

Conditioning Chat-GPT for Information Retrieval: The Unipa-GPT Case Study

Irene Siragusa¹, Roberto Pirrone¹

¹Dipartimento di Ingegneria @ Università degli Studi di Palermo, Viale delle Scienze, Edificio 6, 90128 - PALERMO

Abstract

This paper illustrates the architecture and training of Unipa-GPT, a Large Language Model based chatbot developed for assisting students in choosing a bachelor/master degree course at the University of Palermo. Unipa-GPT relies on gpt-3.5-turbo, it was presented in the context of the European Researchers' Night SHARPER event. In our experiments we adopted both the Retrieval Augmented Generation (RAG) approach and fine-tuning to develop the system. The whole architecture of Unipa-GPT is presented, both the RAG and the fine-tuned systems are compared, and a brief discussion on their performance is reported.

Keywords

Large Language Model, ChatGPT, RAG, Fine-tuning

1. Introduction

Natural Language Processing (NLP), Artificial Intelligence (AI) and Large Language Model (LLM) are highly interesting topics in the scientific community, but they have become also very popular since ChatGPT¹ by OpenAI was released and it was made publicly accessible. Despite the debate about the level of consciousness in ChatGPT and the claims about its standing for a new generation of Artificial General Intelligence (AGI) [2, 3], a plethora of new applications are being developed that are based on either ChatGPT [4] or one of the other LLMs proposed in the last year.

In view of the previous considerations, we investigated the use of ChatGPT for building a virtual assistant for the users of a large public institution like the University of Palermo, and in this paper we propose Unipa-GPT, a chatbot that can guide secondary school students to explore the huge information provided by the University institutional website, when they are choosing a bachelor degree course, by answering to questions related to the available courses, application procedure, tax payments, access to scholarships and more. Unipa-GPT was presented to the wide public in the context of the European Researchers' Night SHARPER event².


Unipa-GPT has been developed mainly as a Retrieval Augmented Generation (RAG) system [5] based on gpt-3.5-turbo that collects the target information from unipa-corpus, a document corpus that has been built purposely by scraping the University website. The aim

NL4AI 2023: Seventh Workshop on Natural Language for Artificial Intelligence, November 6-7th, 2023, Rome, Italy [1]

✉ irene.siragusa02@unipa.it (I. Siragusa); roberto.pirrone@unipa.it (R. Pirrone)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://openai.com/blog/chatgpt>

²<https://www.sharper-night.it/>

of this work is to explore the behaviour and the limitations of LLMs when they are engaged in a Q&A task where precise domain knowledge is required. Obviously, also a fine-tuned version has been built where the corpus has been modified with the aim of saving computational resources i.e. use the less tokens as possible, and a mixed strategy has been adopted where RAG was coupled with fine-tuning to avoid the train step on very detailed information such as the educational objectives of each single class. Both the models have been tested qualitatively by very few students right now, and we present a comparison of their performance based on their judgement on two reference chats along with a discussion of the results.

The paper is arranged as follows: Section 2 illustrates the different corpora we set up for building both the RAG and the fine-tuned Unipa-GPT. The detailed architecture of both systems is reported in Section 3, while the experimental results are reported and discussed in Section 4. Concluding remarks are drawn in Section 5.

2. Corpora

In this section we outline the differences between the versions of the unipa-corpus used for developing the RAG-only and the fine-tuned system.

2.1. unipa-corpus for RAG

The corpus used for Unipa-GPT, called unipa-corpus, is a collection of documents that were collected directly from the website of the University of Palermo. A manual selection of the most interesting pages was made with reference to the target audience of secondary-school students and two main sections of the corpus were identified, that correspond to the *Education* and the *Future Students* sub-trees in the institutional website. Since the expected questions are in Italian, the generated corpus is in Italian.

The *Education* section is the main part of the corpus and it is a collection of all the available courses at the University for the academic year 2023/2024. For each course and each curriculum two files are obtained: *details* is the file that collects all the general details of the course, like name, department of affiliation, typology of course (Bachelor or Master degree), restriction of access and a colloquial description of the course, including its educational objectives and professional opportunities; *course outline* is the second file that collects the course outline divided by year, and the number of credits, the teaching professor, the teaching period and the scientific sector are specified for each class. Three different versions of the *course outline* file were generated, namely `clear`, `full` and `emb`. The `clear` version is the one described above, the `full` version adds a new document for every class in a course and reports its peculiar educational objectives. The `emb` version is a mix of the previous ones where the classes' educational objectives are added directly in the file containing the outline of the course. This distinction led the `clear` corpus and the `emb`³ one to have the same number of files but different information, while the `full` contains the same information of the `emb` corpus but arranged in a different number of documents.

³Despite this corpus is called *embedded*, it does not contains embeddings, the words *embedded* refers to the educational objectives that are inserted in the file with the course outline

Table 1

Number of documents in each version of unipa-corpus.

	<i>Education</i>	<i>Future Students</i>	Total
unipa-corpus-clear	506	104	610
unipa-corpus-full	5794	104	5898
unipa-corpus-emb	506	104	610

The *Future Student* section is the same for the three versions of the corpus, and it is a mix of documents coming from the related section of the University website. The information contained in this files is addressed to the future students of the University, including the academic calendar, the tax rules and reductions, scholarships, University enrolment procedure, and facilities offered to the students.

In Table 1 are reported the statistics of each corpus.

2.2. unipa-corpus for fine-tuning

The unipa-corpus was modified to be in the form required for fine-tuning gpt-3.5-turbo. As already mentioned above, our intent in fine-tuning was lowering the computational resources as much as possible that is using the minimum tokens for training the model. Besides the economic aspect in the case of ChatGPT fine-tuning, this is a crucial topic when dealing with LLMs because also relatively small LLMs like LLama-2-7B [6] require huge computational resources for their fine-tuning or retraining. For our purposes, the chosen corpus was the unipa-corpus-clear since it was the smallest one in terms of tokens to be trained.

The required format for fine-tuning is a sequence of prompt, question and answer: the prompt used is a simple instruction of the chatbot behaviour, while questions and answer were generated differently for each section of the corpus.

For the documents falling in the *Education* section, question-answer pairs were automatically generated by asking gpt-3.5-turbo to describe a specific degree course starting from the corresponding *details* file, and by asking what are the topics of a specific degree course starting from the corresponding *course outline* file. In both cases, the corresponding file was given along with the question, and the given answer was considered as an answer for fine-tuning.

As regards the *Future Students* section, question-answer pairs were extracted directly from the documents already containing a FAQ section, while the other pairs were manually generated. In the second case, a clear question related to a document's section was formulated whenever it was possible, and the answer was either a precise a text or the whole document. Otherwise a generic request was formulated like *parlami di ...*. Some documents were not considered in their entirety since the information contained was highly specific and it was related to non relevant topics.

A validation set was also expunged from the training data by changing questions and/or sampling most important questions. A question-answer pair was randomly picked for each degree course among the *details* and the *course outline* files in the *Education* section. The statistics of the corpus for fine-tuning are reported in Table 2.

Table 2

Number of documents in the unipa-corpus-clear re-arranged for fine-tuning.

	<i>Education</i>	<i>Future Students</i>	Total
Training set	506	269	775
Validation set	253	133	386

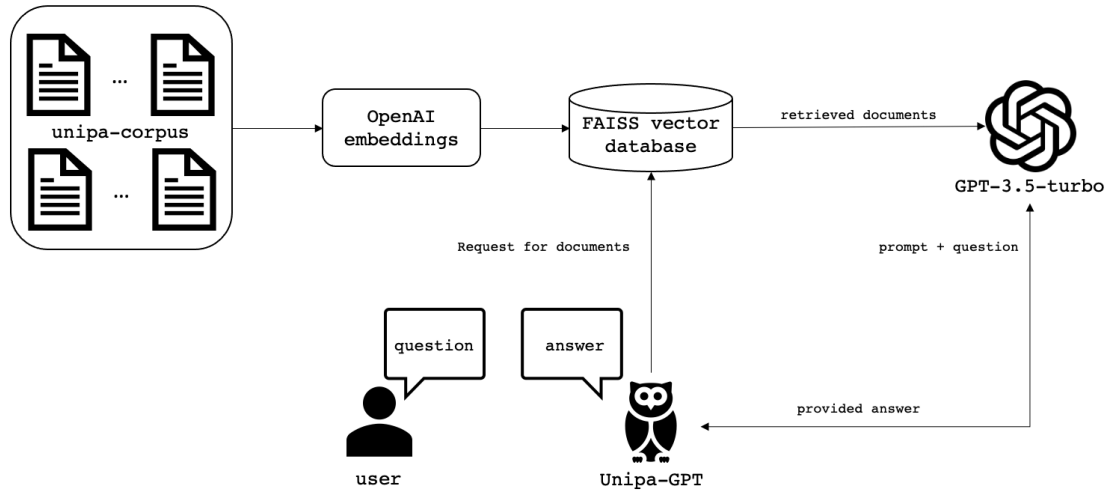


Figure 1: The architecture of the RAG version of Unipa-GPT.

3. System architecture

Unipa-GPT is developed as a RAG architecture [5] made up of two main components, as shown in Figure 1: the retriever and the generator module.

The retriever module consists of a vector database provided by the LangChain library⁴, which makes use of the Facebook AI Similarity Search (FAISS) library [7]. The vector database is filled with the documents in unipa-corpus conveniently divided into chunks of 1000 tokens with an overlap of 50 tokens whose embeddings were extracted using text-embedding-ada-002 by OpenAI⁵.

The generator module consists of an instance of gpt-3.5-turbo [8], a generative Large Language Model based on Transformers [9]. The LLM is inferred with a custom prompt in which the behaviour of the system is explained, and the question of the user is passed along with the most related documents. The expected answer of the system is a reply to the question of the user according to the prompt, that rules the Unipa-GPT overall behaviour, and the domain-knowledge given by the retriever. gpt-3.5-turbo makes inferences using a temperature hyperparameter equal to 0 thus its behaviour is as much deterministic as possible, and the system is prevented to be creative; finally no limits a priori were put on the maximum tokens available for the answer,

⁴https://python.langchain.com/docs/get_started/introduction

⁵<https://openai.com/blog/new-and-improved-embedding-model>

Table 3

The prompts used in the system.

<i>custom prompt</i>	Sei unipa-gpt, il chatbot e assistente virtuale dell'Università degli Studi di Palermo. Rispondi cordialmente e in forma colloquiale alle domande che ti vengono poste. Se ricevi un saluto, rispondi salutando e presentandoti. Se ricevi una domanda riguardante l'università degli studi di Palermo, rispondi in base ai documenti che ti vengono dati insieme alla domanda. Se non sai rispondere, scusati e suggerisci di consultare il sito web, non inventare risposte. Question: {question} Documenti: {context}
<i>condensed prompt</i>	Data la seguente conversazione e la domanda successiva, riformula la domanda successiva in modo tale sia una domanda singola. Conversazione: {chat_history} Domanda successiva: {question} Domanda singola:

in order to prevent broken answers. The chatbot behaviour was implemented via LangChain to keep the chat history and simulate the ChatGPT behavior via the API call of `gpt-3.5-turbo`.

The usage of `gpt-3.5-turbo` and `text-embedding-ada-002` was made via Azure call to the OpenAI API and for `gpt-3.5-turbo` two type of prompting were made, a *custom prompt* and a *condensed prompt*, both in Italian, as shown in Table 3. *Custom prompt* is the explanation of the behaviour of the chatbot where both the previous conversation and the new question are concatenated to the prompt itself. On the contrary, the *condensed prompt* adds to the *custom prompt* another instruction to condense the previous conversation and re-arrange it as a new single question that will be answered accordingly to the *custom prompt*.

In addition to the RAG version illustrated above, a fine-tuned version was implemented with a custom fine-tuned version of `gpt-3.5-turbo` where the `unipa-corpus` explicitly re-arranged, as described in Section 2.2, was used. The same prompt instances mentioned above were used on the fine-tuned model, and the also the RAG architecture was used to compensate for the reduced tokens used in this step. The whole system is shown in Figure 2.

4. Experimental results

The results of the experiments done with both versions of Unipa-GPT are reported in this section along with the main considerations on their performances. Since the chatbot was designed to answer questions posed by secondary school students, we collected the answers from a couple of Italian chats that had not been showed in advance to the system, and were issued by two secondary school students, during an Open Day activity at the University of Palermo. The students were asked to chat with the demo version of Unipa-GPT, and to express a qualitative evaluation of the received answers. The argument of the questions was the same for both students to collect their qualitative judgement on the Unipa-GPT performance.

The first chat (Chat1) consists of 14 questions about subscription procedure, courses at the Department of Engineering, suggestions for the Bachelor course to apply to access a specific Master course, information about University admission test and questions about classes in the

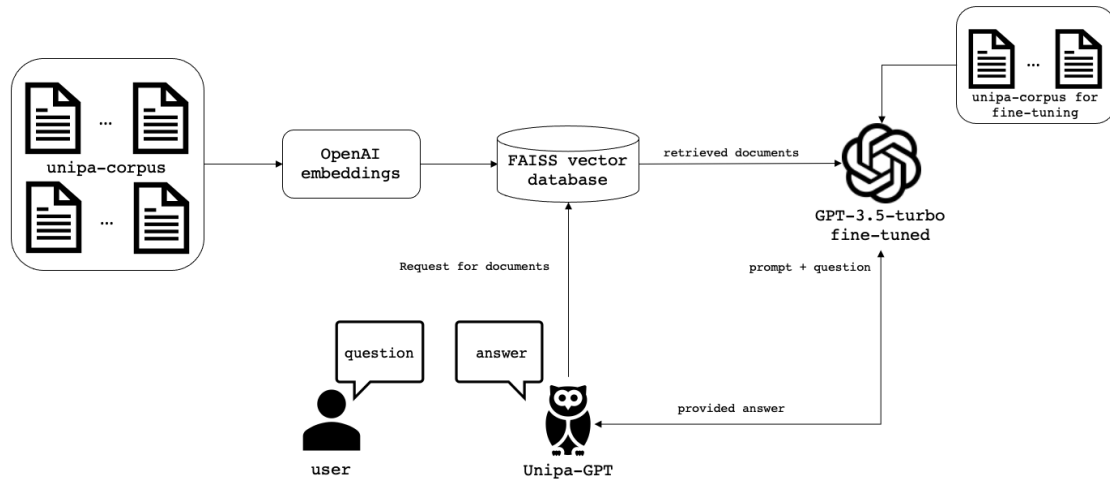


Figure 2: The schema illustrates the architecture of the proposed system where gpt-3.5-turbo-fine-tuned is used after a fine-tuning process over unipa-corpus.

Table 4

Evaluation results of the different runs for both Chat1 and Chat2. Here and in the next tables FT stands for “fine-tuned”.

<i>Custom prompt</i>				<i>Custom prompt FT</i>			
	excellent	good	bad		excellent	good	bad
Chat1-emb	6	1	7	Chat1-emb	4	6	4
Chat1-full	6	4	4	Chat1-full	4	4	6
Chat1-clear	8	4	2	Chat1-clear	6	4	4
Chat2-emb	4	2	4	Chat2-emb	5	1	4
Chat2-full	3	3	4	Chat2-full	6	3	1
Chat2-clear	8	1	1	Chat2-clear	6	1	3
<i>Condensed prompt</i>				<i>Condensed prompt FT</i>			
	excellent	good	bad		excellent	good	bad
Chat1-emb	5	1	8	Chat1-emb	5	2	7
Chat1-full	10	2	2	Chat1-full	4	1	10
Chat1-clear	10	1	3	Chat1-clear	6	0	8
Chat2-emb	6	2	2	Chat2-emb	5	1	4
Chat2-full	8	2	0	Chat2-full	6	1	3
Chat2-clear	9	0	1	Chat2-clear	8	0	2

first year. Conversely, the second one (Chat2) contains 10 questions with typos and grammatical errors, and they are expressed more like a Google search or a list of keywords. The involved topics in the second chat are information about taxes, Erasmus programs, scholarships and Student Desk.

Both Chat1 and Chat2 were issued using with different configurations of the system as regards the corpus used for retrieval clear, full, emb, the prompt *custom prompt*, *condensed prompt*, and the inferred LLM model gpt-3.5-turbo, gpt-3.5-turbo-fine-tuned, for a total of 12 runs. Each student was asked to simply classify the each answer as excellent, good and

Table 5

An extract of Chat2 is reported where the same question is made to systems that use *condensed prompt* and *unipa-corporus-emb* as retrieval corpus. Note that the output of the second model is not entirely reported due its length: in the missing text the model continued generating other cases of taxes reduction.

<i>condensed-emb</i>	<i>condensed-FT-emb</i>
cos'è il contributo onnicomprensivo	
<p>Il contributo onnicomprensivo per l'immatricolazione all'Università degli Studi di Palermo per l'anno accademico 2023/24 è calcolato in base alla classe di reddito ISEE dello studente e alle prestazioni erogate nell'ambito del diritto allo studio universitario. Il pagamento delle tasse e del contributo onnicomprensivo è suddiviso in tre rate, la prima delle quali deve essere pagata entro il 30 settembre 2023 per l'immatricolazione ai corsi di studio ad accesso libero e per le iscrizioni ad anni successivi al primo. Per conoscere i termini e gli importi del pagamento delle tasse e del contributo onnicomprensivo, si consiglia di consultare il Calendario didattico e il Regolamento sulla contribuzione studentesca vigenti. In caso di dubbi o domande specifiche, si suggerisce di contattare direttamente la Segreteria Studenti dell'Università degli Studi di Palermo.</p>	<p>Il contributo onnicomprensivo è il pagamento delle tasse universitarie che gli studenti devono effettuare all'atto dell'immatricolazione o iscrizione. Il pagamento è suddiviso in tre rate: la prima rata, da versare all'atto dell'immatricolazione o iscrizione, nella misura del 40% dell'importo dovuto; la seconda rata, da versare entro il 29 dicembre 2023, nella misura del 30% dell'importo dovuto; la terza rata, da versare entro il 29 febbraio 2024, nella misura del 30% dell'importo dovuto. Gli studenti che si iscrivono per la prima volta alla Laurea triennale o Laurea Magistrale a Ciclo Unico, diplomatisi con voto compreso tra 95 e 100, usufruiscono, per il primo anno, di una riduzione del 25% fino ad un massimo di 350,00 euro. Coloro che si diplomano con voto pari a 100 e lode usufruiscono, per il primo anno, di una riduzione del 50% fino ad un massimo di 500,00 euro.</p> <p>[...]</p>

bad. Good answers are not optimal and precise ones but they contain a partially good answer, and may add non requested pieces of information that are less related to the initial question. Bad answers are either answers that provide wrong information or out of topic replies along with any unexpected/undesired behaviour of the system. Also the replies to questions that the system can not answer are considered bad. In both chats there are questions that are strictly related with the previous answers from the system, where the user asks for further information, e.g. if the user asks for the university admission test for a particular Bachelor course, and system replies that the course is a free admission one, the (expected) question regarding test rules is not applicable thus it is not posed to the system, and the missed reply is considered as bad. Table 4 shows the evaluation results.

Results show that the best models are the ones that use *condensed prompt* with *gpt-3.5-turbo* and *unipa-corporus-clear* as retrieval corpus. The lack of additional information regarding classes' educational objectives in the retrieval corpus, makes it more precise. On the other hand, the use of the *condensed prompt* produces a less degrading conversation compared with the use of the *custom prompt* where the conversation is passed as it is, without refinements. Overall, there are not models that performs bad for the entire conversation.

An unexpected result comes from the comparison between non fine-tuned and fine-tuned models that do not exhibit the expected improvement. Particularly, a high decrease in performance is observed in the *condensed prompt* configurations. The answers of fine-tuned models are generally very long and they tend to generate loops and hallucinations [10]. On one hand, this behaviour may be caused by the training corpus, where an entire document is given as answer. On the other hand, LLMs tend to be repetitive [11]. Answers from fine-tuned models tend also to add extra information that sometimes are pertinent, e.g. the question regarding taxes suggest the possibility for payments reduction as reported in Table 5. In other cases answers are completely out of topic or the chatbot contradicts itself, as shown in Table 6.

With reference to the prompts, *condensed prompt* performs better and the systems that uses *unipa-corporus-full* and *unipa-corporus-clear* as retrieval corpora are the best: they provide very pertinent and similar answers to each other. On the contrary the models with the *custom prompt* tend to be less precise, as reported in Table 7.

Table 6

An extract of Chat1 is reported where the same question is made to systems that use the *custom prompt*. Note that the fine-tuned model contradicts itself.

<i>custom-full</i>	<i>custom-FT-clear</i>
la triennale in ingegneria meccanica è a numero chiuso o a numero aperto?	
La triennale in ingegneria meccanica presso l'Università degli Studi di Palermo è a numero aperto, con un numero di posti in programmazione locale pari a 200. Questo è indicato nel documento relativo all'anno accademico 2023/2024 per il curriculum meccanico. Se hai bisogno di ulteriori informazioni, ti suggerisco di consultare il sito web dell'università	La triennale in ingegneria meccanica è a numero aperto.
devo fare test per entrare alla triennale in ingegneria meccanica?	
Si, per accedere alla triennale in ingegneria meccanica all'Università degli Studi di Palermo è necessario sostenere un test di accesso. Il Centro Orientamento e Tutorato (COT) dell'Università organizza diverse iniziative per affrontare al meglio i test di accesso e superarli con successo, come i workshop "Come affrontare il test di accesso" e la simulazione dei test di accesso. Inoltre, sono disponibili corsi di preparazione alle prove di accesso che inizieranno il 17 luglio 2023 e si svolgeranno esclusivamente online sulla piattaforma Microsoft Teams. Per maggiori informazioni, ti consiglio di consultare i documenti forniti insieme alla domanda.	Si, sono previsti test per entrare alla triennale in ingegneria Meccanica.

Table 7

An extract of Chat1 is reported where the same question is made to system that use *unipa-corpus-clear* as retrieval corpus. Both the answers are correct but the one reported by the RAG-only model is more precise.

<i>condensed-clear</i>	<i>custom-FT-clear</i>
quale corso triennale mi consigli di frequentare per poi passare alla magistrale di aerospaziale?	
Per passare alla laurea magistrale in Ingegneria Aerospaziale all'Università degli Studi di Palermo, il corso triennale consigliato è la Laurea in Ingegneria Meccanica con curriculum Aeronautico. Ti consiglio di consultare il piano di studi del corso di laurea triennale e di contattare il Dipartimento di Ingegneria per ulteriori informazioni.	Ti consiglio di frequentare il corso di Laurea triennale in Ingegneria Meccanica per poi passare alla magistrale in Ingegneria Aerospaziale. Il corso triennale ti fornirà le basi necessarie per affrontare al meglio il percorso magistrale.

All models share the same negative aspects: all the provided answers are well written in Italian⁶ and an uninformed user can easily consider a wrong answer as truthfully by looking at the lexicon and the grammar used. During a conversation, due to the used prompt, the system often answers with an out of context presentation of Unipa-GPT, and a repetition of the question is necessary to obtain the expected answer. Prompt tuning plays a key role in such applications [8, 12] and the reported results are the optimal ones after several tests with different prompt. A test with a English-written prompt was performed but it did not show significant improvements.

Unipa-GPT performs in free Q&A domain where the user can be anyone with different background, thus questions submitted to the system are deeply variable; we found that developed models results more sensible and perform better or worse depending on the specific question. In Chat2, questions for information regarding scholarships and classes, were made as a request and not as a question. This led the system either to answer badly or not to provide answer at all by making its excuses, as reported in Table 8. Questions regarding topics for which knowledge was made available to the system, that are not reported in question format or are written vaguely, limit the capabilities of the system. In Table 8 is reported also the behaviour of the fine-tuned version of gpt-3.5-turbo with the same prompt used in training phase and without the

⁶All the reported experiments where done in Italian, some tests were done also in English but not further analyzed since they were not the primary interest.

Table 8

An extract of Chat2 is reported where the same questions are posed to systems that use the general `gpt-3.5-turbo` and *condensed prompt*, system on the left retrieves documents from the `unipa-corpus-emb`, while the one at the centre uses `unipa-corpus-full`. At the right, the same questions were proposed to the fine-tuned version of `gpt-3.5-turbo`: the prompt given is the same used in training phase and no access was provided to `unipa-corpus`. No fine-tuned systems fails with the no-question but provide the correct answer what the same request is made in a question-like form. Note that the provided answer is not strictly correct since both refer to the first year of medicine MEDIT curriculum. On the other hand, the pure fine-tuned system provides in both cases a credible, but false, answer.

<i>condensed-emb</i>	<i>condensed-full</i>	<i>pure fine-tuned model</i>
lista esami primo anno di medicina		
Gli esami del primo anno del corso di medicina all'Università degli Studi di Palermo per l'anno accademico 2023/24 non sono specificati nei documenti forniti. Ti suggerisco di consultare i bandi di concorso per maggiori informazioni. Inoltre, per quanto riguarda l'immatricolazione, i corsi di studio ad accesso libero possono essere immatricolati senza mora dall'1 agosto [...]	Mi dispiace, ma non ho trovato informazioni riguardanti gli esami del primo anno del corso di medicina all'Università degli Studi di Palermo per l'anno accademico 2023/24. Ti consiglio di consultare i bandi di concorso per maggiori informazioni.	La lista degli esami del primo anno del corso di laurea in Medicina e Chirurgia presso l'Università di Palermo prevede i seguenti esami: Anatomia Umana con Esercitazioni, Biologia Applicata, Lingua Inglese, Chimica Biologica con Elementi di Biologia Molecolare, Fisica, Matematica, Statistica e Informatica, Istologia Umana con Elementi di Biologia Cellulare, Fisiologia Umana con Esercitazioni, e Patologia Generale.
quali sono le materie del primo anno di medicina		
Le materie del primo anno del corso di medicina all'Università degli Studi di Palermo per l'anno accademico 2023/24 sono: Fisica I, Matematica I e Chimica Generale ed Organica C.I. Invece, per il corso di laurea in Scienze Biologiche, le materie del primo anno sono: Matematica con esercitazioni, Citologia e Istologia con esercitazioni, Chimica Generale con esercitazioni e Botanica Generale sistematica con esercitazioni.	Le materie del primo anno del corso di medicina all'Università degli Studi di Palermo per l'anno accademico 2023/24 sono: Fisica I, Matematica I, Chimica Generale ed Organica C.I., Chimica Organica, Biologia, Embriologia ed Istologia C.I., Genetica e Microbiologia C.I., Informatica, Bioinformatica e Inglese Tecnico Scientifico C.I. Ti consiglio di consultare i documenti forniti per maggiori informazioni.	Le materie del primo anno di medicina sono: Anatomia Umana, Istologia, Embriologia, Biochimica, Biologia Cellulare, Genetica, Fisica, Chimica, Lingua Inglese, Metodologia del Lavoro Intellettuale, e Cultura e Valori in Sanità.

access to `unipa-corpus`, to test the effectiveness of fine-tuning: the fine-tuned model provides an answer to both the proposed questions (even to the no-question) but both are wrong, despite they are credible.

The last consideration regards the corpus, in particular the meta-structure of the domain-knowledge. Documents were generated and provided to the vector store without any type of organization, but University documents are likely to be organized in a graph or ontology structure that may provide domain-specific knowledge which is not explicitly reported in the documents. We think that providing the system with a domain-specific ontology coupled to the corpus may lead to better performances and mitigate hallucination e.g. when the system is asked for a Bachelor course, while the answers are referred to the Master course with the same name.

5. Conclusions and future works

In this paper we presented Unipa-GPT, a virtual assistant capable of answering to secondary school students who need information for accessing to and studying at the University of Palermo. The developed system relies on a RAG architecture that uses documents from a corpus purposely scraped from the University institutional website, and comes in two versions that make use of either `gpt-3.5-turbo` or a fine-tuned model `gpt-3.5-turbo-fine-tuned` where the corpus has been reduced to keep the computational resources needed for fine-

tuning low. Significant improvements were not found in the fine-tuned model, and the best performing system was the one that uses the so called *condensed prompt* where the previous conversation and the next question are reformulated to be a unique question. Such a prompt induces gpt-3.5-turbo to summarize the conversation at each question, and then it behaves as instructed using our plain *custom prompt* tailored for the application purposes. Moreover, this system uses the unipa-corpus-clear for retrieval where educational objectives of each class are not reported; we argue that this light version of the corpus provides the information to the LLM in a more compact and precise way, thus generating best answers. Further developments of the systems will cover prompt-tuning, adaptive corpus selection corpus with the integration of a suitable domain ontology, and the development of Unipa-GPT versions that share the same architecture but use different LLMs.

Acknowledgments

We thank all the CHILab team, in particular Dr. Salvatore Contino and PhD students Luca Cruciata, Gaetano Pottino and Paolo Sortino, that contributed in generating the corpus with the scraping and implemented the demonstration interface. This work is supported by the PO FESR 2014-2020 grant n. 086201000543, “SCuSi - Smart Culture in Sicily”.

References

- [1] E. Bassignana, D. Brunato, M. Polignano, A. Ramponi, Preface to the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI), in: Proceedings of the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023), 2023.
- [2] Y. K. Dwivedi, N. Kshetri, L. Hughes, E. L. Slade, A. Jeyaraj, A. K. Kar, A. M. Baabdullah, A. Koochang, V. Raghavan, M. Ahuja, et al., “so what if chatgpt wrote it?” multidisciplinary perspectives on opportunities, challenges and implications of generative conversational ai for research, practice and policy, *International Journal of Information Management* 71 (2023) 102642.
- [3] A. Borji, A categorical archive of chatgpt failures, 2023. [arXiv:2302.03494](https://arxiv.org/abs/2302.03494).
- [4] P. P. Ray, Chatgpt: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope, *Internet of Things and Cyber-Physical Systems* (2023).
- [5] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive nlp tasks, *Advances in Neural Information Processing Systems* 33 (2020) 9459–9474.
- [6] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al., Llama 2: Open foundation and fine-tuned chat models, *arXiv preprint arXiv:2307.09288* (2023).
- [7] J. Johnson, M. Douze, H. Jégou, Billion-scale similarity search with GPUs, *IEEE Transactions on Big Data* 7 (2019) 535–547.

- [8] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [10] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, P. Fung, Survey of hallucination in natural language generation, *ACM Comput. Surv.* 55 (2023). URL: <https://doi.org/10.1145/3571730>. doi:10.1145/3571730.
- [11] A. Holtzman, J. Buys, L. Du, M. Forbes, Y. Choi, The curious case of neural text degeneration, *arXiv preprint arXiv:1904.09751* (2019).
- [12] Z. Zhao, E. Wallace, S. Feng, D. Klein, S. Singh, Calibrate before use: Improving few-shot performance of language models, in: M. Meila, T. Zhang (Eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 12697–12706. URL: <https://proceedings.mlr.press/v139/zhao21c.html>.