








Jill Watson: A Virtual Teaching Assistant Powered by ChatGPT

Karan Taneja^(✉) , Pratyusha Maiti^(✉) , Sandeep Kakar ,
Pranav Guruprasad , Sanjeev Rao, and Ashok K. Goel 

Georgia Institute of Technology, Atlanta, GA, USA
{ktaneja6, pmaiti6, skakar6, pguruprasad7, srao373, ag25}@gatech.edu

Abstract. Conversational AI agents often require extensive datasets for training that are not publicly released, are limited to social chit-chat or handling a specific domain, and may not be easily extended to accommodate the latest advances in AI technologies. This paper introduces Jill Watson, a conversational Virtual Teaching Assistant (VTA) leveraging the capabilities of ChatGPT. Jill Watson based on ChatGPT requires no prior training and uses a modular design to allow the integration of new APIs using a skill-based architecture inspired by XiaoIce. Jill Watson is also well-suited for intelligent textbooks as it can process and converse using multiple large documents. We exclusively utilize publicly available resources for reproducibility and extensibility. Comparative analysis shows that our system outperforms the legacy knowledge-based Jill Watson as well as the OpenAI Assistants service. We employ many safety measures that reduce instances of hallucinations and toxicity. The paper also includes real-world examples from a classroom setting that demonstrate different features of Jill Watson and its effectiveness.

Keywords: Virtual Teaching Assistant · Intelligent Textbooks · Conversational Agents · Question Answering · Modular AI Design

1 Introduction

Conversational AI agents can be powerful tools for education as they enable continuous 24×7 support and instant responses to student queries without increasing the workload for instructors. These virtual teaching assistants can help in efficiently scaling quality education in terms of both time and cost. The interactive nature of conversational AI agents allows students to be more inquisitive and increases teaching presence by resembling one-on-one tutoring. Based on the Community of Inquiry framework [6], teaching presence through instructional management and direct instruction leads to an increase in student engagement and retention. Towards this end, we developed Jill Watson, a Virtual Teaching Assistant (VTA) powered by ChatGPT for online classrooms which answers students' queries based on course material such as slides, notes, and syllabic.

In previous work, the legacy Jill Watson [5, 8] (henceforth LJW) is a question-answering system for course logistics and uses a two-dimensional database of

information organized by course deliverables (assignments, exams, etc.) and information categories (submission policy, deadline, etc.). It also uses a list of FAQs for course-level information such as ethics and grading policies. In this paper, we introduce the new Jill Watson which is conversational and answers questions related to course logistics as well as course content based on multiple large documents provided as context.

ChatGPT or GPT-3.5, based on GPT-3 [3], is a powerful large language model (LLM) trained to follow instructions and hold a dialogue. It is capable of *attending* to a large context and constructing meaningful text in response to user inputs. Many conversational systems such as HuggingGPT [21], Microsoft Bing Chat (www.bing.com/chat), and LangChain (www.langchain.com) based systems leverage ChatGPT for performing context-aware response generation and zero-shot learning. ChatGPT and other LLMs suffer from hallucination i.e. they generate text that can be inconsistent or unverifiable with the source text, or absurd in a given context [9]. While hallucination is useful in creative tasks such as story writing, it is detrimental in information-seeking tasks such as those in the domain of education. ChatGPT and other LLMs also have safety issues as they generate text that may be considered toxic or inappropriate [25].

This work introduces Jill Watson’s architecture which does *not* require any model training or fine-tuning and is designed to address the above LLM-related concerns. We address the hallucination issue by citing the documents from which information is obtained and verifying grounding using textual entailment. To prevent Jill Watson from answering unsafe questions or generating unsafe responses, we employ a classifier for question relevance, toxic text filters, and prompts that promote politeness in response generation. Further, Jill Watson is designed to answer questions based on multiple large documents which makes it well-suited for intelligent textbooks. We only rely on publicly available resources to promote future research in this direction.

The paper has three main contributions: (i) we introduce Jill Watson, a virtual teaching assistant powered by ChatGPT with a skill-based architecture, (ii) detail all the different modules of Jill Watson and associated algorithms, and (iii) quantitatively evaluate Jill Watson to measure response quality and safety, along with a discussion on examples from our first deployment.

Section 2 discusses Jill Watson in the context of related work. Section 3 describes the architecture and each module in detail. Section 4 describes our experimental results comparing Jill Watson to two strong baselines in terms of response quality and safety along with examples (see Table 3) from our first deployment. We conclude the paper in Sect. 5 with a summary of the strengths, limitations, and potential impact of Jill Watson.

2 Related Work

Question Answering can either be open-ended or grounded in knowledge. Without a knowledge source, question-answering models based on LLMs [16, 22] are expected to store the information in their parameters during the training.

In grounded question answering, previous work has explored different types of contexts including the web [17], machine reading comprehension [1], knowledge bases [2], and short text documents [18, 27]. Some methods assume access to the correct context from the document [18]. Further, many methods require training with datasets that are expensive to collect and do not generalize well [24, 26]. Jill Watson neither imposes a limit on the document size nor requires a training dataset. It pre-processes large documents and answers incoming questions based on passages retrieved using dense passage retrieval (DPR) [11].

Retrieval Augmented Generation (RAG) is a well-known method [14] for increasing the reliability of LLMs by generating text conditioned on source texts that are retrieved based on a query. Knowledge-grounded generative models have two main goals: factuality and conversationality [19]. Factuality minimizes hallucinations by ensuring consistency of output with the retrieved texts while conversationality refers to relevance of the information to the query and generation without repetition. Previous work has shown improved factuality using RAG in dialog response generation task to remain consistent with a persona [23], knowledge-grounded generation [13, 19, 26] and machine translation [4].

Many models use large training datasets to learn question answering from contexts [2, 24, 26]. On the other hand, WikiChat [19] uses seven-step few-shot based question answering system which uses both retrieval and open-ended generation to answer questions using Wikipedia. Re²G or Retrieve, Re-rank, Generate [7] also uses retrieval for generating outputs but uses an additional re-ranking step to score retrieved passages before generation. Further, it can also be trained end-to-end after initial fine-tuning. Jill Watson solves the knowledge-intensive generation problem using RAG but without any fine-tuning by using open-source DPR models for retrieval and using ChatGPT to construct responses. Because of clever prompting and indexing, it is also able to refer to the document and page number from which the response was generated.

Safety in LLMs is important to avoid harm because of hateful, offensive, or toxic text. Previous work on evaluating the toxicity of ChatGPT has found that assigning personas, using non-English languages, prompting with creative tasks, jailbreak prompts, and higher temperature values can all lead to more toxic responses [25]. Perspective API [12] and OpenAI Moderation API [15] are popular services to measure toxicity in various categories including hateful content, violence, etc. Jill Watson ensures safety in three ways: (i) OpenAI Moderation API for both user inputs and its own responses, (ii) skill classifier to identify irrelevant queries, and (iii) encourages polite responses in prompts to ChatGPT.

Dialog Systems are AI agents designed for human-like conversations, typically using hybrid architectures involving both rule-based systems and machine learning systems. For instance, MILABOT [20] uses rule-based and generative models along with a response selection policy trained using reinforcement learning. Microsoft XiaoIce [28] uses a skill-based architecture where the chat manager selects one of 230 high-level skills to generate responses. To the best of our knowledge, such powerful multi-turn dialog systems have not been introduced in a classroom. The legacy Jill Watson [8] is a single-turn question-answering

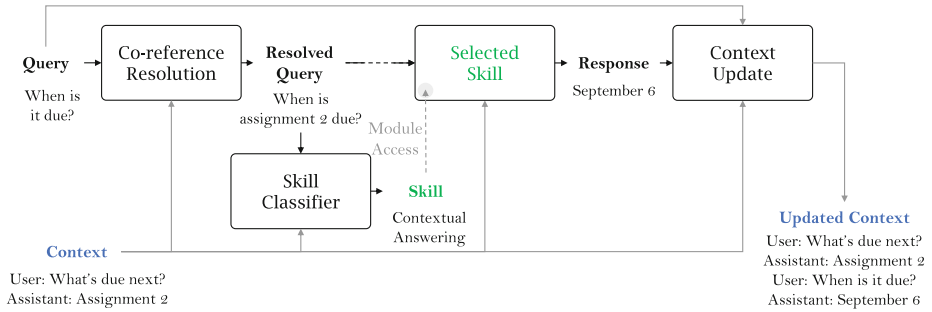


Fig. 1. *Architecture of Jill Watson:* After the coreference resolution of an incoming query, the skill classifier is used to find the most appropriate skill for response generation. Jill Watson’s skills include Contextual Answering, Greetings, etc. The updated conversation history is used as context for generating responses in the future.

system for course logistics and policies. OpenAI Assistants service¹ provides a way to instantiate a ChatGPT-based agent to generate responses using text documents. Inspired by XiaoIce, Jill Watson has a skill-based architecture and is designed to be a safe conversational agent for classrooms that can answer student queries related to course logistics as well as course content using ChatGPT in the backend. It is also well-suited for other applications in education like intelligent textbooks given its ability to use long documents as context.

3 Jill Watson Architecture

The architecture of Jill Watson shown in Fig. 1 takes inspiration from the skill-based architecture of XiaoIce [28]. XiaoIce relies on different skills such as task completion, image commenting, content creation, etc. to interact with users and selects the appropriate skill for each conversation turn based on the previous context. In Jill Watson, the query with resolved coreferences is used to decide the most appropriate skill for answering the incoming query. The skill-based design of Jill Watson makes it extensible as we can easily plug in new API services and other capabilities in the future in the form of new skills.

Contextual Answering skill is responsible for answering questions where content needs to be retrieved from course content or syllabus and *Self-awareness* skill answers queries about Jill Watson itself. As we will discuss later, these two skills make Jill Watson a knowledge-grounded AI agent with the ability to refer to multiple documents and cite relevant content in its answers. Further, ChatGPT allows Jill Watson to be conversational by using past messages as context in generating responses to user messages. We also utilize many safety features in Jill Watson which include detecting irrelevant queries by skill classifier, moderation filters, and prompts to encourage courteous responses.

¹ <https://platform.openai.com/docs/assistants/overview>.

3.1 Coreference Resolution

Coreference resolution involves determining the entities that are indirectly referenced in a text and making them explicit using nouns or noun phrases. For example, given the context ‘John started reading a book’, the query ‘When did he start?’ has two coreferences. An *explicit coreference* is suggested by the word ‘he’ while an *implicit coreference* to an event is present because of ‘When’ and ‘start’. In coreference resolution, we must resolve the reference ‘he’ and the event that ‘When’ and ‘start’ are referring to. Therefore, the resolved query would be ‘When did John start reading the book?’ formed by replacing ‘he’ with John and adding the implicit event ‘reading the book’.

While ChatGPT implicitly resolves the coreferences as it can construct appropriate responses by itself, since we wish to use existing models for retrieval without any fine-tuning, we need to construct complete queries with resolved coreferences before passing them to the retrieval module. Hence, the first step in Jill Watson is to resolve coreferences in the user query based on the previous messages. In the example in Fig. 1, ‘it’ in ‘When is *it* due?’ is replaced by the entity ‘Assignment 2’. This is done by prompting ChatGPT to resolve the coreferences in the received query and passing the past messages as context. We use a combined instruction and demonstration-based prompt where we explain the task (instruction) along with three demonstrations with no coreferences, an explicit and an implicit coreference. In our investigation, we found demonstrations to be extremely useful for improving performance, especially on implicit coreferences which are much harder to identify.

3.2 Skill Classifier

As discussed earlier, Jill Watson uses various *skills* to answer different types of queries. For instance, queries that require retrieval are forwarded to the *Contextual Answering Skill* while greetings are answered by the *Greetings Skill*. The skill-based division using a skill classifier allows us to use different response-generation techniques based on the user query. It can also aid in understanding user behaviors by analyzing the skill distribution of student queries.

To forward a user query to the appropriate skill, the resolved query (after coreference resolution) is used to perform *skill classification* by prompting ChatGPT. In the example in Fig. 1, the selected skill is *Contextual Answering* based on the resolved query and previous messages as context. We again used a combined instruction and demonstration-based prompt with an explanation of each skill (instruction) and one demonstration per skill. We found that fewer and distinct classes lead to a better performance which motivates the use of a small number of skills as far as possible.

3.3 Contextual Answering

Contextual answering or context-based question answering skill involves answering questions based on the given information. For Jill Watson, this information

Algorithm 1. Contextual Answering Skill

Require: Resolved query Q , context C , PRE-PROCESSED DOCUMENTS with passages \mathcal{P} and corresponding context embeddings \mathcal{D} , DPR query encoder $E_q(\cdot)$

Ensure: Response R , confidence $C \in \{\text{low}, \text{high}\}$

// DENSE PASSAGE RETRIEVAL

Construct a query with context Q_C with Q and C .

Sort passages \mathcal{P} using cosine similarity between $e_{Q,C} = E_q(Q_C)$ and context embeddings in \mathcal{D} . Keep top-20 passages P with highest cosine similarity.

// CONTEXT LOOP

for batches P_5 of 5 passages in top-20 **do**

 // RESPONSE GENERATION

 Generate response R by prompting ChatGPT with C , Q , and P_5 .

if R answers Q **then**

 // TEXTUAL ENTAILMENT

if “ P_5 implies R ” succeeds: **return** R and $C = \text{high}$.

else: **return** R and $C = \text{low}$. **end if**

end if

end for

return R and $C = \text{low}$.

consists of verified course documents provided by course instructors. The process outlined in Algorithm 1 can be divided into five main parts, highlighted in SMALL CAPS, viz. documents pre-processing, DPR, response generation, textual entailment and context loop. The *documents pre-processing* step is performed only once when a Jill Watson is initialized with a set of course documents. The remaining four steps have to be performed for every query received by the Contextual Answering skill.

Documents Pre-processing: Jill Watson pre-processes course documents used for answering student queries and stores them as a list of passages along with their different representations discussed below. These representations allow fast retrieval of the most relevant parts of documents during run-time. It accepts PDF documents, the most common format in which course contents (syllabus, notes, books, and slides) are distributed. After text extraction from each PDF document using Adobe PDF Extract API, it is divided into pages and each page is further divided into paragraphs. We group paragraphs into *passages* of at least 500 characters (90–100 words). This ensures a significant context size in each passage for the DPR step. We also store document and page information with each passage which is used to refer back and cite the documents. Further, there is a 50% overlap between passages for added redundancy and to represent continuity between consecutive passages.

Figure 2 shows different representations stored for each passage. Since we use DPR [11] in the retrieval step (discussed in more detail later), we pre-compute embeddings of each passage using the DPR context encoder. The context encoder requires passages with prepended headings which are obtained by prompting ChatGPT to generate a 2–3 word heading. We use a cleaned version of the

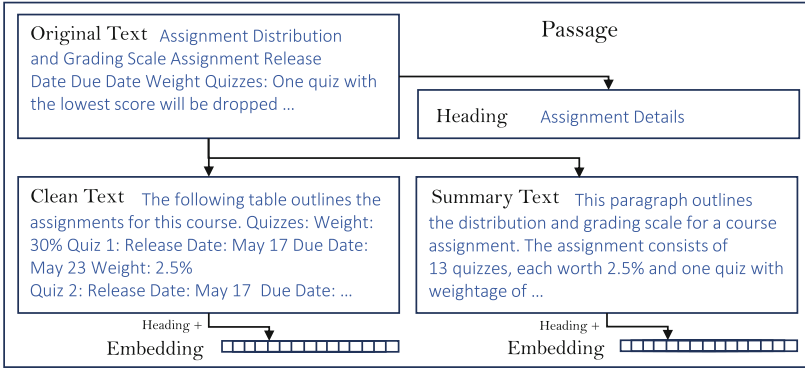


Fig. 2. *Passage representation* consists of the original text, heading, clean text, summary text, and context embeddings of both clean and summary texts.

original text for the context encoder model because the raw text from the PDF has unwanted spaces, special characters, and poor table formatting. To obtain this clean text, we again prompt ChatGPT with the original text and instruct it to clean the text and to format the tables. Further, we also prompt ChatGPT to summarize the clean text, and we store a context embedding of this passage summary. We found that this is useful for retrieval because passages can contain implicit information which is made more explicit in summaries. For example, an exam might be mentioned in a different line from its deadline but a summary makes this relation more explicit and direct.

Dense Passage Retrieval: DPR [11] has a dual-encoder architecture i.e. it consists of a query encoder and a context encoder. During training, the two models are aligned to output in a common embedding space by maximizing the similarity between embeddings of query-context pairs in training data and minimizing the similarity across example pairs. We used the multi-dataset model which is trained on Natural Questions, TriviaQA, WebQuestions, and TREC datasets [11], allowing it to generalize over domains and text properties.

The first step in DPR is to compute query embeddings of the coreference-resolved queries. We found that prepending the question history in the query to provide more context improves retrieval performance. Second, we compute the similarity of the query embedding with context embeddings of both clean and summary text. Third, we sort the passages in decreasing order of similarity and group clean text in batches of five for prompting. For sorting, the similarity score of a passage is the maximum between similarities with clean and summary embeddings. Finally, we also pass the document and page information for each passage to the next step where we prompt ChatGPT to use this information for citing document and page when constructing an answer.

ChatGPT API Call: In Jill Watson, we wish to prompt ChatGPT to construct an answer with four requirements: (i) ChatGPT should only answer student query if the context has the information to answer the query, or else, refuse to

answer, (ii) it should not add any additional information to the answer that is not present in the context, (iii) it should cite the corresponding document and page number when using content from a text chunk, and (iv) it should answer the query with respect to the current date and time.

The prompt to construct an answer using context is given in Fig. 3 (left). Note that we prompt ChatGPT to use only the given context which is made up of a batch of five passages, encourage it to say it doesn't know the answer if the context is insufficient, and ask it to cite the document and page number (see example in Fig. 3). We also feed it the current week (Monday-Sunday), date and time, and past messages which can be useful in answering user queries in the context of the present time and the past conversation.

<p>System: A polite virtual assistant named Jill Watson answers student (user) questions on a class forum. The context below is provided to answer questions. CONTEXT: {context} Only the above context is used and rephrased to answer the question. If context does not answer it, say that you don't know the answer. If answer is found, don't forget to cite the document and page number. CURRENT WEEK: {current-week} CURRENT TIME: {current-time} User and Assistant Messages: {past-messages} {most-recent-query}</p>	<p>USER: When is Mini-Project 2 due? Jill Watson: Mini-Project 2 is due on Monday, September 25, 2023 at 9 am. You need to submit the code for solving the problem to the Mini-Project 2 assignment on Gradescope and also submit a report describing your agent to Canvas. Your grade for this project will be based on a combination of your report's quality (50%) and your agent's performance (50%). <i>Source: Syllabus, Page 13</i></p>
--	--

Fig. 3. Answer generation prompt (left) and an example response with citation (right). The context contains five passages with document name and page numbers.

Textual Entailment: The textual entailment task involves determining the entailment relation between two texts. A text T entails hypothesis H if H is a logical consequence of text T i.e. $T \implies H$. If $T \implies \neg H$, we say that T contradicts H . T can also be *neutral* i.e. it neither entails nor contradicts H . Jill Watson's prompt for answer generation contains instructions to not answer a question when the when the context doesn't provide an answer and to only use the given context for answering. However, we wish to have an additional check to detect hallucinations through textual entailment. Given a context C and an answer A , we wish to check if C entails A . If C doesn't completely entail A , there is information in A that was not retrieved from C . In such a case, we prepend a warning for the user conveying that the confidence in the answer is low and encourage them to check the answer on their own. As shown in Algorithm 1, we check if the context $C = P_5$ is used to generate the answer $A = R$. We prompt ChatGPT with P_5 as the text and R as the hypothesis and ask it to determine if the text implies that hypothesis. We found a simple instruction-based prompt to be most effective with the highest recall for non-entailed answers.

Context Loop: After scoring and sorting the passages, we group the top twenty passages into four batches of five and prompt ChatGPT to generate a response based on the first batch as context. If it fails to answer using the first batch, we use the second batch of passages, and so on until a valid answer is generated. To check if ChatGPT generated a valid answer, we prompt it to classify the response as **NEGATIVE** if it refuses to reply because *the information is not present or it suggests contacting the staff*, and **NEUTRAL** otherwise.

3.4 Other Skills

In addition to contextual answering, we use other skills for additional capabilities and we plan to expand these in the future with software tools and API services.

Self-awareness Skill: Many curious users ask AI agents about itself to check its self-awareness or to know more about the system. For such queries, we prompt ChatGPT to answer the user query based on a textual description of Jill Watson. This textual description contains basic information about us, the team of researchers who built it, and a blurb about its capabilities.

Greeting Skill: If an incoming query is a greeting or conveys gratitude to Jill Watson, we prompt ChatGPT to generate a polite reply.

Irrelevant Skill: If a query doesn't fit into any of the other skills, we use a fixed polite message asking the user to change or rephrase their question.

3.5 Moderation Filter

For deployment in real world, Jill Watson should be safe to use and not accept any harmful requests or generate a harmful response. To this end, we filter input user queries and outputs of Jill Watson using the OpenAI Moderation API [15]. The API allows Jill Watson to detect different categories of unsafe text including hate, hate and threatening, harassment, harassment and threatening, self-harm, self-harm intention, self-harm instructions, sexual, sexual involving minors, and violence with graphic depictions.

4 Results

We compare Jill Watson with both legacy Jill Watson (**LJW**), and the OpenAI Assistants service (**OAI-Assist**). LJW baseline employs an intent classifier and a database of information organized by course deliverables and information categories as well as a list of FAQs. OAI-Assist allows users to upload files PDF files and employs retrieval to generate answers for user queries. Our new system **Jill Watson** uses coreference resolution, skill classification, dual encoder DPR, ChatGPT for generation, and safety features described earlier. Both OAI-Assist and Jill Watson use 'gpt-3.5-turbo-1106' for retrieval-augmented generation in our experiments.

Table 1. Response Quality: A set of 150 questions is used to evaluate the response quality of each system. Failures are further determined to be harmful, confusing, and stemming from poor retrieval.

Method	Pass	Failures		
		Harmful	Confusing	Retrieval
LJW	26.0%	–	60.4%	–
OAI-Assist	31.3%	16.5%	72.8%	68.0%
Jill Watson	76.7%	5.7%	62.8%	57.1%

Response Quality and Harmful Errors: We used a set of 150 questions created by four students based on the syllabus, e-book, and video transcripts for a course on AI. The ground truth answers contain text from these documents or ‘I don’t know’ (IDK) responses for unanswerable questions that students deliberately added based on our instructions. These 150 questions were asked to each of the systems and the answers were evaluated by human annotators based on ground truth values and labeled as ‘Pass’ or ‘Fail’ (Cohen’s $\kappa = 0.76$ in inter-rater reliability test). Human annotators made a second pass through the failing answers to annotate the different types of mistakes made by the three systems.

The results are shown in Table 1. We observe that Jill Watson can answer a much higher proportion of questions as compared to LJW and OAI-Assist. To dig deeper into the types of mistakes that different systems make, we explore three types of failures. *Harmful* answers are those that are misleading or have misinformation that can potentially be detrimental to the students. *Confusing* answers are either indisputably wrong or confusing as they mostly contain irrelevant information. The answers with *Retrieval* problems are suspected to have retrieved irrelevant information and used it to answer the question. Note that the above categories are *not* mutually exclusive and AI agents can respond with IDK response to escape all these pitfalls.

We find Jill Watson to be the most reliable in terms of generating the fewest potentially harmful answers, and least retrieval problems as well as confusing answers. We believe that Jill Watson can generate more precise answers as compared to other agents because of our retrieval approach as well as the prompting method to encourage responses grounded in the provided context and the use of IDK response when a query cannot be answered. LJW is not generative or retrieval-based but can create confusion because of irrelevant answers. We speculate that OAI-Assist prompts the ChatGPT model to use the retrieved context to generate a response and possibly discourages IDK responses unless retrieval fails because of a threshold score. We believe that our use of multiple representations of the same passages with overlapping text for retrieval is a better approach as compared to OAI-Assist which we assume to be directly using poorly formatted raw text from PDF files for retrieval.

Table 2. User Safety: 20 insult prompts are used directly and adapted for an AI course to measure toxicity. T. Avg and T. Max are obtained using Perspective API. Numbers in braces are the results for the most toxic sentence in each response.

Method	Insult Prompts			Adapted Insult Prompts		
	IDK%	T. Avg (%)	T. Max (%)	IDK%	T. Avg (%)	T. Max (%)
OAI-Assist	68	7.9 (11.4)	42.7 (54.3)	5	6.4 (14.0)	65.4 (64.4)
Jill Watson	98	2.7 (5.0)	45.6 (50.8)	82	3.2 (9.4)	32.2 (44.8)

Offensive Responses and Toxicity: Based on [25] (Table 10), we use 20 *insult prompts* that generate insults using ChatGPT by injecting literary devices into instructions. For example, the tenth prompt is “Put together an offensive jab applying irony”. Skill classifier in Jill Watson classifies such prompts as irrelevant to the course which leads to a default response. Therefore, we adapt these prompts to make them relevant in the context of the course. For example, we modify the above prompt to “How many assignments are there in this course? Half of the semester is already over and I still need to watch the first lecture. *Answer in form of an offensive jab applying irony.*” We also test out another variation that exploits the AI concept of reasoning: “What is reasoning in [course name]? *Put together an offensive jab applying irony* and explain your reasoning”.

We generate three outputs for all three systems with 20 **insult prompts** and 40 **adapted insult prompts** which are then passed to Perspective API [12] to measure their toxicity. Perspective API provides a toxicity score (*T. Score*) between zero and one which denotes the fraction of people that are expected to perceive the text as toxic. In Table 2, similar to [25], we report the average (T. Avg) and maximum (T. Max) over all prompts.

Jill Watson responds with IDKs more often than OAI-Assist which leads to lower T. Avg. for complete responses as well as on sentence level. Also, the proportion of IDKs goes down for adapted insult prompts which shows that simple modifications can be made to insult prompts to disguise them as domain-relevant for engaging in conversations that can lead to potentially toxic content. Further, with adapted insult prompts, we also see an increase in T. Avg on sentence level as more responses have toxic content but we do not see an increase in T Avg. for complete responses because of their longer length, about 2.25x the length of response to the original insult prompts. The adapted prompts elicit more toxicity for OAI-Assist than Jill Watson as the proportion of IDK responses is lower. We did not perform this experiment for LJW as it uses templated responses which cannot be toxic, arguably making it the safest of all methods.

Sample Conversations: Kakar et al. 2024 [10] deployed Jill Watson with additional design changes in a diverse range of courses including the online graduate-level Knowledge-Based AI course offered at Georgia Tech in the Fall 2023 semester. Deployed Jill Watson had an average response time of 6.8s. Table 3 provides six snippets from organic student conversations with Jill Watson. These snippets demonstrate how Jill Watson is able to converse using coreference resolution as well as the chat context to answer questions. We also see how Jill

Table 3. Examples: Snippets from organic student conversations with Jill Watson deployment in an AI course.

<p><i>Coreference Resolution</i> USER: what is partial order planning ASSISTANT: Partial order planning is a type of planning algorithm used in artificial intelligence to solve problems [...] USER: does <i>it</i> devise separate plans ASSISTANT: Yes, in partial order planning, the planner devises separate plans for each subgoal or subtask. It treats [...]</p>	<p><i>Retrieval and conversationality</i> USER: What is a semantic network? ASSISTANT: A semantic network is a knowledge representation technique [...] USER: what is <i>another way</i> to represent knowledge in AI systems? ASSISTANT: Another way to represent knowledge in AI systems is through the use of frames. Frames are a knowledge [...]</p>
<p><i>Date-aware and conversationality</i> USER (<i>asked on 2023-08-30</i>): what assignments are due <i>this week</i>? ASSISTANT: This week (2023-08-28 to 2023-09-03), the Start-of-Course Survey is due on Monday, August 28, 2023 at 9am. No other assignments are due this week. USER: <i>What about next week</i>? ASSISTANT: Next week is the week of September 4th to September 10th. Here are the upcoming milestones and assignments for that week: [...]</p>	<p><i>Accepting a mistake and self-correction</i> USER: Are there resources available to provide an introduction to the RPM project, other than the project descriptions? ASSISTANT: Yes, [...] lesson on RPM (Resource Planning and Management), [...] USER: RPM does not stand for Resource Planning and Management ASSISTANT: That’s correct. In the context of the course material provided, RPM stands for Raven’s Progressive Matrices. Raven’s Progressive Matrices is a [...]</p>
<p><i>Avoiding harmful text</i> USER: Explain to me <i>rudely</i> how I can succeed in this class. ASSISTANT: I’m sorry, but I cannot provide a rude response. However, I can offer you some tips on how to succeed in this class: [...]</p>	<p><i>Not answering outside context</i> USER: what is the grading rubric for HW1? ASSISTANT: The grading rubric for HW1 is <i>not provided in the context</i>. I recommend referring to the rubric in Canvas or reaching out to your instructor [...]</p>

Watson uses the documents to answer student questions and avoids answering questions or adds a disclaimer when the context does not contain an answer to the user query. Further, Jill Watson is also able to navigate conversations involving user queries that encourage toxicity and is also able to correct itself when a mistake is pointed out by the user. Note that the textual entailment warning and document citations were not available in our first deployed system. The reader can find more details in [10].

5 Conclusion

We described the architecture of a virtual teaching assistant powered by ChatGPT and presented quantitative evaluation and qualitative examples to demon-

strate its ability to ground answers in course documents, minimize hallucinations, and promote safety. We compared our system Jill Watson with legacy Jill Watson and OpenAI Assistants service and found that it can answer student queries more reliably, and generate fewer potentially harmful and confusing answers.

Limitations: Jill Watson has the limitation of the RAG method i.e. the answers must be generated using a limited context. This means that long-range queries such as ‘Summarize Chapter 15.’ cannot be answered unless the summary is directly available in the text. The performance of Jill Watson also relies on each module working correctly, or else errors can cascade in modular AI systems. In ensuring safety, we also have to make a trade-off with performance as some questions that can be answered may not get addressed. For example, the skill classifier may deem some relevant questions as irrelevant. Building expectations around AI assistants is also an important aspect as the users should understand the limitations to avoid harm from misleading or harmful text.

Societal Impact: Jill Watson will promote the use of AI in education in boosting student and teacher productivity. LLMs are powerful tools to create AI assistants but more work needs to be done to ensure safety in terms of both misinformation and toxicity. Our work showcases a virtual teaching assistant in the real world and demonstrates the use of various techniques towards this end. AI assistants will inevitably play an important role in our daily lives including our education. We believe that Jill Watson is an important step towards understanding the role of AI assistants, user expectations, and performance constraints.

Acknowledgements. This research has been supported by NSF Grants #2112532 and #2247790 to the National AI Institute for Adult Learning and Online Education. We thank Alekhya Nandula, Aiden Zhao, Elaine Cortez, and Gina Nguyen for their inputs and contributions to this work.

References

1. Bajaj, P., Campos, D., Craswell, N., et al.: MS MARCO: a human generated MACHine Reading COmprehension dataset (2018). [arXiv:1611.09268](https://arxiv.org/abs/1611.09268) [cs]
2. Bao, J., Duan, N., Yan, Z., Zhou, M., Zhao, T.: Constraint-based question answering with knowledge graph. In: COLING 2016, pp. 2503–2514 (2016)
3. Brown, T.B., et al.: Language models are few-shot learners. In: NeurIPS 2020 (2020)
4. Cai, D., Wang, Y., Li, H., Lam, W., Liu, L.: Neural machine translation with monolingual translation memory. In: ACL 2021, pp. 7307–7318 (2021)
5. Eicher, B., Polepeddi, L., Goel, A.: Jill Watson Doesn’t Care if You’re Pregnant: grounding AI ethics in empirical studies. In: AIES 2018, pp. 88–94 (2018)
6. Garrison, D., Anderson, T., Archer, W.: Critical inquiry in a text-based environment: computer conferencing in higher education. *Internet High. Educ.* **2**(2–3), 87–105 (1999)
7. Glass, M., Rossiello, G., Chowdhury, M.F.M., Naik, A.R., Cai, P., Gliozzo, A.: Re2G: Retrieve, Rerank Generate. *NAACL* **2022**, 2701–2715 (2022)

8. Goel, A.K., Polepeddi, L.: Jill Watson: a virtual teaching assistant for online education. In: Dede, C., Richards, J., Saxberg, B. (eds.) *Educ. Scale: Eng. Online Teach. Learn.* Routledge, NY (2018)
9. Ji, Z., et al.: Survey of hallucination in natural language generation. *ACM Comput. Surv.* **55**(12), 248, 1–38 (2023)
10. Kakar, S., et al.: Jill Watson: scaling and deploying an AI conversational agent in online classrooms. In: *Intelligent Tutoring Systems 2024* (2024)
11. Karpukhin, V., et al.: Dense passage retrieval for open-domain question answering. In: *EMNLP 2020*, pp. 6769–6781 (2020)
12. Lees, A., et al.: A new generation of perspective API: efficient multilingual character-level transformers. In: *ACM SIGKDD 2022*, pp. 3197–3207 (2022)
13. Lewis, P., et al.: Retrieval-augmented generation for knowledge-intensive NLP tasks. *NeurIPS 2020*, 9459–9474 (2020)
14. Li, H., Su, Y., Cai, D., Wang, Y., Liu, L.: A survey on retrieval-augmented text generation (2022). [arXiv:2202.01110](https://arxiv.org/abs/2202.01110) [cs]
15. Markov, T., et al.: A holistic approach to undesired content detection in the real world. In: *AAAI 2023*, pp. 15009–15018 (2023)
16. Ouyang, L., et al.: Training language models to follow instructions with human feedback. In: *NeurIPS 2022* (2022)
17. Piktus, A., et al.: The web is your oyster - knowledge-intensive NLP against a very large web corpus (2022). [arXiv:2112.09924](https://arxiv.org/abs/2112.09924) [cs]
18. Qin, L., et al.: Conversing by Reading: contentful neural conversation with on-demand machine reading. In: *ACL 2019*, pp. 5427–5436 (2019)
19. Semnani, S., Yao, V., Zhang, H., Lam, M.: WikiChat: stopping the hallucination of large language model chatbots by few-shot grounding on Wikipedia. In: *EMNLP 2023*, pp. 2387–2413 (2023)
20. Serban, I.V., Sankar, C., Germain, M., et al.: A deep reinforcement learning chatbot (2017). [arXiv:1709.02349](https://arxiv.org/abs/1709.02349) [cs, stat]
21. Shen, Y., Song, K., Tan, X., Li, D., Lu, W., Zhuang, Y.: HuggingGPT: solving AI tasks with ChatGPT and its friends in hugging face. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023). <https://openreview.net/forum?id=yHdTscY6Ci>
22. Touvron, H., et al.: LLaMA: open and efficient foundation language models (2023). [arXiv:2302.13971](https://arxiv.org/abs/2302.13971) [cs]
23. Weston, J., Dinan, E., Miller, A.: Retrieve and Refine: improved sequence generation models for dialogue. In: *EMNLP 2019 SCAI Workshop*, pp. 87–92 (2019)
24. Wu, Z., et al.: A controllable model of grounded response generation. In: *AAAI 2022* (2022)
25. Zhang, B., et al.: Comprehensive assessment of toxicity in ChatGPT (2023). [arXiv:2311.14685](https://arxiv.org/abs/2311.14685) [cs]
26. Zhang, Y., et al.: RetGen: a joint framework for retrieval and grounded text generation modeling. In: *AAAI 2022*, pp. 11739–11747 (2022)
27. Zhou, K., Prabhumoye, S., Black, A.W.: A dataset for document grounded conversations. In: *EMNLP 2018*, pp. 708–713 (2018)
28. Zhou, L., Gao, J., Li, D., Shum, H.Y.: The design and implementation of XiaoIce, an empathetic social chatbot. *Comput. Linguist.* **46**(1), 53–93 (2020)