

# NoteChat: A Dataset of Synthetic Doctor-Patient Conversations Conditioned on Clinical Notes

Anonymous EMNLP submission

## Abstract

The detailed clinical records drafted by doctors after each patient’s visit are crucial for medical practitioners and researchers. Automating the creation of these notes with language models can reduce the workload of doctors. However, training such models can be difficult due to the limited public availability of conversations between patients and doctors. In this paper, we introduce NoteChat, a cooperative multi-agent framework leveraging Large Language Models (LLMs) for generating synthetic doctor-patient conversations conditioned on clinical notes. NoteChat consists of Planning, Role-play, and Polish modules. We provide a comprehensive automatic and human evaluation of NoteChat, comparing it with state-of-the-art models, including OpenAI’s ChatGPT and GPT-4. Results demonstrate that NoteChat facilitates high-quality synthetic doctor-patient conversations, underscoring the untapped potential of LLMs in healthcare. This work represents the first instance of multiple LLMs cooperating to complete a doctor-patient conversation conditioned on clinical notes, offering promising avenues for the intersection of AI and healthcare.

## 1 Introduction

Large Language Models (LLMs) are sophisticated artificial intelligence (AI) systems extensively trained on massive amounts of textual data, representing a significant breakthrough in artificial intelligence (AI) (Brown et al., 2020; Longpre et al., 2023). The GPT series by OpenAI (OpenAI, 2023) and other LLMs (Touvron et al., 2023; Taori et al., 2023; Chiang et al., 2023) have demonstrated impressive outcomes and hold significant potential in revolutionizing multiple sectors, such as marketing, education, and customer service. However, the exploration of LLMs in the medical field remains relatively untapped (Gilson et al., 2023), despite their immense potential for transforming healthcare

	Ours-PMC	ChatDoctor	DoctorGLM	Ours-MTS	MTS-Dialog
#dial.	30k	112k	3.4M	20	87
#utt.	579k	224k	11.2M	1.25k	4.79k
Chat	✓	✗	✗	✓	✓
Note	✓	✗	✗	✓	✓
Syn.	AI	✗	✗	AI	Human
Lang	EN	EN	CN	EN	EN
# of utterances in a dialogue					
Avg	19.3	2	3.3	62.5	55.1
Max	61	2	198	112	131
Min	3	2	2	22	7

Table 1: Statistics of our NoteChat dataset and related publicly available resources. OursP and OursM mean our PMC-based and MTS-based datasets. Chat means multi-round question answering or not. Note means whether we can generate a full clinical note from the data. Syn. means whether the data is generated (by annotators or AI) or collected from the real-world.

communication and decision-making (Abacha and Zweigenbaum, 2015). The main reason is the lack of high-quality medical datasets that meet various needs.

To address this issue, Yunxiang et al. (2023) collect 100k real-world doctor-patient conversations from online medical consultation websites and train their ChatDoctor model. And Xiong et al. (2023) converted the ChatDoctor data into Chinese and additionally added relevant Chinese dialogue (Zeng et al., 2020) and question-answering datasets to train their DoctorGLM model. Nonetheless, as indicated in Table 1, the maximum average number of utterances in a dialogue used by ChatDoctor and DoctorGLM is just 3.3, which is typical of online medical consultation websites but markedly less than face-to-face communication between doctors and patients encounter (Drew et al., 2001). Furthermore, the available data consisted almost exclusively of unidirectional question-answering formats, where the doctor only answers the patient’s inquiry. In reality, doctors and patients engage in more dynamic and interactive conversations, where the doctors are expected to not only answer ques-

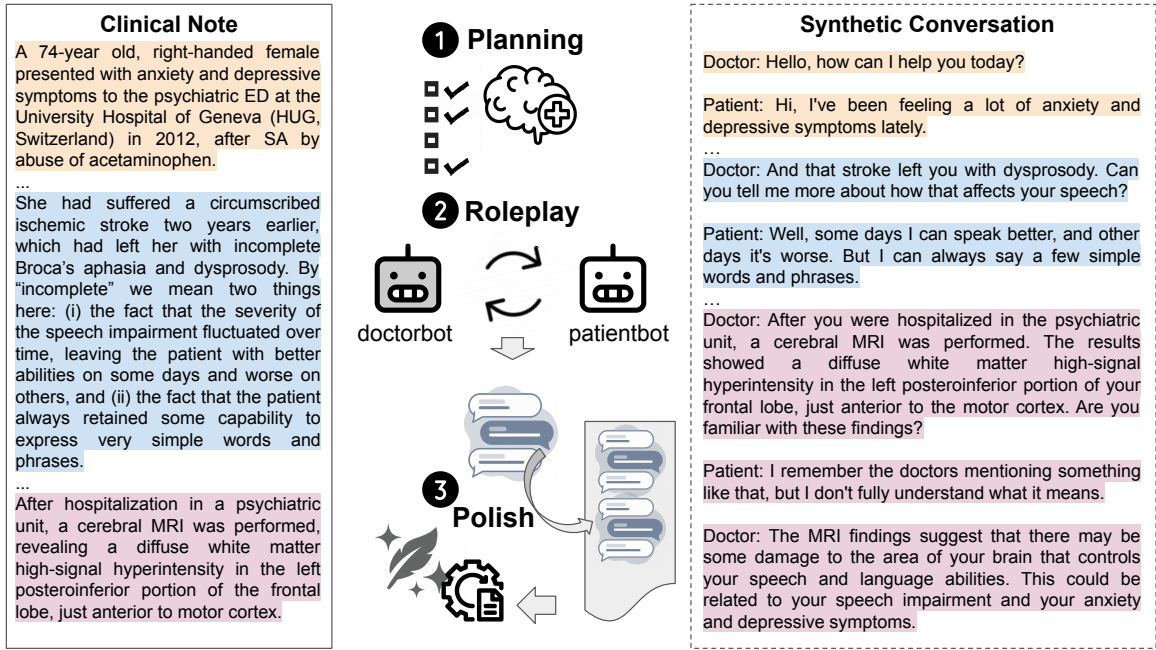


Figure 1: An illustration of NoteChat. **Apricot** indicates that our pipeline can generate smooth doctor-patient conversations. **Blue** shows the characteristics of information seeking, where doctors can actively ask questions to advance the conversation, thanks to ② Roleplay module. In addition, compared with the corresponding note content, the generated utterances are more colloquial, but the key medical concepts are highly overlapped, which reflects NoteChat’s control over factuality (mainly from ① Planning module). **Lavender** means that NoteChat can generate reasonable explanations for patients, and a lot of information in the chat is reasonable imagination instead of hallucination. The two modules of ② Roleplay and ③ Polish can stimulate the imaginative potential of LLMs and reduce unreasonable hallucination through self-examination.

tions but also guide the discussion by asking professional questions and ultimately conclude the conversation after accomplishing the intended objective (Maynard and Heritage, 2005). Typically, the desired objective of a doctor-patient conversation is to arrive at a diagnosis and complete the patient’s clinical note (Krishna et al., 2020; Yim and Yetisgen-Yildiz, 2021), which motivates this work to build a doctor-patient dialogues dataset.

However, it is challenging to collect real-world doctor-patient dialogues or construct synthetic doctor-patient dialogues. The primary difficulty in collecting such data is due to privacy protection and strict data regulations like HIPAA (Rindfleisch, 1997; Annas, 2003). Generating a synthetic doctor-patient dialogue dataset is another potential solution. Manually constructing a large amount of such synthetic data using domain experts is time and effort-consuming (Kelly et al., 2019; Abacha et al., 2023). An alternative is to build synthetic doctor-patient dialogue data by leveraging large language models such as GPT4 (Eysenbach et al., 2023; Li et al., 2023c; Dai et al., 2023).

In the paper, we introduce NoteChat, a coopera-

tive multi-agent framework that generates synthetic doctor-patient conversations conditioning on clinical notes. In the main dialogue generation loop, we first include a Roleplay module into the NoteChat framework where two ChatGPT agents<sup>1</sup> take on the roles of doctor and patient, respectively. Furthermore, we integrate one additional ChatGPT agent into the main loop, namely the Planning module, which is responsible for knowledge planning. Finally, with the feedback and suggestions from doctors and medical students, we added the Polish module to better align with the expert preferences. Extensive automatic and human evaluations demonstrate the efficacy of our cooperative multi-agent framework and show that NoteChat holds great promise for promoting high-quality synthetic doctor-patient conversations.

**In summary, our contributions are as follows:**

- To our knowledge, we are the first instance of multiple LLMs cooperating to complete a doctor-patient conversation conditioned on clinical notes.
- We evaluated the quality of the synthetic doctor-

<sup>1</sup>We use OpenAI’s GPT-3.5 model gpt-3.5-turbo.

patient conversations generated by NoteChat with the state-of-the-art OpenAI’s ChatGPT and GPT-4 using extensive intrinsic and extrinsic evaluation methods.

- Through comprehensive human evaluations, we demonstrate that NoteChat holds promise for high-quality synthetic doctor-patient dialogues.
- We will release the prompts, codes, and synthetic conversations on the 167k PMC-Patient notes to the research community <sup>2</sup>.

## 2 Methods

### 2.1 Data Source

**PMC-Patients** is a comprehensive dataset comprising 167K de-identified patient notes and relations extracted from a diverse range of case reports available in the PubMed Central (PMC) repository (Zhao et al., 2023).

**MTS-Dialog** is a new collection of 1.7k short doctor-patient conversations and corresponding summaries (section headers and contents) Abacha et al. (2023); Ben Abacha et al. (2023); Yim et al. (2023) to foster advancements in the field of automatic clinical note generation from doctor-patient conversations. And this 1.7k short version dataset has 87 corresponding versions of complete dialogues and clinical notes. In this paper, we use the 87 complete dialogues and clinical notes.

### 2.2 NoteChat: Generating doctor-patient dialogues from notes in the GPT Era

To ensure that our synthetic datasets closely resemble authentic dialogues, we first use the prompts in Appendix A.2 to guide the ChatGPT and GPT4 in generating high-quality data as our baselines. In this section, we introduce our NoteChat Framework for this task. All our NoteChat experiments in this paper are based on ChatGPT API (gpt-3.5-turbo), but NoteChat can be used in any model that can handle the instructions.

#### 2.2.1 Preprocessing

**PMC-Patient** PMC-Patient dataset encompasses a vast array of case reports, many of which pertain to rare conditions. To maintain the quality of

the generated dialogue in our study, we instruct ChatGPT to exclude exceptionally rare cases. Furthermore, we also instruct ChatGPT to omit case reports related to animal diseases, as they typically bear less relevance to our objective of focusing on human clinical dialogues.

**MTS-Dialog** MTS-Dialog is constituted by comprehensive and exhaustive clinical notes. Our approach involved deploying the GPT3.5-turbo model for dialogue generation. However, due to the API’s stringent maximum token restriction, incorporating the entire dialogue into a single prompt proved impracticable. Consequently, we implemented a strategy that involved segmenting a clinical note into several sections according to the traditional SOAP format <sup>3</sup>. We used each section header to construct a distinct prompt with the corresponding content in the note, thereby aiding the model in generating individual chats for every section. We add a corresponding postprocessing step for MTS-Dialog with Combine prompt on Table 9, where we concatenated all the small chats from different sections to create a comprehensive dialogue.

#### 2.2.2 Main dialogue generation loop

**Planning module** Typically, a physician’s diagnostic process adheres to a logical sequence, which may be outlined as follows (First et al., 2013; Johnson, 2003; Tsichlis et al., 2021): 1) Eliciting symptoms, such as chest pain, 2) Inquiring about the duration of these symptoms, 3) Obtaining medical history, including personal and familial records, 4) Conducting diagnostic tests, 5) Reaching a conclusion and prescribing appropriate medication. Thus, an effective dialogue dataset should accurately reflect the logical sequence of real-world interactions between doctors and patients. Therefore, before generating dialogues, it is crucial to ensure that the model follows such logic. However, we found models often tend to overlook crucial information or create hallucinations, or often messily skip content that should logically be in the first half of the dialogue and go to generating first with content that should logically appear later.

To circumvent these issues, we first extract clinical domain-specific keywords using CUI (Clinical Uniform Identifier) from the QuickUMLS<sup>4</sup> and require the LLM to build dialogues around these keywords exclusively, where we design the prompt in

<sup>2</sup>For the issue of cost and rate limit, when the paper was submitted, we only have all 167k ChatGPT synthetic data and only generated 10k synthetic data for GPT4 and NoteChat. For a fair comparison, all experiments and statistics in this paper are based on the same 10k PMC-Patient notes. But we will generate and release all 167k data in the future.

<sup>3</sup>SOAP structure details can be found in the Appendix A.1.

<sup>4</sup><https://github.com/Georgetown-IR-Lab/QuickUMLS>

Table 6 with the list of keywords to help the LLM generates the dialogue draft. Note that the output of the Planning module is not this draft but a corresponding checklist, where each CUI in the checklist is extracted in order from the draft. So we actually use LLM’s commonsense to plan the logical order in which CUI should appear in the checklist for use by the roleplay module. Then, the planning module will accompany the entire roleplay module. That is, every time the roleplay module completes a new round of dialogue generation, the planning module will count the newly added CUIs in the dialogue and remove them from the checklist. Therefore, the Planning module not only assumes the responsibility for the correct correlation of the facts but also helps the entire conversation narrow in a more definite direction until the end.

**Roleplay module** The dialogue draft we generated in the Planning module is not high-quality dialogue data. Previous work (Yunxiang et al., 2023) shows that dialogues generated by a single LLM often have problems with language diversity and role homogeneity. Therefore, in order to generate more satisfying dialogues, we use the checklist in the Planning module to generate dialogues round by round, using two LLMs to play the roles of patients and doctors, respectively. This strategy enables us to use distinct prompts so that the doctor’s responses appear more professional and the patient’s dialogue sounds more normal. Furthermore, we can control the direction of each dialogue round by modifying the prompts. More specifically, we determine the keywords covered in each round based on the current checklist, allowing doctor-LLM to guide the dialogue further based on historical dialogues to incorporate keywords. We then let the Planning module update the checklist. This method ensures that the dialogue progression remains logical. Subsequently, we let the patient-LLM respond to the doctor as colloquial a manner as possible, ensuring the patient’s utterance lay language style. All prompts can be found in Table 7.

**Polish module** Although the two modules of Planning and Roleplay bring NoteChat more fine-grained control over LLM, restoring doctor-patient dialogue from clinical notes requires LLM to balance several challenged requirements, including the planning of key information in the clinical note, a reasonable imagination for additional information not occurring in the note, the language style

	NoteChat	ChatGPT	GPT4
total #dial.	10k	10k	10k
avg # in a dialogue			
utterance	20.1	20.5	17.4
word	485	352	390
medical.	53.9	44.5	51.2
avg # of words in an utterance			
doctor	42.2	25.1	33.6
patient	14.0	11.7	9.39
avg medical term density %			
doctor	15.0	15.0	16.9
patient	14.3	13.4	13.0

Table 2: Statistics of three synthetic doctor-patient dialogue datasets conditioned on PMC-Patient notes. In the table, we bifurcated the dialogue into two constituent segments: one representing the doctor and the other the patient, for which we separately computed their corresponding scores. We computed the average count of words in both the doctor and patient utterances across each dialogue in the triad of datasets. Additionally, we derived a metric, indicated as medical term density, which signifies the proportion of the count of Clinical Uniform Identifier (CUI) codes encapsulated within each utterance of doctor and patient to the overall count of words.

characteristics of different roles, and the authenticity after combining everything into one complete dialogue. In the previous Roleplay module, LLM will promote new dialogues based on historical dialogues. Inspired by recent work of rethinking and reranking (Gabriel et al., 2021; Cobbe et al., 2021; Ravaut et al., 2022; Jiang et al., 2022; Shinn et al., 2023), we added the Polish module after the Roleplay module to give LLM another chance for self-reflection and correction. In addition, we invited human experts to summarize 10 special rules based on the preliminary results of NoteChat to help our framework align with experts’ preferences: 1) Make the conversation as colloquial as possible, 2) Increase the number of rounds of interaction, 3) Professional terms and vocabulary should come from the doctors, and patients should be more colloquial, 4) Basic symptoms and medical history should come from the patient, not the doctor, 5) The questions asked by the doctor should be around the case (to avoid hallucination), 6) Physician inquiries should be logical, 7) If there are multiple consultation records, you can split a conversation into multiple ones and then link them with transfer words (e.g., a few days later), 8) Range of rounds of interaction, 9) Must contain the given keywords, 10) Do not generate duplicate information.

Specifically, we converted these requirements

into a part of the Polish prompt one by one. Then the LLM will polish the existing dialogue according to the prompt in Table 8. We found that repeating the polish module multiple times can further improve the quality of the final synthetic dialogue<sup>5</sup>.

### 3 Automatic Evaluation

MTS-Dialog provides the human-annotated ground truth conversation data for every clinical note, but the PMC-Patient dataset only has case reports. So we use intrinsic evaluation for MTS-Dialog synthetic data but extrinsic and human evaluation for PMC-Patient synthetic data.

#### 3.1 Intrinsic Evaluation

We measure this task of note-to-conversation from four aspects of the MTS-Dialog dataset.

**Similarity** We use ROUGE-F1 scores (Lin, 2004) to measure the similarity of the generated conversation and the references.

**Factuality** We use QuickUMLS (Soldaini, 2016) to extract medical concepts from model-generated summaries and ground truth summaries to get two corresponding concept lists. We ask GPT-4 to filter out noise from the concept lists with prompt “Verify whether the concept is important. Only return yes or no”. This process results in two refined lists:  $C_{hypo}$ , and  $C_{ref}$ . We then instruct the GPT-4 to “Verify whether the conversation/note includes or implies the concept I gave you. Only return yes or no.”. Finally, we calculate the Precision, Recall, and F1 scores of  $ref \leftrightarrow hypo$  as our factuality metrics according to GPT4’s decision of “yes” or “no” for every concept in the concept lists.

**Extractiveness** We calculate the ROUGE-F1 of  $src \rightarrow hypo$  (clinical note to model-generated dialogue) as our extractiveness metrics to demonstrate how much information in dialogue is extracted from the clinical note. For AI, a shortcut to improve Factuality is to improve Extractiveness. However, recent work shows increasing the factuality by this way might not be ideal in many scenarios (Ladhak et al., 2022; Goyal et al., 2022).

**Diversity** We use Self-BLEU (SBLEU) (Zhu et al., 2018) to evaluate the diversity of the generated conversation for the patient utterances, doctor utterances, and overall.

<sup>5</sup>After balancing the time, cost, and final performance, we

Similarity	ROUGE1	ROUGE2	ROUGESum
ChatGPT	48.56	16.74	46.36
GPT4	53.29	20.20	50.81
NoteChat	56.48	19.74	53.41
Factuality	Concept-P	Concept-R	Concept-F1
ChatGPT	67.54	35.75	46.23
GPT4	71.46	45.69	55.17
NoteChat	48.23	51.23	49.68
Extractiveness	src->hypo R1	src->hypo R2	src->hypo R-L
ChatGPT	43.73	19.72	40.54
GPT4	52.70	25.70	49.63
NoteChat	37.24	20.83	36.04
Human	35.29	14.38	32.89
Diversity	all-sbleu ↓	doctor-sbleu ↓	patient-sbleu ↓
ChatGPT	0.017	0.006	0.017
GPT4	0.019	0.009	0.019
NoteChat	0.014	0.007	0.014

Table 3: Intrinsic eval results on MTS-dialog<sup>6</sup>.

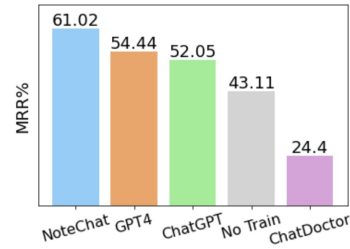


Figure 2: Extrinsic eval results for Medical Chatbot task. LLaMA-7B is fine-tuned on different PMC-Patient synthetic conversations and the ChatDoctor dataset, and then we use MTS-dialog as the evaluation dataset. NoteChat has the highest score, indicating the most preferred by GPT4.

#### 3.2 Extrinsic Evaluation

We used the PMC-Patient synthetic dialogues generated by ChatGPT, GPT4, and NoteChat to fine-tune the LLaMA-7B<sup>7</sup>, where we only use doctor utterances as training labels. Then we evaluate the fine-tuned medical chatbots on the ground truth dialogues from MTS-Dialog.

Recent work shows a higher human evaluation correlation for GPT-4 eval than traditional metrics (Liu et al., 2023b; Gao et al., 2023; Fu et al., 2023; Zheng et al., 2023), so we also use GPT4 preference as measurements to evaluate chatbots’ response quality. Specifically, we instruct GPT4 to give preference ranking<sup>8</sup> based on the conversation history and the real response. We report the Mean Reciprocal Rank (MRR) (Radev et al., 2002) of each model’s final ranking in Figure 2. Generally, a higher MRR implies that evaluators have more preference over an approach.

set the number of repetitions to 2 in our experiments

<sup>6</sup>All experiments are done under the zero-shot setting.

<sup>7</sup><https://huggingface.co/decapoda-research/llama-7b-hf>

<sup>8</sup>Prompts can be found in Appendix 10.

### 3.3 Automatic Evaluation Results

The **intrinsic evaluation** results, as illustrated in Table 3, show that the overall similarity of the conversations generated by NoteChat and Human (MTS-dialog ground truth) is higher than that of GPT4 and ChatGPT baselines. GPT4 is ahead of NoteChat and ChatGPT in both factuality and extractiveness metrics. NoteChat is superior to ChatGPT in factuality but has a lower and closer to human extractiveness score. In Section 4.4, we will discuss the impact of the different factuality and extractiveness scores of the three methods on human expert preferences on our task. Finally, we found that the diversity of NoteChat, especially for patient utterances, is significantly better than the baselines. The **extrinsic evaluation** results are illustrated in Figure 2. In this experiment, LLaMA-7B is first fine-tuned on different PMC-Patient synthetic conversations (or the ChatDoctor dataset). Then we use MTS-dialog as the evaluation dataset. NoteChat-based LLaMA has the highest score, indicating the most preferred by GPT4 when generating MTS-dialog doctor utterances. It is worth noting that this evaluation is also a kind of transfer learning because LLaMA is only trained on different models’ PMC-Patient synthetic dialogue and then tested its zero-shot performance on human-labeled dialogue in MTS-dialog. Therefore, the results of this extrinsic evaluation show from the side that the distribution of NoteChat-generated dialogue is closest to the distribution of human annotations.

## 4 Human Evaluation

In order to assess the quality of synthetic conversations using different methods (ChatGPT, GPT-4, NoteChat), we conducted a human evaluation from both the crowd’s and the experts’ perspectives.

### 4.1 Human Evaluation Settings

The goal of **expert evaluation** is to have human domain experts evaluate whether these machine-generated conversations are comparable to real doctor-patient encounter conversations from a professional perspective (e.g. medical commonsense, knowledge, logic). To do so, we recruited 5 medical practitioners<sup>9</sup>, and their tasks are to read clinical notes and provide qualitative feedback on if

<sup>9</sup>Four licensed physicians and one medical student with hospital internship experience. These experts were not involved in the research, only the human evtion.

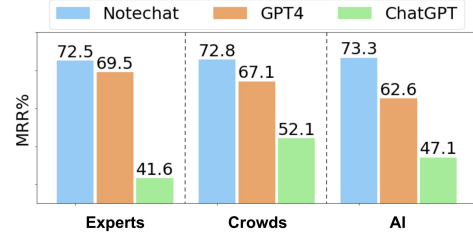


Figure 3: Human&GPT preference.

the machine-generated dialogues can be defined as high-quality doctor-patient interactions in terms of factual accuracy and logical coherence; if not, how should they be improved?

The goal of **crowd evaluation** is to allow the general public to provide ratings for different synthetic conversations based on their lived experience. Since the crowds do not have professional medical knowledge, participants will first read the clinical notes and medical expert annotated conversations as references for high-quality data and then rank different machine-generated conversations for quantitative measurement of their preference. We recruit 10 human evaluators to participate in our crowd evaluation.<sup>10</sup>

### 4.2 Human Evaluation Measurements

We mainly use human preference as measurements to evaluate synthetic conversation quality. Specifically, the participants are provided with the following instructions “*The following three conversations are generated by AI based on this clinical note. Please rank them according to the quality you think, from high to low.*”. We collect the preference ranking from experts, crowds, and GPT4. We report the Mean Reciprocal Rank (MRR) of each model’s final ranking in Figure 3.

### 4.3 Human Evaluation Outcome

All the preference feedback from experts, crowds, and AI are shown in Figure 3. First of all, all preference results are consistent, which is NoteChat > GPT4 > ChatGPT. In the results of experts’ preference, the MRR scores of NoteChat and GPT4 are very close, indicating that from the perspective of experts, the quality of the dialogue data of the two is comparable. In the results of crowds and AI preference, NoteChat is significantly more preferred than GPT4. Finally, NoteChat and GPT4 are preferred over ChatGPT in all three human evls.

<sup>10</sup>All the evaluators have bachelor’s degrees but do not have any medical education background.

1	<p><b>Some information in the note does not come from the conversation, so it is not necessary to mention it in the dialogue, such as dosage information.</b></p> <p><i>Her medications were tab. vildagliptin/metformin (50/1000) (mg) BID, tab. amlodipine/valsartan (5/160) (mg) once daily, and tab. levothyroxine 75 mcg once daily ...</i></p> <p>(GPT4) <b>Patient:</b> I take vildagliptin/metformin 50/1000 mg twice a day, amlodipine/valsartan 5/160 mg once daily, and levothyroxine 75 mcg once daily.</p> <p>(NoteChat) <b>Patient:</b> I have arterial hypertension, diabetes, and hypothyroidism. I take vildagliptin/metformin, amlodipine/valsartan, and levothyroxine.</p>
2	<p><b>Patients sometimes act too professionally.</b></p> <p><i>Two months postoperatively, the patient came to our institution and underwent MRI for reevaluation of the cervical spine in three different positions: neutral, flexion, and extension. While neutral MRI did not show any significant stenosis, flexion and extension MRIs were significant for ...</i></p> <p>(GPT4) <b>Patient:</b> The MRI was done in three different positions: neutral, flexion, and extension. The neutral position didn't show any significant stenosis, but the flexion and extension MRIs showed cervical compression due to spondylosis and paraspinal muscle compression.</p> <p>(NoteChat) <b>Doctor:</b> Thank you. Let me review your medical history records. (Doctor reviews the medical history records) Based on your medical history records, it seems that the previous surgery did not completely resolve your symptoms. We will need to perform a reevaluation of your cervical spine to determine the cause of your symptoms. We will need to perform an MRI in three different positions: neutral, flexion, and extension. The MRI will help us to see if there is any compression in your cervical spine that could be causing your symptoms. <b>Patient:</b> Okay, I understand.</p>
3	<p><b>Highly generalized sentences in notes tend to contain a lot of important information, but in dialogues, they are usually collected by multiple rounds.</b></p> <p><i>Her diagnosis at admission included a gallstone, acute attack of chronic cholecystitis, biliary pancreatitis, septic shock, and multiple organ dysfunction syndrome ...</i></p> <p>(GPT4) <b>Doctor:</b> So, I see from your chart that you were admitted with a gallstone, acute attack of chronic cholecystitis, biliary pancreatitis, septic shock, and multiple organ dysfunction syndrome. Can you tell me more about your symptoms? <b>Patient:</b> Well, I had severe abdominal pain and vomiting.</p> <p>(NoteChat) <b>Doctor:</b> Great, I appreciate your cooperation. Now, could you please provide me with your medical history records? <b>Patient:</b> Sure, here they are. <b>Doctor:</b> Thank you. Based on your medical history records, it seems that you were admitted with a diagnosis of a gallstone, an acute attack of chronic cholecystitis, biliary pancreatitis, septic shock, and multiple organ dysfunction syndrome. Is that correct?</p>
4	<p><b>The logic of the different roles in the dialogue. It is necessary to distinguish which information will be advanced by the doctor and which information will be provided by the patient.</b></p> <p><i>A 66-year-old woman presented at the emergency department of our hospital complaining of sudden onset dizziness and fatigue over the past thirty minutes. Before her arrival, she was at home relaxing and not engaged in any physical activity ...</i></p> <p>(GPT4) <b>Patient:</b> Hi, Doctor. About thirty minutes ago, I suddenly started feeling dizzy and fatigued while I was at home relaxing. <b>I wasn't even engaged in any physical activity.</b> <b>Doctor:</b> I see. Can you please tell me about your past medical history?</p> <p>(NoteChat) <b>Doctor:</b> I see. Can you tell me more about your symptoms? <b>Were you relaxing or engaged in any physical activity before the symptoms occurred?</b> <b>Patient:</b> I was just at home relaxing before the symptoms started.</p>
5	<p><b>NoteChat handles the beginning of the conversation more naturally. So experts prefer NoteChat when other aspects of quality are similar.</b></p> <p><i>A 71-year-old man with a history of atrial fibrillation, acute appendicitis, and early esophageal cancer treated with endoscopic submucosal dissection underwent distal pancreatectomy ...</i></p> <p>(GPT4) <b>Doctor:</b> Hello, I see from your history that you have atrial fibrillation, acute appendicitis, and early esophageal cancer. You were treated with endoscopic submucosal dissection for esophageal cancer, correct? <b>Patient:</b> Yes, that's right.</p> <p>(NoteChat) <b>Patient:</b> Doctor, hello. I have an irregular posterior wall and a submucosal tumor in the anterior wall of my gastric antrum. <b>Doctor:</b> Can you give me your medical records? <b>Patient:</b> Here you go.</p>

Table 4: Expert evaluation case study <sup>11</sup>.

#### 4.4 Heuristic Evaluation with Experts

We interview 5 medical practitioners:

##### Q1) What are the shortcomings of AI syn-

<sup>11</sup>Due to the obvious gap in factuality of ChatGPT, our cases focus on the difference between NoteChat and GPT4.

**thetic conversation compared with real-world doctor-patient encounter conversation?** Experts think that synthetic conversations cover too much information from the clinical note compared to real-world conversations, because some factual information is not provided to note through conversation (such as lab test results). For example, in Table 4 Example 1, the detailed dosage information will be not in the conversation. In Table 4 Example 2, the patient acts too professionally. In the answer, a lot of medical knowledge that doctors will know is described by the patient.

**Q2) What is the difference between ChatGPT, GPT4, and NoteChat synthetic conversations?** All medical practitioners believe that GPT4 and NoteChat lead ChatGPT in terms of factuality. Since our NoteChat is based upon ChatGPT, this human observation shows that our modules successfully inject medical concept knowledge to improve the factuality level from ChatGPT to the level of GPT4. So as shown in Figure 3, ChatGPT is ranked last in most cases.

Regarding the comparison between NoteChat and GPT4, medical practitioners actually believe that the data quality of the two AI-synthetic conversations is generally at a similar level, which aligns with their expert preference in Figure 3. However, the results of preference ranking (as shown in Figure 3) show that when the two kinds of data appear at the same time, crowds and AI much prefer NoteChat. We further conducted a heuristic evaluation to explore the reason here as well as the deficiency of NoteChat and GPT4 synthetic conversations and potential improvement.

First of all, GPT4 prefers to directly copy the information in the note to meet the requirements of factuality, but this will make the conversation unreal. In Table 4 Example 2, the information is highly summarized and put together on the note, but it is unnatural for the same content to appear directly in the dialogue. Compared with the utterance generated by GPT4, a better way is to use multiple rounds of conversation to obtain information one by one. This is a problem common to all AIs in this paper, but GPT4's problem is most obvious.

Second, in reality, doctors are expected to not only answer questions but also advance the discussion by asking professional questions. We observe that the doctor in NoteChat is more likely to advance the conversation compared to the doctor in GPT4 due to our Roleplay module. So we need

to more precisely control the LLMs to distinguish "what information is spoken by the doctor and what information is spoken by the patient." In the Table 4 Example 3, GPT4 directly allowed patients to continue to speculate after answering their own symptoms and further denied physical activities as a possible cause of symptoms. In the same example, we use prompt *"You don't know the patient's medical history and symptoms. You should ask or lead the patient to tell you the symptoms and his medical history and you don't have any information about his medical history and symptoms. All the information of medical history, symptoms, medication history and vaccination history should be told by the patient. You can tell the patient the test results, vital signs and some conclusions."* to instruct the NoteChat Roleplay module (doctor) for this requirement. NoteChat successfully allows doctor and patient to play their roles well and cooperate to advance this layer of logic.

Finally, the beginning style of the generated dialogues should be consistent with real conversations. Physicians emphasize that the initiation of a dialogue is of significant importance, as the initial inquiry often determines subsequent testing and establishes the treatment plan to be pursued. Common symptoms and medical histories should ideally be expressed by the patient or elicited by the doctor through guided questioning. However, in dialogues generated by GPT4, this information is often presented by the doctor, which deviates from conventional clinical practices. For instance, in Example 5 of Table 4, the patient should state symptoms and medical history as the doctor typically does not possess prior knowledge of the patient's history.

## 5 Related Work

A task closely related to our work, but with an inverse direction, is the automatic generation of clinical notes from doctor-patient conversations (Schloss and Konam, 2020; Song et al., 2020; Krishna et al., 2021; Yim and Yetisgen-Yildiz, 2021; Su et al., 2022). This process typically follows a common pipeline. First, clinical conversations are recorded and subsequently transcribed. Next, the utterances are classified into several medical sections, and clusters of utterances containing medically relevant information for each section are predicted from the transcript. Finally, a section-conditioned summarization model is employed to generate concise summaries for each cluster of ut-

terances associated with their respective sections. Recently, the MEDIQA-Chat 2023<sup>12</sup> introduce tasks of both directions (Dialogue2Note Summarization and Note2Dialogue Generation). However, their dataset is either private or limited to less than 2k examples. To the best of our knowledge, we are the first to introduce a large-scale publicly available doctor-patient conversation dataset in English for the purpose of clinical note generation. In addition, our extrinsic eval shows that the NoteChat can be used as auxiliary data for Conversation2Note or Note2Conversation tasks and can also be used as a synthetic medical dialogue dataset alone to help medical LLMs' instruction fine-tuning (Zhang et al., 2023; Li et al., 2023b; Wang et al., 2023; Liu et al., 2023a; Xiong et al., 2023; Zeng et al., 2020).

Our work builds upon the recent advances in deploying two LLMs as cooperative agents (Panait and Luke, 2005) for multi-round conversation generation. In particular, NoteChat is inspired by CAMEL (Li et al., 2023a), which assigns roles to two LLMs (e.g. student and teacher) in order to facilitate conversation between the two agents for a particular task (e.g. teaching). Similar to CAMEL's findings, we found that roleplay by itself may hallucinate or generate fake replies that repeat most of the previous utterances. To solve this issue, we proposed a novel Planning module to ground agents to certain keywords. Cho et al. (2023) also addresses the challenges of using LLM to craft a dialogue dataset with specified personas. They emphasize the importance of grounding and context in conversation generation. Similarly, NoteChat relies on structured clinical notes segmented using the SOAP format to provide context for our dialogue synthesis to diagnose a patient. However, their work is limited to generating open-domain dialogue, while we focus on task-oriented dialogue.

## 6 Conclusion

In this study, we present *NoteChat*, a cooperative multi-agent framework leveraging LLMs for generating synthetic doctor-patient conversations conditioned on clinical notes. NoteChat consists of Planning, Roleplay, and Polish modules. Extensive evaluations demonstrate that NoteChat facilitates high-quality synthetic doctor-patient conversations, underscoring the untapped potential of LLMs in healthcare and offering promising avenues for the intersection of AI and healthcare.

<sup>12</sup><https://sites.google.com/view/mediqa2023>

## 7 Limitations and Ethical Considerations

This study offers valuable insights, but with a few limitations, we would like to note.

**Privacy Implications.** LLMs can present privacy concerns in using clinical notes to generate patient-doctor conversation, potentially violating HIPAA regulations. However, in this study, all experiments were sourced from publicly available real patient data collected from research articles with at least CC BY-NC-SA license. We also present an approach for generating synthetic conversations from case reports in the PubMed Central repository.

**Biases.** LLMs trained on vast amounts of text data may inadvertently capture and reproduce biases present in the data. For example, they may prefer certain questions related to Metformin or link particular health conditions to specific populations. Thus the doctor bot trained from our synthetic data may perpetuate incorrect information or provide inaccurate answers. Moreover, the case reports used to generate synthetic conversations usually focus on unusual observations and rare conditions. Thus the doctor bot may hallucinate or overtreat patients with common diseases.

**Broader Impacts.** We have performed a preliminary study to generate synthetic conversation from case reports within research articles indexed from January of the year 2002 to July of the year 2022 by PubMed Central. The credibility of these case reports is ensured as they are peer-reviewed and published in academic journals. Moreover, the type of disease is diverse as they are sourced from various hospital departments and are not limited to intensive care units (such as MIMIC). Thus, models trained using our synthetic data may inherit the benefits from these characteristics.

## References

- Asma Ben Abacha, Wen-wai Yim, Yadan Fan, and Thomas Lin. 2023. An empirical study of clinical note generation from doctor-patient encounters. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2283–2294.
- Asma Ben Abacha and Pierre Zweigenbaum. 2015. Means: A medical question-answering system combining nlp techniques and semantic web technologies. *Information processing & management*, 51(5):570–594.
- George J Annas. 2003. Hipaa regulations: a new era of medical-record privacy? *New England Journal of Medicine*, 348:1486.
- Asma Ben Abacha, Wen-wai Yim, Griffin Adams, Neal Snider, and Meliha Yetisgen. 2023. Overview of the mediqua-chat 2023 shared tasks on the summarization and generation of doctor-patient conversations. In *ACL-ClinicalNLP 2023*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. *Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality*.
- Won Ik Cho, Yoon Kyung Lee, Seoyeon Bae, Ji-Hwan Kim, Sangah Nancy Park, Moosung Kim, Sowon Hahn, and Nam Soo Kim. 2023. When crowd meets persona: Creating a large-scale open-domain persona dialogue corpus. *ArXiv*, abs/2304.00350.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. *Training verifiers to solve math word problems*. *ArXiv preprint*, abs/2110.14168.
- Haixing Dai, Zheng Liu, Wenxiong Liao, Xiaoke Huang, Zihao Wu, Lin Zhao, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, Hongmin Cai, Quanzheng Li, Ding-gang Shen, Tianming Liu, and Xiang Li. 2023. Chataug: Leveraging chatgpt for text data augmentation. *ArXiv*, abs/2302.13007.
- Paul Drew, John Chatwin, and Sarah Collins. 2001. Conversation analysis: a method for research into interactions between patients and health-care professionals. *Health Expectations*, 4(1):58–70.
- Gunther Eysenbach et al. 2023. The role of chatgpt, generative language models, and artificial intelligence in medical education: A conversation with chatgpt and a call for papers. *JMIR Medical Education*, 9(1):e46885.
- Lewis R First, Humayun J Chaudhry, and Donald E Melnick. 2013. Quality, cost, and value of clinical skills assessment. *New England Journal of Medicine*, 368(10):963–964.

- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Saadia Gabriel, Antoine Bosselut, Jeff Da, Ari Holtzman, Jan Buys, Kyle Lo, Asli Celikyilmaz, and Yejin Choi. 2021. [Discourse understanding and factual consistency in abstractive summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 435–447, Online. Association for Computational Linguistics.
- Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. 2023. Human-like summarization evaluation with chatgpt. *arXiv preprint arXiv:2304.02554*.
- Aidan Gilson, Conrad W Safranek, Thomas Huang, Vimig Socrates, Ling Chi, Richard Andrew Taylor, David Chartash, et al. 2023. How does chatgpt perform on the united states medical licensing examination? the implications of large language models for medical education and knowledge assessment. *JMIR Medical Education*, 9(1):e45312.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Dongfu Jiang, Bill Yuchen Lin, and Xiang Ren. 2022. Pairreranker: Pairwise reranking for natural language generation. *arXiv preprint arXiv:2212.10555*.
- Hillary Johnson. 2003. A critical review of standardized patient examinations as part of the usmle. *AMA Journal of Ethics*, 5(12):426–429.
- Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, 17:1–9.
- Kundan Krishna, Sopan Khosla, Jeffrey P Bigham, and Zachary C Lipton. 2020. Generating soap notes from doctor-patient conversations using modular summarization techniques. *arXiv preprint arXiv:2005.01795*.
- Kundan Krishna, Amy Pavel, Benjamin Schloss, Jeffrey P Bigham, and Zachary C Lipton. 2021. Extracting structured data from physician-patient conversations by predicting noteworthy utterances. In *Explainable AI in Healthcare and Medicine*, pages 155–169. Springer.
- Faisal Ladhak, Esin Durmus, He He, Claire Cardie, and Kathleen McKeown. 2022. [Faithful or extractive? on mitigating the faithfulness-abstractiveness trade-off in abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1410–1421, Dublin, Ireland. Association for Computational Linguistics.
- G. Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023a. Camel: Communicative agents for "mind" exploration of large scale language model society. *ArXiv*, abs/2303.17760.
- Jianquan Li, Xidong Wang, Xiangbo Wu, Zhiyi Zhang, Xiaolong Xu, Jie Fu, Prayag Tiwari, Xiang Wan, and Benyou Wang. 2023b. [Huatu-26m, a large-scale chinese medical qa dataset](#).
- Zongjie Li, Chaozheng Wang, Pingchuan Ma, Chaowei Liu, Shuai Wang, Daoyuan Wu, and Cuiyun Gao. 2023c. On the feasibility of specialized ability stealing for large language code models. *arXiv preprint arXiv:2303.03012*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Hongcheng Liu, Yusheng Liao, Yutong Meng, Yu Wang, and Yanfeng Wang. 2023a. Medicalgpt-zh. <https://github.com/MediaBrain-SJTU/MedicalGPT-zh>.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023b. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. [The flan collection: Designing data and methods for effective instruction tuning](#).
- Douglas W Maynard and John Heritage. 2005. Conversation analysis, doctor–patient interaction and medical communication. *Medical education*, 39(4):428–435.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Liviu Panait and Sean Luke. 2005. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11:387–434.
- V Podder, V Lew, and S Ghassemzadeh. 2021. Soap notes.[updated 2021 sep 2]. *StatPearls [Internet]*. StatPearls Publishing. Available from: <https://www.ncbi.nlm.nih.gov/books/NBK482263>.
- Dragomir R Radev, Hong Qi, Harris Wu, and Weiguo Fan. 2002. Evaluating web-based question answering systems. In *LREC*. Citeseer.
- Mathieu Ravaut, Shafiq Joty, and Nancy Chen. 2022. [SummaReranker: A multi-task mixture-of-experts re-ranking framework for abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4504–4524, Dublin, Ireland. Association for Computational Linguistics.

793	Thomas C Rindfleisch. 1997. Privacy, information technology, and health care. <i>Communications of the ACM</i> , 40(8):92–100.	848
794		849
795		
796	Benjamin Schloss and Sandeep Konam. 2020. Towards an automated soap note: classifying utterances from medical conversations. In <i>Machine Learning for Healthcare Conference</i> , pages 610–631. PMLR.	850
797		851
798		852
799		853
800	Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. <a href="#">Reflexion: Language agents with verbal reinforcement learning</a> .	854
801		855
802		856
803		857
804	Luca Soldaini. 2016. Quickumls: a fast, unsupervised approach for medical concept extraction.	858
805		859
806		860
807	Yan Song, Yuanhe Tian, Nan Wang, and Fei Xia. 2020. Summarizing medical conversations via identifying important utterances. In <i>Proceedings of the 28th International Conference on Computational Linguistics</i> , pages 717–729, Barcelona, Spain (Online).	861
808		862
809		863
810		864
811	Jing Su, Longxiang Zhang, Hamidreza Hassanzadeh, and Thomas Schaaf. 2022. Extract and abstract with bart for clinical notes from doctor-patient conversations. In <i>Interspeech</i> .	865
812		866
813		867
814		868
815	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. <a href="https://github.com/tatsu-lab/stanford_alpaca">https://github.com/tatsu-lab/stanford_alpaca</a> .	869
816		870
817		871
818		872
819		873
820	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	874
821		875
822		876
823		
824		
825		
826		
827	Jason T Tsichlis, Andrew M Del Re, and J Bryan Carmody. 2021. The past, present, and future of the united states medical licensing examination step 2 clinical skills examination. <i>Cureus</i> , 13(8).	877
828		878
829		879
830		880
831	Haochun Wang, Chi Liu, Nuwa Xi, Zewen Qiang, Sendong Zhao, Bing Qin, and Ting Liu. 2023. Huatuo: Tuning llama model with chinese medical knowledge. <i>arXiv preprint arXiv:2304.06975</i> .	
832		
833		
834		
835	Honglin Xiong, Sheng Wang, Yitao Zhu, Zihao Zhao, Yuxiao Liu, Qian Wang, and Dinggang Shen. 2023. Doctorglm: Fine-tuning your chinese doctor is not a herculean task. <i>arXiv preprint arXiv:2304.01097</i> .	
836		
837		
838		
839	Wen-wai Yim, Yujuan Fu, Asma Ben Abacha, Neal Snider, Thomas Lin, and Meliha Yetisgen. 2023. Acibench: a novel ambient clinical intelligence dataset for benchmarking automatic visit note generation. <i>Submitted to Nature Scientific Data</i> .	
840		
841		
842		
843		
844	Wen-wai Yim and Meliha Yetisgen-Yildiz. 2021. Towards automating medical scribing: Clinic visit dialogue2note sentence alignment and snippet summarization. In <i>Proceedings of the Second Workshop on</i>	
845		
846		
847		
	<i>Natural Language Processing for Medical Conversations</i> , pages 10–20.	
	Li Yunxiang, Li Zihan, Zhang Kai, Dan Ruilong, and Zhang You. 2023. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. <i>arXiv preprint arXiv:2303.14070</i> .	
	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 <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 9241–9250.	
	Hongbo Zhang, Junying Chen, Feng Jiang, Fei Yu, Zhihong Chen, Jianquan Li, Guiming Chen, Xiangbo Wu, Zhiyi Zhang, Qingying Xiao, Xiang Wan, Benyou Wang, and Haizhou Li. 2023. Huatuoqpt, towards taming language models to be a doctor. <i>arXiv preprint arXiv:2305.15075</i> .	
	Zhengyun Zhao, Qiao Jin, Fangyuan Chen, Tuorui Peng, and Sheng Yu. 2023. <a href="#">Pmc-patients: A large-scale dataset of patient summaries and relations for benchmarking retrieval-based clinical decision support systems</a> .	
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. <a href="#">Judging llm-as-a-judge with mt-bench and chatbot arena</a> .	
	Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Texus: A benchmarking platform for text generation models. <i>SIGIR</i> .	

## A Example Appendix

### A.1 SOAP Structure

The SOAP (Subjective, Objective, Assessment, and Plan) structure is commonly used by providers (Podder et al., 2021).

1. Chief Complaint includes a brief description of a patient’s conditions and the reasons for the visit.
2. The Subjective section is a detailed report of the patient’s current conditions, such as source, onset, and duration of symptoms, mainly based on the patient’s self-report. This section usually includes a history of present illness and symptoms, current medications, and allergies.
3. The Objective section documents the results of physical exam findings, laboratory data, vital signs, and descriptions of imaging results.
4. The Assessment section typically contains medical diagnoses and reasons that lead to medical diagnoses. The assessment is typically based on the content from the chief complaint, and the subjective and objective sections.
5. The Plan section addresses treatment plans based on the assessment.

### A.2 Prompts for ChatGPT&GPT4

We use the following prompts to instruct ChatGPT and GPT4 to generate the synthetic doctor-patient dialogue based on the provided clinical note.

*Generate the conversation between doctor and patient. But for some cases, if the patient eventually dies (according to the clinical note), you can add the patient’s family at the end of the conversation to make it more reasonable. The conversation should include all the information in the following note, especially paying attention to those numbers and medical concepts. The conversation can be more colloquial. When the doctor is speaking, the patient can have many modal particles (e.g. hmm, yes, okay) to increase interaction. All the numbers and medical concepts that appear in the note should be mentioned by the doctor. Professional medical terms and numbers should more likely occur in the doctor’s utterances but not in the patient’s answer. The doctor may describe and explain professional judgment to the patient and instruct the patient on follow-up requirements but not ask questions that require professional medical knowledge to an-*

*swer. The patient’s answer should be succinct and accurate in a colloquial lay language style.*

### A.3 Experimental Settings

In our study on generating conversation datasets using ChatGPT and GPT-4, we adopted a temperature setting of 0.7. This setting was consistently applied across our methodologies. For each round of dialogue, we set the max tokens for doctor role-play as 200 tokens and the patient role-play as 100 tokens. For the intrinsic evaluation phase, we selected a subset of 20 data points from the MT-Dialog dataset and randomly chose 100 datasets from the pmc dataset for testing. In terms of external evaluation, we selected three random data points from each model’s output on the pmc dataset to use as few-shot examples. These were inputted into GPT-4, which then generated dialogues from clinical notes or clinical notes from conversations based 20 data sets from the MT-Dialog dataset. During the external chatbot evaluation, we used 10k datasets generated by ChatGPT, GPT-4, and NoteChat-ChatGPT to fine-tune Llama-7b on two A100-40g gpus. During the fine-tuning process, we used DeepSpeed Zero-2 for training, with a learning rate of 1.5e-5, a batch size of 16, max tokens of 6144 and 3 training epochs.

### A.4 Color for Polish Promopt

We have used consistently different colors to indicate in the polish prompt, as shown in Table 8, which parts of our prompt have achieved these ten different functions.

1. **Yellow**: Make the conversation as colloquial as possible
2. **Orchid**: Increase the number of rounds of interaction
3. **Pink**: Professional terms and vocabulary should come from the doctors, and patients should be more colloquial
4. **Gray**: Basic symptoms and medical history should come from the patient, not the doctor
5. **BrickRed**: The questions asked by the doctor should be around the case (to avoid hallucination)
6. **SkyBlue**: Physician inquiries should be logical
7. **Emerald**: If there are multiple consultation records, you can split a conversation into multiple ones and then link them with transfer

words (e.g., a few days later)

8. **BurntOrange**: Range of rounds of interaction

9. **Thistle**: Must contain the given keywords

10. **Periwinkle**: Do not generate duplicate information

Note that there are some similar and repeated parts in the prompt, which are because we found that mentioning a certain point multiple times in different places in the prompt is more helpful for LLM to avoid certain problems.

Section	Subsection	Definition
Subjective	Chief Complaint	Patient’s primary motivation for the visit and type of visit
	Review of Systems	Patient’s report of system related health and symptoms
	Past Medical History	Patient’s reported diagnoses/conditions (when and what, excluding laboratory and imaging results and surgeries)
	Past Surgical History	Patient’s reported prior surgeries (what, when, where)
	Family Medical History	Conditions affecting patient’s close genetic relatives
	Social History	Patient’s alcohol, tobacco, and drug related behaviors
	Medications	Patient’s list of medications (not prescribed during visit)
	Allergies	Patient’s list of allergies (primarily medicinal)
	Miscellaneous	Patient’s clinically relevant social and other circumstances
Objective	Immunizations	Vaccination record (not frequently discussed)
	Laboratory and Imaging Results	Clinician’s discussion of laboratory/imaging results
Assessment	Assessment	Synthesis of reason for visit and pertinent diagnosis
Plan	Diagnostics & Appointments	Plan for future tests, appointments, or surgeries
	Prescriptions & Therapeutics	Plan for medications and therapeutics

Table 5: Details of the SOAP structure using in our CC and CCUser datasets.

<b>Planning Module</b>	<p>Apply the Doctor and Patient prompt to generate the beginning and lead the Doctor LLM to ask about the medical record. Continue to generate 20 to 40 utterances conversations between doctor and patient to ask or tell the patient regarding the case(you must follow up the history conversation). The conversations you generate must cover all the keywords I gave you. You cannot revise or eliminate any keywords and you cannot use synonyms of the keywords. Your conversation should also include all information. If it’s difficult to include all the information and key words, you can use the original sentences in the clinical note.</p> <p>The Clinical Note: <a href="#">Clinical Note</a></p> <p>The Key Words: <i>key<sub>1</sub>, key<sub>2</sub>,...</i></p> <p>Your conversations must include all the keywords I provided to you, and if it’s not possible to include them all, you can make slight modifications based on the original wording in the notes. You cannot revise or eliminate any key words and you cannot use synonyms of the keywords. Your conversation should also include all information. If it’s difficult to include all the information and key words, you can use the original sentences in the clinical note. Your generation must follow the logical sequence of a doctor’s inquiry. Your conversations must follow the logical sequence of a doctor’s inquiry. For example, the general logical order of the conversation is: first discussing symptoms, then discussing the medical history, followed by discussing testing and results, and finally discussing the conclusion and treatment options, etc. The doctor didn’t know any information of medical history or symptoms. This information should be told by the patient</p>

Table 6: Planning Module prompt.

<b>Doctor Prompt</b>	<p>Please role-play as a doctor and further generate questions or conclusion, or the test result(such as medication test result or vital signs) based on the above dialogue and clinical note(after mentioned examination, you have to know test results and vital signs so you shouldn't ask the patient about a test result or vital signs). Add 'Doctor:' before each round. Your question, answer or conclusion(tell the patient the test result) should be around the keywords (I gave you) corresponding to the clinical note(finally, the whole conversation should include all the keywords). the answer of your questions can be found on the clinical note. You cannot modify these key words or use synonyms. You need to ensure the treatment plan, medication, and dosage you give to the patient must also be totally consistent with the clinical note. Do not ask questions which answers cannot be found in the clinical note. You may describe and explain professional judgment to the patient and instruct the patient on follow-up requirements, but not ask questions that require professional medical knowledge to answer. The order of the questions you ask must match the order of the keywords I provided. If it's not possible to include them all, you can make slight modifications based on the original wording in the notes. If the history conversation has included the keywords, there is no need to include them again. The treatment plan and conclusions you provide must align completely with the clinical notes. Do not add treatment plans that is not present in the clinical notes. You don't know the patient's medical history and symptoms. You should ask or lead the patient to tell you the symptoms and his medical history, and you don't have any information about his medical history and symptoms. All the information of medical history, symptoms, medication history, and vaccination history should be told by the patient. You can tell the patient the test results, vital signs, and some conclusions.</p> <p>The Clinical Note: <a href="#">Clinical Note</a></p> <p>The Key Words: <i>key<sub>1</sub>, key<sub>2</sub>,...</i></p> <p>The History Conversation: <a href="#">History Dialogue</a></p> <p>You should only generate one utterance based on history conversation. Remember, you are the doctor, not the patient. Don't mention the information that has been mentioned in history conversation. If you feel that the patient's information is incomplete, you can supplement it based on the clinical note and include relevant keywords. However, please refrain from saying, 'based on medical record or clinical note.'</p> <p>Instead, you should say, 'I guess...'</p>
<b>Patient Prompt</b>	<p>Act as a patient to reply to the doctor. Add 'Patient:' before each round. Your answer should align with the clinical notes. You are just an ordinary person. Your response should be made as colloquial as possible. Don't mention any experimental results, conclusions, or medical dosage. because you're just an ordinary person and may not understand the meaning of these results. But you could tell the doctor your medical history, medication history, or vaccination history (medical history, medication history, or vaccination history are all long to medical history). Your response should revolve around the doctor's words and avoid adding information that was not mentioned.</p> <p>The Clinical Note: <a href="#">Clinical Note</a></p> <p>The History Conversation: <a href="#">History Dialogue</a></p> <p>Your reply should be succinct and accurate in a colloquial lay language style and must be aligned with clinical notes. Don't generate the part which should be said by the doctor. Do not say all the information unless the doctor asks about it. You cannot say any information about your test result or vital signs. Your medical history, vaccination history, and medication history all belong to medical history. Your reply must be completely aligned with the clinical note. But you cannot say any examination or test results because you are not a doctor. You must not be able to use highly specialized terms or medical terminology. You can only describe limited common symptoms. You shouldn't use the abbreviation if you know the full name(you should use the full name, not the abbreviation, such as D9 must be day 9, D7 must be day 7</p>

Table 7: Roleplay module prompt for doctor role and patient role.

Polish Prompt	<p>Expand the conversation. The conversation for patient parts can be more colloquial. When the doctor is speaking, the patient can have many modal particles (e.g. hmm, yes, okay) to increase interaction. All the numbers and medical concepts that appear in the note should be mentioned by the doctor. Professional medical terms and numbers should always occur in the doctor's utterances but not in the patient's answer. The doctor may describe and explain professional judgment to the patient and instruct the patient on follow-up requirements, but not ask questions that require professional medical knowledge to answer and the question must be around the clinical note(the patient could find the answer on the clinical note). All the information of medical history, symptoms and medication history should be told by patient. The patient's answer should be succinct and accurate in a colloquial lay language style. The answer should align with the clinical notes and as colloquial as possible. You can add some transitional phrases to make the conversation more logical.</p> <p>For example:  Example 1:  Patient: I understand, please go ahead.  (After examination)  Doctor: The result shows....  Example 2:  Patient: Thank you for the diagnosis, doctor.  (After two years)  Doctor: Hi...  Example 3:  Patient: Okay, I understand.  (Few days latter)  Doctor: Hi...</p> <p>Your conversations must follow the logical sequence of a doctor's inquiry. For example, the general logical order of the conversation is: first discussing symptoms, then discussing the medical history, followed by discussing testing and results, and finally discussing treatment options, conclusion etc." If you find this conversation to be incoherent, you can try dividing it into two separate coherent conversations. Patients should not say too much information at once.</p> <p>The Clinical Note: <b>Clinical Note</b>  The Key Words: <i>key<sub>1</sub>, key<sub>2</sub>,...</i>  The History Conversation: <b>Conversation</b></p> <p>There are only one patient and one doctor and just return the conversation. Your conversation must include all the key words I gave you.  Your conversation should also include all information. if it's difficult to include them all, you can use the original sentences in the notes.  The common symptoms and common medical history should be told by the patient.  Some specific symptoms and medical history should be added by the doctor after the patient has finished describing his symptoms and medical history.</p> <p>For example:  Doctor: Can you give me your medical history record?  Patient: Here you are.  Doctor: Based on your medical history record...  Because after the patient has finished describing common symptoms or medical history, he will give doctor his medical history records.  After patient gives the doctor his medical history record, the doctor could know medical history record. Otherwise he didn't know any information of the medical history.  Some results should not come from history clinical note they should come from the examination.  All the examination results, history examination results, vital sign and medical number must be told by doctor.  The revised conversation should be at least around 30 to 40 utterances (the doctor or patient should say too much information at once).  The conversation must include all the information on the clinical note.  You must include all the key words I gave you. If it is difficult to include all the key words you could use original the sentences of clinical note.  You cannot revise or eliminate any key words and you cannot use synonyms of the key words.  You shouldn't use the abbreviation if you know the full name(you should use full name not abbreviation, such as D9 must be day 9, D7 must be day 7. If both the full name and the abbreviation appear, it's better to use the full name rather than the abbreviation.  Patients must not say any highly specialized terms, medical terminology or medical dosage.  They can only describe limited common symptoms.</p> <p>The doctor should supplement the remaining information based on test results.  Don't repeat the same information in long paragraphs. The utterance of the dialogue needs to be expanded as much as possible.</p>

Table 8: Polish prompt.

<b>Combine Prompt</b>	<p>The above two paragraphs were extracted from a complete conversation.  Please concatenate the two dialogues together. Add 'Doctor:' before the doctor's words  and 'Patient:' before the patient's words for easier differentiation.  Please combine these two dialogues.  It means that your generation should include all the information  such as dosage of the medication which is mentioned in the clinical note  if the dosage is not mentioned in the clinical not  you should not mention it and the length should be longer than  both of these two conversations even longer than the sum of them.  You should try to ensure that the dialogue is smooth,  and don't use any greetings such as 'Hi there', 'how are you feeling today?',  'Hey', 'Hello' or any farewells in the dialogue.  The entire conversation takes place at the same time and place,  and revolves around the same patient and doctor.  Try to make the conversation smoother. Try to make these two dialogues into one dialogue  that takes place at the same time and place. Modify this conversation  by deleting all greeting sentences  such as 'Hi', 'Hey', 'Hi there', 'How are you feeling today', and 'Good Morning'.  The conversation must include these key words:<i>key<sub>1</sub>, key<sub>2</sub>, ...</i>  and you should also eliminate the repeat parts.</p>
-----------------------	---

Table 9: Combine prompt.

---

In this task, we ask for your expertise in annotating the quality of system-generated replies by machine learning models. Mainly we provide the history dialogue along with system-generated replies and ask for your preference.

Output your ranking for system-generated replies.  
Use the following format, and do not add any other text.

Some examples:

$a > b > c > d > e$   
 $e > d > c > b > a$

History Conversation:  
[*History Conversation*]

Conversation snippet:  
[*utterance*]

System-generated summaries:

1. [*Utterance1*]
2. [*Utterance2*]
3. [*Utterance3*]
4. [*Utterance4*]
5. [*Utterance5*]

Now, output your ranking:

---

Table 10: GPT-4 Prompt for preference ranking in extrinsic evaluation.

---

In this task, we ask for your expertise in annotating the quality of the system-generated dialogues by machine learning models. Mainly we provide the ground truth dialogue and the clinical note along with system-generated dialogues and ask for your preference.

Output your ranking for system-generated dialogues.  
Use the following format, and do not add any other text.

Some examples:

$a > b > c$   
 $c > b > a$

Clinical Note:  
[*Clinical Note*]

Ground Truth Dialogue:  
[*dialogue*]

System-generated summaries:

1. [*dialogue1*]
2. [*dialogue2*]
3. [*dialogue3*]

Now, output your ranking:

---

Table 11: GPT-4 Prompt for preference ranking in human evaluation.