

Evaluation of ChatGPT on Biomedical Tasks: A Zero-Shot Comparison with Fine-Tuned Generative Transformers

Israt Jahan^{†,§}, Md Tahmid Rahman Laskar^{‡,§,§}, Chun Peng[†], Jimmy Xiangji Huang^{‡,§}

[†]Department of Biology, York University

[‡]School of Information Technology, York University

[§]Information Retrieval and Knowledge Management Research Lab, York University

[§]Dialpad Canada Inc.

Toronto, Ontario, Canada

{israt18, tahmid20, cpeng, jhuang}@yorku.ca

Abstract

ChatGPT is a large language model developed by OpenAI¹. Despite its impressive performance across various tasks, no prior work has investigated its capability in the biomedical domain yet. To this end, this paper aims to evaluate the performance of ChatGPT on various benchmark biomedical tasks, such as relation extraction, document classification, question answering, and summarization. To the best of our knowledge, this is the first work that conducts an extensive evaluation of ChatGPT in the biomedical domain. Interestingly, we find based on our evaluation that in biomedical datasets that have smaller training sets, zero-shot ChatGPT even outperforms the state-of-the-art fine-tuned generative transformer models, such as BioGPT and BioBART. This suggests that ChatGPT's pre-training on large text corpora makes it quite specialized even in the biomedical domain. Our findings demonstrate that ChatGPT has the potential to be a valuable tool for various tasks in the biomedical domain that lack large annotated data.

1 Introduction

The rapid growth of language models (Rogers et al., 2021; Zhou et al., 2023) in the field of Natural Language Processing (NLP) in recent years has led to significant advancements in various domains, including the biomedical domain (Kalyan et al., 2022). Although specialized models (e.g., BioBERT (Lee et al., 2020), BioBART (Yuan et al., 2022a), BioGPT (Luo et al., 2022), etc.) have shown promising results in the biomedical domain, they require fine-tuning using domain-specific datasets. This fine-tuning process can be time-consuming due to the requirement of task-specific large annotated datasets. In contrast, zero-shot learning enables models to perform tasks without the need for fine-tuning on task-specific datasets. ChatGPT, a large language model, has

¹<https://openai.com/blog/chatgpt>

demonstrated impressive zero-shot performance across various tasks (Laskar et al., 2023). However, its performance in the biomedical domain remains to be thoroughly investigated. In this regard, this paper presents a comprehensive evaluation of ChatGPT on four key biomedical tasks: relation extraction, question answering, document classification, and summarization.

In this paper, our primary objective is to explore the extent to which ChatGPT can perform these tasks without fine-tuning and assess its performance by comparing with state-of-the-art generative fine-tuned models, BioGPT and BioBART. To our best knowledge, this is the first work that evaluated ChatGPT on benchmark biomedical datasets. Our evaluation of ChatGPT can have a profound impact on the biomedical domain that lacks domain-specific datasets by exploring its zero-shot learning capabilities. To ensure the reproducibility of our evaluation and to help facilitate future research, we will release all the ChatGPT-generated responses along with our evaluation code here: <https://github.com/tahmedge/chatgpt-eval-biomed>.

2 Related Work

The effective utilization of transformer-based (Vaswani et al., 2017) NLP models like BERT (Devlin et al., 2019) have also led to significant progress in the biomedical domain (Lee et al., 2020; Alsentzer et al., 2019; Beltagy et al., 2019; Gu et al., 2020; Peng et al., 2019) in recent years. BERT leverages the encoder of the transformer architecture, while GPT leverages the decoder of the transformer. In addition to these models, sequence-to-sequence models like BART (Lewis et al., 2019) that leverage both the encoder and the decoder of transformer have also emerged as a powerful approach in various text generation tasks.

It has been observed that domain-specific pre-training of these models on the biomedical text corpora followed by fine-tuning on task-specific

Dataset	Type	Data Split (Train / Valid / Test)	Prompt
BC5CDR	Chemical-Disease Relation Extraction	500 / 500 / 500	Identify each pair of drugs and the drug-induced side-effects (e.g., diseases) in the following passage: [PASSAGE]
KD-DTI	Drug-Target Relation Extraction	12K / 1K / 1.3K	Identify the drug-target interactions in the following passage (along with the interaction type among the following: 'inhibitor', 'agonist', 'modulator', 'activator', 'blocker', 'inducer', 'antagonist', 'cleavage', 'disruption', 'intercalation', 'inactivator', 'bind', 'binder', 'partial agonist', 'cofactor', 'substrate', 'ligand', 'chelator', 'downregulator', 'other', 'anti-body', 'other/unknown'): [PASSAGE]
DDI	Drug-Drug Relation Extraction	664 / 50 / 191	Identify the pairs of drug-drug interactions in the passage given below based on one of the following interaction types: (i) mechanism: this type is used to identify drug-drug interactions that are described by their pharmacokinetic mechanism. (ii) effect: this type is used to identify drug-drug interactions describing an effect. (iii) advice: this type is used when a recommendation or advice regarding a drug-drug interaction is given. (iv) int: this type is used when a drug-drug interaction appears in the text without providing any additional information. [PASSAGE]
HoC	Document Classification	9972 / 4947 / 4947	The 10 hallmark cancer taxonomy with their definitions MeQSum given below: (i) Sustaining proliferative signaling: Cancer cells can initiate and maintain continuous cell division by producing their own growth factors or by altering the sensitivity of receptors to growth factors. (ii) Evading growth suppressors: Cancer cells can bypass the normal cellular mechanisms that limit cell division and growth, such as the inactivation of tumor suppressor genes. (iii) Resisting cell death: Cancer cells develop resistance to apoptosis, the programmed cell death process, which allows them to survive and continue dividing. (iv) Enabling replicative immortality: Cancer cells can extend their ability to divide indefinitely by maintaining the length of telomeres, the protective end caps on chromosomes. (v) Inducing angiogenesis: Cancer cells stimulate the growth of new blood vessels from existing vessels, providing the necessary nutrients and oxygen to support their rapid growth. (vi) Activating invasion and metastasis: Cancer cells can invade surrounding tissues and migrate to distant sites in the body, forming secondary tumors called metastases. (vii) Cellular energetics: Cancer cells rewire their metabolism to support rapid cell division and growth, often relying more on glycolysis even in the presence of oxygen (a phenomenon known as the Warburg effect). (viii) Avoiding immune destruction: Cancer cells can avoid detection and elimination by the immune system through various mechanisms, such as downregulating cell surface markers or producing immunosuppressive signals. (ix) Tumor promoting inflammation: Chronic inflammation can promote the development and progression of cancer by supplying growth factors, survival signals, and other molecules that facilitate cancer cell proliferation and survival. (x) Genome instability and mutation: Cancer cells exhibit increased genomic instability, leading to a higher mutation rate, which in turn drives the evolution of more aggressive and drug-resistant cancer cells. Classify the following sentence in one of the above 10 hallmark cancer taxonomy. If cannot be classified, answer as "empty": [SENTENCE]
PubMedQA	Question Answering	450 / 50 / 500	For the question, the reference context, and the answer given below, is it possible to infer the answer for that question from the reference context? Only reply as either Yes or No or Maybe. Question: [QUESTION] Reference context: [REFERENCE CONTEXT] Answer: [ANSWER]

Table 1: Datasets Details with our Prompts for Relation Extraction, Document Classification, and Question Answering Tasks.

biomedical datasets have helped these models to achieve state-of-the-art performance in a variety of BioNLP tasks (Gu et al., 2021). However, one major limitation of using such fine-tuned models is that they require task-specific large annotated datasets, which is significantly less available in the BioNLP domain in comparison to the general NLP domain. In this regard, having a strong zero-shot model could potentially alleviate the need for large annotated datasets, as it could enable the model to perform well on tasks that it was not trained on.

Recently, large autoregressive language models like GPT-3 (Brown et al., 2020) have demonstrated impressive few-shot learning capability. More recently, a new variant of GPT-3, called the Instruct-GPT model (Ouyang et al., 2022) has been proposed that leverages the reinforcement learning from human feedback (RLHF) mechanism. The resulting InstructGPT models (in other words, GPT-3.5) are much better at following instructions than the original GPT-3 model, resulting in an impressive zero-shot performance across various tasks. ChatGPT, a very recent addition to the GPT-3.5 series, has been trained using dialog-based instructional data alongside its regular training phase. Though ChatGPT has demonstrated strong zero-

shot performance across various NLP tasks (Laskar et al., 2023; Qin et al., 2023; Bang et al., 2023; Yang et al., 2023), it is yet to be investigated in the biomedical domain. To this end, this paper aims to evaluate the zero-shot capability of ChatGPT in the biomedical domain.

3 Our Methodology

For a given test sample X , we prepare a task instruction T and concatenate the text in the test sample with the task instruction to construct the prompt P . Then the prompt P is given as input to ChatGPT (*gpt-3.5-turbo*) to generate the response R . In this paper, we evaluate ChatGPT on 4 biomedical tasks across 11 benchmark datasets. Below, we describe these tasks, the datasets we use for evaluation, and the prompt P that we construct for each task depending on the respective dataset.

(i) Relation Extraction: Given a text sequence S , the biomedical relation extraction task aims to extract relations between entities mentioned in the text by identifying all possible relation triplets. In this paper, we evaluate drug-target-interaction in the KD-DTI dataset (Hou et al., 2022), chemical-disease-relation in the BC5CDR dataset (Li et al.,

Dataset	Type	Data Split (Train / Valid / Test)	Prompt
iCliniq	Dialog Summarization	24851 / 3105 / 3108	Write a very short and concise one line summary of the following dialogue as an informal question in a healthcare forum: [DIALOGUE]
HealthCare Magic	Dialog Summarization	181122 / 22641 / 22642	Write a very short and concise one line summary of the following dialogue as a question in a healthcare forum: [DIALOGUE]
MeQSum	Question Summarization	500 / - / 500	Rewrite the following question in a short and concise form: [QUESTION]
MEDIQA-QS	Question Summarization	- / 50 / 100	Rewrite the following question in a short and concise form: [QUESTION]
MEDIQA-MAS	Answer Summarization	- / 50 / 80	For the following question, some relevant answers are given below. Please write down a short concise answer by summarizing the given answers. Question: [QUESTION] Answer 1: [ANSWER1] Answer 2: [ANSWER2]
MEDIQA-ANS	Answer Summarization	- / - / 552	Write a very short and concise summary of the following article based on the question given below: [ARTICLE]

Table 2: Datasets Details with our Prompts for Abstractive Summarization tasks.

2016), and drug-drug-interaction in the DDI dataset (Herrero-Zazo et al., 2013). Our prompts for these datasets are demonstrated in Table 1.

(ii) Document Classification: Given a text document S , the goal is to classify the type of the document. For this task, we use the HoC (the Hallmarks of Cancers corpus) dataset (Baker et al., 2016) that consists of 1580 PubMed abstracts. This dataset was annotated at the sentence level by human experts among ten currently known hallmarks of cancer. Our prompt is shown in Table 1.

(iii) Question Answering: For the question-answering task, we evaluate the performance of ChatGPT on the PubMedQA dataset (Jin et al., 2019). Here, the objective is to determine whether the answer to a given question can be inferred from the reference context. We give the question, the reference context, and the answer as input to ChatGPT to determine whether the answer to the given question can be inferred from the given reference context, with ChatGPT being prompted to reply either as *yes*, *no*, or *maybe* (see Table 1 for details).

(iv) Abstractive Summarization: Given a text sequence S , the goal is to generate a concise abstractive summary of S . To this end, we evaluate ChatGPT on various biomedical summarization tasks, such as healthcare question summarization (we used *MeQSum* (Abacha and Demner-Fushman, 2019) and *MEDIQA-QS* (Abacha et al., 2021) datasets), medical answer summarization (we used *MEDIQA-ANS* (Savery et al., 2020) and *MEDIQA-MAS* (Abacha et al., 2021) datasets), and dialogue summarization (we used the *iCliniq* and *HealthCareMagic* datasets (Zeng et al., 2020; Mrini et al., 2021) for doctor-patient dialogue summariza-

tion to generate short queries for healthcare forums describing patient’s medical conditions). We show our prompts for this task in Table 2.

4 Experiments

Since ChatGPT is a generative model, we consider two state-of-the-art generative transformers as our baselines. Below, we first present these baselines, followed by presenting the results.

4.1 Fine-tuned Baselines

BioGPT: The backbone of BioGPT (Luo et al., 2022) is GPT-2 (Radford et al., 2019), which is a decoder of the transformer. The BioGPT model was trained over PubMed titles and abstracts via leveraging the standard language modeling task. We compare zero-shot ChatGPT with BioGPT models fine-tuned on relation extraction, document classification, and question-answering tasks.

BioBART: BioBART is a sequence-to-sequence model that was pre-trained over PubMed abstracts (Yuan et al., 2022a). The pre-training process involves reconstructing corrupted input sequences. We compare the zero-shot ChatGPT with BioBART fine-tuned on abstractive summarization datasets.

4.2 Results & Discussion

We first compare the performance of ChatGPT with BioGPT on relation extraction, document classification, and the question-answering task (see Table 3). Then we compare its performance with BioBART on summarization datasets (see Table 4). More evaluation details are given in Appendix A.1.

Relation Extraction Evaluation: We observe that in the BC5CDR and KD-DTI datasets for re-

Model	BC5CDR			KD-DTI			DDI			HoC	PubMedQA
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	F1	Accuracy
ChatGPT	36.20	73.10	48.42	19.19	66.02	29.74	47.11	45.77	46.43	59.14	51.60
BioGPT (Luo et al., 2022)	49.52	43.25	46.17	40.00	39.72	38.42	41.70	44.75	40.76	85.12	78.20

Table 3: Performance on Relation Extraction, Document Classification, and Question Answering datasets.

Model	iCliniq				HealthCareMagic				MeQSum				MEDIQA-QS				MEDIQA-MAS				MEDIQA-ANS			
	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S
ChatGPT	30.5	12.8	25.4	89.3	28.1	9.8	2.4	88.9	30.0	12.3	26.2	89.0	30.6	11.6	26.7	89.0	38.9	14.6	22.1	87.9	28.7	10.4	24.4	89.0
BioBART-Base	61.1	48.5	59.4	94.1	46.7	26.0	44.1	91.8	53.8	36.5	51.3	92.9	30.1	11.3	27.4	89.8	32.9	11.3	29.3	86.1	19.0	7.5	16.8	85.0
BioBART-Large	60.3	48.0	58.7	94.0	46.5	26.1	44.2	91.9	55.6	38.1	53.2	93.3	32.0	12.4	29.7	90.3	30.6	10.4	27.0	86.1	21.6	9.3	19.2	85.7

Table 4: Performance on Abstractive Summarization datasets. Here, ‘R-1’, ‘R-2’, ‘R-L’ and ‘B-S’ denote ‘ROUGE-1’, ‘ROUGE-2’, ‘ROUGE-L’, and ‘BERTScore’, respectively. For BERTScore (Zhang et al., 2019), we use the RoBERTa-Large model (Liu et al., 2019). The results for BioBART models are taken from Yuan et al. (2022a).

lation extraction, ChatGPT led to higher recall scores but much lower precision scores compared to the fine-tuned BioGPT model. This is because ChatGPT tends to generate long and descriptive responses, leading to many inaccurate relation extractions. Though in terms of F1, it outperforms fine-tuned BioGPT in the BC5CDR dataset, it fails to outperform in the KD-DTI dataset. More importantly, it outperforms BioGPT in the DDI dataset in all metrics: Precision, Recall, and F1.

While analyzing the results in different datasets, we observe that in both BC5CDR and DDI datasets where ChatGPT outperforms BioGPT, the training set is small, only 500 and 664 instances, respectively. On the other hand, in the KD-DTI dataset where ChatGPT fails to outperform BioGPT, the training set contains 12000 instances. This gives us a strong indication that even in the biomedical domain, zero-shot ChatGPT can outperform fine-tuned biomedical models in smaller-sized datasets.

We also observe that more descriptive prompts may help ChatGPT to obtain better Precision scores. Contrary to the KD-DTI dataset, we describe the definition of each interaction type in the DDI dataset (see Table 1) where ChatGPT performs the best. To further investigate the effect of prompts in relation extraction, we evaluate the performance in BC5CDR with a new prompt:

i. Identify the chemical-disease interactions in the passage given below: [PASSAGE].

We observe that the Precision, Recall, and F1 scores are decreased by 16.07%, 10.3%, and 14.29%, respectively, with this prompt variation.

Document Classification Evaluation: We observe that in the HoC dataset, the zero-shot ChatGPT achieves an F1 score of 59.14, in comparison to its counterpart fine-tuned BioGPT which

achieves an F1 score of 85.12. We also investigate the effect of prompt tuning by evaluating with two new prompts that are less descriptive (see Appendix A.2 for more details):

i. Prompting without explicitly mentioning the name of 10 HoC classes, drops F1 to 38.20.

ii. Prompting with the name of each HoC class is given without providing the definition of each class, drops the F1 score to 46.93.

Question Answering Evaluation: We observe that in the PubMedQA dataset, the zero-shot ChatGPT achieves much lower accuracy than BioGPT (51.60 by ChatGPT in comparison to 78.20 by BioGPT). However, the BioGPT model was fine-tuned on about 270K QA-pairs in various versions of the PubMedQA dataset for this task. While ChatGPT achieves more than 50% accuracy even without any few-shot examples in the prompt.

Summarization Evaluation: We observe that in terms of all ROUGE scores (Lin, 2004), ChatGPT performs much worse than BioBART in datasets that have dedicated training sets, such as iCliniq, HealthCareMagic, and MeQSum. Meanwhile, it performs on par with BioBART in the MEDIQA-QS dataset. More importantly, it outperforms BioBART in both MEDIQA-ANS and MEDIQA-MAS datasets. Note that MEDIQA-ANS, MEDIQA-MAS, and MEDIQA-QS datasets do not have any dedicated training data and ChatGPT achieves comparable or even better performance in these datasets compared to the BioBART model fine-tuned on other related datasets (Yuan et al., 2022a). This further confirms that zero-shot ChatGPT is more useful than domain-specific fine-tuned models in biomedical datasets that lack large training data.

5 Conclusions and Future Work

In this paper, we evaluate ChatGPT on 4 benchmark biomedical tasks to observe that in datasets that have large training data, ChatGPT performs quite poorly in comparison to the fine-tuned models (BioGPT and BioBART), whereas it outperforms fine-tuned models on datasets where the training data size is small. These findings suggest that ChatGPT can be useful in low-resource biomedical tasks. We also observe that ChatGPT is sensitive to prompts, as variations in prompts led to a noticeable difference in results.

Though in this paper, we mostly evaluate ChatGPT on tasks that require it to generate responses by only analyzing the input text, in the future, we will investigate the performance of ChatGPT on more challenging tasks, such as named entity recognition and entity linking (Yadav and Bethard, 2018; Yan et al., 2021; Yuan et al., 2022b; Laskar et al., 2022a,b,c), as well as problems in information retrieval (Huang et al., 2005; Huang and Hu, 2009; Yin et al., 2010; Laskar et al., 2020, 2022d). We will also explore the ethical implications (bias or privacy concerns) of using ChatGPT in the biomedical domain.

Limitations

Since the training datasets of ChatGPT are unknown, some data used for evaluation may or may not exist during the training phase of ChatGPT. Also, a new version called the GPT-4 model has been released that may ensure higher accuracy. Nonetheless, GPT-4 is very costly to use, around 60x more expensive than ChatGPT. Meanwhile, even using the paid ChatGPT Plus² subscription, it is available for just a limited use (allows evaluation of only 25 samples in 3 hours). Another limitation of this research is that the results mentioned in this paper for ChatGPT may not be reproducible, as ChatGPT may generate different responses for the same input prompt. Although the experimental results may change over time, this work will still give a concrete direction for future research using ChatGPT like large language models in the biomedical domain.

Ethics Statement

The paper evaluates ChatGPT on 4 benchmark biomedical tasks that require ChatGPT to generate

²<https://openai.com/blog/chatgpt-plus>, last accessed June 4th, 2023

a response based on the information provided in the input text. Thus, no data or prompt was provided as input that could lead to ChatGPT generating any responses that pose any ethical or privacy concerns. This evaluation is only done in some academic datasets that already have gold labels available and so it does not create any concerns like humans relying on ChatGPT responses for sensitive issues, such as disease diagnosis. Since this paper only evaluates the performance of ChatGPT and investigates its effectiveness and limitations, conducting this evaluation does not lead to any unwanted biases. Only the publicly available academic datasets are used that did not require any licensing. Thus, no personally identifiable information has been used.

Acknowledgements

We would like to thank all the anonymous reviewers for their detailed review comments. This work is done at York University and supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada and the York Research Chairs (YRC) program.

References

- Asma Ben Abacha and Dina Demner-Fushman. 2019. On the summarization of consumer health questions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2228–2234.
- Asma Ben Abacha, Yassine M'rabet, Yuhao Zhang, Chaitanya Shivade, Curtis Langlotz, and Dina Demner-Fushman. 2021. Overview of the mediqua 2021 shared task on summarization in the medical domain. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 74–85.
- Emily Alsentzer, John R Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. *arXiv preprint arXiv:1904.03323*.
- Simon Baker, Ilona Silins, Yufan Guo, Imran Ali, Johan Högberg, Ulla Stenius, and Anna Korhonen. 2016. Automatic semantic classification of scientific literature according to the hallmarks of cancer. *Bioinformatics*, 32(3):432–440.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert: A pretrained language model for scientific text. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-specific language model pretraining for biomedical natural language processing. *arXiv preprint arXiv:2007.15779*.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- María Herrero-Zazo, Isabel Segura-Bedmar, Paloma Martínez, and Thierry Declerck. 2013. The ddi corpus: An annotated corpus with pharmacological substances and drug–drug interactions. *Journal of biomedical informatics*, 46(5):914–920.
- Yutai Hou, Yingce Xia, Lijun Wu, Shufang Xie, Yang Fan, Jinhua Zhu, Tao Qin, and Tie-Yan Liu. 2022. Discovering drug–target interaction knowledge from biomedical literature. *Bioinformatics*, 38(22):5100–5107.
- Xiangji Huang and Qinmin Hu. 2009. A bayesian learning approach to promoting diversity in ranking for biomedical information retrieval. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 307–314.
- Xiangji Huang, Ming Zhong, and Luo Si. 2005. **York university at TREC 2005: Genomics track**. In *Proceedings of the Fourteenth Text REtrieval Conference, TREC 2005, Gaithersburg, Maryland, USA, November 15-18, 2005*, volume 500-266 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*.
- Katikapalli Subramanyam Kalyan, Ajit Rajasekharan, and Sivanesan Sangeetha. 2022. Ammu: a survey of transformer-based biomedical pretrained language models. *Journal of biomedical informatics*, 126:103982.
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Xiangji Huang. 2023. **A systematic study and comprehensive evaluation of chatgpt on benchmark datasets**. *arXiv preprint arXiv:2305.18486*.
- Md Tahmid Rahman Laskar, Cheng Chen, Xue-yong Fu, and Shashi Bhushan Tn. 2022a. Improving named entity recognition in telephone conversations via effective active learning with human in the loop. In *Proceedings of the Fourth Workshop on Data Science with Human-in-the-Loop (Language Advances)*, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Md Tahmid Rahman Laskar, Cheng Chen, Jonathan Johnston, Xue-Yong Fu, Shashi Bhushan TN, and Simon Corston-Oliver. 2022b. An auto encoder-based dimensionality reduction technique for efficient entity linking in business phone conversations. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3363–3367.
- Md Tahmid Rahman Laskar, Cheng Chen, Aliaksandr Martsinovich, Jonathan Johnston, Xue-Yong Fu, Shashi Bhushan Tn, and Simon Corston-Oliver. 2022c. **BLINK with Elasticsearch for efficient entity linking in business conversations**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 344–352, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Md Tahmid Rahman Laskar, Enamul Hoque, and Jimmy Xiangji Huang. 2022d. **Domain adaptation with pre-trained transformers for query-focused abstractive text summarization**. *Computational Linguistics*, 48(2):279–320.
- Md Tahmid Rahman Laskar, Xiangji Huang, and Enamul Hoque. 2020. Contextualized embeddings based transformer encoder for sentence similarity modeling in answer selection task. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5505–5514.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.

- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciak, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. Biogpt: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6).
- Khalil Mrini, Franck Deroncourt, Seunghyun Yoon, Trung Bui, Walter Chang, Emilia Farcas, and Ndapandula Nakashole. 2021. A gradually soft multi-task and data-augmented approach to medical question understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1505–1515.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 58–65.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. Is chatgpt a general-purpose natural language processing task solver? *arXiv preprint arXiv:2302.06476*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Max Savery, Asma Ben Abacha, Soumya Gayen, and Dina Demner-Fushman. 2020. Question-driven summarization of answers to consumer health questions. *Scientific Data*, 7(1):1–9.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Vikas Yadav and Steven Bethard. 2018. A survey on recent advances in named entity recognition from deep learning models. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2145–2158.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various ner subtasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5808–5822.
- Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023. Exploring the limits of chatgpt for query or aspect-based text summarization. *arXiv preprint arXiv:2302.08081*.
- Xiaoshi Yin, Jimmy Xiangji Huang, Xiaofeng Zhou, and Zhoujun Li. 2010. A survival modeling approach to biomedical search result diversification using wikipedia. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 901–902.
- Hongyi Yuan, Zheng Yuan, Ruyi Gan, Jiaying Zhang, Yutao Xie, and Sheng Yu. 2022a. Biobart: Pretraining and evaluation of a biomedical generative language model. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 97–109.
- Hongyi Yuan, Zheng Yuan, and Sheng Yu. 2022b. Generative biomedical entity linking via knowledge base-guided pre-training and synonyms-aware fine-tuning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4038–4048.
- Guangtao Zeng, Wenmian Yang, Zeqian Ju, Yue Yang, Sicheng Wang, Ruisi Zhang, Meng Zhou, Jiaqi Zeng, Xiangyu Dong, Ruoyu Zhang, et al. 2020. Med-dialog: Large-scale medical dialogue datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9241–9250.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. 2023. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *arXiv preprint arXiv:2302.09419*.

A Appendix

A.1 Evaluating ChatGPT on Different Tasks

Since ChatGPT generated responses can be lengthy, and may sometimes contain unnecessary information while not in a specific format, especially in tasks that may have multiple answers (e.g., Relation Extraction), it could be quite difficult to automatically evaluate its performance in such tasks by comparing with the gold labels by using just an evaluation script. Thus, for some datasets and tasks, we manually evaluate the ChatGPT generated responses by ourselves and compare them with the gold labels. Below we describe our evaluation approach for different tasks:

- **Relation Extraction:** The authors manually evaluated the ChatGPT generated responses in this task by comparing them with the gold labels. To ensure the reproducibility of our evaluation, we will release the ChatGPT generated responses.
- **Document Classification:** We created an evaluation script and identifies if the gold label (one of the 10 HoC classes) is present in the ChatGPT generated response. For fair evaluation, we lowercase each character in both the gold label and the ChatGPT generated response. Our evaluation script will be made publicly available to ensure reproducibility of our findings.
- **Question Answering:** Similar to Document Classification, we also evaluated using an evaluation script that compares the gold label and the ChatGPT generated response (here, we also convert each character to lowercase). The evaluation script will also be made public.
- **Abstractive Summarization:** We used the HuggingFace’s Evaluate³ library (Wolf et al., 2020) to calculate the ROUGE scores and the BERTScore for the Abstractive Summarization task evaluation.

³<https://huggingface.co/docs/evaluate/index>

A.2 Effects of Prompt Variation

We investigate the effects of prompt tuning in the HoC dataset by evaluating the performance of ChatGPT based on the following prompt variations:

- Prompting with explicitly defining the 10 HoC classes achieves an F1 score of 59.14 (see Row 1 in Table 5).
- Prompting without explicitly mentioning the name of 10 HoC classes, drops F1 to 38.20 (see Row 2 in Table 5).
- Prompting with the name of each HoC class is given without providing the definition of each class, drops the F1 score to 46.93 (see Row 3 in Table 5).

Our findings demonstrate that more descriptive prompts yield better results.

A.3 Sample ChatGPT Generated Responses

Some sample prompts with the ChatGPT generated responses for Relation Extraction, Document Classification, and Question Answering tasks are given in Table 6 and for the Abstractive Summarization task are given in Table 7.

#	Prompt	F1
1.	<p>The 10 hallmark cancer taxonomy with their definitions MeQSum given below:</p> <p>(i) Sustaining proliferative signaling: Cancer cells can initiate and maintain continuous cell division by producing their own growth factors or by altering the sensitivity of receptors to growth factors.</p> <p>(ii) Evading growth suppressors: Cancer cells can bypass the normal cellular mechanisms that limit cell division and growth, such as the inactivation of tumor suppressor genes.</p> <p>(iii) Resisting cell death: Cancer cells develop resistance to apoptosis, the programmed cell death process, which allows them to survive and continue dividing.</p> <p>(iv) Enabling replicative immortality: Cancer cells can extend their ability to divide indefinitely by maintaining the length of telomeres, the protective end caps on chromosomes.</p> <p>(v) Inducing angiogenesis: Cancer cells stimulate the growth of new blood vessels from existing vessels, providing the necessary nutrients and oxygen to support their rapid growth.</p> <p>(vi) Activating invasion and metastasis: Cancer cells can invade surrounding tissues and migrate to distant sites in the body, forming secondary tumors called metastases.</p> <p>(vii) Cellular energetics: Cancer cells rewire their metabolism to support rapid cell division and growth, often relying more on glycolysis even in the presence of oxygen (a phenomenon known as the Warburg effect).</p> <p>(viii) Avoiding immune destruction: Cancer cells can avoid detection and elimination by the immune system through various mechanisms, such as downregulating cell surface markers or producing immunosuppressive signals.</p> <p>(ix) Tumor promoting inflammation: Chronic inflammation can promote the development and progression of cancer by supplying growth factors, survival signals, and other molecules that facilitate cancer cell proliferation and survival.</p> <p>(x) Genome instability and mutation: Cancer cells exhibit increased genomic instability, leading to a higher mutation rate, which in turn drives the evolution of more aggressive and drug-resistant cancer cells.</p> <p>Classify the following sentence in one of the above 10 hallmark cancer taxonomy. If cannot be classified, answer as "empty":</p> <p>[SENTENCE]</p>	59.14
2.	<p>Is it possible to classify the following sentence in one of the 10 categories in the Hallmarks of Cancer taxonomy? If possible, write down the class.</p> <p>[SENTENCE]</p>	38.20
3.	<p>Classify the sentence given below in one of the 10 categories (i. activating invasion and metastasis, ii. tumor promoting inflammation, iii. inducing angiogenesis, iv. evading growth suppressors, v. resisting cell death,vi. cellular energetics, vii. genomic instability and mutation, viii. sustaining proliferative signaling, ix. avoiding immune destruction, x. enabling replicative immortality) in the Hallmarks of Cancer taxonomy? If cannot be classified, answer as "empty".</p> <p>[SENTENCE]</p>	46.93

Table 5: Effects of Prompt Variation in the HoC dataset for the Document Classification Task.

Dataset	Type	Prompt	ChatGPT Response	Gold Label
BC5CDR	Chemical-Disease Relation Extraction	Identify each pair of drugs and the drug-induced side-effects (e.g., diseases) in the following passage: Famotidine-associated delirium. A series of six cases. Famotidine is a histamine H2-receptor antagonist used in inpatient settings for prevention of stress ulcers and is showing increasing popularity because of its low cost. Although all of the currently available H2-receptor antagonists have shown the propensity to cause delirium, only two previously reported cases have been associated with famotidine. The authors report on six cases of famotidine-associated delirium in hospitalized patients who cleared completely upon removal of famotidine. The pharmacokinetics of famotidine are reviewed, with no change in its metabolism in the elderly population seen. The implications of using famotidine in elderly persons are discussed.	Drug: Famotidine Drug-induced disease: Delirium Side effect: Increased risk of delirium in elderly patients	chemical: famotidine disease: delirium
KD-DTI	Drug-Target Relation Extraction	Identify the drug-target interactions in the following passage (along with the interaction type among the following: 'inhibitor', 'agonist', 'modulator', 'activator', 'blocker', 'inducer', 'antagonist', 'cleavage', 'disruption', 'intercalation', 'inactivator', 'bind', 'binder', 'partial agonist', 'cofactor', 'substrate', 'ligand', 'chelator', 'downregulator', 'other', 'antibody', 'other/unknown'): Analogue inhibitors by modifying oseltamivir based on the crystal neuraminidase structure for treating drug-resistant H5N1 virus. The worldwide spread of H5N1 avian influenza and the increasing reports about its resistance to the existing drugs have made a priority for the development of the new anti-influenza molecules. The crystal structure of H5N1 avian influenza neuraminidase reported recently by Russell et al. [R.J. Russell, L.F. Haire, D.J. Stevens, P.J. Collins, Y. P. Lin, G.M. Blackburn, A.J. Hay, S.J. Gamblin, J.J. Skehel, The structure of H5N1 avian influenza neuraminidase suggests new opportunities for drug design, Nature 443 (2006) 45-49] have provided new opportunities for drug design in this regard. It is revealed through the structure that the active sites of the group-1 neuraminidases, which contain the N1 subtype, have a very different three-dimensional structure from those of group-2 neuraminidases. The key difference is in the 150-loop cavity adjacent to the conserved active site in neuraminidase. Based on these findings and by modifying oseltamivir, six analog inhibitors were proposed as candidates for developing inhibitors against H5N1 virus, particularly against the oseltamivir-resistant H5N1 virus strain.	Drug-target interactions: - Oseltamivir and neuraminidase - Analog inhibitors and neuraminidase Interaction type: Inhibitor	drug: Oseltamivir target: Neuraminidase interaction: inhibitor
DDI	Drug-Drug Relation Extraction	Identify the pairs of drug-drug interactions in the passage given below based on one of the following interaction types: (i) mechanism: this type is used to identify drug-drug interactions that are described by their pharmacokinetic mechanism. (ii) effect: this type is used to identify drug-drug interactions describing an effect. (iii) advice: this type is used when a recommendation or advice regarding a drug-drug interaction is given. (iv) int: this type is used when a drug-drug interaction appears in the text without providing any additional information. Interactions between Leukine and other drugs have not been fully evaluated. Drugs which may potentiate the myeloproliferative effects of Leukine, such as lithium and corticosteroids, should be used with caution.	1. Leukine and lithium - effect 2. Leukine and corticosteroids - effect 3. Lithium and Leukine - advice 4. Corticosteroids and Leukine - advice	drug: Leukine target: lithium interaction: effect drug: Leukine target: corticosteroids interaction: effect
HoC	Document Classification	The 10 hallmark cancer taxonomy with their definitions are given below: (i) Sustaining proliferative signaling: Cancer cells can initiate and maintain continuous cell division by producing their own growth factors or by altering the sensitivity of receptors to growth factors. (ii) Evading growth suppressors: Cancer cells can bypass the normal cellular mechanisms that limit cell division and growth, such as the inactivation of tumor suppressor genes. (iii) Resisting cell death: Cancer cells develop resistance to apoptosis, the programmed cell death process, which allows them to survive and continue dividing. (iv) Enabling replicative immortality: Cancer cells can extend their ability to divide indefinitely by maintaining the length of telomeres, the protective end caps on chromosomes. (v) Inducing angiogenesis: Cancer cells stimulate the growth of new blood vessels from existing vessels, providing the necessary nutrients and oxygen to support their rapid growth. (vi) Activating invasion and metastasis: Cancer cells can invade surrounding tissues and migrate to distant sites in the body, forming secondary tumors called metastases. (vii) Cellular energetics: Cancer cells rewire their metabolism to support rapid cell division and growth, often relying more on glycolysis even in the presence of oxygen (a phenomenon known as the Warburg effect). (viii) Avoiding immune destruction: Cancer cells can avoid detection and elimination by the immune system through various mechanisms, such as downregulating cell surface markers or producing immunosuppressive signals. (ix) Tumor promoting inflammation: Chronic inflammation can promote the development and progression of cancer by supplying growth factors, survival signals, and other molecules that facilitate cancer cell proliferation and survival. (x) Genome instability and mutation: Cancer cells exhibit increased genomic instability, leading to a higher mutation rate, which in turn drives the evolution of more aggressive and drug-resistant cancer cells. Classify the following sentence in one of the above 10 hallmark cancer taxonomy. If cannot be classified, answer as "empty": Heterogeneity in DNA damage within the cell population was observed as a function of radiation dose.	genomic instability and mutation.	genomic instability and mutation.
PubMedQA	Question Answering	For the question, the reference context, and the answer given below, is it possible to infer the answer for that question from the reference context? Only reply as either Yes or No or Maybe. Question: Do some U.S. states have higher/lower injury mortality rates than others? Reference context: This article examines the hypothesis that the six U.S. states with the highest rates of road traffic deaths (group 1 states) also had above-average rates of other forms of injury such as falling, poisoning, drowning, fire, suffocation, homicide, and suicide, and also for the retail trade and construction industries. The converse, second hypothesis, for the six states with the lowest rates of road traffic deaths (group 2 states) is also examined. Data for these 12 states for the period 1983 to 1995 included nine categories of unintentional and four categories of intentional injury. Seventy-four percent of the group 1 states conformed to the first hypothesis, and 85% of the group 2 states conformed to the second hypothesis. Answer: Group 1 states are likely to exhibit above-average rates for most other categories of injury death, whereas group 2 states are even more likely to exhibit below-average rates for most other categories of injury death.	Yes.	Yes.

Table 6: Sample prompts with ChatGPT generated responses for Relation Extraction, Document Classification, and Question Answering Tasks.

Dataset	Type	Prompt	ChatGPT Response	Gold Summary
iCliniq	Dialog Summarization	Write a very short and concise one line summary of the following dialogue as an informal question in a healthcare forum: Patient: Hello doctor, I am depressed. What can I do? My parents are not supporting for a surgery which is important for my self-confidence and self-esteem. My father is not taking care of my health or well-being. Please help. Doctor: Hello. Any related incident or cause for having depression is unclear in this query. Details of surgery and symptoms which are not mentioned are appreciated to know the complete picture. Depressive feelings are generally linked with some loss, self-related issues, stress, or environmental or biological factors. It can be managed by: 1. Medicines. 2. Psychotherapy to identify the causing agents and try resolving the issue. 3. Modifying thoughts related to concerns. 4. Physical exercises like walking, jogging, and workout. 5. Diet changes - Reduce the intake of too much salt and junk food and reduce sugar intake. Consume healthy foods and fruits having serotonin and omega-3 can reduce depression.	What are some ways to manage depression without surgery and what details should I provide to my doctor?	I am depressed, and my parents are not supporting me to overcome it. Please help
HealthCare Magic	Dialog Summarization	Write a very short and concise one line summary of the following dialogue as a question in a healthcare forum: Patient: Dear Doctor,I have a 1.5 years old sick daughter. I am looking for the written file to understand some clinical numbers of a normal young child(1-2 years old):1: what is the diameter of the hepatic portal vein? 2: what is the frequency of the hepatic portal flow?I will appreciate if you could help me with these information.Kelvin Lu Doctor: 4-8 mm is range for diameter for hepatic portal vein for 1.5 year of children , there is no thing like frequency of hepatic portal flow. please make it clear	What is the diameter of the hepatic portal vein in a 1-2 year old child, and what is the frequency of the hepatic portal flow?	What should be the diameter of hepatic portal vein of a 1.5 years old?
MeQSum	Question Summarization	Rewrite the following question in a short and concise form: Question: Williams' syndrome . I would like to have my daughter tested for William's syndrome. Could you please tell me where I would go or who does it in my area? Thank you!!	Where can I get my daughter tested for Williams' syndrome in my area?	Where can I get genetic testing for william's syndrome?
MEDIQA-QS	Question Summarization	Rewrite the following question in a short and concise form: Question: 23 surgeries and counting.....lower lip birthmark, have tried all options out there and guess what still have it, continues to grow back.....any suggestions? Is there a cure coming in the next few years hopefully?	Any cure for a growing lower lip birthmark after 23 surgeries?	How can I get rid of a lower lip birthmark permanently?
MEDIQA-ANS	Answer Summarization	Write a very short and concise summary of the following article based on the question given below: Question: Bad Breath I have very bad breath and at times it can make myself and others sick. I need some advice as to what I need to do. Article: Bad Breath Halitosis Summary There are many reasons why you might have bad breath. You can get it if you don't brush and floss regularly. Bacteria that build up in your mouth and between your teeth produce the bad odor. Other problems in your mouth, such as gum disease, dry mouth or cavities, may also cause it. Sinusitis or problems with your nose may be to blame. You can also have bad breath if you eat some foods, like raw onions, garlic or cabbage. And of course smoking causes its own bad aroma. Some diseases and medicines are associated with a specific breath odor. Having good dental habits, like brushing and flossing regularly, help fight bad breath. Mouthwashes, mints or chewing gum may make your breath fresher. If you have an underlying disorder, treating it may help eliminate the breath odor.	The article discusses the causes and considerations of bad breath, which is usually related to poor dental hygiene. Other disorders and diseases can also produce distinct breath odors, such as fruity breath being a sign of ketoacidosis in diabetes. Proper dental hygiene, avoiding smoking, and following healthcare provider's instructions are recommended for treating bad breath. If bad breath persists or is accompanied by other symptoms, medical attention may be necessary.	Bad breath could be caused by poor oral hygiene, dental and gastroesophageal diseases, infections, smoking, alcoholism, and certain foods.
MEDIQA-MAS	Answer Summarization	For the following question, some relevant answers are given below. Please write down a short concise answer by summarizing the given answers. Question: Can I use a Power Plate machine if I have a pacemaker? Answer 1: Most machines and devices will not interfere with your pacemaker. But some with strong magnetic fields may. Always ask your provider about any specific device that you need to avoid. DO NOT put a magnet near your pacemaker. Most appliances in your home are safe to be around. This includes your refrigerator, washer, dryer, toaster, blender, computers and fax machines, hair dryer, stove, CD player, remote controls, and microwaves. You should keep several devices at least 12 inches (30 centimeters) away from the site where the pacemaker is placed under your skin. These include: Battery powered cordless tools (such as screwdrivers and drills) Plug-in power tools (such as drills and table saws) Electric lawnmowers and leaf blowers Slot machines Stereo speakers Tell all providers that you have a pacemaker before any tests are done. Some medical equipment may interfere with your pacemaker. Stay away from large motors, generators, and equipment. DO NOT lean over the open hood of a car that is running. Also stay away from: Radio transmitters and high-voltage power lines Products that use magnetic therapy, such as some mattresses, pillows, and massagers Large electrical- or gasoline-powered appliances If you have a cell phone: DO NOT put it in a pocket on the same side of your body as your pacemaker. When using your cell phone, hold it to your ear on the opposite side of your body. Be careful around metal detectors and security wands. Handheld security wands may interfere with your pacemaker. Show your wallet card and ask to be hand searched. Most security gates at airports and stores are OK. But DO NOT stand near these devices for long periods. Your pacemaker may set off alarms. After any operation, have your provider check your pacemaker. Answer 2: You will probably be able to go home after 1 day or even the same day in some cases. You should be able to return to your normal activity level quickly. Ask your provider how much you can use the arm on the side of your body where the pacemaker was placed. You may be advised not to: Lift anything heavier than 10 to 15 pounds (4.5 to 6.75 kilograms) Push, pull, and twist your arm for 2 to 3 weeks. Raise your arm above your shoulder for several weeks. When you leave the hospital, you will be given a card to keep in your wallet. This card lists the details of your pacemaker and has contact information for emergencies. You should always carry this wallet card with you. You should try to remember the name of the pacemaker manufacturer if you can in case you lose your card.	It is important to ask your provider about specific devices that may interfere with a pacemaker, but most appliances in the home are safe. It is recommended to stay away from large motors, generators, and equipment, products that use magnetic therapy, and to be cautious around metal detectors and security wands. Additionally, after getting a pacemaker, some activity restrictions may apply, such as not lifting heavy objects and limiting arm movements. It is important to carry a wallet card with pacemaker details and emergency contact information.	Most machines are safe to use with a pacemaker. Some devices should be kept at least 12 inches away. These include battery powered cordless tools (such as screwdrivers and drills) Plug-in power tools (such as drills and table saws) Electric lawnmowers and leaf blowers Slot machines Stereo speakers Some medical equipment may interfere with your pacemaker. Stay away from large motors, generators, and equipment. Radio transmitters and high-voltage power lines Products that use magnetic therapy, such as some mattresses, pillows, and massagers Large electrical- or gasoline-powered appliances. Do not put your cell phone in a pocket or to the ear on the same side as your pacemaker.

Table 7: Sample prompts with ChatGPT generated responses for Abstractive Summarization tasks.