlhf), 2023.
- 333 [16] Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang.  
334 Glm: General language model pretraining with autoregressive blank infilling. *arXiv preprint*  
335 *arXiv:2103.10360*, 2021.
- 336 [17] Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher,  
337 and Dragomir Radev. Summeval: Re-evaluating summarization evaluation. *Transactions of the*  
338 *Association for Computational Linguistics*, 9:391–409, 2021.
- 339 [18] Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training  
340 quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
- 341 [19] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason  
342 Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse  
343 text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- 344 [20] Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. Samsun corpus:  
345 A human-annotated dialogue dataset for abstractive summarization. *arXiv preprint*  
346 *arXiv:1911.12237*, 2019.
- 347 [21] Peter Hase and Mohit Bansal. When can models learn from explanations? a formal framework  
348 for understanding the roles of explanation data. In *ACL Workshop on Learning with Natural*  
349 *Language Supervision*, 2022.
- 350 [22] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu  
351 Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference*  
352 *on Learning Representations*, 2021.
- 353 [23] Yebowen Hu, Tim Ganter, Hanieh Deilamsalehy, Franck Dernoncourt, Hassan Foroosh, and  
354 Fei Liu. Meetingbank: A benchmark dataset for meeting summarization. *arXiv preprint*  
355 *arXiv:2305.17529*, 2023.
- 356 [24] McKinsey Global Institute. How hybrid work has changed the way people work,  
357 live, and shop, 2023. URL [https://www.mckinsey.com/mgi/our-research/](https://www.mckinsey.com/mgi/our-research/empty-spaces-and-hybrid-places-chapter-1)  
358 [empty-spaces-and-hybrid-places-chapter-1](https://www.mckinsey.com/mgi/our-research/empty-spaces-and-hybrid-places-chapter-1). Accessed: 2023-09-30.
- 359 [25] Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara  
360 Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al. The icsi meeting corpus.  
361 In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003.*  
362 *Proceedings.(ICASSP'03).*, volume 1, pages I–I. IEEE, 2003.
- 363 [26] Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and  
364 Robert McHardy. Challenges and applications of large language models. *arXiv preprint*  
365 *arXiv:2307.10169*, 2023.
- 366 [27] Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin  
367 Choi, Noah A Smith, and Daniel S Weld. Genie: Toward reproducible and standardized human  
368 evaluation for text generation. *arXiv preprint arXiv:2101.06561*, 2021.
- 369 [28] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
370 language models are zero-shot reasoners. *Advances in neural information processing systems*,  
371 35:22199–22213, 2022.
- 372 [29] Vijay Anand Korthikanti, Jared Casper, Sangkug Lym, Lawrence McAfee, Michael Ander-  
373 sch, Mohammad Shoeybi, and Bryan Catanzaro. Reducing activation recomputation in large  
374 transformer models. *Proceedings of Machine Learning and Systems*, 5, 2023.
- 375 [30] Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. Evaluating the  
376 factual consistency of abstractive text summarization. *arXiv preprint arXiv:1910.12840*, 2019.
- 377 [31] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga,  
378 Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of  
379 language models. *arXiv preprint arXiv:2211.09110*, 2022.



- 380 [32] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization*  
381 *branches out*, pages 74–81, 2004.
- 382 [33] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically  
383 ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In  
384 *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*  
385 *(Volume 1: Long Papers)*, pages 8086–8098, 2022.
- 386 [34] Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task gen-  
387 eralization via natural language crowdsourcing instructions. In *60th Annual Meeting of the*  
388 *Association for Computational Linguistics, ACL 2022*, pages 3470–3487. Association for  
389 Computational Linguistics (ACL), 2022.
- 390 [35] Anna Nédoluzhko, Muskaan Singh, Marie Hledíková, Tirthankar Ghosal, and Ondřej Bojar.  
391 ELITR Minuting Corpus: A novel dataset for automatic minuting from multi-party meetings  
392 in English and Czech. In *Proceedings of the 13th International Conference on Language*  
393 *Resources and Evaluation (LREC-2022)*, Marseille, France, June 2022. European Language  
394 Resources Association (ELRA). In print.
- 395 [36] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin,  
396 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to  
397 follow instructions with human feedback. *Advances in Neural Information Processing Systems*,  
398 35:27730–27744, 2022.
- 399 [37] Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli,  
400 Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb  
401 dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv*  
402 *preprint arXiv:2306.01116*, 2023.
- 403 [38] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena,  
404 Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified  
405 text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- 406 [39] Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai,  
407 Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari,  
408 Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan  
409 Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang,  
410 Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang,  
411 Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries,  
412 Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush.  
413 Multitask prompted training enables zero-shot task generalization, 2021.
- 414 [40] Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. Answers unite!  
415 unsupervised metrics for reinforced summarization models. *arXiv preprint arXiv:1909.01610*,  
416 2019.
- 417 [41] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy  
418 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.  
419 [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- 420 [42] MosaicML NLP Team. Introducing mpt-7b: A new standard for open-source, commercially  
421 usable llms, 2023. URL [www.mosaicml.com/blog/mpt-7b](http://www.mosaicml.com/blog/mpt-7b). Accessed: 2023-05-05.
- 422 [43] Falcon Team Technology Innovation Institute. Falcon llm, 2023. URL [https://](https://falconllm.tii.ae/i)  
423 [falconllm.tii.ae/i](https://falconllm.tii.ae/i). Accessed: 2023-10-01.
- 424 [44] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timo-  
425 thée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open  
426 and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 427 [45] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei,  
428 Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open  
429 foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

- 430 [46] Oleg Vasilyev, Vedant Dharnidharka, and John Bohannon. Fill in the blanc: Human-free quality  
431 estimation of document summaries. *arXiv preprint arXiv:2002.09836*, 2020.
- 432 [47] Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive  
433 Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>,  
434 May 2021.
- 435 [48] Yiming Wang, Zhuosheng Zhang, and Rui Wang. Element-aware summarization with large  
436 language models: Expert-aligned evaluation and chain-of-thought method. *arXiv preprint*  
437 *arXiv:2305.13412*, 2023.
- 438 [49] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi,  
439 and Hannaneh Hajishirzi. Self-instruct: Aligning language model with self generated instruc-  
440 tions. *arXiv preprint arXiv:2212.10560*, 2022.
- 441 [50] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei,  
442 Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al.  
443 Super-naturalinstructions:generalization via declarative instructions on 1600+ tasks. In *EMNLP*,  
444 2022.
- 445 [51] Albert Webson and Ellie Pavlick. Do prompt-based models really understand the meaning of  
446 their prompts? In *Proceedings of the 2022 Conference of the North American Chapter of the*  
447 *Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344,  
448 2022.
- 449 [52] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan  
450 Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In  
451 *International Conference on Learning Representations*, 2021.
- 452 [53] JD Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. Why johnny  
453 can’t prompt: how non-ai experts try (and fail) to design llm prompts. In *Proceedings of the*  
454 *2023 CHI Conference on Human Factors in Computing Systems*, pages 1–21, 2023.
- 455 [54] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen,  
456 Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained  
457 transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- 458 [55] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore:  
459 Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
- 460 [56] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B  
461 Hashimoto. Benchmarking large language models for news summarization. *arXiv preprint*  
462 *arXiv:2301.13848*, 2023.
- 463 [57] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use:  
464 Improving few-shot performance of language models. In *International Conference on Machine*  
465 *Learning*, pages 12697–12706. PMLR, 2021.
- 466 [58] Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan  
467 Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, et al. Qmsum: A new benchmark for  
468 query-based multi-domain meeting summarization. *arXiv preprint arXiv:2104.05938*, 2021.
- 469 [59] Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. Adapting language models for zero-shot  
470 learning by meta-tuning on dataset and prompt collections. In *Findings of the Association for*  
471 *Computational Linguistics: EMNLP 2021*, pages 2856–2878, 2021.

## 472 A Appendix

### 473 A.1 More Related works

474 Evaluating LLMs remain a challenge. Promising benchmarks such as HELM [31] have been  
475 developed to evaluate the LLM abilities. While news summarization features in the benchmark,

476 conversational summarization is missed out. In this we also argue for inclusion of conversational  
477 summarization in the benchmark. Automatic evaluation of summarization remains a challenge [17].  
478 ROUGE [32] and BERTScore [55] still remain the most commonly reported reports for benchmarking  
479 purposes. Other metrics compute the quality of the summaries based on Factual consistency [30].  
480 Given, the subjective nature of the summary and the difficulty of developing a single or even multiple  
481 target summaries, reference free metrics have been developed in recent times to compute the quality  
482 of the summaries [46]. Human evaluations still remain central to computing the quality of the  
483 summaries [27]. Other metrics rely on question-answering approaches to measure the quality of the  
484 summaries [40].

## 485 A.2 Other LLMs

486 **MPT-7b** MPT [42] is a family of generative LLM from MosaicML NLP team trained on 1T tokens  
487 of text and code. The performance of the model is comparable to Llama [44]. The instruct version of  
488 the model is further trained on for instruction following from a dataset [15] derived from [14] and [5].  
489 **StableLM-7b** is a decoder only LLM based on GPT-NeoX [3] from Stability AI. The model is  
490 trained partly on RefinedWeb extract [37], RedPajama data [13] and, The Pile dataset [19].

491 **Dolly** Dolly is a class of LLMs based on Pythia [6] trained on Databricks 15k dolly dataset [14] and  
492 The Pile [19] dataset. From the authors claim Dolly is not a state-of-the-art model, however has good  
493 instruction following capabilities.

494 **T0pp** T0 is a class of encoder-decoder models based on T5 [38] trained in a multi-task prompted  
495 setting consisting of various NLP tasks with the goals of improving zero-shot generalizational  
496 abilities [39]. T0pp exhibits strong generalizational abilities on various tasks including summarization.

497

## 498 A.3 Dataset resources

- 499 1. Dialogsum: <https://huggingface.co/datasets/knkarthick/dialogsum>  
500 2. Samsum: <https://huggingface.co/datasets/samsum>  
501 3. Promptsources: [https://github.com/bigscience-workshop/promptsources/  
502 tree/main/promptsources/templates/samsum](https://github.com/bigscience-workshop/promptsources/tree/main/promptsources/templates/samsum)

## 503 A.4 Model hyper-parameters

504 We num train epochs=25, train batch size=2, warmup steps=500, weight decay=0.01, learning rate=5e-  
505 10, AdamW optimizer for Flan-T5 models and Dolly models. We used FP32 for Flan-T5 and Fp16  
506 precision for Dolly model weights.

507 For Lama-7b-chat finetuning we used LORA [22] adapters for finetuning.

```
508 lora_r = 8  
509 lora_alpha = 16  
510 lora_dropout = 0.05  
511 lora_target_modules= [  
512 "q_proj",  
513 "v_proj",  
514 ] with ,  
515 num_train_epochs=10,  
516 per_device_train_batch_size=2,  
517 per_device_eval_batch_size=2,  
518 warmup_steps=500,  
519 weight_decay=0.01,  
520
```

521 When inferencing with GPT3.5 we used temperature=0.2, max tokens=500, top p=1, frequency  
522 penalty=0, presence penalty=0.

Model	Samsam							Dialogsum						
	R-1	R-2	R-L	B-Score	LR	N1	N2	R-1	R-2	R-L	B-Score	LR	N1	N2
Llama-7b	0.14	0.02	0.10	0.701	56.58	47.06	85.70	0.16	0.02	0.11	0.736	40.24	31.22	61.87
Falcon-7b	0.14	0.02	0.10	0.703	58.80	48.74	89.68	0.16	0.02	0.11	0.737	36.32	26.61	54.37
MPT-7b	0.14	0.02	0.10	0.693	58.30	45.26	83.00	0.16	0.02	0.11	0.734	44.80	34.93	68.13
Stablelm-7b	0.14	0.01	0.10	0.691	56.45	45.86	83.52	0.14	0.02	0.10	0.732	42.77	33.50	66.29
MPT-7b-inst*	0.15	0.03	0.11	0.703	60.57	49.44	88.52	0.16	0.02	0.11	0.729	44.57	40.49	74.74
MPT-7b-chat*	0.17	0.04	0.12	0.713	63.64	49.92	92.34	0.19	0.04	0.13	0.756	45.32	34.38	70.01
Alpaca*	0.18	0.04	0.13	0.721	35.19	27.82	52.80	0.17	0.03	0.13	0.740	21.78	15.06	32.29
Falcon-7b-inst*	0.21	0.05	0.15	0.739	33.96	27.82	52.63	0.22	0.05	0.15	0.761	24.99	18.10	38.80
Flan-T5-xxl*	0.45	0.20	0.36	0.851	<b>17.77</b>	9.69	<b>24.08</b>	0.32	0.10	0.25	0.794	<b>12.82</b>	6.92	18.10
Dolly*	0.17	0.03	0.12	0.723	57.42	42.03	80.99	0.18	0.03	0.13	0.755	44.52	30.14	62.91
Dolly**	0.12	0.01	0.09	0.727	53.59	43.41	79.20	0.14	0.02	0.10	0.740	51.06	35.95	70.97
Flan-T5-small*	0.42	0.19	0.33	0.844	24.38	8.71	27.25	0.22	0.07	0.17	0.742	17.11	5.54	<b>15.62</b>
Flan-T5-small**	0.43	0.20	0.33	0.845	24.45	<b>8.70</b>	27.33	0.23	0.07	0.18	0.741	17.16	5.63	15.86
Flan-T5-base*	0.46	0.23	0.36	0.856	24.29	9.14	28.01	0.31	0.11	0.24	0.785	18.82	<b>4.80</b>	15.93
Flan-T5-base**	0.47	0.23	0.37	0.857	24.23	9.09	27.92	0.32	0.12	0.25	0.787	18.58	4.82	15.97
Flan-T5-large*	0.44	0.19	0.35	0.849	18.38	10.24	25.25	0.29	0.09	0.23	0.793	13.23	7.59	18.75
Flan-T5-large**	0.46	0.20	0.36	0.851	18.38	10.24	24.22	0.30	0.10	0.24	0.793	13.24	7.52	18.72
Flan-T5-xl*	0.48	0.26	0.39	0.862	24.37	9.87	28.21	0.37	0.16	0.29	0.824	18.10	5.71	18.13
Flan-T5-xl**	0.49	0.26	0.40	0.863	24.41	9.92	28.27	0.38	0.17	0.31	0.827	18.17	5.71	18.13
Llama-7b-chat*	0.18	0.04	0.13	0.721	32.41	25.65	48.58	0.21	0.05	0.15	0.765	31.46	22.54	47.39
Llama-7b-chat**±	0.18	0.05	0.14	0.721	31.66	25.03	47.45	0.21	0.06	0.15	0.760	30.54	22.17	46.17
T0pp-11b**	0.46	0.20	0.36	0.853	18.91	10.15	25.69	0.33	0.11	0.26	0.804	15.19	7.58	20.45
GPT-3.5 turbo	0.14	0.06	0.10	0.746	189.27	80.86	195.89	0.14	0.05	0.11	0.741	140.31	62.33	150.97

Table 2: Zero-shot evaluation of models on conversational summarization task on all the models (including those which were found to be not suitable for conversational summarization). This table is a superset of Table 1. The \* indicates the models are instruction tuned. \*\* Indicates that the models were further instructed tuned for conversational summarization task. ± indicates that the model was trained using adapters.

Model	Samsam			Samsam - FT			dialogsum			dialogsum -FT		
	R1	RL	B-Score	R1	RL	B-Score	R1	RL	B-Score	R1	RL	B-Score
Flan-T5-small	0.42	0.33	0.846	0.43	0.34	0.862	0.24	0.20	0.751	0.24	0.20	0.753
Flan-T5-base	0.46	0.37	0.857	0.47	0.38	0.857	0.33	0.26	0.793	0.34	0.26	0.794
Flan-T5-large	0.45	0.36	0.852	0.49	0.40	0.886	0.31	0.25	0.802	0.30	0.25	0.801
Flan-T5-xl	0.49	0.39	0.863	0.49	0.40	0.863	0.39	0.31	0.824	0.39	0.32	0.836
Llama-7b-chat*	0.32	0.24	0.808	0.33	0.24	0.809	0.26	0.21	0.789	0.29	0.21	0.792
GPT-3.5	0.21	0.14	0.748	-	-	-	0.24	0.20	0.775	-	-	-

Table 3: Average best performing prompt result.

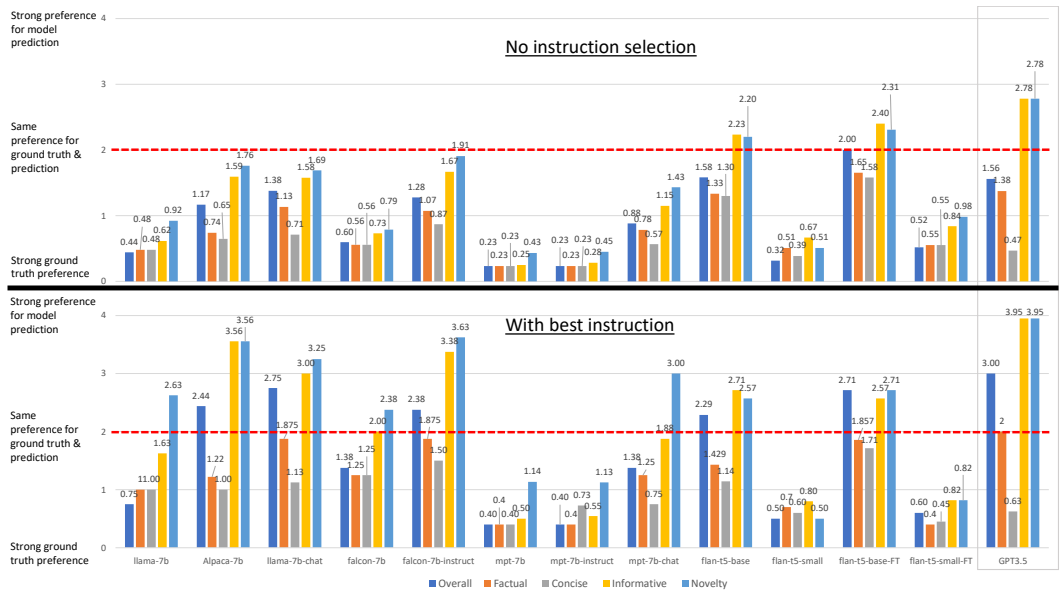


Figure 4: Shows human evaluation of summaries generated by the models

1	Generate a summary for this dialogue	Used
2	Summarize:	Used
3	To sum up this dialog	Used
4	Given the above dialogue write a summary	Used
5	Summarize this dialogue:	Used
6	Sum up the following dialogue	Used
7	Write a dialogue that match this summary	Not Used

Table 4: Prompts from the Promptsources dataset and if they're used for generating the summary in this work.

Model	Samsam					Dialogsum				
	R1(p/r)	R2(p/r)	RL(p/r)	N1	N2	R1(p/r)	R2(p/r)	RL(p/r)	N1	N2
Flan-T5-small*	0.72/0.30	0.37/0.15	0.58/0.24	9.15	28.41	0.80/0.20	0.50/0.12	0.63/0.16	5.75	16.10
Flan-T5-base*	0.74/0.31	0.37/0.15	0.59/0.25	9.60	29.21	0.84/0.23	0.55/0.15	0.70/0.19	4.97	16.41
Flan-T5-large*	0.68/0.20	0.24/0.07	0.51/0.15	10.75	26.34	0.70/0.13	0.28/0.05	0.53/0.10	7.86	19.32
Llama-7b	0.39/0.21	0.10/0.05	0.28/0.14	49.66	89.49	0.43/0.21	0.11/0.05	0.29/0.13	23.49	48.84
Falcon-7b	0.34/0.27	0.07/0.17	0.20/0.17	51.32	93.59	0.46/0.21	0.11/0.05	0.29/0.12	27.64	56.09
MPT-7b	0.34/0.07	0.10/0.07	0.22/0.17	47.53	86.49	0.42/0.24	0.11/0.07	0.25/0.14	36.22	70.24
Stablelm-7b	0.31/0.24	0.06/0.05	0.19/0.15	48.30	87.20	0.42/0.22	0.11/0.06	0.25/0.13	34.87	68.43
MPT-7b-inst*	0.31/0.26	0.08/0.07	0.19/0.07	51.96	92.31	0.54/0.18	0.14/0.05	0.37/0.12	41.91	76.98
Dolly*	0.38/0.30	0.10/0.08	0.24/0.19	44.26	84.53	0.49/0.27	0.16/0.09	0.29/0.16	31.35	64.89
Alpaca*	0.39/0.21	0.10/0.05	0.28/0.14	29.33	55.03	0.47/0.20	0.14/0.06	0.34/0.14	15.61	33.22
Falcon-7b-inst*	0.41/0.20	0.08/0.05	0.28/0.13	29.27	54.88	0.54/0.18	0.14/0.05	0.37/0.05	18.81	40.02
Flan-T5-xl*	0.73/0.31	0.37/0.15	0.57/0.24	10.36	29.42	0.82/0.22	0.48/0.13	0.67/0.18	5.90	18.68
Flan-T5-xxl*	0.69/0.19	0.25/0.07	0.52/0.15	10.18	25.11	0.74/0.14	0.32/0.06	0.57/0.11	7.01	18.54
MPT-7b-chat*	0.32/0.29	0.09/0.08	0.20/0.19	52.65	96.40	0.45/0.27	0.13/0.08	0.26/0.16	35.80	72.21
Llama-7b-chat*	0.38/0.21	0.09/0.21	0.28/0.06	27.05	50.68	0.50/0.14	0.14/0.06	0.33/0.14	23.49	48.84
Flan-T5-small**	0.72/0.30	0.38/0.15	0.58/0.25	9.13	28.51	0.80/0.20	0.49/0.12	0.62/0.16	5.83	16.35
Flan-T5-base**	0.74/0.30	0.38/0.15	0.58/0.25	9.55	29.12	0.84/0.23	0.54/0.15	0.69/0.19	4.99	16.45
Flan-T5-large**	0.68/0.20	0.24/0.07	0.51/0.15	10.77	26.36	0.70/0.13	0.28/0.05	0.53/0.10	7.79	19.29
Flan-T5-xl**	0.73/0.31	0.37/0.15	0.57/0.24	10.41	29.48	0.82/0.22	0.48/0.13	0.67/0.13	5.90	18.67
Dolly**	0.31/0.22	0.06/0.04	0.22/0.13	45.58	82.58	0.43/0.23	0.08/0.05	0.24/0.13	33.69	67.67
T0pp-11b**	0.69/0.25	0.25/0.08	0.52/0.16	10.65	26.80	0.74/0.16	0.32/0.07	0.56/0.12	7.84	21.06
Llama-7b-chat**±	0.38/0.21	0.09/0.06	0.28/0.14	26.34	49.43	0.50/0.13	0.13/0.06	0.33/0.13	22.94	47.60

Table 5: Table shows the Rouge and Novel 1-gram and Novel 2-gram scores of the generated summary compared to the . High Rouge of copying mechanisms in the summary generated

transcription\_input

Summarize this dialogue:  
 #Person1#: Honey, I think you should quit smoking.  
 #Person2#: Why? You said I was not when smoking.  
 #Person1#: But I want you to be fit.  
 #Person2#: Smoking is killing, I know.  
 #Person1#: Check out this article. It says smoking can lead to lung cancer.  
 #Person2#: I don't believe it.  
 #Person1#: But you know that smoking does harm to health, right?  
 #Person2#: Of course I know it, but you know it's hard to quit smoking. ...  
 #Person1#: Stop beating around the bush. Will you quit or not?  
 #Person2#: Yes, ma'am. Whatever you say.

summary\_1

In this dialogue, Person1 wants Person2 to quit smoking because it is harmful to their health. Person2 is initially resistant, but eventually agrees to quit after Person1 presents an article that highlights the dangers of smoking.

summary\_2

Person1 asks Person2 to quit smoking for health. Person2 thinks it's hard but agrees.

Overall  
 My overall preference for the given conversation.

Prefer summary 1  Slightly prefer summary 1  Both  Slightly prefer summary 2  Prefer summary 2  Neither

Factual  
 Preference for factually consistent with the input

Prefer summary 1  Slightly prefer summary 1  Both  Slightly prefer summary 2  Prefer summary 2  Neither

Concise  
 Preference for conciseness for the given input.

Prefer summary 1  Slightly prefer summary 1  Both  Slightly prefer summary 2  Prefer summary 2  Neither

Informative  
 Preference for informativeness for the given input.

Prefer summary 1  Slightly prefer summary 1  Both  Slightly prefer summary 2  Prefer summary 2  Neither

Novelty  
 Preference for the novelty of given the input.

Prefer summary 1  Slightly prefer summary 1  Both  Slightly prefer summary 2  Prefer summary 2  Neither

Record and get next

Figure 5: Interface for humans rating the summaries comparing two summaries.

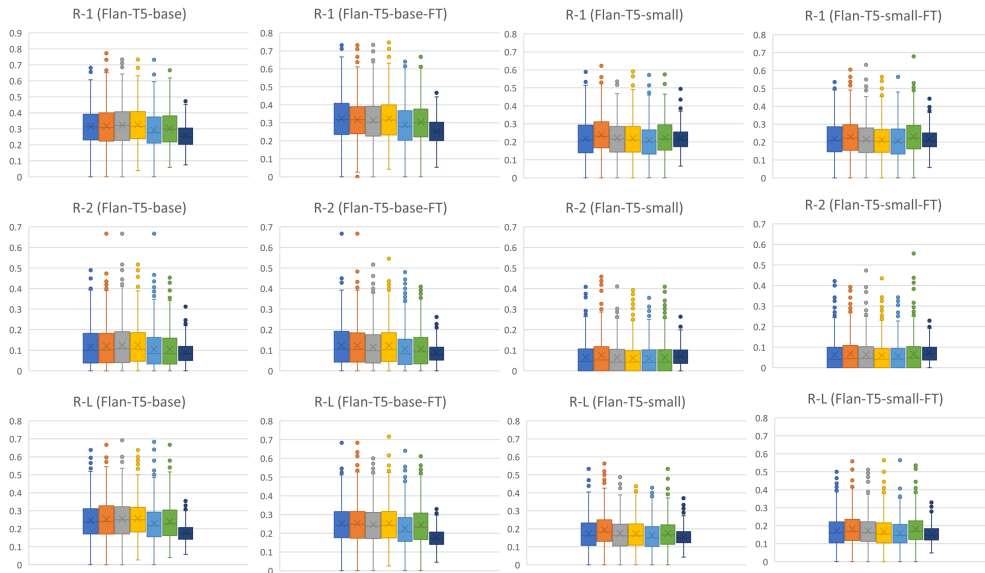


Figure 6: Box plots with Rouge scores per prompt from Prompsource dataset and various LLMs. Flan models (instruction tuned) models seem to be more robust to different prompts. ROUGE scores for all the prompts in Table 4.

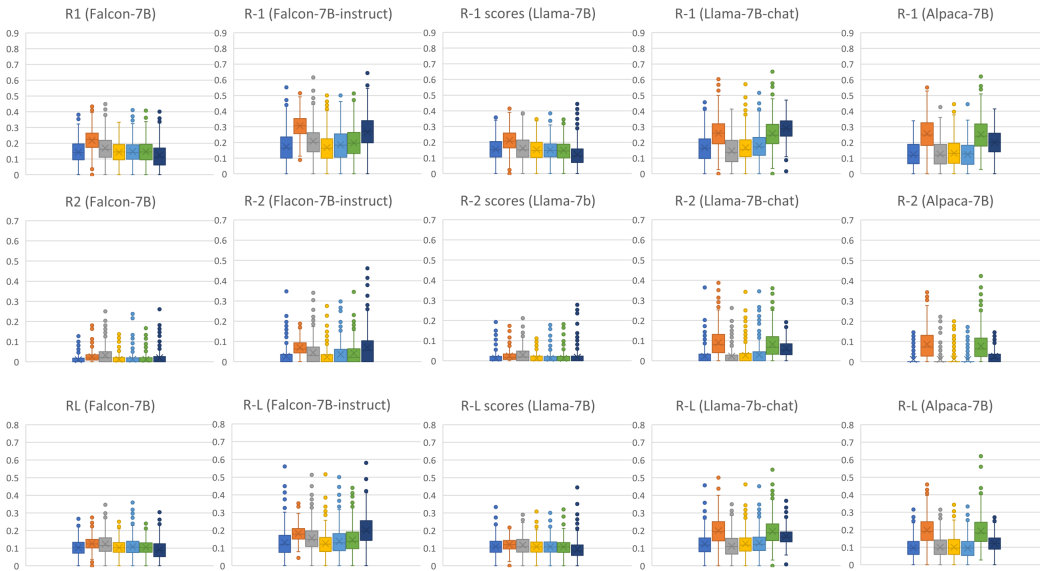


Figure 7: Box plots with Rouge scores per prompt from Prompsource dataset for various LLMs. Instruction tunings seems to improve the ROUGE scores across prompts, however the performance on different prompts seems to be an issue. ROUGE scores for all the prompts in Table 4.

Table 6: Example (not handpicked). Chosen randomly from the dataset.

(Conversation)	(Ground truth)
<p>Person1: Hello, how are you doing today?                      Person2: I ' Ve been having trouble breathing lately.                      Person1: Have you had any type of cold lately?                      Person2: No, I haven ' t had a cold.                      I just have a heavy feeling in my chest when I try to breathe.                      Person1: Do you have any allergies that you know of?                      Person2: No, I don ' t have any allergies that I know of.                      Person1: Does this happen all the time or mostly when you are active?                      Person2: It happens a lot when I work out.                      Person1: I am going to send you to a pulmonary specialist who can run tests on you for asthma.                      Person2: Thank you for your help, doctor.</p>	<p>Person2 has trouble breathing.                      The doctor asks Person2 about it and will send Person2 to a pulmonary specialist.</p>
<p>(Flan-T5-small)                      I am going to send you to a pulmonary specialist who can run tests on you for asthma.                      Do you have any allergies that you know?</p>	<p>(Flan-T5-base)                      Person1 is going to send Person2 to a pulmonary specialist who can run tests on Person2 for asthma.                      He will send Person1 and Person2 a message.</p>
<p>(Flan-T5-small***)                      Person2: Hi, how are you doing today?                      I ' ve been having problems breathing lately.                      I haven ' t had a cold. I just have a heavy feeling in my chest when I try to breathe                      Person1: Do you have any allergies that you know?</p>	<p>(Flan-T5-base***)                      Person2 has been having trouble breathing lately. She has a heavy feeling in her chest when she try to breathe. Person1 is going to send her to a pulmonary specialist who can run tests on her for asthma.</p>
<p>(Llama-7b)                      1. Describe your dialogue:                      2. How should you greet the patient? (1 Marks)                      3. What should you ask about the lifestyle? (1 Marks)                      4. What symptoms would you suspect? (1 Marks)                      5. List three diagnostic tests that your friend doctor would run on a patient suffering from your partner patient disease. (2 Marks)</p>	<p>(Alpaca-7b)                      The patient has been having difficulty breathing and it is a recurring issue. The doctor notices that it happens mostly when the patient is active, so they suggest sending the patient to a pulmonary specialist for tests to rule out asthma. The patient is thankful for the doctor's help.</p>
<p>(Llama-7b-chat)                      Person 1 (doctor) is interacting with Person 2 (patient) during a consultation. Person 2 has been experiencing difficulty breathing, and the doctor asks questions to gather more information and determine the cause. The doctor suspects asthma and recommends a pulmonary specialist for further tests.</p>	<p>(Llama-7b-chat***)                      Summary: The two people in the dialogue are discussing symptoms related to breathing difficulties. Person two explains that they have been experiencing a heavy feeling in their chest when trying to breathe, but have not had any colds or allergies. The first person suggests that they may have asthma and sends them to a pulmonary specialist for further testing.</p>
<p>(Falcon-7b)                      Write a summary of Person1:                      'Person1 has a heavy feeling in the chest when he/she tries to breathe.' Person1 goes to the primary care physician who diagnoses him/her with a lung infections. Person1: and the doctor sends him/her to another doctor for further tests. The doctor has been breathing difficulties. The doctor runs tests like a lung disease. A specialist helps Person2: 's diagnose</p>	<p>(Falcon-7b-instruct)                      After discussing symptoms with Person1, an appointment is arranged with a pulmonary specialist to examine person2 and diagnose any underlying respiratory issues.</p>
<p>(Flan-T5-XL)                      Person2 has been having trouble breathing lately. Person1 is going to send him to a pulmonary specialist who can run tests on him for asthma.</p>	<p>(Flan-T5-XXL)                      Person2 has been having breathing problems lately. They happen primarily when he exercises, so he might have Asthma.</p>

## Example (contd. from Table 6)

<p>(Flan-T5-Large)          Person2 has trouble breathing.          She has not had any cold, and her chest is heavy when she tries to breath.          At Person1 she will have tests for asthma.</p>	<p>(Flan-T5-Large**)          Person2 has trouble breathing because of allergies. She may have asthma and a specialist will run tests on her.</p>
<p>(MPT-7b)          Person2: I been having difficulty in breathing lately. Person2: Never, I have just a heavy feeling in my chest when I try to breathe.          Person2: I have no allergies that I know of. The D.P.R.I.V.E.R.S.          Copyright©          コスメをしくできるガウ! ,          2016 All Rights Reserved. We're all set for our launch!</p>	<p>(MPT-7b-chat)          Person1 and Person2 are speaking with each other. Person1 is listening to Person2's concerns about their breathing difficulties and suspects that Person2 may have asthma. Person1 suggests that Person2 should see a pulmonary specialist and arranges an appointment for them. Person1 suspects that Person2 may have asthma because it happens when Person2 is active which is a common symptom of asthma. Person1 also asks about allergies, colds and whether it happens all the time.</p>
<p>(MPT-7b-instruct)          A person that has breathing problems should consult a doctor that can perform a pulmonary function test for asthma. The World Economic Forum has created a unique and exclusive partnership with the International Monetary Fund, the World Bank and the United Nations to tackle the main issues and challenges of our time. The U.S. Department of State's Bureau of East Asian and Pacific Affairs is seeking young professionals for the Regional Internship Program (RIP) at the U.S. Embassy in Beijing.</p>	<p>(GPT-3.5)          I appreciate it. Person 1 is a doctor who is asking Person 2 how they are feeling. Person 2 responds that they have been having trouble breathing lately and a heavy feeling in their chest. Person 1 asks if they have had a cold or any allergies, and Person 2 responds that they have not. Person 1 asks if this happens all the time or mostly when they are active, and Person 2 responds that it happens a lot when they work out. Person 1 decides to send Person 2 to a pulmonary specialist to run tests for asthma. Person 2 thanks Person 1 for their help.</p>



Figure 8: Rouge, Bert novel n-grams and LR on Dialog sum dataset (contd. in next page)

<b>Llama-7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.155097	0.012483	0.109163	0.741976	0.39952	0.295743	0.599184
To sum up this dialog	0.161713	0.029629	0.116964	0.738562	0.410761	0.32151	0.649726
Sum up the following dialogue	0.148922	0.012307	0.107055	0.741244	0.407716	0.314225	0.620105
Summarize:	0.153161	0.012565	0.108862	0.743285	0.410667	0.300399	0.596294
Given the above dialogue write a summary	0.123213	0.017752	0.093643	0.709805	0.379939	0.33302	0.637694
Summarize this dialogue:	0.151364	0.012279	0.107938	0.741257	0.405815	0.308598	0.609301
<b>Alpaca-7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.123043	0.01104	0.096235	0.715468	0.169501	0.132597	0.262098
Summarize:	0.122846	0.011024	0.09479	0.71513	0.178707	0.138414	0.269087
To sum up this dialog	0.250834	0.077504	0.192573	0.77986	0.301346	0.181213	0.427168
Given the above dialogue write a summary	0.256705	0.086362	0.19838	0.784792	0.323217	0.196489	0.476503
Summarize this dialogue:	0.131366	0.013207	0.100833	0.719379	0.176512	0.136996	0.270987
Sum up the following dialogue	0.12561	0.012196	0.099977	0.715172	0.157391	0.117977	0.2313
<b>llama-7b-chat-FT-LORA</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.14481	0.021173	0.110217	0.731245	0.23102	0.170951	0.34261
Generate a summary for this	0.165926	0.02286	0.120241	0.746968	0.274295	0.203458	0.409811
Summarize this dialogue:	0.148658	0.019937	0.11203	0.738389	0.281885	0.211636	0.419865
Given the above dialogue write a summary	0.286605	0.088575	0.200529	0.794326	0.354381	0.23352	0.538818
Summarize:	0.170485	0.025389	0.12577	0.751375	0.341037	0.260788	0.512327
To sum up this dialog	0.259702	0.084096	0.198623	0.788785	0.349657	0.250022	0.552924
<b>llama-7b-chat**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.166289	0.027592	0.126325	0.748326	0.276779	0.199318	0.4106
Summarize:	0.177038	0.029221	0.132263	0.752538	0.347118	0.268846	0.529419
Given the above dialogue write a summary	0.258367	0.086615	0.199377	0.792302	0.355037	0.235759	0.543415
Summarize this dialogue:	0.155213	0.021733	0.116584	0.743092	0.269732	0.198723	0.3972
To sum up this dialog	0.285702	0.083828	0.19695	0.789559	0.345861	0.245569	0.542865
Sum up the following dialogue	0.148853	0.022171	0.111753	0.732514	0.242253	0.172483	0.346875
<b>falcon 7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.146674	0.012244	0.104962	0.740799	0.386214	0.275641	0.562292
Summarize this dialogue:	0.14644	0.012668	0.108058	0.738016	0.363445	0.267092	0.540104
Given the above dialogue write a summary	0.12315	0.02002	0.091768	0.70866	0.337709	0.274743	0.541918
Summarize:	0.14311	0.011729	0.102972	0.740492	0.364294	0.257496	0.524667
Generate a summary for this	0.142773	0.010107	0.102845	0.737669	0.331336	0.237425	0.48786
To sum up this dialog	0.169428	0.031676	0.123349	0.744305	0.396157	0.284468	0.605376
<b>falcon 7b instruct</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.166967	0.023472	0.124576	0.745567	0.236715	0.172025	0.357036
Summarize this dialogue:	0.185693	0.036116	0.137477	0.752301	0.24014	0.176336	0.371843
Generate a summary for this	0.172575	0.026236	0.13056	0.744003	0.231244	0.158482	0.340374
Given the above dialogue write a summary	0.26566	0.070478	0.200301	0.785344	0.231605	0.173266	0.380553
To sum up this dialog	0.208224	0.045168	0.155561	0.760238	0.306331	0.215654	0.472222

Figure 9: Rouge, Bert novel n-grams and LR on Dialog sum dataset (contd. in next page)

Sum up the following dialogue	0.197336	0.041285	0.145594	0.75609	0.253452	0.190242	0.405983
<b>flan t5 small</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.220849	0.06469	0.174849	0.739101	0.166336	0.057042	0.160337
Summarize this dialogue:	0.218634	0.061893	0.172445	0.737129	0.16738	0.063482	0.170796
Summarize:	0.208724	0.060363	0.164737	0.734579	0.171026	0.070176	0.190372
Generate a summary for this	0.217678	0.065163	0.174395	0.738971	0.165249	0.054691	0.155188
Given the above dialogue write a summary	0.239075	0.074692	0.194387	0.751289	0.164287	0.034602	0.109689
To sum up this dialog	0.222722	0.063573	0.174229	0.748869	0.192317	0.052654	0.150685
<b>flan t5 small**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.206299	0.055989	0.159021	0.729974	0.169743	0.072055	0.190081
To sum up this dialog	0.231967	0.06588	0.182556	0.751114	0.194985	0.052697	0.151882
Sum up the following dialogue	0.214377	0.062978	0.17244	0.736924	0.16832	0.057129	0.160164
Given the above dialogue write a summary	0.229277	0.068427	0.181946	0.750092	0.167528	0.037606	0.118998
Generate a summary for this	0.218365	0.063444	0.171822	0.738577	0.163012	0.055599	0.159295
Summarize this dialogue:	0.211452	0.060423	0.16837	0.733221	0.165853	0.062515	0.171191
<b>flan t5 base</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.322372	0.124192	0.255298	0.790596	0.181411	0.050651	0.160012
Given the above dialogue write a summary	0.317451	0.120006	0.252177	0.788916	0.176301	0.043973	0.14672
Summarize:	0.292114	0.104867	0.230782	0.778736	0.197628	0.053353	0.168494
Summarize this dialogue:	0.32667	0.125102	0.25722	0.792915	0.183676	0.050118	0.166641
To sum up this dialog	0.302969	0.104503	0.241155	0.784637	0.210806	0.039846	0.145457
Generate a summary for this	0.315809	0.116995	0.247073	0.785773	0.179514	0.049935	0.168245
<b>flan t5 base**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.323974	0.122721	0.251531	0.787683	0.175363	0.049676	0.16524
Summarize this dialogue:	0.323559	0.122226	0.253564	0.791757	0.182332	0.051762	0.168762
Summarize:	0.290381	0.105061	0.226647	0.778846	0.198275	0.055769	0.171603
To sum up this dialog	0.304067	0.107208	0.244015	0.784353	0.204823	0.038294	0.141856
Given the above dialogue write a summary	0.317706	0.12313	0.251897	0.789862	0.174227	0.043076	0.145844
Sum up the following dialogue	0.313669	0.114791	0.248405	0.788237	0.179662	0.050824	0.165103
<b>flan t5 large</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.298818	0.098571	0.240091	0.79512	0.137105	0.078751	0.197492
To sum up this dialog	0.300911	0.104235	0.245926	0.790678	0.13389	0.074518	0.186834
Summarize this dialogue:	0.296258	0.095401	0.235329	0.797522	0.127828	0.074517	0.180357
Given the above dialogue write a summary	0.300012	0.101514	0.241985	0.789	0.134489	0.076606	0.194473
Sum up the following dialogue	0.308708	0.102489	0.246222	0.802357	0.137219	0.076345	0.191035
Summarize:	0.257798	0.074306	0.214406	0.785214	0.123281	0.074951	0.174774
<b>flan t5 large-FT</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.271149	0.079915	0.221449	0.787357	0.120611	0.072525	0.170887
Sum up the following dialogue	0.299363	0.098003	0.240935	0.798535	0.138283	0.07732	0.192509
To sum up this dialog	0.299735	0.101324	0.24312	0.78918	0.130116	0.071952	0.184161
Generate a summary for this	0.299255	0.098783	0.241988	0.793976	0.135931	0.075357	0.192023

Figure 10: Rouge, Bert novel n-grams and LR on Dialog sum dataset (contd. in next page)

Given the above dialogue write a summary	0.297744	0.098212	0.239776	0.788899	0.133343	0.075389	0.190052
Summarize this dialogue:	0.301441	0.095856	0.243957	0.800948	0.136137	0.078929	0.193592
<b>flan t5 xl</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.382845	0.167842	0.308651	0.823862	0.174189	0.053869	0.176848
Sum up the following dialogue	0.375802	0.163135	0.306933	0.823262	0.182485	0.053901	0.176824
To sum up this dialog	0.376861	0.162373	0.306246	0.827145	0.186537	0.056169	0.180301
Given the above dialogue write a summary	0.375544	0.167272	0.306846	0.826753	0.176331	0.054801	0.174932
Summarize this dialogue:	0.379134	0.169924	0.310737	0.826006	0.183926	0.057328	0.181918
Summarize:	0.377441	0.164603	0.308148	0.837087	0.182768	0.066485	0.197083
<b>flan t5 xxl</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.295254	0.09226	0.238475	0.78514	0.119622	0.068781	0.171903
Given the above dialogue write a summary	0.315609	0.097389	0.257521	0.795451	0.124431	0.072778	0.185828
Summarize this dialogue:	0.32805	0.115463	0.265209	0.794181	0.13637	0.066337	0.184973
Sum up the following dialogue	0.339075	0.116157	0.278561	0.797664	0.129232	0.06729	0.181505
To sum up this dialog	0.331992	0.115183	0.26942	0.798283	0.137369	0.071532	0.188815
Generate a summary for this	0.32412	0.120772	0.264038	0.796111	0.122485	0.068155	0.173301
<b>GPT-35</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Given the above dialogue write a summary	0.135999	0.054377	0.105135	0.736754	1.809941	0.782086	1.843037
Generate a summary for this	0.236238	0.098671	0.191451	0.774839	1.021735	0.440279	1.100863
Summarize this dialogue:	0.111463	0.041679	0.084602	0.741608	2.151012	0.975895	2.400766
Summarize:	0.139401	0.057176	0.110111	0.749084	1.947661	0.772832	1.880909
Sum up the following dialogue	0.102011	0.038773	0.077755	0.737168	2.221335	0.945967	2.342505
To sum up this dialog	0.097661	0.039338	0.073692	0.733833	2.204351	0.934519	2.185292
<b>T0pp</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this	0.326715	0.113065	0.261189	0.803355	0.152048	0.074924	0.204107
Summarize this dialogue:	0.328963	0.112688	0.260172	0.804439	0.160057	0.078188	0.21322
Given the above dialogue write a summary	0.331677	0.116408	0.264242	0.80328	0.142965	0.071972	0.196545
Summarize:	0.331093	0.105155	0.259262	0.805496	0.156579	0.078341	0.207284
Sum up the following dialogue	0.337288	0.119383	0.270275	0.806868	0.153429	0.076063	0.203969
To sum up this dialog	0.331773	0.117043	0.26753	0.802943	0.146162	0.07551	0.201669
<b>stablelm</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
To sum up this dialog	0.146775	0.023737	0.108277	0.729948	0.443051	0.349731	0.701991
Summarize:	0.146498	0.011435	0.10403	0.739341	0.445244	0.333723	0.66178
Summarize this dialogue:	0.145053	0.012268	0.101972	0.739322	0.442735	0.346407	0.678268
Generate a summary for this	0.145766	0.011632	0.10538	0.73953	0.430457	0.327275	0.65164
Given the above dialogue write a summary	0.129337	0.019264	0.095513	0.714057	0.375939	0.323324	0.624799
Sum up the following dialogue	0.142083	0.012869	0.10139	0.731902	0.428777	0.329807	0.658729
<b>mpt 7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Write a dialogue that match this summary	0.231065	0.026577	0.12765	0.738155	2.027707	1.594844	2.54711
Sum up the following dialogue	0.131568	0.015097	0.096129	0.726031	0.407433	0.300765	0.58201

To sum up this dialog	0.157214	0.029741	0.113231	0.735217	0.456145	0.371393	0.737388
Given the above dialogue write a summary	0.146025	0.027024	0.106522	0.724074	0.431499	0.340296	0.660324
Generate a summary for this	0.153545	0.015142	0.107567	0.744884	0.455223	0.338262	0.671552
Summarize:	0.145686	0.0138	0.103283	0.734998	0.475127	0.375062	0.724554
Summarize this dialogue:	0.147687	0.013721	0.104138	0.737296	0.462702	0.369859	0.712154
<b>mpt 7b instruct</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Write a dialogue that match this summary	0.248541	0.034401	0.137344	0.748172	2.036852	1.682254	2.657518
Sum up the following dialogue	0.136118	0.016041	0.097606	0.731028	0.410583	0.331181	0.628588
Given the above dialogue write a summary	0.139063	0.024377	0.103418	0.718878	0.441112	0.420767	0.773514
Summarize:	0.123223	0.008994	0.087282	0.716617	0.464403	0.452329	0.801463
To sum up this dialog	0.170395	0.035432	0.122733	0.744242	0.456188	0.405947	0.783691
Generate a summary for this	0.1445	0.01309	0.102604	0.735782	0.44275	0.370925	0.704518
Summarize this dialogue:	0.123662	0.009847	0.089732	0.714568	0.459108	0.448355	0.792561
<b>mpt 7b chat</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Write a dialogue that match this summary	0.27604	0.0547	0.15366	0.775746	2.023749	1.442607	2.411826
Generate a summary for this	0.164312	0.02179	0.11691	0.749668	0.438559	0.335884	0.668655
Sum up the following dialogue	0.160771	0.022286	0.112933	0.746871	0.437798	0.337811	0.670768
Given the above dialogue write a summary	0.216199	0.068369	0.16386	0.769369	0.438509	0.305399	0.67679
Summarize this dialogue:	0.164424	0.020946	0.115544	0.747723	0.46588	0.359704	0.709164
Summarize:	0.162509	0.017333	0.112602	0.743885	0.480345	0.364861	0.724643
To sum up this dialog	0.200555	0.055429	0.145139	0.758801	0.458418	0.359136	0.751067

Figure 11: Rouge, Bert novel n-grams and LR on SAMsum dataset (contd. in next page)

<b>Llama-7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this dialogue	0.138151	0.011232	0.096096	0.696381	0.532532	0.430856	0.79213
Summarize this dialogue:	0.140075	0.011797	0.097608	0.692316	0.560422	0.459221	0.831045
Given the above dialogue write a summary	0.134712	0.018687	0.095484	0.700108	0.566197	0.484294	0.884038
Summarize:	0.142032	0.013697	0.100249	0.694947	0.564884	0.450905	0.819776
Sum up the following dialogue	0.138351	0.012247	0.098149	0.693184	0.537648	0.440747	0.800536
To sum up this dialog	0.158298	0.025	0.110279	0.723472	0.6331	0.557317	1.01443
<b>Alpaca-7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
To sum up this dialog	0.274779	0.087774	0.202459	0.784598	0.488296	0.365464	0.725447
Summarize this dialogue:	0.126125	0.016269	0.09287	0.689864	0.34206	0.294787	0.522726
Sum up the following dialogue	0.120929	0.017911	0.092474	0.682573	0.25215	0.198631	0.35941
Given the above dialogue write a summary	0.299811	0.100909	0.221014	0.802135	0.441332	0.313617	0.677838
Generate a summary for this dialogue	0.116115	0.01407	0.087	0.681406	0.254707	0.217611	0.38721
Summarize:	0.128742	0.014729	0.096472	0.687265	0.332946	0.279145	0.495593
<b>llama-7b-chat</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize this dialogue:	0.126199	0.026178	0.097576	0.689695	0.294159	0.243503	0.434528
Summarize:	0.146753	0.032657	0.109718	0.701594	0.371549	0.296411	0.534962
Generate a summary for this dialogue	0.100954	0.016443	0.079728	0.674158	0.206717	0.177299	0.310414
Sum up the following dialogue	0.101414	0.016671	0.080353	0.674374	0.217001	0.181127	0.317643
Given the above dialogue write a summary	0.26926	0.089237	0.199548	0.775598	0.399962	0.293658	0.610027
To sum up this dialog	0.320837	0.112322	0.242416	0.807502	0.419169	0.310027	0.639226
<b>llama-7b-chat**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize this dialogue:	0.125293	0.022487	0.094775	0.688839	0.301221	0.251131	0.448525
Generate a summary for this dialogue	0.106675	0.017424	0.082528	0.677814	0.228789	0.195356	0.341239
Given the above dialogue write a summary	0.265568	0.089776	0.202422	0.779068	0.403039	0.292503	0.614073
Sum up the following dialogue	0.100072	0.014085	0.076629	0.667446	0.210996	0.180551	0.310914
Summarize:	0.150035	0.032711	0.110513	0.70207	0.370832	0.303188	0.548543

Figure 12: Rouge, Bert novel n-grams and LR on SAMsum dataset (contd. in next page)

To sum up this dialog	0.317069	0.109511	0.237412	0.807368	0.429941	0.316229	0.651681
<b>falcon 7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.138088	0.013231	0.097488	0.697743	0.575662	0.480943	0.877054
Summarize:	0.136126	0.01174	0.095757	0.699371	0.622519	0.521622	0.942095
Summarize this dialogue:	0.137002	0.012657	0.094589	0.698954	0.60495	0.511499	0.92654
Generate a summary for this dialogue	0.137287	0.011518	0.097851	0.696043	0.538723	0.436403	0.80921
To sum up this dialog	0.17606	0.034745	0.122639	0.731453	0.619681	0.503324	0.9497
Given the above dialogue write a summary	0.133605	0.020958	0.094634	0.700101	0.566262	0.470778	0.876445
<b>falcon 7b instruct</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.163407	0.027373	0.119564	0.71579	0.381512	0.307338	0.572282
Summarize this dialogue:	0.210885	0.045688	0.154768	0.742217	0.320168	0.255625	0.487716
Generate a summary for this dialogue	0.142276	0.021434	0.107446	0.702387	0.2483	0.203035	0.381988
Sum up the following dialogue	0.190806	0.041909	0.142071	0.728361	0.275203	0.223624	0.420916
Given the above dialogue write a summary	0.275482	0.072237	0.201047	0.78784	0.3364	0.293131	0.556635
To sum up this dialog	0.223003	0.048269	0.158714	0.758777	0.475864	0.386221	0.73841
<b>flan t5 small</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.424134	0.19869	0.331925	0.84593	0.240866	0.087103	0.270402
Summarize:	0.416116	0.191786	0.324276	0.843437	0.255013	0.096081	0.287791
Summarize this dialogue:	0.423168	0.196865	0.329597	0.844902	0.244977	0.089552	0.274155
Generate a summary for this dialogue	0.421951	0.194399	0.329904	0.844547	0.236166	0.083252	0.264308
To sum up this dialog	0.418872	0.189302	0.324848	0.841838	0.256304	0.090344	0.279999
Given the above dialogue write a summary	0.416849	0.191247	0.32746	0.843206	0.229195	0.076071	0.25811
<b>flan t5 small**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Given the above dialogue write a summary	0.4207	0.194494	0.330137	0.844591	0.229464	0.074112	0.2559
Generate a summary for this dialogue	0.425009	0.199445	0.331532	0.845756	0.237111	0.082875	0.264925
Summarize:	0.420493	0.193945	0.32469	0.844425	0.256145	0.096569	0.288803
To sum up this dialog	0.416721	0.187867	0.322595	0.841248	0.258522	0.093436	0.285204
Summarize this dialogue:	0.424686	0.196242	0.33153	0.846229	0.244996	0.088037	0.273978

Figure 13: Rouge, Bert novel n-grams and LR on SAMsum dataset (contd. in next page)

Sum up the following dialogue	0.424029	0.19749	0.332784	0.845635	0.24094	0.086688	0.271283
<b>flan t5 base</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
To sum up this dialog	0.456785	0.225835	0.364757	0.855416	0.241282	0.093333	0.277662
Generate a summary for this dialogue	0.458924	0.229698	0.364091	0.856017	0.239277	0.090211	0.278621
Given the above dialogue write a summary	0.455562	0.226962	0.363869	0.855499	0.230164	0.080242	0.267037
Summarize this dialogue:	0.459947	0.23078	0.366927	0.856481	0.248915	0.094768	0.28515
Summarize:	0.454209	0.223392	0.359889	0.854749	0.255367	0.097744	0.293865
Sum up the following dialogue	0.45732	0.227539	0.362455	0.855595	0.242161	0.092254	0.27837
<b>flan t5 base**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.458306	0.228776	0.365684	0.855244	0.253586	0.097256	0.291026
To sum up this dialog	0.456886	0.226381	0.365502	0.855112	0.240872	0.092336	0.275326
Given the above dialogue write a summary	0.456669	0.228126	0.365725	0.855943	0.229177	0.079639	0.265848
Generate a summary for this dialogue	0.457649	0.228695	0.363815	0.855589	0.240278	0.08956	0.279615
Summarize this dialogue:	0.458968	0.228204	0.366502	0.856239	0.248801	0.096055	0.286384
Sum up the following dialogue	0.46044	0.230491	0.365305	0.856088	0.241238	0.090595	0.277022
<b>flan t5 large</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Given the above dialogue write a summary	0.447581	0.192464	0.353051	0.849431	0.175731	0.091765	0.241429
Generate a summary for this dialogue	0.441133	0.182931	0.347351	0.84891	0.181865	0.100838	0.25219
Summarize:	0.431042	0.176629	0.342292	0.846037	0.184177	0.104794	0.253661
Summarize this dialogue:	0.4471	0.191483	0.355392	0.850469	0.188625	0.107861	0.259229
Sum up the following dialogue	0.454356	0.201503	0.362175	0.851893	0.184824	0.104233	0.252943
To sum up this dialog	0.448531	0.19508	0.355392	0.85194	0.18742	0.10486	0.255467
<b>flan t5 large**</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
To sum up this dialog	0.445391	0.19212	0.352128	0.850234	0.187317	0.107694	0.257889
Given the above dialogue write a summary	0.449893	0.195342	0.354392	0.85051	0.179205	0.092971	0.248061
Summarize:	0.4368	0.181819	0.348932	0.847873	0.177966	0.101775	0.243194
Summarize this dialogue:	0.450291	0.196271	0.360784	0.850779	0.190799	0.108495	0.261573

Figure 14: Rouge, Bert novel n-grams and LR on SAMsum dataset (contd. in next page)

Sum up the following dialogue	0.451589	0.195899	0.353651	0.85128	0.186073	0.102902	0.253479
Generate a summary for this dialogue	0.456443	0.195864	0.362597	0.852474	0.181571	0.100598	0.250304
<b>flan t5 xl</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this dialogue	0.482227	0.253183	0.385786	0.861587	0.23712	0.095679	0.276221
Summarize:	0.481108	0.249328	0.384194	0.861054	0.253999	0.104758	0.292021
Summarize this dialogue:	0.487758	0.257244	0.390022	0.862574	0.249201	0.103033	0.287512
Sum up the following dialogue	0.484681	0.256303	0.386974	0.86174	0.242142	0.100413	0.280825
To sum up this dialog	0.482373	0.254279	0.388943	0.861261	0.244367	0.10072	0.283744
Given the above dialogue write a summary	0.48473	0.257693	0.389706	0.862628	0.23538	0.087678	0.272229
<b>flan t5 xxl</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Sum up the following dialogue	0.445267	0.19799	0.356423	0.848548	0.173269	0.096662	0.235865
Summarize:	0.438043	0.182785	0.34837	0.847752	0.178606	0.100618	0.241043
Generate a summary for this dialogue	0.455566	0.204244	0.364262	0.852396	0.17007	0.092868	0.231381
Summarize this dialogue:	0.455063	0.201777	0.361122	0.852339	0.182798	0.101024	0.247952
Given the above dialogue write a summary	0.451885	0.199967	0.357116	0.851314	0.179026	0.091815	0.244427
To sum up this dialog	0.457944	0.207235	0.366748	0.853512	0.182481	0.098383	0.244232
<b>GPT3.5</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this dialogue	0.147895	0.054254	0.116342	0.726377	1.704756	0.81086	1.840422
Summarize this dialogue:	0.126942	0.053637	0.098358	0.730847	2.660202	1.219457	2.95443
Given the above dialogue write a summary	0.122386	0.04909	0.097204	0.725427	2.703784	1.259391	2.946073
To sum up this dialog	0.176719	0.073031	0.135415	0.756771	1.882938	1.055004	2.355559
Sum up the following dialogue	0.109183	0.043131	0.086664	0.725799	2.764192	1.190182	2.943563
Summarize:	0.098294	0.027807	0.074608	0.688649	2.79711	1.339082	3.007391
<b>T0pp</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize this dialogue:	0.46031	0.205191	0.364192	0.853453	0.194585	0.105112	0.261555
Sum up the following dialogue	0.454903	0.198567	0.360058	0.852486	0.186436	0.102332	0.254516
To sum up this dialog	0.45764	0.197935	0.359596	0.852282	0.184857	0.100857	0.252031



Generate a summary for this dialogue	0.455786	0.193037	0.358573	0.85201	0.190025	0.103923	0.262268
Summarize:	0.452621	0.197417	0.354793	0.851454	0.200065	0.106945	0.266188
Given the above dialogue write a summary	0.456897	0.198228	0.359995	0.85337	0.178658	0.089705	0.24508
<b>mpt 7b</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize this dialogue:	0.130743	0.017367	0.094772	0.687118	0.579671	0.441913	0.798518
Sum up the following dialogue	0.12737	0.014797	0.092596	0.672495	0.543793	0.423902	0.764869
Given the above dialogue write a summary	0.148239	0.0303	0.109428	0.69497	0.589931	0.442074	0.834862
Summarize:	0.138481	0.012449	0.098907	0.693921	0.601138	0.472018	0.853063
To sum up this dialog	0.172531	0.037449	0.118778	0.722212	0.642891	0.525815	0.975826
Generate a summary for this dialogue	0.129616	0.012423	0.094945	0.685504	0.540657	0.409823	0.753422
<b>mpt 7b instruct</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Generate a summary for this dialogue	0.13026	0.012304	0.091362	0.687481	0.582291	0.480074	0.858761
Summarize this dialogue:	0.132523	0.017338	0.095159	0.688021	0.602581	0.493965	0.86798
To sum up this dialog	0.197526	0.04441	0.138587	0.741606	0.634858	0.524643	0.959402
Summarize:	0.133676	0.013195	0.095182	0.687833	0.629766	0.521584	0.913392
Sum up the following dialogue	0.127604	0.016235	0.091713	0.678287	0.567648	0.439284	0.774835
Given the above dialogue write a summary	0.18664	0.04545	0.133084	0.736619	0.616821	0.506748	0.93708
<b>mpt 7b chat</b>	R1	R2	RL	BS-F1	LR-Pred	N1	N2
Summarize:	0.126538	0.016198	0.090059	0.682858	0.653942	0.510196	0.905745
To sum up this dialog	0.21497	0.062735	0.155566	0.760399	0.673022	0.543196	1.036007
Sum up the following dialogue	0.142378	0.022415	0.099596	0.690534	0.607876	0.475325	0.856632
Summarize this dialogue:	0.134201	0.018483	0.09398	0.687583	0.630332	0.501101	0.906677
Given the above dialogue write a summary	0.227252	0.071722	0.166757	0.764069	0.637268	0.474886	0.940555
Generate a summary for this dialogue	0.141838	0.021354	0.098994	0.694801	0.615857	0.490256	0.894808