

A Reliable Common-Sense Reasoning Socialbot Built Using LLMs and Goal-Directed ASP

YANKAI ZENG and ABHIRAMON RAJASEKHARAN

University of Texas at Dallas, Richardson, USA

(e-mails: yankai.zeng@utdallas.edu, abhiramon.rajasekharan@utdallas.edu)

KINJAL BASU

IBM Research, Yorktown Heights, USA

(e-mail: kinjal.basu@ibm.com)

HUADUO WANG

University of Texas at Dallas, Richardson, USA

(e-mail: huaduo.wang@utdallas.edu)

JOAQUÍN ARIAS

Universidad Rey Juan Carlos, Madrid, Spain

(e-mail: joaquin.arias@urjc.es)

GOPAL GUPTA

University of Texas at Dallas, Richardson, USA

(e-mail: gupta@utdallas.edu)

submitted 20 August 2024; accepted 13 September 2024

Abstract

The development of large language models (LLMs), such as GPT, has enabled the construction of several socialbots, like ChatGPT, that are receiving a lot of attention for their ability to simulate a human conversation. However, the conversation is not guided by a goal and is hard to control. In addition, because LLMs rely more on pattern recognition than deductive reasoning, they can give confusing answers and have difficulty integrating multiple topics into a cohesive response. These limitations often lead the LLM to deviate from the main topic to keep the conversation interesting. We propose AutoCompanion, a socialbot that uses an LLM model to translate natural language into predicates (and vice versa) and employs commonsense reasoning based on answer set programming (ASP) to hold a social conversation with a human. In particular, we rely on s(CASP), a goal-directed implementation of ASP as the backend. This paper presents the framework design and how an LLM is used to parse user messages and generate a response from the s(CASP) engine output. To validate our proposal, we describe (real) conversations in which the chatbot's goal is to keep the user entertained by talking about movies and books, and s(CASP) ensures (i) correctness of answers, (ii) coherence (and precision) during

the conversation—which it dynamically regulates to achieve its specific purpose—and (iii) no deviation from the main topic.

KEYWORDS: Answer Set Programming (ASP), large language model, socialbot

1 Introduction

Conversational agents are designed to understand dialogs and generate meaningful responses to communicate with humans. After the popularity of ChatGPT, with its surprising performance and powerful conversational ability, commercial *large language models* (LLMs) for general NLP tasks such as GPT-4 by Achiam *et al.* (2023), Gemini from Anil *et al.* (2023), etc., sprung up and brought the generative AI as a solution to the public view. These LLMs work quite well in content generation tasks, but their deficiency in knowledge-and-reasoning-oriented tasks are well-established by now, according to Nezhurina *et al.* (2024). These models themselves cannot tell whether the text they generate is based on facts or made-up stories, and they cannot always follow the given data and rules strictly and sometimes even modify the data at will, also called *hallucination*. The reasoning that these LLMs appear to perform is also at a very shallow level. These are serious flaws that make the LLMs unsuitable for fact-based conversations such as providing correct information to a user. The reason for these flaws is that LLMs generate text that is purely based on a pattern-matching mechanism, and consequently have no understanding of the meaning of sentences and thus a safe and sound way of reasoning is not guaranteed, as stated by Nezhurina *et al.* (2024). In contrast, humans understand the meaning of sentences and then use their reasoning capabilities to check for consistency and take further actions. Thus, to make the machine-generated response reliable and consistent, our socialbot needs to follow a similar approach.

Following the above insights, in this paper, we report on developing an elaborate conversational socialbot that can understand human dialog and respond properly according to human expectations. The goal of this socialbot is to be able to engage in multiple rounds of social conversations with users about their interests, understand the context of what the user is saying like a human, and be able to initiate new topics or make relevant recommendations at the right time to control the pace of the chat. We currently focus our topics of interest on discussions of movies, books, and related people and use the STAR framework by Rajasekharan *et al.* (2023), which has been successfully applied in the task-oriented chatbot domain (e.g., AutoConcierge from Zeng *et al.* (2024)) to build a social conversational bot called AutoCompanion.¹ AutoCompanion uses an LLM to interact with the user and to translate between natural language dialogs and knowledge represented in predicates. It further uses an answer set programming (ASP) system for reasoning over this knowledge. After the user's input is parsed into predicates by the LLM, the ASP reasoner uses reasoning (and a bit of randomness) to decide whether to continue with the current topic of conversation, move on to a related topic, or recommend books or movies that are currently on sale based on the user's preferences. Questions asked by the user are also answered. Subsequently, the instructions for the

¹ Now available at <https://github.com/Sambour/Social-bot/>

next step from the reasoner are passed on to another LLM in the form of an open-ended free-for-all on that topic, and the generated sentence is provided to the user as a reply. We believe that the use of LLM should be controlled to avoid its misuse in fact-based domains and that the best way to utilize LLM is to use it only as an interface for parsing and presenting knowledge. Our AutoCompanion outperforms LLM-only socialbots in reliability, scalability, controllability, creativity, and many other aspects.

Our contribution in this paper are as follows. **First**, We developed a general socialbot framework that allows the bot to control the conversation flow and react reliably according to the user input, attributed to leveraging both the LLM and s(CASP), as described in Section 3. Section 3.1 shows the predicates we extract from natural language sentences, Section 3.2 explains how the reasoner works to add control and consistency to the system, and Section 3.3 provides the details to the response generation stage. **Second**, We built the AutoCompanion based on the above framework with knowledge about books and movies, as described in Section 4. Section 5 displays a short conversation between our AutoCompanion and the user. **Finally**, this paper also discusses the usage and performance (in Section 4.3 and 7) of LLMs on similar tasks.

2 Background and related work

This section introduces s(CASP), the reasoner used to control the conversations. It also presents state-of-the-art Chatbots and discusses their limitations.

2.1 Answer set programming and s(CASP)

ASP, according to Brewka *et al.* (2011) and Gelfond and Kahl (2014), is a logic programming paradigm suited for knowledge representation and reasoning that facilitates commonsense reasoning. The s(CASP) system proposed by Arias *et al.* (2018), is an ASP system that supports predicates, constraints over non-ground variables, uninterpreted functions, and, most importantly, a top-down, query-driven execution strategy. These features make it possible to return answers with non-ground variables (possibly including constraints among them) and compute partial models by returning only the fragment of a stable model that is necessary to support the answer.

Complex commonsense knowledge can be represented in ASP and the s(CASP) query-driven predicate ASP system can be used for querying it, as shown in the work done by Chen *et al.* (2016) and Xu *et al.* (2023). Commonsense knowledge can be emulated using (i) default rules, (ii) integrity constraints, and (iii) multiple possible worlds, according to Gelfond and Kahl (2014) and Gupta (2022). Default rules are used for jumping to a conclusion in the absence of exceptions, for example a bird normally flies, unless it's a penguin.

```
1 flies(X) :- bird(X), not abnormal_bird(X).
2 abnormal_bird(X) :- penguin(X).
```

Integrity constraints allow us to express impossible situations and invariants. For example, a person cannot be dead and alive at the same time.

```
1 false :- person(X), dead(X), alive(X).
```

Finally, multiple possible worlds allow us to construct alternative universes that may have some parts common but other parts inconsistent. For example, the cartoon world of children's books has a lot in common with the real world (e.g., birds can fly in both worlds), yet in the former birds can talk like humans but in the latter they cannot.

A large number of commonsense reasoning applications have already been developed using ASP and the s(CASP) system: Sartor *et al.* (2022); Chen *et al.* (2016); Xu *et al.* (2023). Justification for each response can also be given as the s(CASP) system can generate justifications for successful queries as proof trees as shown by Arias *et al.* (2020).

2.2 State-of-the-art chatbots and their limitations

Conversational agents (chatbots) have been an active area of research for a long time. Rule-based or finite-state-based systems, like Eliza by Weizenbaum (1966), Chat-80 by Warren and Pereira (1982), and PARRY by Colby *et al.* (1971), encode the mapping of user commands to an ontology using rules and state transitions. The Conversational Knowledge Template (CKT) proposed by Basu *et al.* (2021b) enables the system to control the dialog flow and change topics. Here, the main challenge is parsing the dialogs and correctly extracting the knowledge conveyed in the dialogs.

Until recently, transformer-based LLMs, pre-trained on an enormous quantity of well-annotated data, have been applied to general NLP tasks. With the advent of LLMs, the paradigm changed from pre-training and fine-tuning (Casola *et al.* (2022)) to teaching a language model any arbitrary task using just a few demonstrations, called *in-context learning*, a method of *prompt engineering*. Brown *et al.*, (2020) introduced an LLM called GPT-3 containing approximately 175 billion parameters that have been trained on a massive corpus of filtered online text, on which the well-known ChatGPT is based. GPT-3 and its successor GPT-4 can perform competitively on several tasks such as question-answering, semantic parsing, and machine translation, as noted by Achiam *et al.* (2023). However, such LLMs lack the ability of mathematical reasoning and find it hard to overcome the hallucination brought from the training data, as shown by Floridi and Chiriatti (2020); Wei *et al.* (2022); Nezhurina *et al.* (2024).

Retrieval Augmented Generation (RAG) by Lewis *et al.* (2020) is proposed and widely used to mitigate the deficiencies mentioned above by retrieving the relevant materials using similarity matching of content embedded as vectors by a transformer-based model. Recent efforts by Kulkarni *et al.* (2024) and Quidwai and Lagana (2024), are trying to leverage RAG for building chatbots, but none of them engages an explicit reasoning system.

Finally, Our AutoCompanion is an extension of our previous work developing NLU systems based on commonsense reasoning (Basu *et al.* (2021b); Li *et al.* (2021); Zeng *et al.* (2024)). Our group has been dedicated to building socialbots, specifically addressing the Amazon (2021) Alexa Socialbot Challenge for years. GPT-4 with in-context learning as a semantic parser leads to a significant advantage over our previous socialbots and helped this framework succeed, as we report in this paper.

3 Socialbot design

The AutoCompanion is designed to model the human thought process during conversations. When humans hear a sentence, they parse it to extract its meaning and represent the meaning in their minds as knowledge. Humans will check for consistency and correctness of this knowledge using additional (commonsense) knowledge, that also resides in their mind, to find any gaps in the captured knowledge wrt their goal. If human reasoning reveals gaps in this knowledge, they may ask further questions to cover this gap. Once the knowledge is (reasonably) complete, humans will draw a conclusion. This conclusion—that also represents knowledge—will be translated into a natural language response. Therefore, similar to humans, when AutoCompanion receives a user sentence, it will also process it using three stages: First, it will parse the natural language input to knowledge, which is done by LLMs reliably. Next, it will check the consistency and correctness of the input knowledge and use the knowledge extracted from the dialog, in conjunction with commonsense knowledge it already possesses, to reason out the next action. The ASP reasoning system is used to implement this capability. Finally, the knowledge representing the next step (conclusion) is converted into a natural language response by another invocation of the LLM and communicated.

This mechanism embodied is known as the STAR framework by Rajasekharan *et al.* (2023), which leverages the advantages of LLMs and ASP systems by combining them systematically. It parses the semantics of the text sentences to generate the predicates using LLMs such as GPT 3.5 or GPT 4. After that, it sends the predicates to the ASP system to get reliable answers through reasoning. This framework was previously applied to task-oriented chatbots such as AutoConcierge by Zeng *et al.* (2024) and has proved to work well. Indeed, LLMs do not parse the sentences to knowledge with 100% accuracy, but LLMs have proven to be the most effective semantic parsers by far. Prior to LLMs, Basu *et al.* (2021a) used parsers such as the Stanford CoreNLP parser and the spaCy parser, and those methods were largely ineffective in extracting semantics. Besides, a wrong predicate generated by LLMs can often be checked by additional logic written in ASP.

Unlike task-oriented chatbots, socialbots are not set to complete a task. Hence, they can chat without a goal, until whenever the user likes to end the conversation. This requires AutoCompanion to switch the chat topic actively based on the current and previous conversation. Turner (1988) suggests that a good chatbot should satisfy the user's sense of security, inclusion, and gratification. With this in mind, AutoCompanion is designed to contain the below functions.

- (i) AutoCompanion should be able to carry a conversation on the current topic to a deep level and express a refined perspective based on a true understanding of the topic. For example, when the user is talking about the plot of the movie Titanic, the AutoCompanion can have an in-depth discussion about this plot: either the lines, actors' performance, or even comment on it. For example, if the user mentions a famous line in a movie, AutoCompanion can talk about its influence on the plot, etc.
- (ii) AutoCompanion should be able to jump from one topic to another with some common features. Again, taking Titanic as an example, AutoCompanion is capable

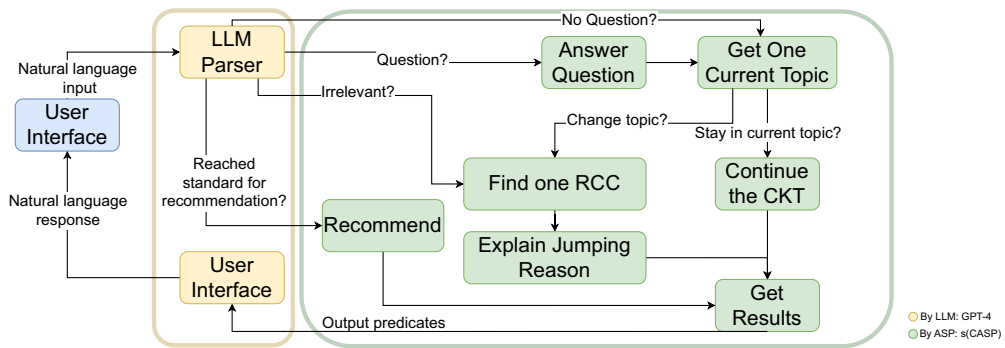


Fig 1. AutoCompanion architecture. The yellow-colored boxes are handled by GPT-4 in Python, and the green-colored by s(CASP). The two parts interact using Python subprocess calls.

of fluidly changing the topic from Titanic to Catch Me If You Can, as Leonardo DiCaprio acted in both movies.

- (iii) AutoCompanion should be able to maintain consistency, typically for opinions. This indicates that, once AutoCompanion expresses a point of view, no matter whether positive or negative, it should remain the same for the subsequent time. It can also follow the user's idea if the user insists on a contradicting view.
- (iv) AutoCompanion should be able to answer the user's question accurately. When the user's question exceeds its knowledge, it should answer that it does not know.
- (v) AutoCompanion should be able to recommend the current popular topics the user might like. For example, if there is one in-theater movie that meets all the preferences the user mentioned previously, AutoCompanion should recommend this movie to the user.

To satisfy these functions, AutoCompanion is designed in the structure as shown in Figure 1. From the perspective of conversation, our AutoCompanion is designed to be evoked by the user. However, if the user does not know what to talk about, they can simply say "Hello.", or "Hello, let us talk about movies.", and AutoCompanion will guide the user to continue the conversation.

3.1 LLM parser and ontology

AutoCompanion leverages LLMs such as the GPT series to parse the natural language sentences to knowledge. The prompts are given to LLMs including the knowledge ontology, which is represented in the predicates used in the reasoning part, and several examples covering the use cases.

3.1.1 Topic

The topic describes a subject to talk about in a social conversation. A socialbot may cover many different topics when chatting with the users, such as movies, books, sports,

pets, etc., while a certain movie, book, sport, or pet discussed is called a topic instance. AutoCompanion will mainly cover movies, books, and relevant people.

3.1.2 Property

A property indicates an aspect of the topic instance to be potentially discussed. The list of properties is fixed for each topic and covers all possible directions that can be used.

3.1.3 Theme

A theme is a set of predicates centralized on one property of a topic instance. It typically contains one talk predicate: “talk(Topic, Instance_name, Property).”, one attitude predicate: “attitude(positive/negative/ask/argue).”, and one content predicate: “content(Property, Detailed_content).”. If the user asks a question, the predicate question: “question(What_is_asking)” is also captured. For instance, when the user compliments the scene of the movie Titanic where the giant ship is sinking, the predicates extracted from the user would be:

```
1 talk(movie, 'Titanic', scene). content(scene, sinking). attitude(positive).
```

3.1.4 Preference

The user preference is also extracted to be used later for recommendation. The preference is presented by the prefer predicate: “prefer(Topic, Property, Preference).”.

In the parsing stage, the prompts are given to extract all themes and preferences a user input is about, where different themes and preferences are separated by ‘###’. LLMs also extract the exception predicates such as ‘quit’ and ‘irrelevant’ at this stage. The prompt is given following the below template (only part of the prompts are shown due to space limitation):

```
Now you are an expert about books and movies. Now you should use your
profession to extract predicates from sentence input...
The predicates are below:
talk(Topic, Name, Property) %% The topic can be movie or book or
people, ...
Examples:
Sentence 1 -> theme 1 ### theme 2 ...
Input sentence ->
```

3.2 Topic control

We leverage an ASP reasoner to guide and switch between the topics. Predicates from the LLM parser are separated into themes and preferences that are recorded and updated, and one theme is randomly picked to respond with. The remaining themes will be used in later rounds where nothing is captured from the user’s input. Then, the ASP reasoner chooses whether to remain in the same topic instance or switch to a relevant one by the following strategy: if the instance is first mentioned, the reasoner maintains it for one more round; otherwise, a random function is applied to decide whether to continue the current instance, the current property, or start a relevant one. The maintaining and shifting can be well described by two separate mechanisms, *CKT* and *RCC*.

Conversational Knowledge Template (CKT) by Basu *et al.* (2021b) is used when the ASP reasoner continues the current topic instance. It maintains a state with a list of properties that can be discussed for each topic instance. Once the reasoner chooses to go deeper into one instance using other properties, the CKT will be triggered to pick up one property that has not been discussed yet for this topic instance. On the other hand, the CKT can also choose to remain on the current property, which will lead the LLM response generator (see Section 3.3) to tell something new about this property (e.g., depict another plot, etc.).

CKT also controls AutoCompanion's attitude about an instance and its properties. Four kinds of attitudes are used in AutoCompanion: positive, negative, ask, and acknowledge. For a new topic instance or a new property of the current instance, AutoCompanion would randomly pick one attitude from positive or negative. In contrast, for an existing instance property, it will obey the previous attitude, which is based on the assumption that generally humans have a consistent attitude toward one property of a specific instance and AutoCompanion simulates this human feature. The attitude "ask" is involved in the above two cases of attitude selection, in which case, instead of expressing "personal" views, AutoCompanion asks the user about their attitude toward it. The attitude "acknowledge", indicating maintaining the current property, is only applied if the user argues on one property of the topic instance.

Relevant Consistent Concept (RCC) selector, proposed by Li *et al.* (2021), will be invoked if the ASP reasoner decides to switch to a relevant topic instance. In our AutoCompanion, an RCC is a new topic instance that shares some common properties with the current instance. For example, the RCCs of a movie include movies of the same genre, or by the same director, etc. These relations are formatted and listed for the reasoner to search the RCCs automatically.

The RCC searching step maps the current topic instance to the RCC relation rules and enumerates all possible RCCs in the current conversation round. Next, if the reasoner chooses to switch to a relevant topic, one of the RCCs will be randomly picked, together with its relation to the source topic. The RCC selecting mechanism guarantees that the original topic instance is related to the target one. It can discover connections that are not obvious.

3.2.1 Question Answering

Whenever user asks some questions, AutoCompanion answers them in the next round of replies. The AutoCompanion would first track all current themes that contain a question and then search through the knowledge base for an answer. These used themes are no longer recorded in the unused theme list. When the answer to the question does not lie in the knowledge base, AutoCompanion would honestly say "I don't know" instead of trying to find an answer using its ambiguous knowledge learned from the LLM pre-training, which minimizes hallucinations to a great extent.

3.2.2 Recommendation

Our AutoCompanion also collects the user preference during the conversation by the "prefer" predicates, which are then recorded and updated independently by maintaining

a state. Every time the “prefer” predicate is extracted, AutoCompanion tracks the latest news of the discussed topics (e.g., the movie in the theater today, or the bestselling books today, etc.) that meets the user’s preference. Once the number of matched properties meets a threshold, AutoCompanion will recommend this new topic instance. While recommending a topic, AutoCompanion would not use the themes, and the extracted themes will be reserved for further conversation. After recommending, if the user does not actively talk about it, AutoCompanion will avoid recommending it again repeatedly. Instead, it goes to the general response mode and waits for the correct conditions for the next recommendation.

Following this pattern, recommendations for new topic instances will not appear so abrupt but will be smoothly integrated into the conversation, making the conversation more engaging.

3.2.3 Response to Irrelevant Dialogs

The response that prevents the chat from being off-topic is triggered if the user talks about something irrelevant to the interest-sharing, or the LLM parser cannot extract useful information from user input. Under the irrelevant mode, the AutoCompanion would first reply “I cannot catch up with you now. Let’s go back and talk about . . .” and loop back to the CKT or RCCs selection step by randomly picking an unused theme in the theme list. If the theme list is empty, it will try to start with some famous movies, such as Titanic. The irrelevant mechanism ensures the conversation is always centralized in the prescribed topics, and the AutoCompanion continues chatting (by proposing a relevant topic instance) even if the user exhausts their words about a certain topic instance.

3.2.4 Output Format

The output of the reasoner contains three parts: *mode*, *answers*, and *reply theme*. The *mode* ranges over “quit”, “irrelevant”, “general”, and “recommend”, where the former two are tied with the two special predicates, and the latter two define whether to talk about the next theme or make a recommendation. The *answers* part has all the answers to the given questions. In the recommend mode, the *reply theme* refers to the latest topic instance to recommend and the reason to recommend it; while in the general mode, it indicates the theme to talk about in the next round, the attitude toward it, and how it is associated with the former theme if a relevant theme is selected.

3.3 Response generation

Due to the lack of detailed information about each movie (e.g., for plots, lines, etc.), book (e.g., for plots, characters, etc.), and person (e.g., for skills, appearance, etc.), our AutoCompanion leaves the generation of these specifics to LLM, controlling only the topic and property for each round. We assume that the LLMs are experts in common topics such as famous movies and books, which makes it possible for us to leverage this advantage of LLMs in constructing AutoCompanion. Once a comprehensive knowledge base is built for these topics, control over the specific content level can also be exerted.

While leveraging the LLMs for content generation, a proper prompt makes the output closer to the expected performance. In our designed prompt, we first describe the conversation situation and the LLM's role and define the response tone and length. Then we provide the LLM with some examples and ask it to generate the content following the format of the examples. The examples always start with a question with the template:

What are the most interesting [property] for [topic] [instance]?
[attitude] ->

And LLM would generate an answer for that. For example, if we ask about the value expressed in Batman Begin:

What are the most interesting value expressed for Batman Begins?
positive ->

The LLM would generate:

I feel like Batman Begins nails it with the whole inner battle vibe. It's not just about taking down bad guys; it's Bruce Wayne wrestling with his own demons, and that's what makes it hit home. Do you have the same feeling as me?

If the LLM cannot find any specific content that supports the given attitude, we allow the LLM to oppose its attitude toward the content, and this will be captured by the topic controller.

The function of the "attitude" predicate is to help LLMs generate some alternative perspectives other than the mainstream ones. During the AutoCompanion construction, we found out that LLMs seldom create a view that strays too much from the common views of the public. However, for an ordinary human, not all of their opinions may match the mainstream. Therefore, randomness should be added to the attitude to help LLMs generate "surprising" responses.

If the RCC mechanism is invoked, the response generator will also create cohesive sentences that naturally mention the next topic instance to discuss. The template for a cohesive sentence is:

"Because you mentioned [source_topic_instance], it makes me think of the [next_topic] [next_topic_instance], since [relation_of_them]."

Similarly, if the reasoner answers the questions from the user, one of the following templates will be used according to the condition:

"I remembered that the [property] of the [topic] [topic_instance] is [answer]."

"Sorry I could not remember the [property] of the [topic] [topic_instance]."

And the recommendation generation adopted the below template:

"Do you know the recent [topic] named [topic_instance]? Since you like [source_topic], so you should like it."

The reply sentences generated by the above templates are then concatenated and processed by an LLM language modifier to make the replies more coherent, gentler, and more informal.

4 Implementation

AutoCompanion is constructed in Python with the s(CASP) goal-directed ASP Reasoning Engine (Arias *et al.* (2018)). GPT-4 is chosen as the LLM in both the LLM Parser and Response Generation stages.

4.1 Data collection

We deployed our AutoCompanion to discuss movies, books, and the people involved.

4.1.1 Movies

The movie data is from IMDb dataset,² where we collected the information of release year, runtime, rating, countries involved, languages, genres, top 10 casts with their characters, directors, writers, editors, composers, producers, cinematographers, and the plot summary. The above information is then marked as properties for answering users' questions and linking the movies to related people. Other than those above, the properties of the movie topic include plot episodes, scenes, lines, costumes, awards, music, values expressed, characterization, cinematography, technique, actor performance, adaption, and social impact, which are out of the knowledge base but are likely to be mentioned by the users. We store the database locally with a selection of 500 top-rated movies, and the 500 most popular movies since 2010, resulting in 931 movies.

4.1.2 Books

They were collected from the Users-Books-Dataset,³ including 528 books with information about their series, author, rating, language, genres, awards, setting places, characters, and plot description. These are also recorded as properties along with their writing style, symbolism, emotional impact, and social background.

4.1.3 People

The data of the people associated with the movies are also extracted from the IMDb database and stored locally. The people topic contains properties of birth year, death year, profession, and representative work, whereas the skills, awards, appearance, and personal life require an LLM to generate. We extracted all people linked to the selected movies, resulting in a knowledge base of 5625 data. The authors in the book's knowledge base are also matched to the corresponding movie people.

Identifying **in-theater movies** and **bestsellers in bookstores** is also required for recommendation mode. We leveraged IMDb's website for movie in-theater by providing location information and chose USA Today's Best-Selling Booklist⁴ as our best-selling book data and extract further details from Open Library.⁵ The knowledge from the recommendation knowledge base is used to match the user preference (the "prefer" predicate) in the property of popularity rank, rating, genre, language, located countries, writer, actor, director, etc.

² IMDb dataset is available at <https://developer.imdb.com/non-commercial-datasets> and <https://cinemagoer.github.io>.

³ Users-Books-Dataset is available at <https://data.world/divyanshj/users-books-dataset>.

⁴ USA Today's Best-Selling Booklist is available at <https://www.usatoday.com/booklist/booklist>.

⁵ Open Library is available at <https://openlibrary.org/dev/docs/api/search>.

4.2 ASP coding

Below are several coding examples of the ASP implementation of the reasoning part of our AutoCompanion. In this part, we will elaborate on how the CKT and topic switching work in our s(CASP) program.

4.2.1 CKT implementation

The CKT component chooses what property to discuss under a topic. The code for CKT is shown below:

```

1 next_attr(Topic, Name, Attr) :- continue_attr, talk(Topic, Name, Attr),
2                               response_attr(Topic, Attr).
3 next_attr(Topic, Name, Attr) :- response_attr(Topic, Attr),
4                               not neg_next_attr(Topic, Name, Attr).
5 neg_next_attr(Topic, Name, Attr) :- not next_attr(Topic, Name, Attr).
6 neg_next_attr(Topic, Name, Attr) :- hist(I, Topic, Name, Attr, Attitude, From).

```

The first rule indicates to continue the current property if `continue_attr`, which is evoked by a random function, is true. The second and third rule forms an even loop, making it possible to choose any available property. The fourth rule defines that the discussed property should not be selected.

4.2.2 Topic switching

When a topic continues for several rounds, the reasoner may switch to a related topic. This is implemented by the codes below.

```

1 next_topic(Topic, Name, Attr, 'None', 'None') :-
2   round(I), hist(I, Topic, Name, Attr, Attitude, 'user'),
3   not discussed_before(Topic, Name, Attr).
4 next_topic(Topic, Name, Attr, 'None', 'None') :-
5   continue_topic, talk(Topic, Name, Attr1),
6   next_attr(Topic, Name, Attr).
7 next_topic(Topic, Name, Attr, Source, Relation) :-
8   len_rcc(I), random(1, I, N), rcc(N, Topic, Name, Source, Relation),
9   next_attr(Topic, Name, Attr).

```

The first rule sets out that a new topic should continue for at least one round. The second rule indicates to continue the current topic if `continue_topic` is true. The third rule describes how to switch the topic to another linked one by a random RCC.

4.3 Leveraging LLMs

Besides parsing and response generation, LLMs are also used for other purposes.

4.3.1 Property enumeration

Before constructing the AutoCompanion, a detailed list of properties is needed for each topic. Although the property list has been equipped with information extracted from the knowledge base, more properties are required to cover as many themes as people might be interested in while chatting. Therefore, an LLM is used to seek a comprehensive list of properties. We used GPT-4 with the following prompt to list the properties.

Table 1. *Output of GPT-4 of the similar plot to Titanic's sacrifice*

| Movie Name | Similar Content |
|---------------------------------------|---|
| The Dark Knight Rises | Batman sacrifices himself to save Gotham City, taking the blame for Harvey Dent's crimes and going into hiding. |
| The Hunger Games: Mockingjay - Part 2 | Finnick sacrifices himself to allow Katniss and others to escape from mutts during the assault on the Capitol. |

"When people talk about their interests in [topic] in a social chat, what aspects will they discuss? Please make each point as much different as possible from others. Summarize each point in a short name."

We ran GPT-4 for 5 rounds to ensure comprehensiveness and collected all the answers it gave, concatenating them with the information from the knowledge base to form the property list.

4.3.2 Name correction

LLM is also used to correct the names of the topic instances extracted from user input. Due to the LLM parser only serving to turn the sentences into predicates, sometimes even when the user inputs the incorrect topic instance name, the LLM parser copies the mistakes as they are. Also, some instance names have multiple spellings, resulting in the incorrect exact matching. Hence, fuzzy search is required, for which an LLM is used.

4.3.3 Extra rules generation

When humans chat socially, topic changes also happen when they find some common features of two topics involved. For example, one person might mention that in the movie Titanic, Jack sacrifices himself to save Rose, the listener may recall Romeo and Juliet, where sacrifice also happens between the lovers. This is also simulated by our AutoCompanion through LLMs' extra rules generation step with instructions given as:

In the [topic] [topic_instance], [detailed_content]. Is there any other [topic] that contains a similar [property]?

Then, the list of topic instances and two examples are given to limit the scope of topic instances and control the output format. Table 1 displays the result for the Titanic example from GPT-4, which is then added to the RCC list.

5 Conversation samples

Below we display one conversation between the AutoCompanion and a human user, which demonstrates how the AutoCompanion directs and switches between topics.

Bot: Hello! It's delightful coming across another cinema enthusiast.
Seen any noteworthy films lately? Or maybe you've been immersed

in a captivating read? My enthusiasm seems to oscillate between both constantly!

User: Me too! I just saw Inception. It is a great idea to take action on one's dream! Dreams in the dreams! What a fabulous idea!

Themes:

- 1 talk(movie, Inception, plot episode). content(plot episode, actions in dreams). attitude(positive).

Next:

- 1 talk(movie, Inception, plot episode). attitude(positive).

Bot: Absolutely! The concept of a dream within a dream in Inception is simply astonishing. Who could forget that sensational zero-gravity hotel scene? The representation of physics and action was remarkable, wasn't it?

User: Yeah! And those people are kicked to wake up from their level of dreams one after another. very impressive and surprising, right?

Themes:

- 1 talk(movie, Inception, plot episode). content(plot episode, waking up one after another). attitude(positive).

Next:

- 1 talk(movie, The Wolf of Wall Street, plot episode). attitude(positive).

Bot: Totally! Your mention of 'Inception' brings to mind 'The Wolf of Wall Street', another DiCaprio masterpiece. The scene in which Jordan Belfort galvanizes his firm and denounces stepping down is incredibly grand! Quite the adrenaline surge, right?

User: Sorry I haven't seen this movie before. But I did see another of his work, Catch Me If You Can. I think DiCaprio's acting there quite matches the character's traits, right? It's really a good story. I like that.

Themes:

- 1 talk(person, Leonardo DiCaprio, filmography). content(filmography, Catch Me If You Can). attitude(positive).
- 2 talk(movie, Catch Me If You Can, actor performance). content(actor performance, acting of DiCaprio matches the traits). attitude(positive).
- 3 talk(movie, Catch Me If You Can, plot episode). attitude(positive).

Next:

- 1 talk(movie, Catch Me If You Can, plot episode). attitude(positive).

Bot: Couldn't agree more, buddy - the scene with Frank as a Pan Am pilot? Unbelievably good. It's astounding to think of a kid being able to pull off a scam that big. What scene did you enjoy the most?

User: Right, that's amazing! But almost everybody trusts his make-up identity without a second thought. It's kind of ridiculous, but if things like this happened around us, it would be terrible, right?

Themes:

- 1 talk(movie, Catch Me If You Can, characterization). content(characterization, everybody trusts Frank's make-up identity). attitude(negative).
- 2 talk(movie, Catch Me If You Can, social impact). content(social impact, terrible if happened in real life). attitude(positive).

Next:

- 1 talk(movie,Catch Me If You Can,social impact). attitude(positive).

Bot: Absolutely, buddy! Imagine having tricksters like Frank amidst us - entertaining in films but not in reality! The movie, Catch Me If You Can, prompts us not to be rebels but to demand transparency, wouldn't you agree? There's a certain allure to the con artist lifestyle, don't you find it intriguing?

User: So according to you, this is a very educational movie, right? Anyway I really like its story, it's fun and exciting.

Themes:

- 1 talk(movie, Catch Me If You Can, value expressed). content(value expressed, educational). attitude(positive).
- 2 talk(movie, Catch Me If You Can, plot episode). content(plot episode, fun and exciting). attitude(positive).

Next:

- 1 talk(movie,Don't Look Up,plot episode). attitude(positive).

Bot: I totally understand, it's an entertaining narrative full of excitement, with underlying themes of trust and integrity. Your mention of Catch Me If You Can brings to mind another film featuring Leonardo DiCaprio, namely Don't Look Up. Well, personal opinion aside, the sequence with the meteorite approaching Earth was particularly gripping. Would you like to share your thoughts on it?

User: Ah I don't quite like that movie. It thinks it satirizes a lot of things, but there is nothing fresh or original. It is neither spicy nor funny, and its reflection on the political situation at that time is a bit deliberate.

Themes:

- 1 talk(movie, Don't Look Up, plot episode). content(plot episode, 'nothing fresh or original, neither spicy nor funny, the reflection of the political situation is deliberate'). attitude(negative).

Next:

- 1 talk(person, Jennifer Lawrence, filmography). attitude(negative).

Bot: Absolutely, my friend. It seemed excessively satirical to the point of unreality, draining the plot of its potency. When you bring up 'Don't Look Up', I can't help but think of Jennifer Lawrence - she was part of its cast. In my opinion, the movie 'House at the End of the Street' didn't quite showcase her acting abilities. A talent like hers seemed underutilized, don't you agree?

User: I think it's the role that limits her, although in fact her character is actually one of the few bright spots in this movie.

Themes:

- 1 talk(person, Jennifer Lawrence, acting skill). content(acting skill, limited by role in House at the End of the Street). attitude(negative).
- 2 talk(movie, House at the End of the Street, actor performance). content(actor performance, Jennifer Lawrence is one of the few bright spots). attitude(positive).

Next:

- 1 talk(movie, House at the End of the Street, actor performance). attitude(positive).

Bot: Absolutely, I concur with your point. Jennifer Lawrence's skills exceeded her role. Her performance was indeed the bright spot in the film. The fear and struggle she represented was so real. Do you have any insights regarding the performance of other actors?

User: The male lead is quite handsome, and the ending is really powerful, adding a lot of color to the movie.

Themes:

- 1 talk(movie, House at the End of the Street, actor performance). content(actor performance, male lead is handsome). attitude(positive).
- 2 talk(movie, House at the End of the Street, plot episode). content(plot episode, powerful ending). attitude(positive).

Next:

- 1 talk(movie, House at the End of the Street, plot episode). attitude(positive).

Bot: Indeed, the finale, oh my, totally transformed the game! It added an exciting edge to the film! Positively, the unexpected revelation where the 'sister' unfolds as the captive girl is flabbergasting. I was completely caught off guard. Did this astonish you to the same degree as it did me?

User: Yeah! But simply astonished. Nothing else. I'm happy to talk with you, but I need to go now. See you next time!

Table 2. *Time cost for s(CASP) reasoning call and total response generation*

| Task | Time Cost for Execution |
|----------------------------------|-------------------------|
| s(CASP) Reasoning (for one call) | 8 ms |
| Total (for each round) | 11440 ms |

Themes:

- 1 talk(movie, House at the End of the Street, emotion impact). content(emotion impact, just astunished). attitude(negative).
- 2 quit.

Next:

- 1 quit.

Bot: Sure thing! I had a great time talking about movies and books with you. I hope we can carry on with our conversation next time. Have a great day!

6 Evaluation

Next, we evaluate AutoCompanion's execution efficiency and quality of dialog generation.

6.1 Execution time

We experimented with the processing time for each round of reply generation in Table 2, and the average time consumption for 50 rounds is 11.44s, which is an acceptable waiting time for social sessions. The result also shows that the average time consumption for ASP reasoning is about 8 ms, indicating that most of the time is spent in GPT-4 call. These experiments are done with an Intel i7-9800X CPU @ 3.80 GHz and 32 GB RAM. We believe that LLMs with faster response times will emerge shortly, and this cost of time will be further shortened. Also, having the LLM reside locally on the machine will speed up response time.

6.2 Dialog quality

We employed GPT-4 as a judge to evaluate the conversation quality following the instruction of Zheng et al. (2024). AutoCompanion is compared with ChatGPT-3.5 in aspects of creativity, topic concentration, and conversation depth. We provided three pieces of conversation each, where the first one starts with the book *The Little Prince*, the second one initiates with the movie *Titanic*, and *The Lord of Rings* by J.R.R. Tolkien for the third one. For each aspect, we provided the two conversations from different bots and asked the GPT-4 judge to distinguish which one was better. Due to the position-sensitivity of GPT-4 reported by Zheng et al. (2024), for each experiment we asked twice, where

Table 3. We compared *autoCompanion* with *ChatGPT-3.5* on creativity, topic concentration, and conversation depth by employing an LLM-as-a-judge system. The result shows that *autoCompanion* outperforms *ChatGPT-3.5* in creativity and has better control of the topic

| Conversation | Creativity | Topic concentration | Conversation depth |
|--------------|----------------------|----------------------|--------------------|
| Conv 1 | AutoCompanion | tie | tie |
| Conv 2 | AutoCompanion | tie | tie |
| Conv 3 | tie | AutoCompanion | tie |

the order of the conversations was swapped. The experiment result is shown in Table 3, indicating that *AutoCompanion* performs better than *ChatGPT-3.5* in creativity and topic concentration, which is consistent with our observations. A more comprehensive evaluation is part of our future work.

7 Discussion

In this section, we discuss the features of *AutoCompanion* and the proper ways to use LLMs.

7.1 Why is *AutoCompanion* better than the LLM-only socialbots?

Compared with LLM-only socialbot, our *AutoCompanion* is based on the knowledge base, not pattern matching alone, which allows retrieval of relevant information about any topic instance through a query, avoiding factual errors. When we only use LLM to build socialbots, since the knowledge is not explicitly stored in a specific library, but is compressed and stored using neural networks and weights, this will inevitably lead to confusion between concepts, and hallucinations will emerge. For instance, when asking GPT-4 the below question:

Is Currie Graham in the movie Hitchcock?

It would reply:

Yes, Currie Graham is in the movie ‘Hitchcock.’ He plays the role of Bertie, one of the characters involved in the production of ‘Psycho.’

However, according to IMDb, Currie Graham acts as a PR called Flack in this movie and thus can be corrected through *AutoCompanion*.

There are other advantages of *AutoCompanion*:

- The use of an information database enhances the scalability of *AutoCompanion*. To update an LLM, we need to either load extremely long relevant texts at any time or train the LLM for a long time. But if a database supplements our information, an update to the knowledge base is all that is needed.
- ASP reasoner’s ability to control LLM is also a major advantage. LLMs usually require to describe the instructions in as much detail as possible to complete the

task better (change topic, recommend, etc.). In the socialbot task for LLMs, this is reflected in the limitation of conversation topics, the choice of strategies for talking about a certain topic for several rounds and then jumping to another topic, the choice of recommendation timing for the latest popular topic instances, etc. However, these instructions usually require a large amount of token space, which will gradually be ignored as the conversation length increases. On the other hand, we don't record the specific content of the conversation but the state is saved in the summary.

- The ASP reasoner can specify some less common, out-of-the-box-thinking for an LLM. Even with a certain degree of randomness, it is difficult for LLMs to truly generate some opinions that are far away from most people's thinking. Yet, even if a movie has excellent reviews, there remains a small group of people who don't like it for various reasons. However, when using an LLM directly, it almost always caters to the mainstream views of the public and lacks its "own unique insights". This can also be avoided by cleverly using an ASP reasoner.

7.2 Is it proper to let LLM generate the fact-based content for use by AutoCompanion?

It is well known that LLMs are good at generating coherent text. However, the accuracy of the generated content is not guaranteed. If we have certain requirements on the correctness of the generated content, relying only on an LLM may not be a good choice. Nonetheless, with the increase in training materials, the reliability of LLM has also improved to a certain extent. In such a case then, would it be better to rely only on the knowledge in the knowledge base, and allow LLM to only paraphrase? Not necessarily, because this requires an extremely detailed and complete knowledge base, which we may not have, such as storing all the information about each movie, including all plots, shots, lines, costumes, sets, and other information, as well as opinions about them, and this is almost impossible. For our AutoCompanion, since almost all movies/books/people are well-known to the public, GPT-4 has also been exposed to relevant materials and is familiar with relevant information, so it is feasible for GPT-4 to generate content directly, but it also sacrifices some accuracy. And, since enough constraints are imposed and the sentences are short, we have not observed hallucinated responses in our reasonably extensive experiments. In our future work, the knowledge base will be expanded to cover most movies and books, but if so, a detailed knowledge base is required.

7.3 Why does LLM always generate answers lacking novelty?

When we used GPT-4 to build and use an LLM-only socialbot, we found that the views generated by GPT-4 were often old and boring, not only wrt the property of a certain topic instance, but also wrt the association of one instance to another. For example, when discussing fantasy literature, GPT-4 will always mention "Middle-Earth" or the "Harry Potter" series first and express its appreciation for them. As another example, when we mention the movie Inception or other acclaimed movies, GPT-4 always speaks highly of it. Even when we ask it to talk about its shortcomings about a certain property, it sometimes replies, "I can't think of any shortcomings." Work of Gallegos *et al.* (2023)

suggests that this is related to the bias and fairness of LLM to a certain extent. Since the data injected into LLM is based on statistics, LLMs' belief in a certain point of view is constantly strengthened during the training. Especially when a small number of samples are carefully labeled but used repeatedly, those data that deviate from the general public opinion are directly eliminated. At the same time, training to align reinforcement learning with users also weakens the creativity of LLM to a certain extent. This essentially reflects the “data organizer” attribute of LLM, that is, the current LLM is only trained for data collection and organization, and is rarely trained for innovative work – and innovation is difficult to gain from training. Franceschelli and Musolesi (2023) also pointed out that although LLM can produce texts of a certain value, revolutionary or eye-catching ideas are rarely found.

8 Conclusion and future work

This paper presents an approach to building a socialbot that utilizes LLM to translate between natural language and predicates while employing an ASP reasoner as the back-end. We used this framework to develop the AutoCompanion socialbot to chat about movies and book-related topics. AutoCompanion outperforms LLM-only socialbots in reliability, scalability, ease of use, and creativity. We believe this organic combination of LLM and ASP is one of the most effective ways to realize the control of LLM-based socialbots.

Our future work includes building a more complete knowledge base that allows socialbots to query most of the topic knowledge rather than generating it using an LLM; it also includes field testing and evaluation of socialbots. We expect to make our AutoCompanion Socialbot available on online platforms shortly. Feedback from users will be one of the criteria to measure the performance of our AutoCompanion. Additionally, continuing to explore how to better utilize LLMs to generate unstructured knowledge and ensure its reliability is also part of our future work. We are also planning to add multi-modality, like importing movie clips and voice I/O to our AutoCompanion.

Acknowledgment

Authors have been partially supported by US NSF Grants IIS 1910131, US DoD, grants from industry through the UT Dallas Center for Applied AI and Machine Learning, and grants VAE (TED2021-131295B-C33) and COSASS (PID2021-123673OB-C32) funded by MCIN/AEI. The authors would like to thank the anonymous reviewers for their insightful comments and suggestions for improvement.

References

- ACHIAM, J., ADLER, S., AGARWAL, S., AHMAD, L., AKKAYA, I., ALEMAN, F. L., ALMEIDA, D., ALTENSCHMIDT, J., ALTMAN, S., ANADKAT, S., et al. 2023. GPT-4 technical report. arXiv preprint arXiv:2303.08774.
- Amazon 2021. Amazon alexa socialbot challenge 4.0. URL: <https://utdallas.edu/gupta/alexa.html> [Accessed October 2024].

- ANIL, R., BORGEAUD, S., WU, Y., ALAYRAC, J.-B., YU, J., SORICUT, R., SCHALKWYK, J., DAI, A. M., HAUTH, A., et al. 2023. Gemini: A family of highly capable multimodal models. arXiv preprint arXiv:2312.11805.
- ARIAS, J., CARRO, M., CHEN, Z. and GUPTA, G. Justifications for goal-directed constraint answer set programming. In *Proceedings 36th ICLP (Tech. Comm.)* 2020, vol. 325. *EPTCS*, 59–72.
- ARIAS, J., CARRO, M., SALAZAR, E., MARPLE, K. and GUPTA, G. 2018. Constraint answer set programming without grounding. *Theory and Practice of Logic Programming*, 18, 3–4, 337–354.
- BASU, K., VARANASI, S., SHAKERIN, F., ARIAS, J. and GUPTA, G. Knowledge-driven natural language understanding of english text and its applications. In *Proceedings of the AAAI Conference on Artificial Intelligence* 2021a, 12554–12563.
- BASU, K., WANG, H., DOMINGUEZ, N., LI, X., LI, F., VARANASI, S. C. and GUPTA, G. 2021b. CASPR: a commonsense reasoning-based conversational socialbot. arXiv preprint arXiv:2110.05387, *abs/2110.05387b*.
- BREWKA, G., EITER, T. and TRUSZCZYNSKI, M. 2011. Answer set programming at a glance. *Commun. ACM*, 54, 92–103.
- BROWN, T. B., MANN, B. et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*. URL: <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- CASOLA, S., LAURIOLA, I. and LAVELLI, A. 2022. Pre-trained transformers: an empirical comparison. *Machine Learning with Applications*, 9, 100334.
- CHEN, Z., MARPLE, K., SALAZAR, E., GUPTA, G. and TAMIL, L. 2016. A physician advisory system for chronic heart failure management based on knowledge patterns. *Theory and Practice of Logic Programming*, 16, 5–6, 604–618.
- COLBY, K. M., WEBER, S. and HILF, F. D. 1971. Artificial paranoia. *Artificial Intelligence*, 2, 1, 1–25.
- FLORIDI, L. and CHIRIATTI, M. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30, 1–14.
- FRANCESCHELLI, G. and MUSOLESI, M. 2023. On the creativity of large language models. arXiv preprint arXiv:2304.00008.
- GALLEGOS, I. O., ROSSI, R. A., BARROW, J., TANJIM, M. M., KIM, S., DERNONCOURT, F., YU, T., ZHANG, R. and AHMED, N. K. 2023. Bias and fairness in large language models: A survey. *arXiv preprint arXiv:2309.00770*.
- GELFOND, M. and KAHL, Y. 2014. *Knowledge Representation, Reasoning, and the Design of Intelligent Agents: Answer Set Programming Approach*. Cambridge University Press.
- GUPTA, G. 2022. Automating common sense reasoning with ASP and s(CASP). Technical Report, UT Dallas. URL: <https://utdallas.edu/gupta/csr-scasp.pdf>.
- KULKARNI, M., TANGARAJAN, P., KIM, K. and TRIVEDI, A. 2024. Reinforcement learning for optimizing rag for domain chatbots. arXiv preprint arXiv:2401.06800.
- LEWIS, P., PEREZ, E., PIKTUS, A., PETRONI, F., KARPUKHIN, V., GOYAL, N., KÜTTLER, H., LEWIS, M., Yih, W.-t., Rocktäschel, T., et al. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- LI, F., WANG, H., BASU, K., SALAZAR, E. and GUPTA, G. Discasp: A graph-based ASP system for finding relevant consistent concepts with applications to conversational socialbots. In *Proc 37th ICLP (Tech. Comm.)* 2021, vol. 345. *EPTCS*, 205–218.
- NEZHURINA, M., CIPOLINA-KUN, L., CHERTI, M. and JITSEV, J. 2024. Alice in wonderland: Simple tasks showing complete reasoning breakdown in state-of-the-art large language models. arXiv preprint arXiv:2406.02061.

- QUIDWAI, M. A. and LAGANA, A. 2024. A rag chatbot for precision medicine of multiple myeloma. *medRxiv*, 2024–03.
- RAJASEKHARAN, A., ZENG, Y., PADALKAR, P. and GUPTA, G. Reliable natural language understanding with large language models and answer set programming. In *Proceedings 39th International Conference on Logic Programming, ICLP 2023*, vol. 385. *EPTCS*, 274–287.
- SARTOR, G., DAVILA, J., BILLI, M., PISANO, G., CONTISSA, G. and KOWALSKI, R. Integration of logical english and s(CASP). In *Proc. ICLP Workshops: GDE'22 2022*, vol. 2970. *CEUR Workshop Proceedings*. [CEUR-WS.org](https://ceur-ws.org).
- TURNER, J. H. 1988. *A Theory of Social Interaction*. Stanford University Press.
- WARREN, D. H. and PEREIRA, F. C. 1982. An efficient easily adaptable system for interpreting natural language queries. *American Journal of Computational Linguistics*, 8, 3–4, 110–122.
- WEI, J., WANG, X., SCHUURMANS, D., BOSMA, M., ICHTER, B., XIA, F., CHI, E. H., LE, Q. V. and ZHOU, D. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, vol. 35, 2022.
- WEIZENBAUM, J. 1966. Eliza—A computer program for the study of natural language communication between man and machine. *CACM*, 9, 1, 36–45.
- XU, Z., ARIAS, J. and Others. Jury-trial story construction and analysis using goal-directed answer set programming. In *Proceedings of PADL 2023*, vol. 13880. *LNCS*, 261–278. Springer.
- ZENG, Y., RAJASEKHARAN, A., PADALKAR, P., BASU, K., ARIAS, J. and GUPTA, G. Automated interactive domain-specific conversational agents that understand human dialogs. In *International Symposium on Practical Aspects of Declarative Languages 2024*, 204–222. Springer.
- ZHENG, L., CHIANG, W.-L., SHENG, Y., ZHUANG, S., WU, Z., ZHUANG, Y., LIN, Z., LI, Z., LI, D., XING, E., et al. 2024. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. *Advances in Neural Information Processing Systems*, vol. 36.