

---

# LEVERAGING LARGE LANGUAGE MODEL AS SIMULATED PATIENTS FOR CLINICAL EDUCATION

---

Yanzeng Li<sup>1</sup>, Cheng Zeng<sup>2,3</sup>, Jialun Zhong<sup>1</sup>, Ruoyu Zhang<sup>1</sup>, Minhao Zhang<sup>1</sup>, Lei Zou<sup>1\*</sup>

<sup>1</sup>Wangxuan Institute of Computer Technology, Peking University

<sup>2</sup>School of Computer Science, Wuhan University. <sup>3</sup>CureFun Co.

## ABSTRACT

Simulated Patients (SPs) play a crucial role in clinical medical education by providing realistic scenarios for student practice. However, the high cost of training and hiring qualified SPs, along with the heavy workload and potential risks they face in consistently portraying actual patients, limit students' access to this type of clinical training. Consequently, the integration of computer program-based simulated patients has emerged as a valuable educational tool in recent years. With the rapid development of Large Language Models (LLMs), their exceptional capabilities in conversational artificial intelligence and role-playing have been demonstrated, making them a feasible option for implementing Virtual Simulated Patient (VSP). In this paper, we present an integrated model-agnostic framework called CureFun that harnesses the potential of LLMs in clinical medical education. This framework facilitates natural conversations between students and simulated patients, evaluates their dialogue, and provides suggestions to enhance students' clinical inquiry skills. Through comprehensive evaluations, our approach demonstrates more authentic and professional SP-scenario dialogue flows compared to other LLM-based chatbots, thus proving its proficiency in simulating patients. Additionally, leveraging CureFun's evaluation ability, we assess several medical LLMs and discuss the possibilities and limitations of using LLMs as virtual doctors from the perspective of their diagnostic abilities.

## 1 Introduction

Clinical medical education plays a pivotal role in training aspiring healthcare professionals by equipping them with the necessary skills, knowledge, and providing practical experience to deliver high-quality patient care. One essential component of clinical education is Simulated Patients (SPs), who are individuals trained to portray specific medical conditions, symptoms, and behaviors in a standardized and consistent manner. By interacting with SPs, students can practice their clinical skills, communication, and decision-making in a controlled and realistic environment [1, 2, 3, 4, 5]. While SPs have been reliable and invaluable in clinical education, its universality and popularity often be limited due to high costs associated with training and hiring qualified individuals [6, 7]. Meanwhile, SPs themselves face risks, including physical discomfort and psychological stress in consistently portraying actual patients [8, 9]. As a result, the Virtual Simulated Patient (VSP) [10, 11, 12] has emerged as an innovative approach to address these limitations. VSP systems leverage comprehensive technology, such as natural language processing and conversational artificial intelligence [13, 14, 15], to simulate actual patient encounters and provide students with realistic clinical scenarios for conversation and diagnosis. One recently advancement is the rise of large language models (LLMs), powerful artificial intelligence systems capable of processing and generating human-like text and conversations [16, 17]. LLMs, such as GPT [18], Llama [19], PaLM [20], have demonstrated remarkable capabilities in natural language understanding and generation, leading to their application in various domains, including digital medicine applications [21, 22, 23, 24, 25]. Recently, researchers have begun to explore the potential of LLMs as SPs, e.g., some of contemporaneous works have explored utilizing LLMs like ChatGPT in simulation-based training [26, 27, 28, 29].

Undoubtedly, the integration of *LLMs as SP* holds great promise for clinical education. Benefiting from large-scale pre-training and aligning with human preferences, LLMs enable remarkable abilities such as following instructions, analyzing text content, and recalling existed information from the context [30]. These fundamental capabilities are essential for a qualified SP. However, the existing approaches, including LLM-based methods, face certain

---

\*Corresponding Author

challenges [25]. One notable issue is hallucinations, which can result in the generation of fictional information and factual errors, thereby reducing the realism of clinical training. While existing LLMs primarily aim to align with the perspective of healthcare advisors to provide helpful responses in addressing users' concerns [31, 32], they often struggle to accurately portray an actual patient, leading to problems such as role flipping in SP conversations. Additionally, there are several general challenges in conversation AI, such as instruction leakage, flake replies, infinite repetition, and toxic responses [33, 34]. These issues need to be addressed to enhance the performance and reliability of VSP. Another challenging aspect in the utilization of SPs for clinical education is the assessment of students' medical dialogues. In traditional clinical examinations involving SPs, the student-patient encounter is typically evaluated by the SPs and supervisors by a checklist, which assesses various aspects of the student's performance [35, 36]. Traditional manual or rule-based grading methods face a series of issues, including raters' imperfect subjective judgments, the inability to dynamically and comprehensively evaluate the whole conversation, and difficulty in scaling up to large-scale assessments [37, 38, 39].

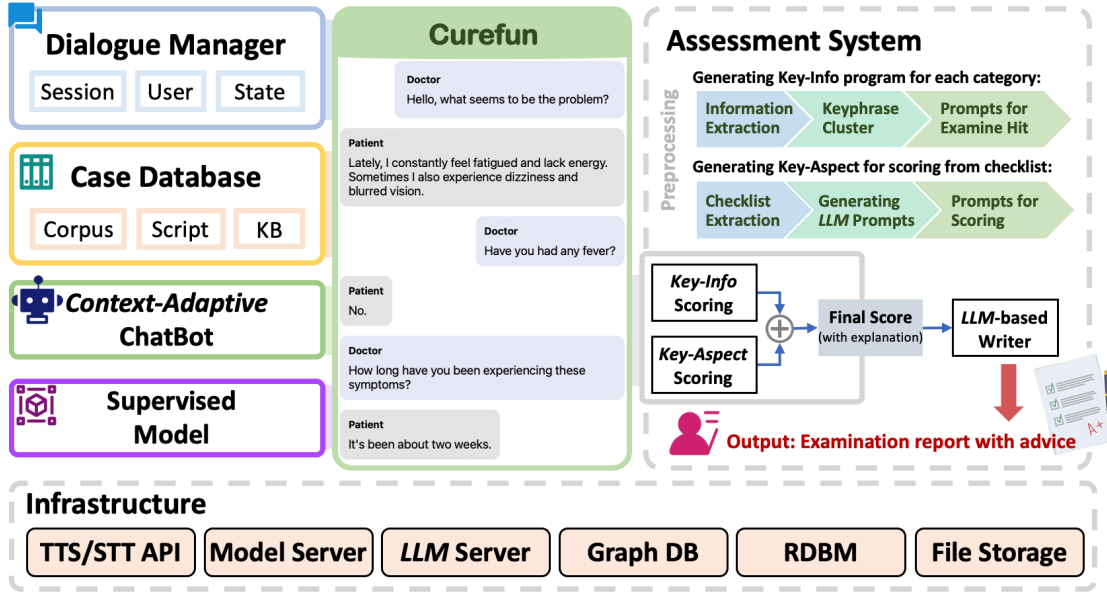


Figure 1: The overview diagram of this study. Curefun integrates LLMs to simulate patient roles, enhancing the dialogue flow via structured graph memory, and providing automatic assessment for student-patient conversations.

In this study, we present an integrated model-agnostic framework called CureFun that utilizes the comprehensive abilities of LLMs. Specifically, in the process of preparing the SP template, implementing student-patient conversation, and evaluating SP dialogue history, we utilize instructions such as Prompt and Chain-of-Thought to control the LLM behaviors according to our predetermined procedure. This allows us to construct a structured SP case graph, enhancing and controlling the dialogue flow in the form of retrieval-augmented generation (RAG) [40, 41]. Finally, through the generated multi-granularity assessment items, supervised by multiple LLMs and their voting, we can grade an appropriate score to the student, and provide suggestions. In summary, this framework aims to address existing issues and develop a productive VSP application for clinical education. Through comprehensive evaluations, we demonstrate that our approach enables more authentic and professional dialogue flows in SP scenarios compared to other LLM-based chatbots, and shields most of the aforementioned flaws. Furthermore, we enhance CureFun's assessment capability by automatically converting traditional SP evaluation checklists into LLM-executable programs. We then ensemble multiple LLMs to collaborate and provide a comprehensive and reliable assessment of students' medical dialogues. This enables large-scale and efficient SP-involved assessment in clinical education. Moreover, leveraging CureFun's evaluation capability, we assess several LLMs and discuss the possibilities and limitations of using LLMs as Virtual Doctors (VDs) [22]. In conclusion, our exploration highlights the potential of LLMs as VSPs for more efficient clinical education. And, it provides in-depth insights into the development of medical LLMs for intelligent diagnosis and treatment.

## 2 Result

In this section, we present the results of our comprehensive evaluation of the proposed framework, CureFun, in the context of acting SPs, automatic assessment, and evaluating LLMs as VDs in our SP system.

### 2.1 The Preparation of SP Cases

We have meticulously selected 8 cases in the Chinese language, with each case representing an individual SP script. These cases cover a range of diseases, including gastric disorders, diabetes, chronic obstructive pulmonary disease (COPD), COVID, pneumonia, and bronchiectasis. Moreover, these cases span multiple medical specialties such as pulmonology, endocrinology, and gastroenterology. The dataset used in this study was obtained from Wuhan Talent Information Technology Co., Ltd. After performing data cleaning, preprocessing, and manual data selection, we identified these eight cases as representative and comprehensive high-quality data for our subsequent experiments. An example of one of our SP cases is illustrated in Figure 5.

### 2.2 Generation Quality

We conducted evaluation experiments on 8 SP cases. We employed experts who are proficient in SP education assessment. They engaged in conversations with six mature LLM chatbots, with all SPs being given the same role-playing prompt to start with. During the interaction with the SPs, the experts asking questions according to the context and necessity of diagnose, to ensure a smooth conversation. The conversations were terminated either when the maximum number of rounds ( $N=20$ ) was reached or when the expert deemed it appropriate to end the dialogue. During the evaluation, the experts were unaware of which model was behind the tested VSP. In this study, we uses open-source LLMs with large parameter-scale and well-established commercial models as backbone models, owing to their advanced capacity to comprehend instructions and participate in dialogue [42, 43]. Specifically, we include the following LLMs in our study: GPT-3.5-turbo, PaLM [20], ERNIE-4 [44], Mixtral-8x7B [45], and Qwen-72B [46]. All of the included LLMs are capable of conversion in Chinese language, and chat versions of open-source LLMs are specially utilized in our investigation.

In the LLM as SP experiment, we observed that the different LLMs exhibited varying characteristics based on the enlighten of the same prompt. Figure 2 presents the statistics of the response lengths generated by each LLM when playing the role of the patient and responding to the doctor’s inquiries. It can be observed that PaLM often provides concise replies, whereas GPT-3.5-turbo tends to respond with more detailed content. After incorporating we proposed framework, there was no significant change in the distribution of response lengths, indicating that our framework would not visibly impact the inherent characteristics and personalities of the underlying LLMs.




Figure 2: Violin plot of the token length distribution of responses from different underlying LLMs when acting as SP and answering inquiries from doctors. The statistics on the left side of each violin represent the vanilla models’ generation, while the right sides represent the responses generated with our proposed framework.

In order to further investigate the effectiveness of *LLM as SP*, we conducted pairwise comparisons of various models following the setup of Chatbot Arena [47]. We employed GPT-4 [18], which currently demonstrates the best performance, as the judge and selected the superior performer among each pair of SP dialogues. The results of the win rates are presented in Figure 3. Furthermore, we employed Bootstrap ELO rating (B-ELO) to calculate the stable capacity ratings of each model, considering both unordered and sufficient competition [43]. We computed vanilla ELO [48] ratings using the configuration of  $initialrate = 1600$ ,  $K = 100$ , and then recomputed ELO scores 1,000 times in the

randomly shuffled comparison order. The median of the 1000 ELO scores was adopted as the final score, as depicted in Table 1.

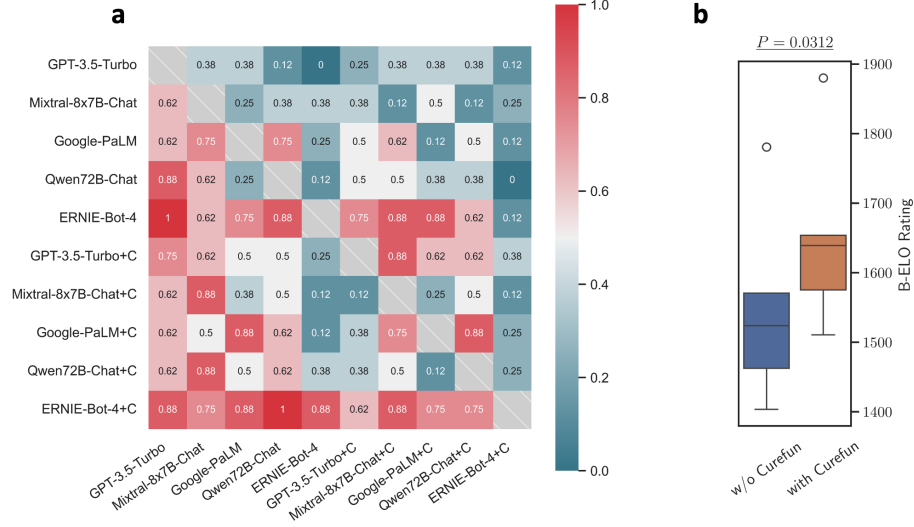


Figure 3: (a) Heatmap of pairwise comparison for LLMs in acting SPs. The left-side labels represent the LLM used as “player”, while the bottom labels represent the “opponents”. “+C” denotes the corresponding model is collaborating with our framework. (b) The B-ELO score distribution with or without our framework,  $P < 0.05$ , one-sided Wilcoxon’s rank-sum test.

From Figure 3 and Table 1, it can be observed that ERNIE-Bot-4 combined our framework achieved the best performance on the *LLM as SP* task. This could be attributed to the fact that our case scripts are all in Chinese language, and ERNIE bot is one of the earliest LLMs specifically designed for Chinese corpora. Conversely, Mixtral is primarily built on datasets of Indo-European languages, hence the performance lags a bit. Disregarding the individual idiosyncrasies of these LLMs, it is noteworthy that our framework consistently and significantly improved performance across all backbone LLMs (Figure 3b;  $P < 0.05$ , one-sided Wilcoxon’s rank-sum test). Specifically, integrating our framework into GPT-3.5-Turbo resulted in a 250.18-point increase in B-ELO score for the SP role-playing capacity, signifying a substantial advancement.

Table 1: B-ELO Ratings for LLM’s capacity to play the role of SP.

Model	B-ELO	
	w/o Ours	with Ours
Mixtral-8x7B	1462.40	1510.60(+48.20)
Qwen72B	1523.93	1575.20 (+51.27)
PaLM	1570.91	1639.07 (+68.16)
GPT-3.5-Turbo	1403.54	1653.72 (+250.18)
ERNIE-Bot 4	1780.88	1880.15 (+99.27)

Indeed, all of these LLMs have exhibited strong capabilities in role-playing. However, there have been some discrepancies in the role of SP. For example, GPT-3.5 demonstrates its high-level ethics and frequently flips the role as doctor following lengthy dialogues. Whereas our framework can consistently remind the model of the current task and the attributes it should consider. This feature is one of the key reasons why our framework has significantly enhanced the performance of GPT-3.5. Furthermore, the preferences of the judge, GPT-4, may introduce some biases. For instance, while the patient portrayed by PaLM is vivid and professional, GPT-4 often perceives it as lacking in detail due to PaLM’s tendency to generate shorter responses. Overall, these comparative experiments consistently demonstrate the effectiveness of our framework in the scenario of *LLM as SP*.

### 2.3 Evaluate Assessment

To validate the efficacy of automated assessment module within our proposed framework, we applied Curefun to evaluate the previously collected dialogue histories, consisting of a total of 80 records in 8 specific cases. Furthermore, an expert human evaluator was hired to independently fill out checklists and provide the grade for the aforementioned medicine dialogues, based on traditional scoring standard. We compared the scores obtained from both the automated scoring program and the human evaluators, sought to evaluate the consistency between these two assessment methods. To assess the degree of agreement between the human evaluators and our program, we employed two correlation measures: Spearman’s rank correlation and Pearson correlation. These measures allow us to quantify the strength and direction of the relationship between the two sets of scores.

Table 2: Comparison of correlation measures between human evaluators’ scores and program evaluators’ scores.

Item	Spearman’s Rank		Pearson’s	
	Correlation	p-value	Correlation	p-value
Case1	0.954	0.000	0.927	0.000
Case2	0.655	0.040	0.943	0.000
Case3	0.822	0.004	0.903	0.000
Case4	0.810	0.004	0.820	0.004
Case5	0.803	0.005	0.765	0.010
Case6	0.758	0.011	0.772	0.009
Case7	0.832	0.003	0.809	0.005
Case8	0.820	0.004	0.832	0.003

Table 2 shows Spearman’s rank correlation and Pearson correlation between human evaluators’ and our program’s assessment scores. Overall, both Spearman’s rank correlation and Pearson correlation coefficients are consistently closed to 1 (on average 0.81 and 0.85, respectively), suggesting a high degree of agreement between the two sets of scores. And the p-values associated with both correlation tests were found to be less than 0.05 for each group of cases. Those obtained correlation coefficients indicate a strong positive relationship between the human evaluators’ scores and the scores generated by our program. The high correlations and statistically significant p-values suggest that our program’s scores align closely with the judgments made by human evaluators. The score distribution depicted in Figure 4 further supports our analysis. The distribution appears to be symmetrical and concentrated around a central value, indicating a consistent and reliable scoring mechanism. This indicates that our automated assessment method produces reliable and accurate assessments, making it a suitable alternative to human evaluators in SP tests.




Figure 4: The distribution of scores from program evaluator and human evaluator.

### 2.4 Evaluate LLMs as VDs

Based on our comprehensive VSP, we have the capability to evaluate various public-available LLMs and test their standardized diagnostic interviewing abilities [49], which is a novel perspective due to its dynamic character. We integrate Curefun into the to-be-evaluated LLMs’ chat flow as the role of users, and through multi-turn dialogues in predefined diagnosis scenarios, we stimulate LLMs to conduct medicine interviews with our VSP. Ultimately, the entire dialogue histories are scored using aforementioned assessment method, which records and analyzes various indicators during the conversation process. Except for the scores used to directly evaluate the VDs’ diagnostic abilities, we also

Table 3: Model Comparison

Model	Information Density	Emotional Tendency	Response Length	Turn Number	Overall Score
Llama2-70B	0.02	0.93	2996.17	5.39	0.34
Mixtral-8x7B	0.11	0.69	647.50	5.88	0.33
BianQue-2 (6B)	0.14	0.69	178.00	4.25	0.25
DISC-MedLLM (13B)	0.15	0.66	361.87	4.57	0.43
ERNIE-4-Bot	0.13	0.69	662.09	5.70	0.37
ChatGPT (3.5-turbo)	0.15	0.70	371.74	7.74	0.51
Human (Non-medical background)	0.15	0.71	125.00	8.45	0.45
Human (Clinical medicine student)	0.19	0.48	212.99	23.11	0.72
Human (Expert)	0.27	0.56	135.75	13.38	0.78

define several non-scoring indicators to provide more reference information for analyzing the VDs’ performance. These indicators include:

- **Information Density:** Measures the density of entities and medical entities contained in the dialogue, indicating the professionalism and information density of the conversation.
- **Emotional Tendency:** Measures the emotional inclination in the dialogue, judging the emotional polarity of the conversation through an emotion analysis model, indicating the friendliness of the doctor role in the dialogue.
- **Response Length:** The length of the doctor’s response in the dialogue, indicating the level of detail in the doctor’s response.
- **Turn Number:** The total number of turns in the dialogue.

These indicators provide a comprehensive evaluation of the VDs’ diagnostic abilities and help to analyze their performance in detail.

We employ several kinds of LLMs as VDs, including SOTA business LLMs (ChatGPT and ERNIE-4-Bot), open-sourced general LLMs (Llama2-70B [19], Mixtral-8x7B [45]), and medical-specific LLMs (BianQue-2 (6B) [50], DISC-MedLLM (13B) [51]), to evaluate their diagnostic abilities. Those LLMs are evaluated in the same 8 cases as the SPs, and each experiment is repeated 5 times to ensure the stability and reliability of the results. We also recruited human evaluators with different medical backgrounds to participate in the evaluation as a reference group. The human evaluators are classified into three categories: non-medical background folks, clinical medicine students, and experts, each of group contains 3 individuals. The non-medical background human evaluators have no medical knowledge, the clinical medicine students are second-year or third-year medical student, and the experts are senior physicians with rich clinical experience. The evaluation was conducted in a controlled environment, with each evaluator interacting with our SP chatbot.

Table 3 presents the evaluation results of the performance and metrics of LLMs acting as VDs. The results show that ChatGPT obtains the highest score among all LLMs. DISC-MedLLM achieves the second-highest score, indicating that medical-specific LLMs could exhibit better diagnostic abilities than general LLMs. However, BianQue-2 achieves a relative lower score, which may attribute to its smallest parameter scale thus lack of multi-turn dialogue ability. Besides, human evaluators, especially the experts, outperform all LLMs in terms of diagnostic ability, indicating that LLMs still have room for improvement in simulating real-world medical scenarios.

From the non-scoring indicators, we can observe the diversified preferences of LLMs and human-beings. For example, LLMs, especially Llama2-70B, tend to generate long responses while human evaluators like to talk in relative shorter length. All LLMs tend to end the conversation in fewer rounds, while human evaluators prefer more rounds of conversation, especially clinical medicine students have an average of 23 rounds of conversation. This may be because human evaluators tend to communicate with patients to obtain information through in-depth conversation, while LLMs prefer to solve problems in fewer rounds to reduce costs. Additionally, it can be noted that LLMs consistently perform well in emotional tendency, or the degree of friendliness in communication with patients. This is because LLMs are usually given a certain degree of friendliness during alignment training. The friendliness of human evaluators in this indicator varies, which may be due to the fact that human evaluators are influenced by more factors during the conversation, such as empathy for patients and their individual career experience<sup>2</sup>.

<sup>2</sup>An interesting finding is that human doctors tend to shift towards negative sentiment polarity over time during the entire test, whereas the LLMs remain positive throughout. This suggests the potential of LLMs for greater stability than humans in medical consultation.

In conclusion, although LLM falls short of human experts in our SP testing, it has already reached a level comparable to ordinary individuals without medical backgrounds. Additionally, LLM consistently maintains a superior emotional indicator when interacting with patients compared to humans. This result indicates that LLMs can assist human doctors to some extent in medical consultations, highlighting the potential of LLMs as pre-diagnostic and triage tools.

It’s worth noting that SPs and VDs, considered complementary tasks, do not align perfectly in real-world scenarios, as confirmed by our diagnostic evaluation of various LLMs. Highly performing LLMs in diagnostic QA benchmarks do not meet expectations in our VSP setting. Firstly, SPs are primarily designed for education and rely on standardized examinations, prioritize comprehensive and standardized communication, while real patient interactions prioritize efficiency and flexibility. This disparity exists not only between VSPs and VDs but also in real-world clinical education and practice. Therefore, aspiring students are advised to gain practical experience to align with real-world medical scenarios. Based on these observations, we can develop an LLM training process integrating VSP and VD tasks to facilitate self-improvement, as explored in a recent study [52].

### 3 Method

In this section, we describe the detailed methodology and main components of our proposed framework. We developed Curefun from the ground up, including the data processing pipeline, the graph-driven context-adaptive SP chatbot, the LLM-based automatic assessment methods, and some infrastructure for supporting our proposed framework and improving the user experience. The overview of our framework is illustrated in Figure 1.

#### 3.1 Data Processing

For each SP case, we utilize Named Entity Recognition (NER) [53] and relation extraction [54] models to extract entities and relationships between entities from the SP script, forming the skeleton graph of the corresponding case. Then, we employ open information extraction method [55, 56] to extract the attributes and the corresponding attribute values associated with entities in the skeleton graph, creating a case graph, which can describe the internal relationships and various attributes of the case. Specifically, we trained a BERT-based NER model [57, 58] on both general corpus and medical datasets. This model is used to extract entities in the general domain and the medical domain [59], including medications, symptoms, and diseases, etc. Then we employed a simple BERT-based relation classification model [60] to detect relationships between various entities [61]. An open attribute extraction model [62] was introduced to supplement the attributes of medical entities. Some literal values that cannot be contained in the entity-relation skeleton, such as body temperature and blood pressure, will be extracted as attributes and linked to the skeleton graph. Figure 5 presents a constructed case graph through aforementioned processes.

The purpose of using aforementioned information extraction methods to construct a case graph is to utilize the RAG manner. When a student mentions relevant symptoms or relationships, the RAG method stimulate model to retrieve relevant subgraphs of entities, relationships, and attributes from the whole case graph. This effectively reduces the length of the case scripts occupying the LLM input buffer, alleviating the degradation of generation quality caused by excessively long form dialogue. It enhances the quality and coherence of dialogue generation in VSP’s conversation [63, 64].

Except constructing the case graph, we also preprocess the corresponding SP checklist for automatic assessment, which is detailed in Section 3.3.

#### 3.2 Graph-Driven Context-Adaptive SP Chatbot

To enhance the dialogue quality and coherence of the VSP, we propose the integration of a graph-driven context-adaptive mechanism into the LLM-based chatbot. This mechanism aims to dynamically adjust the flow of the dialogue by utilizing a controlled and structured graph memory. The process is illustrated in Figure 6 (a), which consists of four main steps: **Extract**, **Retrieve**, **Rewrite**, and **Generate** (ERRG). In the “Extract” step, the chatbot extract the core entities and possible relations from the user’s input and the ongoing dialogue context. These extracted mentions are then linked to the corresponding entity nodes in the case graph through substring matching and string proximity. In the “Retrieve” step, the chatbot utilizes the LLM to generate a formal query based on the extracted entities and relations, then executes the query on the case graph to retrieve the relevant subgraph. In the “Rewrite” step, the chatbot transforms the retrieved subgraph into natural language form, thereby providing context-aware evidence for generating the response. Finally, in the “Generate” step, the chatbot generates an optimal response by considering the retrieved evidence, the current dialogue context, and the user’s input.







It is worth noting that our proposed framework, Curefun, serves as a model-agnostic shell capable of accommodating a wide range of chat-oriented LLMs, as demonstrated in Table 1. However, in light of considerations pertaining to usability, cost, and efficacy, we employ a Llama-13B based checkpoint [65] as the backbone model in the production. Subsequently, the automatic assessment and the corresponding experiments in Section 2.3 and Section 2.4 were conducted using this model as the foundation.

### 3.3 LLM-based Automatic Assessment

In order to comprehensively and dynamically evaluate students’ performances and expand to large-scale assessments, we utilize LLM to analyze, extract, and summarize the various rating items in the original checklist, and eventually forming two elements for scoring:

- *Aspects*: These aspect items necessitate proactive inquiry by the student in order to accrue scoring credit. E.g., whether the student asked about the patient’s genetic medical history.
- *Information*: These key information items must be elicited from patient through inquiry to obtain scoring credit. E.g., specific symptoms and important physiological indicators.

The scoring elements at granularity hierarchical level simplify the complex checklist into single tasks compatible with LLMs, facilitating more accurate assessment. An example of a structured assessment checklist is illustrated in Figure 7. This approach allows us to seamlessly transform the complex checklist into an automated evaluation procedure.

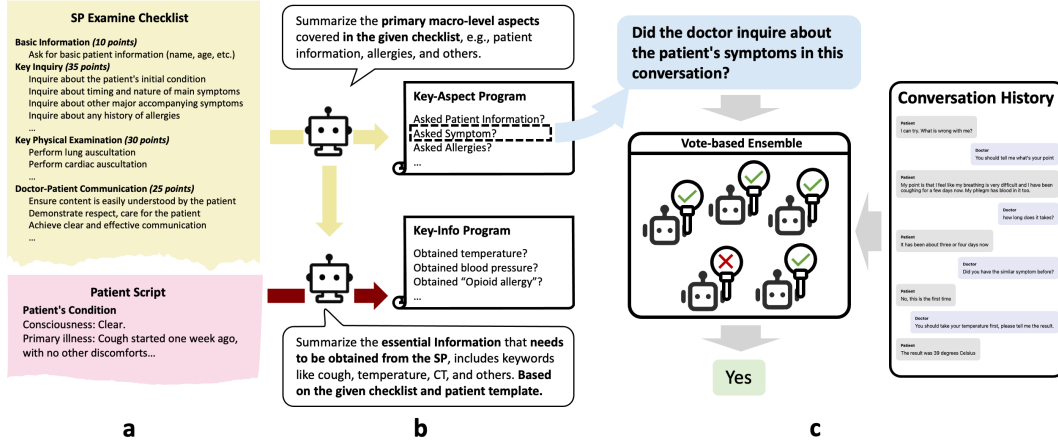


Figure 7: A sample of an assessment. (a) Original case checklist. (b) Processed automatic scoring program. (c) Vote-based LLMs ensemble.

After generating the automatic scoring program, we employed a series of LLMs to perform scoring for each item and ensemble their opinion through voting, as shown in Figure 7(c). After voting for all items in the program, we can obtain a reliable and stable score for the clinical history. In practice, we used 5 different Llama-13B checkpoints for ensemble. The proportions of *Aspect* and *Information* in the total score are 0.3 and 0.7, respectively, according to the checklist in the actual SP examination. The scores obtained through this approach allow for a comprehensive assessment of students’ mastery of various aspects of the SP case in the examination. Furthermore, it effectively prevents attacks targeting specific exam items, as obtaining a valid score for “Information” requires eliciting it through inquiry and cannot be achieved by exhaustively listing it in endless questions. In addition to assessing the mastery of “Aspects” and “Key Information”, we have also defined several metrics, including the emotional inclination of the dialogue, the entity-related information density, the dialogue length, and the number of dialogue turns. These metrics, as shown in Table 3 in Section 2.4, are difficult to quantify and may reduce the robustness of assessment. Therefore, they serve as non-scoring indicators to provide reference information for educators and students. The experimental results presented in Section 2.3 demonstrate that our proposed assessment programs yield scores consistent with human grading.

### 3.4 Auxiliary Modules

We introduce several auxiliary modules that enhance the performance of our proposed framework.

**TTS/STT Models.** To facilitate natural conversations between students and our proposed system, we integrate a Text-to-Speech (TTS) [66] and Speech-to-Text (STT) [67] module into our framework. The TTS module converts the

generated patients' responses into realistic speech, providing a more immersive experience for students. Conversely, the STT module transcribes the students' spoken inquiries into text format, enabling the LLMs to process and respond accordingly. By incorporating TTS and STT technologies, our framework bridges the gap between written dialogue and spoken interaction, enhancing the authenticity and effectiveness of the VSP experience.

**Graph Database.** To enable efficient storage and retrieval of character settings and LLM-fabled attributes, we employ the graph database [68] as a knowledge repository in our framework. The graph database can efficiently organize medical concepts, conditions, and symptoms as nodes, while the relationships between them are represented as edges [69]. Hereby we introduce the standardized Resource Description Framework (RDF) <sup>3</sup> format to represent the structured patient information, then use SPARQL [70] queries to retrieve the relevant information in graph database. This structure allows for flexible querying and navigation of patient information, facilitating the generation of accurate and contextually relevant responses by the LLMs.

**LLM Server.** To support the computational requirements of our framework, we introduce the dedicated and high-performance LLM server [71]. The server hosts the LLMs, adopting acceleration technique like page attention [71] and speculative decoding [72], enabling efficient and parallel processing of student inquiries and SPs' responses. The LLM server plays a critical role in ensuring the overall performance and scalability of our framework.

## References

- [1] Robert C Pascucci, Peter H Weinstock, Brigid E O'Connor, Kristina M Fancy, and Elaine C Meyer. Integrating actors into a simulation program: a primer. *Simulation in Healthcare*, 9(2):120–126, 2014.
- [2] Jill S Sanko, Ilya Shekhter, Richard R Kyle Jr, Stephen Di Benedetto, and David J Birnbach. Establishing a convention for acting in healthcare simulation: merging art and science. *Simulation in Healthcare*, 8(4):215–220, 2013.
- [3] David M Gaba. The future vision of simulation in healthcare. *Simulation in Healthcare*, 2(2):126–135, 2007.
- [4] Alessandra R Mesquita, Divaldo P Lyra Jr, Giselle C Brito, Blcie J Balisa-Rocha, Patricia M Aguiar, and Abilio C de Almeida Neto. Developing communication skills in pharmacy: a systematic review of the use of simulated patient methods. *Patient education and counseling*, 78(2):143–148, 2010.
- [5] Amitai Ziv, Paul Root Wolpe, Stephen D Small, and Shimon Glick. Simulation-based medical education: an ethical imperative. *Simulation in Healthcare*, 1(4):252–256, 2006.
- [6] Maureen Hillier, Tony L Williams, and Tiffani Chidume. Standardization of standardized patient training in medical simulation. 2020.
- [7] Heidi M. Felix and Leslie V. Simon. Types of standardized patients and their recruitment in medical simulation. 2019. URL <https://www.ncbi.nlm.nih.gov/books/NBK549907/>.
- [8] Lonneke Bokken, Jan Van Dalen, and Jan-Joost Rethans. Performance-related stress symptoms in simulated patients. *Medical education*, 38(10):1089–1094, 2004.
- [9] Elizabeth T Newlin-Canzone. The effect of improvisations and observations on standardized patient encounters, subjective workload and stress. 2011.
- [10] Elham Mousavinasab, Nahid Zarifsanaiy, Sharareh R. Niakan Kalhori, Mahnaz Rakhshan, Leila Keikha, and Marjan Ghazi Saeedi. Intelligent tutoring systems: a systematic review of characteristics, applications, and evaluation methods. *Interactive Learning Environments*, 29(1):142–163, 2021.
- [11] Víctor López, Eduardo M Eisman, and Juan Luis Castro. A tool for training primary health care medical students: The virtual simulated patient. In *2008 20th IEEE International Conference on Tools with Artificial Intelligence*, volume 2, pages 194–201. IEEE, 2008.
- [12] Raffaello Furlan, Mauro Gatti, Roberto Menè, Dana Shiffer, Chiara Marchiori, Alessandro Giaj Levra, Vincenzo Saturnino, Enrico Brunetta, and Franca Dipaola. A natural language processing-based virtual patient simulator and intelligent tutoring system for the clinical diagnostic process: simulator development and case study. *JMIR medical informatics*, 9(4):e24073, 2021.
- [13] Carlos Rodrigues, Arsénio Reis, Rodrigo Pereira, Paulo Martins, José Sousa, and Tiago Pinto. A review of conversational agents in education. In *International Conference on Technology and Innovation in Learning, Teaching and Education*, pages 461–467. Springer, 2022.
- [14] Joao Luis Zeni Montenegro, Cristiano André da Costa, and Rodrigo da Rosa Righi. Survey of conversational agents in health. *Expert Systems with Applications*, 129:56–67, 2019.

<sup>3</sup><http://www.w3.org/TR/rdf-sparql-query/>

- [15] Lijia Chen, Pingping Chen, and Zhijian Lin. Artificial intelligence in education: A review. *Ieee Access*, 8: 75264–75278, 2020.
- [16] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- [17] Bonan Min, Hayley Ross, Elinor Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, 56(2):1–40, 2023.
- [18] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmann, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [19] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [20] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [21] Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Anthony B Costa, Mona G Flores, et al. A large language model for electronic health records. *NPJ Digital Medicine*, 5(1):194, 2022.
- [22] Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. Large language models in medicine. *Nature medicine*, 29(8):1930–1940, 2023.
- [23] Bertalan Meskó and Eric J Topol. The imperative for regulatory oversight of large language models (or generative ai) in healthcare. *NPJ digital medicine*, 6(1):120, 2023.
- [24] Tiffany H Kung, Morgan Cheatham, Arielle Medenilla, Czarina Sillos, Lorie De Leon, Camille Elepaño, Maria Madriaga, Rimel Aggabao, Giezel Diaz-Candido, James Maningo, et al. Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models. *PLoS digital health*, 2(2):e0000198, 2023.
- [25] Alaa Abd-Alrazaq, Rawan AlSaad, Dari Alhuwail, Arfan Ahmed, Pdraig Mark Healy, Syed Latifi, Sarah Aziz, Rafat Damseh, Sadam Alabed Alrazak, Javaid Sheikh, et al. Large language models in medical education: Opportunities, challenges, and future directions. *JMIR Medical Education*, 9(1):e48291, 2023.
- [26] Siyuan Chen, Mengyue Wu, Kenny Q Zhu, Kunyao Lan, Zhiling Zhang, and Lyuchun Cui. Llm-empowered chatbots for psychiatrist and patient simulation: Application and evaluation. *arXiv preprint arXiv:2305.13614*, 2023.
- [27] Trista M Benítez, Yueyuan Xu, J Donald Boudreau, Alfred Wei Chieh Kow, Fernando Bello, Le Van Phuoc, Xiaofei Wang, Xiaodong Sun, Gilberto Ka-Kit Leung, Yanyan Lan, et al. Harnessing the potential of large language models in medical education: promise and pitfalls. *Journal of the American Medical Informatics Association*, page ocad252, 2024.
- [28] Friederike Holderried, Christian Stegemann-Philipps, Lea Herschbach, Julia-Astrid Moldt, Andrew Nevins, Jan Griewatz, Martin Holderried, Anne Herrmann-Werner, Teresa Festl-Wietek, Moritz Mahling, et al. A generative pretrained transformer (gpt)-powered chatbot as a simulated patient to practice history taking: Prospective, mixed methods study. *JMIR Medical Education*, 10(1):e53961, 2024.
- [29] Neil Sardesai, Paolo Russo, Jonathan Martin, and Anand Sardesai. Utilizing generative conversational artificial intelligence to create simulated patient encounters: a pilot study for anaesthesia training. *Postgraduate Medical Journal*, page qgad137, 2024.
- [30] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.
- [31] Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*, 2023.
- [32] Tessa Han, Aounon Kumar, Chirag Agarwal, and Himabindu Lakkaraju. Towards safe and aligned large language models for medicine. *arXiv preprint arXiv:2403.03744*, 2024.
- [33] Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. Camel: Communicative agents for "mind" exploration of large language model society. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

- [34] Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274, 2023.
- [35] R Ladyshevsky. Simulated patients and assessment. *Medical Teacher*, 21(3):266–269, 1999.
- [36] Cathy M Smith, Carol C O’Byrne, and Debra Nestel. Simulated patient methodology and assessment. *Simulated patient methodology: theory, evidence and practice*, pages 85–92, 2014.
- [37] Leonora Kaldaras and Kevin C Haudek. Validation of automated scoring for learning progression-aligned next generation science standards performance assessments. In *Frontiers in Education*, volume 7, page 968289. Frontiers Media SA, 2022.
- [38] Jinhao Wang and Michelle Stallone Brown. Automated essay scoring versus human scoring: A correlational study. *Contemporary Issues in Technology and Teacher Education*, 8(4):310–325, 2008.
- [39] Marcus Messer, Neil CC Brown, Michael Kölling, and Miaojing Shi. Automated grading and feedback tools for programming education: A systematic review. *arXiv preprint arXiv:2306.11722*, 2023.
- [40] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2023.
- [41] Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.
- [42] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [43] Haodong Duan, Jueqi Wei, Chonghua Wang, Hongwei Liu, Yixiao Fang, Songyang Zhang, Dahua Lin, and Kai Chen. Botchat: Evaluating llms’ capabilities of having multi-turn dialogues. *arXiv preprint arXiv:2310.13650*, 2023.
- [44] Yu Sun, Shuohuan Wang, Shikun Feng, Siyu Ding, Chao Pang, Junyuan Shang, Jiaxiang Liu, Xuyi Chen, Yanbin Zhao, Yuxiang Lu, et al. Ernie 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation. *arXiv preprint arXiv:2107.02137*, 2021.
- [45] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- [46] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- [47] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023.
- [48] Arpad E Elo. The proposed uscf rating system, its development, theory, and applications. *Chess Life*, 22(8): 242–247, 1967.
- [49] Sandeep Reddy. Evaluating large language models for use in healthcare: A framework for translational value assessment. *Informatics in Medicine Unlocked*, 41:101304, 2023. ISSN 2352-9148. doi:<https://doi.org/10.1016/j.imu.2023.101304>. URL <https://www.sciencedirect.com/science/article/pii/S2352914823001508>.
- [50] Yirong Chen, Zhenyu Wang, Xiaofen Xing, Zhipei Xu, Kai Fang, Junhong Wang, Sihang Li, Jieling Wu, Qi Liu, Xiangmin Xu, et al. Bianque: Balancing the questioning and suggestion ability of health llms with multi-turn health conversations polished by chatgpt. *arXiv preprint arXiv:2310.15896*, 2023.
- [51] Zhijie Bao, Wei Chen, Shengze Xiao, Kuang Ren, Jiaao Wu, Cheng Zhong, Jiajie Peng, Xuanjing Huang, and Zhongyu Wei. Disc-medllm: Bridging general large language models and real-world medical consultation. *arXiv preprint arXiv:2308.14346*, 2023.
- [52] Tao Tu, Anil Palepu, Mike Schaekermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Tomasev, et al. Towards conversational diagnostic ai. *arXiv preprint arXiv:2401.05654*, 2024.

- [53] Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering*, 34(1):50–70, 2020.
- [54] Hailin Wang, Guoming Lu, Jin Yin, and Ke Qin. Relation extraction: A brief survey on deep neural network based methods. In *2021 The 4th International Conference on Software Engineering and Information Management*, pages 220–228, 2021.
- [55] Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. Clinical information extraction applications: a literature review. *Journal of biomedical informatics*, 77:34–49, 2018.
- [56] Yanzeng Li and Lei Zou. gbuilder: A scalable knowledge graph construction system for unstructured corpus. *arXiv preprint arXiv:2208.09705*, 2022.
- [57] Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1):86, 2021.
- [58] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi:10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- [59] Priyanka Bose, Sriram Srinivasan, William C Sleeman IV, Jatinder Palta, Rishabh Kapoor, and Preetam Ghosh. A survey on recent named entity recognition and relationship extraction techniques on clinical texts. *Applied Sciences*, 11(18):8319, 2021.
- [60] Arpita Roy and Shimei Pan. Incorporating medical knowledge in BERT for clinical relation extraction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5357–5366, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. doi:10.18653/v1/2021.emnlp-main.435. URL <https://aclanthology.org/2021.emnlp-main.435>.
- [61] Udo Hahn and Michel Oleynik. Medical information extraction in the age of deep learning. *Yearbook of medical informatics*, 29(01):208–220, 2020.
- [62] Yanzeng Li, Bingcong Xue, Ruoyu Zhang, and Lei Zou. Attgen: Attribute tree generation for real-world attribute joint extraction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2139–2152, 2023.
- [63] Chen Wei, Zhichen Yu, and Simon Fong. How to build a chatbot: chatbot framework and its capabilities. In *Proceedings of the 2018 10th international conference on machine learning and computing*, pages 369–373, 2018.
- [64] Fangyuan Xu, Weijia Shi, and Eunsol Choi. Recomp: Improving retrieval-augmented lms with compression and selective augmentation. *arXiv preprint arXiv:2310.04408*, 2023.
- [65] Jiaxing Zhang, Ruyi Gan, Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, Yongfeng Huang, Xiayu Li, Yanghan Wu, Junyu Lu, Xinyu Zhu, Weifeng Chen, Ting Han, Kunhao Pan, Rui Wang, Hao Wang, Xiaojun Wu, Zhongshen Zeng, and Chongpei Chen. Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence. *CoRR*, abs/2209.02970, 2022.
- [66] Xu Tan, Tao Qin, Frank Soong, and Tie-Yan Liu. A survey on neural speech synthesis. *arXiv preprint arXiv:2106.15561*, 2021.
- [67] Mishaim Malik, Muhammad Kamran Malik, Khawar Mehmood, and Imran Makhdoom. Automatic speech recognition: a survey. *Multimedia Tools and Applications*, 80:9411–9457, 2021.
- [68] Lei Zou, Jinghui Mo, Lei Chen, M Tamer Özsu, and Dongyan Zhao. gstore: answering sparql queries via subgraph matching. *Proceedings of the VLDB Endowment*, 4(8):482–493, 2011.
- [69] Renzo Angles and Claudio Gutierrez. Survey of graph database models. *ACM Computing Surveys (CSUR)*, 40(1): 1–39, 2008.
- [70] Yanzeng Li, Zilong Zheng, Wenjuan Han, and Lei Zou. Vgstore: A multimodal extension to SPARQL for querying RDF scene graph. In Anastasia Dimou, Armin Haller, Anna Lisa Gentile, and Petar Ristoski, editors, *Proceedings of the ISWC 2022 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 21st International Semantic Web Conference (ISWC 2022), Virtual Conference, Hangzhou, China, October 23-27, 2022*, volume 3254 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2022. URL <https://ceur-ws.org/Vol-3254/paper377.pdf>.
- [71] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626, 2023.

- [72] Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*, 2023.