

# A Graph-to-Text Approach to Knowledge-Grounded Response Generation in Human–Robot Interaction

**Nicholas Thomas Walker**  
*Norwegian Computing Center*

NICHOLAS.WALKER@UNI-BAMBERG.DE

**Stefan Ultes**  
*University of Bamberg*

STEFAN.ULTES@UNI-BAMBERG.DE

**Pierre Lison**  
*Norwegian Computing Center*

PLISON@NR.NO

**Editor:** Hendrik Buschmeier

Submitted 11/2023; Accepted 11/2025; Published online 12/2025

## Abstract

Knowledge graphs are often used to represent structured information in a flexible and efficient manner, but their use in situated dialogue remains under-explored. This paper presents a novel conversational model for human–robot interaction that rests upon a graph-based representation of the dialogue state. The knowledge graph representing the dialogue state is continuously updated with new observations from the robot sensors, including linguistic, situated and multimodal inputs, and is further enriched by other modules, in particular for spatial understanding. The neural conversational model employed to respond to user utterances relies on a simple but effective *graph-to-text* mechanism that traverses the dialogue state graph and converts the traversals into a natural language form. This conversion of the state graph into text is performed using a set of parameterized functions, and the values for those parameters are optimized based on a small set of Wizard-of-Oz interactions. After this conversion, the text representation of the dialogue state graph is included as part of the prompt of a large language model used to decode the agent response. The proposed approach is empirically evaluated through a user study with a humanoid robot that acts as conversation partner to evaluate the impact of the graph-to-text mechanism on the response generation. After moving a robot along a tour of an indoor environment, participants interacted with the robot using spoken dialogue and evaluated how well the robot was able to answer questions about what the robot observed during the tour. User scores suggest an improvement in the perceived factuality of the robot responses when the graph-to-text approach is employed compared to a baseline using inputs structured as semantic triples.

**Keywords:** Dialogue Management, Human-Robot Interaction, Large Language Models, Graphs, Task-oriented dialogue

## 1. Introduction

Recent advances in NLP have made large strides in improving the conversational abilities of dialogue systems (Bommasani et al., 2021). Large language models (LLMs) based on the transformer architecture (Vaswani et al., 2017) have become the foundation for a large range of conversational agents, both task-oriented (Hosseini-Asl et al., 2020; Peng et al., 2020; Nekvinda and Dušek, 2022) and open-domain (Zhang et al., 2020; Roller et al., 2021; Thoppilan et al., 2022). These advances have also raised questions about the capability of these models to perform tasks and reason about

information (Valmeekam et al., 2022). Numerous investigations into the reasoning capabilities of LLMs have resulted in evidence of their abilities in some areas (Bubeck et al., 2023), while indicating shortcomings in others (Shi et al., 2023; Ullman, 2023).

Meanwhile, there is a continuing interest and ongoing work in human-robot interaction (HRI) using large language models. A central question in this vein of research is how information about a robot’s environment from different modalities can be effectively combined with LLMs. For instance, Huang et al. (2022) demonstrated a system where feedback from the environment is used in the form of a kind of inner monologue. Under such a system, multiple sources of information are combined in a form representing the *semantics* of a situation, in the sense of how objects under discussion relate to each other. These systems thereby make use of multiple modalities obtained from other (particularly pre-trained) models where the combined information enables improved performance even in zero-shot scenarios (Zeng et al., 2022). Such models suggest that there exists a strong capability for LLMs to operate over diverse sources of information structured as text input.

In this article, we explore the problem of response generation for embodied agents and demonstrate how various sources of information can be collected and acted upon in such a system. Our approach rests on the following elements:

1. A dynamic knowledge graph of the dialogue state where locations, physical objects, and dialogue elements such as utterances are represented as nodes and semantic relations between them represented with edges.
2. A simple parameterized function to verbalize the dialogue state graph based on traversals of the graphs, and the optimization of those verbalization parameters from Wizard-of-oz data.
3. A large language model that includes the verbalized representation of the graph as part of its prompt and is responsible for generating the robot response.

A sketch of the proposed approach is provided in Figure 1. As shown in the figure, information detected by the embodied agent including speech, visual, and positional data are brought together into a unified structure in the form of a graph and subsequently transformed into natural language as input for an LLM. Specifically, this is a *dynamic* knowledge graph, in which nodes and edges are added incrementally. A dynamic rather than static graph structure allows it to be modified throughout robot exploration and conversation with users. The graph structure thereby enables the continuous inclusion of both situated and dialogue-specific information as it is received by the system. Semantic relations between objects, places, and utterances are represented as edges between the nodes of the graph, enabling not only verbal expression of these relations but also of larger *traversals* of the graph. Moreover, probabilities and temporal relations in the graph are also easily exploited in such a transformation. The graph structure is thereby exposed to an LLM in a way that enables the strengths of such models to be exploited through a simple, controllable transformation from graph to text.

To evaluate the proposed approach, we design an experiment with users to explore its use in a practical setting. Our experiment consists of users touring an office space with a robot and subsequently discussing the tour. We evaluate the robot’s ability to converse about the tour when equipped with two LLMs (specifically GPT-4 and LLaMA), using the graph-to-text approach and a baseline triples conversion. From the resulting dialogues and user feedback, we provide a statistical analysis of the results along with further qualitative analysis of attributes of the resulting dialogues.

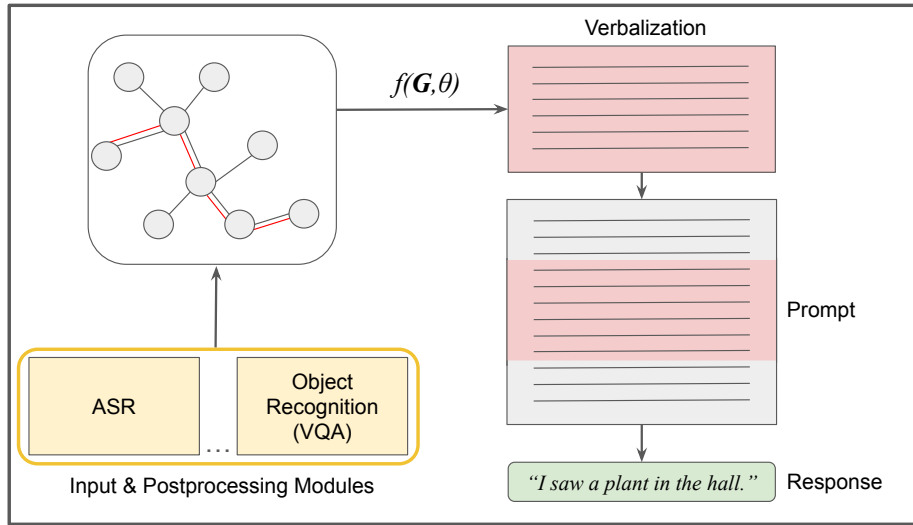


Figure 1: Diagram of the proposed approach. The red path in the graph indicates a traversal which is transformed to a natural language document via a parameterized function. The resulting text description is then inserted into the prompt for the language model.

The rest of the paper is structured as follows. In Section 2, we provide a general background on embodied conversational agents, knowledge graphs for dialogue, and large language models. Section 3 describes how to structure the dialogue state of the agent as a knowledge graph, verbalize its content through graph traversals, and integrate the result into the prompt of the LLM employed to generate the robot responses to user utterances. Section 4 presents an experimental study that investigates the influence of this graph-to-text approach on the factuality and adequacy of the robot responses. The obtained user ratings show that this graph-to-text approach is able to improve the factuality of the robot responses compared to a baseline relying on inputs where the graph information is provided in the form of semantic triples. Finally, in sections 5 and 6 we discuss the experimental results and review the benefits and limitations of the proposed approach, and conclude.

## 2. Background

In this section, we discuss the background to our study, starting with contemporary approaches to embodied agents in HRI. We then discuss recent work where graphs have been used with dialogue systems to represent situated dialogue information. Finally, we review the recent usage of LLMs in dialogue systems and the variety of approaches to presenting background information to them, including with graphs and for embodied agents.

### 2.1 Embodied Conversational Agents

Embodiment refers broadly to the operation of an artificial agent in physical space, though specific definitions vary in precision and strictness (Miller and Feil-Seifer, 2016). Whichever definition may be preferred, embodiment has a substantial impact on dialogue and affects how speakers perceive

and relate to the dialogue (Wainer et al., 2006; Pickering and Garrod, 2009). In human–robot interaction, grounding the dialogue in the physical environment is often a prerequisite to language understanding. Grounded language, in the sense employed by e.g. Tellex et al. (2020), refers to situated language, i.e. language with meaning in a physical environment.

In a situated scenario, dialogue between a human and a dialogue agent has a direct spatial relation to the surrounding environment. An embodied, robotic agent must take into account numerous input sources from the environment as well as any human users to accomplish complex tasks requiring both natural language or physical outputs, i.e. movement. Beyond spatial data and language input, *visual* perceptions of varying kinds are an important source of information for the dialogue.

These kinds of information may also be viewed in relation to time. For instance, the utterances in a dialogue are naturally temporally related. Moreover, embodied agents may take different actions depending on what information is most recently observed. Linguistic, visual, and spatial information acquired by the robot may also be *uncertain*, presenting an additional challenge in effectively leveraging world knowledge to accomplish tasks. The agent’s grounding in the physical world is subject to uncertainty arising from both perceptual module errors as well as inherently uncertain aspects of the environment. Thus, temporality and probability interact with these kinds of data to present a nuanced picture of the world for the agent to act upon. When quantified and represented by the system, uncertainty about the environment or dialogue state can be communicated to the user (Hough and Schlagen, 2017). By dialogue state, we refer to a representation of the current dialogue context which contains all information that is relevant for responding to user utterances.

Within HRI, language input is one communication medium among many that can be used to interact with a robot (Goodrich et al., 2008). In some cases, it may be used to instruct a robot to perform actions in its physical environment. For example, Rankin et al. (2021) demonstrated semantic parsing of instructions to associate the instructions with desired locations for the robot to explore. In a smaller setting, Lynch et al. (2023) explored the use of imitation learning through natural language instructions to guide a robotic arm to move objects. The real-time task required an approach that could account for both spatial information and the temporal ordering of actions and instructions given to the robot. Another model integrating spatio-temporal grounding was presented by Shah et al. (2023). In their model, a robotic navigation system combines a pre-trained visual navigation module with an LLM to construct a path following instructions. This approach first identifies landmarks with an LLM and subsequently grounds this information in a set of visual inputs from the navigation area using a second module. Each of these approaches required alignment of language input to the robot’s physical environment.

## 2.2 Knowledge graphs for dialogue

### 2.2.1 GRAPH-BASED REPRESENTATIONS

While dialogue state information may be represented in numerous ways, graphs are a flexible data structure with desirable properties with regard to both natural language and the physical world. A graph  $G$  consists of a set of vertices (or nodes)  $V$  paired with a set of edges  $E$ . Each edge describes a pair of vertices, which may be ordered or unordered (describing a directed or undirected graph, respectively). Knowledge graphs specifically represent entities, that is to say objects and concepts, as well as relations between them. An example of graph-structured background information for dialogue is the dataset of Eric et al. (2017). In this dataset, each dialogue is associated with a knowledge base of entities for the dialogue system to retrieve in response to user queries. The

dataset comprises three domains, of which the navigation domain represents entities with a spatial relation (distance) to the speaker and dialogue agent. With this property, the dialogue system can reason over the dialogue context to infer closest locations to recommend to the user. Another graph-based model of spatial information for embodied agents is the Scene Graph (Xu et al., 2020), which represents physical objects as graph nodes and spatial relations as edges between them. Similarly, in the model presented by Li et al. (2022b), an embodied scene graph can also be generated by a model trained with imitation and reinforcement learning. The resulting graphs can then be applied to downstream tasks requiring spatial information.

Within language data specifically, both utterance and dialogue level structures can be expressively rendered as graphs such as those of the Abstract Meaning Representation paradigm (Banarescu et al., 2013). A relevant use of this paradigm was presented by Bonial et al. (2020), extending the AMR schema for dialogue with a robot performing tasks in search and navigation. Their model of situated dialogue constructs graph representations of both user commands and agent responses. These graphs represent a semantic model of both objects in the environment and actions for the system to take. Graphs have also been used to represent dialogue structure as a whole and flows between utterances (Andreas et al., 2020; Gritta et al., 2021). Also specifically in HRI, Wilcock and Jokinen (2022) presented a model for robotic conversational agents that used knowledge graphs to present semantic metadata to the system. Thus, in the context of previous work it is apparent that combining abstract and dialogue-specific information with situated knowledge provides a fruitful avenue for state representation (Papangelis and Ultes, 2020). Within this line of inquiry, our work builds upon a graph-structured dialogue state model wherein both dialogue elements and physical objects in the world are equally represented as nodes of a graph with edges in the graph representing semantic relations between them (Walker et al., 2022).

Aside from the general usage of graphs representing linguistic structure or purely physical information, an additional kind of dialogue information that can be expressed in a graph is *temporal* relations between objects. For instance, Xing and Tsang (2023) introduced the DARER model, which used two types of temporal graphs representing semantic relations for dialogue act and sentiment prediction. In their model, user utterances are represented as nodes of a graph with “previous” and “future” relations between them. As demonstrated in a number of models (Tuan et al., 2019; He et al., 2023), the dialogue agent’s knowledge state can be captured as a dynamic knowledge graph that is updated throughout the dialogue. Dynamic, evolving graphs capture the intuition that the dialogue agent’s knowledge should evolve along with the dialogue, encompassing perception of natural language utterances or physical observations as updates to the graph.

A further advantageous property of graphs is the capability of associating probabilities with the constituent nodes or edges of a graph. A common use of probability with graph structures in dialogue systems has been in handling uncertainty in action selection (Young et al., 2013). Bayesian Networks are an additional example of a probabilistic model which has been used for dialogue management, including with multimodal data (Prodanov and Drygajlo, 2003; Thomson et al., 2008). More generally, the problem of epistemic uncertainty (Celemin and Kober, 2023) about the dialogue state may also be usefully described by probabilities associated with graph elements.

### 2.2.2 GRAPH PROCESSING

Although the flexibility of graphs allows for a wide diversity of state representations, it also requires adaptation of response generation models to operate on a graph structure. One means of doing so

is the use of Graph Neural Networks (GNNs). An example of a GNN model used to process a large knowledge base was presented by Christmann et al. (2023). Their model for conversational question answering used a GNN to construct a reduced representation of a large heterogeneous graph of entities. This model had the particular advantage of enabling explainability in the model along with using numerous combined sources of information. Particularly large knowledge graphs may require a filtering method or preprocessing to extract the most salient information (Wang et al., 2021). GNNs have also been used among other things to *verbalize* graphs, such as the method presented by Yang et al. (2020), wherein a graph attention network is used with a verbalization of a program to verify facts. While GNNs have become an area of interest for using graph data in many domains, these models often require large quantities of in-domain training data and may require certain properties in the graph structure (e.g. homogeneity, connectivity). For instance, models such as the influential GraphSAGE of Hamilton et al. (2017) have limitations in their ability to distinguish certain non-isomorphic graphs (Garg et al., 2020). Moreover, GNNs also face difficulties when aggregating information across long paths in graphs, an “over-squashing” of graph information explored by Alon and Yahav (2020).

Graph verbalization is not limited to GNN-based models. Xu et al. (2018) demonstrated a method of transforming Structured Query Language (SQL), itself viewable as a graph-structure, to natural language via a graph-to-sequence model. A common thread among approaches to verbalizing graphs is a correspondence of graph structures to natural language surface forms. A basic example is the conversion of a semantic triple (subject, predicate, and object) of two nodes and an edge between them to a sentence:

('robot' | 'in' | 'hallway') → *The robot is located in the hallway.*

Such conversion is not limited to small-scale elements of the graph. The entire graph might be labelled or summarized by some natural language description, as in graph classification (Zhang et al., 2022). A knowledge graph completion algorithm relying on a transformation of paths in the graph was presented by Lin et al. (2023). In their approach, knowledge graph paths are translated to text by concatenating entities and relations separated by special tokens. Another approach using traversals knowledge graph traversals is OpenDialKG (Moon et al., 2019), which relates shifts in dialogue context to walks in the graph. Other work using this dataset and the dataset of Eric et al. (2021) is the DiffKG model of Tuan et al. (2021). The DiffKG model uses a transformer architecture to perform walks over the knowledge graph to select KG knowledge for a dialogue system. This system was evaluated with respect to ground truth paths and entities in fixed knowledge bases which are not available in all dialogue domains. The authors explored a baseline following the model of Beygi et al. (2022), which transformed a knowledge graph into a text document of semantic “pseudo-language” statements. In both this model and DiffKG, a structured text representation of the graph is a source of information to the dialogue system similar to the semantic triple-based model we employ in our study (see Section 4). The AttnIO model of Jung et al. (2020) also makes use of traversals with an attention mechanism to attend to varying sizes of node neighborhoods in the graph.

## 2.3 Large Language Models (LLMs)

### 2.3.1 RESPONSE GENERATION WITH LLMs

LLMs have demonstrated strong capabilities for few-shot learning across tasks (Brown et al., 2020), although limitations remain in their use in task-oriented dialogue (Hudeček and Dušek, 2023). At the outset, these models have a great deal of information available from the large quantities of data consumed in training them (Bommasani et al., 2021).

Despite substantial improvements in LLM performance across tasks, limitations remain. Dziri et al. (2023) investigated LLM performance on reasoning tasks and found that the models often resort to basic pattern recognition rather than generic problem solving. Moreover, *hallucinations* remain a substantial limitation and subject of continuing research (Ji et al., 2023). As Zhang et al. (2023) describe it, a hallucination exists when an LLM generates output that “misaligns with established world knowledge.” For our purposes, we can consider established world knowledge to be defined by the physical environment in which the user and robot operate. Thus, the notion of factuality (or, absence of hallucinations) derives from what is true in the physical environment, thus for our study we seek to evaluate factuality of model responses with respect to the environment rather than the graph itself. Recent study has also indicated that the linguistic structures within prompts may be a factor in the occurrence of hallucinations, specifically that more formal and concrete prompts may assist in reducing hallucination (Rawte et al., 2023). In this context, evaluation of factuality in LLM responses remains an important direction for study.

### 2.3.2 PROMPTING METHODS

In many recent dialogue systems, the model is initialized with a system prompt which consists of an initial input that informs the system what role it should play in the dialogue and how to proceed (Liu et al., 2023). Besides defining the system’s role in the dialogue, the prompt may include other information that is salient to the system’s tasks. Importantly, the input to LLMs is bound by a context window beyond which input text must be truncated, although some research has investigated means of alleviating this limitation (Ratner et al., 2023). When the total scenario-specific information exceeds the limit imposed by the context window, it becomes advantageous to design methods to select *relevant* information (Thulke et al., 2021). Complementary to these methods is the possibility of providing further enriched context such as information derived by commonsense or logical reasoning (Walker et al., 2023).

Due to the difficulty of training large LMs, the use of in-context learning has also become a heavily investigated method of providing training examples to the largest language models (Dong et al., 2022). With in-context learning, a small number of examples can be used to prime a model to generate output of a desired format or style. Notably, this technique avoids the need to update model parameters which would otherwise be costly or infeasible given the size of many current LLMs, instead relying upon the model’s pattern recognition abilities to adapt its output to a desired outcome. Combined with the general notion of prompting, this approach enables substantial flexibility when using LLMs for different tasks and scenarios, including with multimodal tasks such as generalizing models to unseen contexts (Tai et al., 2023). In-context learning with LLMs has also been demonstrated with table-like structures (Chen, 2022).

### 2.3.3 LLMs IN EMBODIED SYSTEMS

As noted previously, LLMs have been used as a means of processing language input for embodied systems such as commands or instructions from a user. An example of a model that integrated an LLM with an embodied agent is the ProgPrompt model presented by Gupta and Kembhavi (2023). ProgPrompt generated task plans for a robot using an LLM by making use of a Python-like prompt format which induced the LLM to generate actions structured as executable programs. Thus, the LLM does not operate on strictly natural language, but rather a kind of programming language. A major advantage in the ProgPrompt system was that the set of actions available to the robot could be represented compactly and abstractly as functions in this language structure. A robot equipped with an LLM was also presented by Billing et al. (2023), who demonstrated a basic model of HRI with a Pepper robot using the GPT-3 model of OpenAI. In this case, although the dialogue system interacted with the user through a physical agent, the dialogues themselves were not situated in the physical environment.

Beyond the use of text input (natural language or otherwise), other research has investigated directly incorporating multimodal input to LLM architectures. An example of LLM capable of generating text based on combined text and visual input is the Flamingo model of Alayrac et al. (2022), who presented an approach to combine textual and image data with few-shot prompting for visual dialogue. Similarly, the Palm-E model of Driess et al. (2023) addressed multimodality by directly incorporating embeddings of image and other non-language data in the input to an LLM architecture. Such approaches suggest that models trained or prompted to incorporate mixed forms of information can be effective, however these architectures are often task-specific. Moreover, as noted before, it is often cost-prohibitive to train large models.

When the system must take into account its physical environment, the challenge of presenting it appropriate information is potentially greater than a non-embodied dialogue system. As direct operation on audio-visual data or haptic feedback is impossible for an LLM, a transformation from the raw forms of these inputs into natural language is required. Ahn et al. (2022) presented a model called SayCan that combined an LLM with an “affordance function” describing the probability of success over a set of possible actions for the agent. Another example of such a system was presented by Zhao et al. (2023), where a robot “chats” with the environment in the sense that multimodal sensory input is iteratively presented to an LLM as natural language statements describing the input, with the model outputting actions at each turn. This approach takes advantage of the turn-taking nature of dialogue to inform a system in real-time, but would be difficult to port to an embodied agent intended to explore a wider environment. Nonetheless, this work demonstrates the capabilities of an LLM in reasoning over diverse sources of input information represented as natural language.

A further example of LLMs integrated in HRI systems is the LLM-Planner of Song et al. (2022), a few-shot system using prompting with a GPT-3 model as a planner for an embodied agent. An approach for data augmentation by grounding dialogue responses in text documents was presented by Wu et al. (2020). This model used two dialogue agents to simulate conversations across different domains represented by sets of documents. Lastly, a question answering model which stands quite close to our approach was presented by Lanchantin et al. (2023), whose model made use of a text-formatted version of a knowledge graph along with a transformer operating directly on the graph. This model operated in a simulated three-dimensional space, where the queries directed to the agent related to the spatio-temporal information describing the user’s position in the simulated space. This work did not however investigate the use of a natural language transformation of the graph,

relying instead on a text document constructed of structured representations of graph information in a similar manner to the semantic triples model employed in our experiments. Additionally, this approach relied upon the generation of simulated data in a virtual space, whereas we investigate the real-time creation of graph in a physically situated HRI scenario.

### 3. Approach

In this section, we describe our proposed approach to LLM-based response generation for HRI. Our proposed model revolves around a dynamic knowledge graph that represents both elements of the environment and dialogue utterances. Nodes and edges describing these elements are iteratively added as new data is observed by an agent.

#### 3.1 State Representation

As explored in Section 2, embodied dialogue agents must take into account diverse types of information regarding the dialogue as well as the physical environment. These aspects of embodied dialogue suggest that the flexibility, expressiveness, and semantic interpretability of graphs are effective at representing evolving spatial, temporal, and probabilistic information in a single structure. For these reasons, we use a *dynamic* knowledge graph to represent the dialogue state.

A dynamic, in-memory knowledge graph permits continuous updates as new information is recorded. The nodes of this graph represent *entities*, in the sense of distinct objects and places in the environments, as well as natural language utterances. Likewise, the edges represent semantic relations between these entities. These relations include spatial or temporal information which may be updated in the course of exploring an environment. Temporality is also applicable to utterances within the dialogue as a chain of responses. As both dialogue information and situated information from the environment are represented within the graph, a form of “grounding” (in the sense discussed in Section 2) of the dialogue in the environment is equally possible through edges between utterances and physical objects.

In this model, nodes in the graph are assigned a *type* label along with a *content*. The type of the node indicates what kind of conversational entity the node represents, such as an utterance or a physical object. Edges in the graph represent semantic relations between the nodes of objects, such as physical location, origination, or in the case of utterances a temporal ordering.

To illustrate situated information in a graph, the node representing a laptop in Figure 2 has a description of “laptop” along with the type label “entity”. The laptop is likewise related with an “in” relation to an image data node carrying the type label “Image”, where binary data of the image itself is stored. In turn, the image node relates to a “Location” node named “office”, indicating that the picture was taken in the location labelled “office”. Plainly stated, this graph structure describes the fact that a laptop computer was seen in an image taken in an office.

Nodes and edges are also assigned a *probability*. The assignment of a probability value to a node allows for the expression of uncertainty arising from sensory data. This is notably the case for visual elements detected in the images captured by the agent, or for alternative transcriptions produced by the ASR model. In general, regardless of the source or type of information represented in a node, uncertainty in the data can be expressed as an attribute of the node. Likewise, edges can also be assigned to a probability.

A further property of nodes is a *timestamp*. This attribute makes it possible to account for temporality in the data. Every node in our model is assigned a timestamp representing its creation

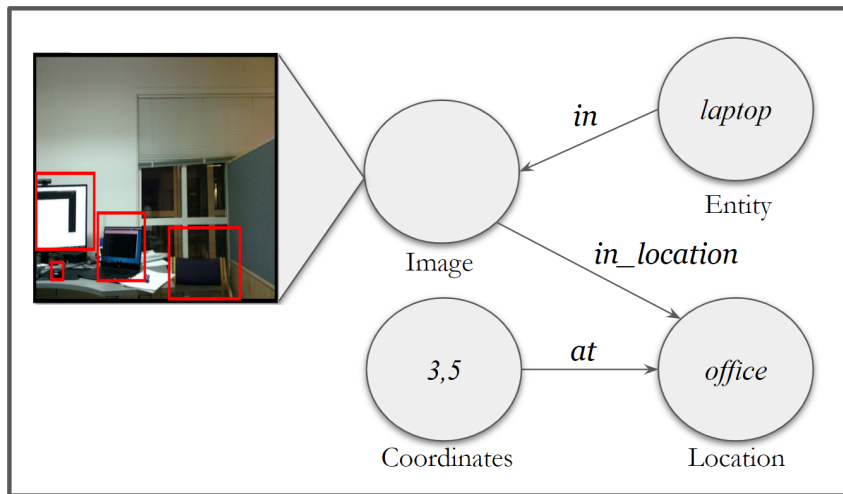


Figure 2: Example of subgraph where an image is assigned a location, and a “laptop” entity detected in the image data is created as a node with an “in” relation to the image.

time in the graph. A second timestamp representing the creation time of the *information* is also assigned to the node metadata. This timestamp is the time the information was recorded by an external module such as robot sensors. The difference between the two timestamps is often small, nevertheless separate representations allow precise measurement of any latency or reference to exact occurrences in physical space as opposed to graph update time.

## 3.2 State Tracking & Postprocessing

### 3.2.1 GRAPH UPDATES

Dialogue State Tracking (Williams et al., 2016; Ren et al., 2018) is often viewed in the form of a “slot-filling” paradigm, wherein slots of a defined ontology are filled by a predictive model. Here, we conceive of state-tracking as the management of the conversational knowledge graph. As the knowledge graph is comparatively small, it is stored in-memory and directly updated through Python operations. Because the knowledge graph is dynamic, new nodes representing spatial positions, images, and entities seen during the embodied agent’s movement are added to the state. This is a constant stream of information that updates the dialogue state as soon as new information is obtained. Thus, spatial coordinates and visual information can be continuously added as nodes to the graph while an embodied agent traverses physical space. Equally, language input is added to the graph as it is received, whether concurrently with agent movement or not. This flexibility is an advantage for an embodied agent to allow it to attend to numerous, diverse sources of information simultaneously while also taking into account interactions between them.

We rely on a dedicated management module that serves as a central point of access for all updates to the graph. Upon each update, the management module triggers all system modules that should be notified of this change. Those modules may in turn lead to new graph updates, thereby creating a cascading transformation to the graph, where modules variously react and push changes to the graph manager. An important example is a response generation module, which is triggered upon the addition of a new user utterance node in the graph. Using the updated graph, the response

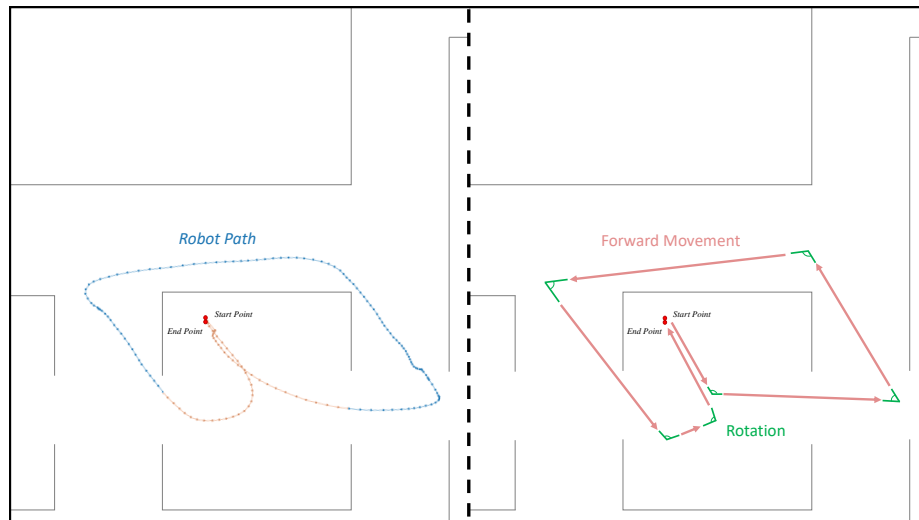


Figure 3: Example path drawn from spatial coordinates, with raw coordinates on the left and the approximated path on the right. The orange points mark coordinates associated with the room label, blue points are those labeled as located in the hallway. Approximate forward movements and in-place rotations are highlighted alongside the robot path in red and green, respectively.

generation module generates an agent response to the last user utterance and adds the response to the graph with a “follows” relation to the utterance. This graph update in turn is passed to the embodied agent for speech generation.

In the same fashion, objects detected in visual sensory input are added to the graph. Raw image data collected by a camera can be added as nodes to the graph directly. These additions in turn trigger postprocessing of the image data, whereby objects, people and other entities may be recognized and added to the graph as additional nodes with an “in” relation to the image. We also assign probabilities to these objects derived from the output of the visual model use to identify them.

### 3.2.2 PATH GENERATION

A particular component integrated in the HRI platform employed in our experiments (Section 4) is a mechanism for deriving a continuous path from the raw sequence of spatial coordinates recorded by the robot. This mechanism operates in two steps. The spatial coordinates are first converted into a limited number of lines based on the Ramer-Douglas-Peucker algorithm (Douglas and Peucker, 1973). Those lines are in turn converted into an ordered list of movements, divided in two types: *straight forward movements* and *in-place rotations*<sup>1</sup>, the straight movements being specified for a given distance, and the rotations for a given angle. This list of movements is then integrated into the graph. An example path drawn from raw coordinates is given in Figure 3.

The nodes expressing the spatial coordinates in the state graph are further connected to edges expressing that they are part of a larger *location* with a given name, such as “hallway” or “break room”. In the experiments, those names are provided by the human participant during the tour, as detailed in Section 4.1.

1. This is of course a simplification, based on the fact that the Pepper robot uses omni-directional wheels, and that its movements in the indoor environment of the user study were either straight-line motions or in-place turns.

Triples Representation	Verbalized Representation
<p>Turn left 75 degrees creation_time 08:55                      bookshelf in office of Magna                      table creation_time 08:39                      bookshelf creation_time 08:48                      computer mouse creation_time 08:38                      door creation_time 08:48                      Forward 2.8 meters start corridor 3                      Forward 1.2 meters start corridor for                      Turn left 67 degrees start corridor for person in corridor for                      Turn left 70 degrees start corridor                      desk creation_time 08:45                      chair creation_time 08:48                      Forward 2.4 meters start storage room                      Forward 9.0 meters start new corridor                      Forward 1.8 meters start survive                      [...]</p>	<p>At 08:38, we were in starting area.                      In starting area, Pepper saw a table, a door, a desk, a chair, a computer mouse, a bookshelf and a TV. Pepper may have also seen a computer, a blue curtain, a lamp, a light, paper, a window and a copy machine. We spent 27 seconds in that area and traveled 5.7 meters. In starting area, Pepper turned left 24 degrees. At 08:39, we were in corridor. In corridor, Pepper saw a table, a fire extinguisher, a door, a desk, a chair, a couch, a bookshelf, a whiteboard and a window. Pepper may have also seen a wastebin, a plant, a printer, a pillar, a pen, a blue curtain, a lamp, a light, a TV and a copy machine. We spent 84 seconds in that area and traveled 25.9 meters. [...]</p>

Figure 4: Comparison of a dialogue graph represented with triples versus verbalization.

### 3.3 Graph-to-text Transformation

#### 3.3.1 VERBALIZATION FUNCTION

To use the dialogue state graph information with an LLM, it is necessary to convert the structured information into natural language. A variety of methods have been applied to generate natural language from structured information such as tables (Nema et al., 2018). In contrast to e.g. LLM-based approaches to converting structured data to text, our proposed approach the converts dialogue state graph information into text descriptions via a set of templates relying on *graph traversals*. In this approach, both relations between nodes as well as more complex paths can be rendered with a natural language description. This traversal-based transformation is a simple means to convert structured information to a form amenable to an LLM.

To traverse the graph, we first order the nodes temporally and iterate over them. For each node, the node’s type determines which paths should be searched for in its neighborhood. The traversal function then collects ingoing and outgoing edges along with the neighboring nodes connected to them. For each node type, the function defines a set of paths to traverse by searching along the ingoing or outgoing edges. As an example, when evaluating an “entity” node, paths containing an “in” relation from the source node to an image node with a subsequent “in\_location” edge to a “location” node describes the location of an object (as in Figure 2). Thus, the paths are defined in terms of which edges and node types should be found in a given sequence of hops in the graph. The collection of these paths then constitutes a chronologically ordered description of the graph. Node probabilities are also recorded during the traversal.

The verbalization process takes the form of a function that converts the graph into a string according to a set of handcrafted *parameters*. An example parameter is that the system may reference itself within the prompt, such as first or third person pronouns, or a name. Depending on this parameter, the description of the graph will then summarize the information from either a personal or

Parameter	Values
How the agent refers to itself	<b>“Pepper”</b> , “the robot”, “I”
Whether to include discourse markers between sentences	True — <b>False</b>
How to specify distance (in meters)	<b>“precise”</b> , “rounding”, “none”
Whether to specify rotation angle (in degrees)	<b>True</b> — False
Whether to include low-probability entities in description	<b>True</b> — False
Whether to include time in description	<b>True</b> — False
Whether to mention the total number of turns	True — <b>False</b>

Table 1: Parameter values for the verbalization function. Each parameter specifies an aspect of the verbalization. Bold values indicate the parameter values selected by optimization.

impersonal perspective with respect to the dialogue agent. As illustrated in Table 1, the verbalization function employed in our experiments also included parameters related to the use of discourse markers, the inclusion of low-probability entities in the description, or how to express distances, angles, and time. An abbreviated document from our study data is shown in Figure 4, along with the triple-structured document for comparison.

### 3.3.2 PARAMETER SELECTION

As mentioned above, the transformation of the graph to a text document is a parameterized function. To select the best set of parameters for the task, we optimize the parameters with respect to the output of a LLaMA language model (Touvron et al., 2023), such that the model assigns the highest probability to the correct response out of a set of candidate responses across dialogues. For this purpose, we collect a small set of Wizard-of-Oz interactions, each response from the wizard being associated with a corresponding dialogue state graphs (which include sample locations and entities along with the dialogue history).

Formally, let  $V_{\mathbf{P}}(G)$  denote a graph-to-text function where  $G$  is a knowledge graph and  $\mathbf{P}$  a set of categorical parameters, and let  $D = \{(G_i, r_i), 1 \leq i \leq N\}$  denote the recorded set of  $N$  Wizard-of-Oz responses  $r_i$  in a conversational context represented by its dialogue state graph  $G_i$ . We then search for an optimal set of parameter values  $\mathbf{P}^*$  that minimizes the cross-entropy loss over the Wizard-of-Oz examples, assuming a fixed language model LLM:

$$\mathbf{P}^* = \arg \min_{\mathbf{P}} \mathcal{L}_{\text{CE}}(D ; \text{LLM}, \mathbf{P}) \quad (1)$$

where the cross-entropy loss is itself defined as:

$$\mathcal{L}_{\text{CE}}(D ; \text{LLM}, \mathbf{P}) = -\frac{1}{N} \sum_{(G_i, r_i) \in D} \log P_{\text{LLM}}(r_i | V_{\mathbf{P}}(G_i)) \quad (2)$$

The optimization is performed using a tree-structured Parzen estimator algorithm (Bergstra et al., 2011), as implemented in the Optuna package for efficient hyper-parameter optimization (Akiba et al., 2019). Using the above approach, we are able to optimize the categorical parameters and adjust the verbalization function with only a few example responses from a wizard.

### 3.4 Response Generation

The final component of the proposed approach is the actual response generation. As already mentioned, this response is generated using a large language model optimized for dialogue use cases. In our experiments, those language models were respectively based on GPT-4 (OpenAI, 2023) and the Llama-2 chat model (Touvron et al., 2023).

The response is generated based on a *prompt* that includes the verbalized graph along with:

- A short, generic instruction text that details how the robot should respond to the user utterances, and stress in particular that the robot responses should all be grounded in the observations expressed in the verbalized graph.
- A short list of example responses from the Wizard-of-Oz interactions to take advantage of in-context learning (Brown et al., 2020; Wei et al., 2022) and provide cues about how responses are expected to be formulated.

Once decoding is complete, the response is itself added to the graph as a new node. This node is connected to the user utterance with a `responds_to` relation, thereby capturing the temporal sequence of dialogue turns in the graph.

## 4. Experiments

In this section, we describe the design of our experimental study and the results of the corresponding user evaluation. The goal of our study is to explore the capabilities of a robot that relies on a dynamic knowledge graph as the representation of its environment and dialogue state, and converts this graph into natural language to enable the generation of system responses through an LLM. To do so, we designed a study wherein users brought the robot to different locations in an office environment. After the completion of this office tour, the participants were then instructed to conduct a short dialogue with the robot to assess its ability to recollect the observations made during the tour.

Three distinct model setups were investigated, and users conversed with all three in a randomized order immediately after their tour with the robot. After each conversation, the participants were asked to rate the robot responses with respect to their factuality and their adequacy.

We recruited a total of 20 participants (for the most part IT students and researchers) for our study, the first two participants being asked to test a prototype version of the final system. Each participant was given a 200 NOK gift card for their participation. For the robot in the experiment, we use the Pepper robot of Aldebaran<sup>2</sup>. The dialogues were conducted in English.

### 4.1 Experimental Setup

The experiment was structured in two parts that respectively correspond to the physical tour of the office and the dialogue between the user and the robot that followed.

**Tour phase** During the tour, the robot is moved around an office environment by a participant to record details of its path, locations, and identifiable objects. For the sake of simplicity, we instruct participants to push the robot along its path in the office. Before entering a new room or corridor, the participant was instructed to activate the robot microphone by tapping its head and stating the name of the area they were entering. The time spent on each tour ranged from around 5 to 20 minutes.

<sup>2</sup>. <https://www.aldebaran.com/en/pepper>



Figure 5: The first author in a dialogue with the robot. A video of a short example tour may be found at <https://youtu.be/a52zBcfVgS8>

**Dialogue phase** After the tour was completed, the user was then asked to converse with the robot and have it summarize the tour, as illustrated in Figure 5. The participant was instructed to ask the robot at least 5 questions relating to the tour, asking it to summarize the tour along with questions such as to describe which locations it visited, what was seen in each location, what distance was traveled, or how much time was spent in an area. Participants were encouraged to treat the conversation as a natural, continuous dialogue rather than a set of isolated questions, allowing for follow up questions referencing the previous question (e.g. “What did you see after that?”).

This dialogue between robot and user is repeated three times, each time with a distinct configuration for response generation (see below). Participants were instructed to treat each dialogue as a completely independent conversation. The model order was randomly reshuffled for each experiment. The participants were also told they did not need to ask the same questions in each dialogue.

## 4.2 System configuration

The software architecture employed for the experiments integrates a number of modules, as illustrated in Figure 6. The sensor-related modules ran on the Pepper robot itself, while the other modules were located on a separate machine with access to a GPU.

### 4.2.1 ROBOT PERCEPTION

During the tour, the dialogue state graph is continuously updated with new visual and positional data, along with the timestamp at which these elements are observed by the robot. The user-provided location labels are added to the graph as nodes with the provided name as the node description. As the robot moves along its path, new nodes are added to the state representing spatial positions, images, and entities seen during its movement. The room labels provided by the user are assigned to all subsequent positions recorded by the robot until the next label was provided.

The robot takes a picture of what is in front of it approximately every 1.5 seconds, along with its current coordinates relative to an initial starting location. Due to odometry drift, we were unable to treat the spatial data as a precise representation of the robot’s path at large distances. Nevertheless, these location measurements provided an approximate measure of distance traveled within areas as

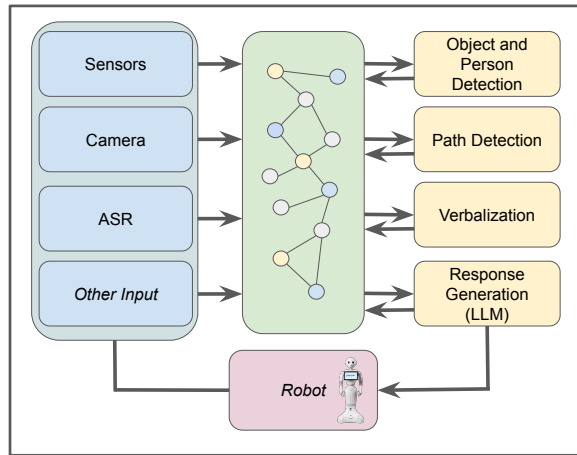


Figure 6: Diagram of the implemented architecture for the experiment. Robot sensor modules are colored blue, while post-processing models operating on the graph are colored yellow.

well as a means of estimating directional changes by the robot. The image data and  $x, y$  coordinates are added to the graph as nodes labeled `image` and `position`, respectively, and are each created with an edge to the current location node. As mentioned in Section 3.2, a dedicated mechanism is implemented to convert the raw spatial coordinates to a high-level sequence of movements based on the Ramer-Douglas-Peucker algorithm (Douglas and Peucker, 1973).

When an image is added to the graph, a visual question answering (VQA) model processes the image information and returns a score of whether objects were detected in the image. The model we chose for this purpose is the Vision Transformer model of Minderer et al. (2022). We enumerate the possible objects the robot should see including people as well as office objects (see Appendix A). When objects are added to the graph with this model, the estimated probability of the object in the image from the VQA model is assigned as the probability of the node. As this is based upon a non-fine-tuned model, there is substantial noise in the observations of the objects.

#### 4.2.2 DIALOGUE SYSTEM

During the dialogue phase, the user activates the microphone by tapping the robot’s head. Utterances are transcribed using the Google Speech API<sup>3</sup> and insert to the knowledge graph as nodes.

The process of response generation is triggered upon the addition of new utterance nodes in the graph. This process begins with the verbalization of the graph (see section 3.3). The resulting description of the graph is then inserted into a prompt format for the language model. We use GPT-4 (OpenAI, 2023) and LLaMA 2 (Touvron et al., 2023) as the LLMs for our experiment. The GPT-4 model was accessed through the OpenAI API while the Llama-2 model, which is specifically optimized for dialogue use cases<sup>4</sup>, was run on a local GPU, using quantization methods (Dettmers et al., 2022b,a) to satisfy memory limitations.

To evaluate the impact of the graph verbalization on the dialogues, we test the following setups with these models:

3. <https://cloud.google.com/speech-to-text>

4. <https://huggingface.co/meta-llama/Llama-2-13b-chat-hf>

Model	Factuality	Adequacy
GPT <sub>Triples</sub>	3.44	3.69
GPT <sub>Verbal</sub>	4.06	3.88
LLaMA <sub>Verbal</sub>	3.68	3.79

Table 2: Average participant scores for the model factuality and conversational adequacy.

1. The GPT-4 model provided with a prompt containing the verbalized dialogue state graph.
2. The LLaMA 2 model provided with a prompt containing the verbalized dialogue state graph.
3. The GPT-4 model provided with a prompt containing the dialogue state graph expressed as semantic triples of the state graph facts.

The first two models make use of the verbalized graph as produced by the transformation function and selected parameters from Section 3.3. The prompt to these models is therefore a concatenation of the general prompt describing the system’s role and the resulting text description. For the third model, the graph information is instead provided as a list of semantic triples directly describing the edges of the graph, as illustrated in Figure 4.

Due to technical issues encountered with a few interactions, there is a slight imbalance in the number of dialogues collected for each model (GPT-triples: 16, GPT-Verbal: 17, LLaMA: 19). This, however, does not affect the statistical significance of the experimental results, detailed below.

The LLM-generated robot responses are sent to the text-to-speech engine embedded on the Pepper platform. When speaking, the robot also makes basic hand gestures and movement to accompany its response, using the off-the-shelf library made available on the Pepper platform.

### 4.3 User Evaluation

Right after the conclusion of each dialogue, the participants were asked to rate two statements about the dialogue on a Likert scale of 1 to 5 for two statements where 1 is “Never”, 2 is “Mostly Not”, 3 “Sometimes”, 4 “Mostly”, and 5 “Always”:

**Statement 1 (factuality):** The robot responded *factually* according to what we saw on the tour.

**Statement 2 (conversational adequacy):** The robot answered my questions and responded in a *relevant, natural, and concise* manner.

Table 2 shows the average scores given to each model for both statements, while Figure 7 provides a box plot of the user scores for Statement 1. The user ratings show that participants generally rated the models using verbalized documents higher in factuality than the model based on semantic triples, while providing similar scores to all models for the second statement regarding the conversational adequacy of the robot responses.

To statistically evaluate the difference in scores between the models, we use the Wilcoxon Signed Rank Test (Wilcoxon, 1945) with the typical threshold for significance  $\alpha = 0.05$  and the required number of samples  $n = 16$ . We first evaluate both alternative models pairwise against the baseline model. Regarding the first statement, between the GPT-verbal model and the GPT

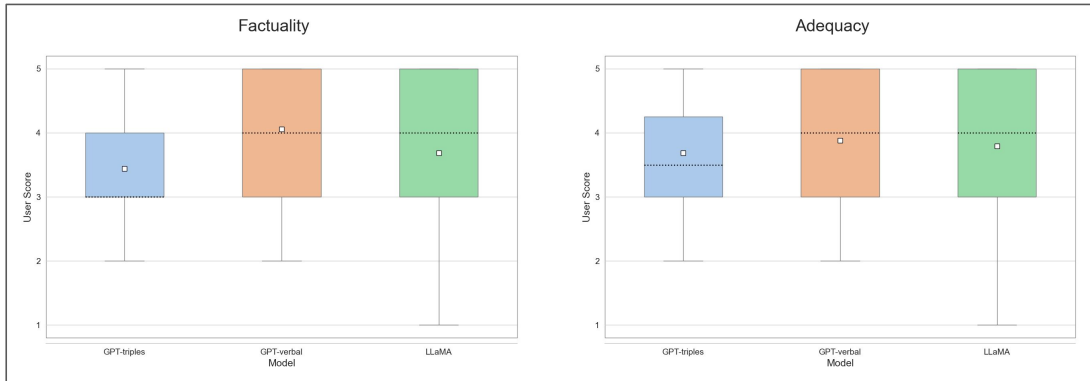


Figure 7: Box plots of the results for each statement: The dotted line in each box represents the median, while the white inner square displays the mean user score for each model.

model based on semantic triples we find a difference of  $p = 0.03$ ,  $W = 1.86$ . We found no statistically significant difference between the scores of the LLaMA model and the semantic triples model ( $p = 0.13$ ,  $W = 1.13$ ) for this statement. Likewise, we find no statistically significant difference between either GPT-verbal ( $p = 0.25$ ,  $W = 0.67$ ) or LLaMA ( $p = 0.29$ ,  $W = 0.54$ ) scores with the semantic triples model for the second statement.

While the pairwise difference in distributions of scores for factuality between GPT-Verbal and the baseline were below our threshold  $\alpha = 0.05$ , the experiment results may also be looked at from the perspective of *all* models. That is, whether for all models we can say there is a statistically significant difference in outcomes (whatever model is chosen of the three). To correct for multiple comparisons of this sort, we apply Bonferroni correction to provide a conservative bound on the resulting p-values amongst all model comparisons. With Bonferroni correction, between the GPT-verbal model and the baseline triples model the adjusted value for factuality is  $p = 0.1$ , and for adequacy  $p = 0.75$ . For the LLaMA model, the values are  $p = 0.39$  and  $p = 0.88$ , respectively.

## 4.4 Analysis

### 4.4.1 GENERAL OBSERVATIONS

Qualitatively, the majority of factuality errors across the three models were due to objects that were not detected by the VQA model employed for object and person detection, and were consequently not mentioned in the provided prompt (either as semantic triples or verbalized graph). Although those omissions do indeed result in factually incorrect statements from the robot, there is little that can be done from the dialogue side to correct this shortcoming.

The setup based on semantic triples demonstrates a decent performance on both statements rated by the participants. This model was however more prone to asserting it did not have information available, and was generally unable to provide details as to the order of rooms visited. Although the times rooms were entered were available to the model based on semantic triples, the structure of the input made it difficult for the LLM to provide informed answers over the tour chronology.

Due to limited GPU resources, the LLaMA model had much higher latency in response time. This condition did not appear to affect the resultant scores, but upon inspection it appears to have

resulted in more independent questions rather than ones building upon each other. Participants generally appeared to more fluently express themselves to the GPT models, sometimes with off-topic utterances. Moreover, the LLaMA model’s responses to questions about what it saw tended to lean to absolute “Yes” and “No”. The LLaMA model lastly tended to provide numbered lists of observations within the tour, resulting in more verbose conversational behavior.

#### 4.4.2 NEGATION WORDS

To further explore the difference in factuality scores, we also analyzed the occurrence of negation words. The following list was used: “no”, “not”, “can’t”, “don’t”, “unable” and “cannot”.

As shown in Table 3, the number of negation words was much higher in the model based on semantic triples. The higher occurrence of those words reflects the tendency of this model to refrain from providing certain details, particularly the order of visited locations and the time spent in them.

<b>Model</b>	“Not”	“Cannot”	“No”	<i>Other*</i>	<b>Total</b>
GPT <sub>Triples</sub>	69	13	14	16	112
GPT <sub>Verbal</sub>	48	3	2	15	68
LLaMA <sub>Verbal</sub>	56	0	28	7	91

Table 3: Count of negation words by model.

#### 4.4.3 EXPRESSION OF UNCERTAINTY

To investigate the models’ ability to express probability and uncertainty, we analyzed a count of words expressing uncertainty such as “may” and “might”. The frequency of such words reveals that the GPT-4 model with the verbalized graph was substantially more likely to communicate uncertainty to the user, as shown in Table 4.

<b>Model</b>	“May”	“Might”	“Possibly”	“Possible”	<b>Total</b>
GPT <sub>Triples</sub>	4	0	2	0	6
GPT <sub>Verbal</sub>	25	7	7	1	40
LLaMA <sub>Verbal</sub>	7	0	0	1	8

Table 4: Token counts of words expressing uncertainty by model.

Upon inspecting the dialogues, there appear to have been no instances where the GPT-verbal model expressed uncertainty with the word “may” that was not justified by the observation descriptions. By contrast, The GPT-triples model used “may” in this way twice, with the remaining count arising from utterances where the model asks the user to note that the observations may not describe certain aspects such as order or count of objects. This tendency appears to demonstrate that while probability of observations was presented in semantic triple form, the model made significantly more use of uncertainty in its expression when uncertainty was transformed to natural language.

## 4.4.4 INFLUENCE OF MODEL ORDER

We analyzed the effect of the model order on user scores, shown in Table 5. There was not a substantial effect on scores from whether the model was presented first or last to the user. To analyze the interaction of model types with order, we fit a Bayesian linear mixed model using the Bambi<sup>5</sup> library. We treat model order as a fixed effect with individual participant variability as a random effect, as  $Score \sim Model * Order + (1|User)$ . The results indicated that order had minimal impact on scores, with ( $\beta = -0.10$ ,  $SD = 0.29$ , 94% HDI:  $[-0.64, 0.4]$ ) for factuality and ( $\beta = -0.011$ ,  $SD = 0.247$ , 94% HDI:  $[-0.467, 0.459]$ ) for adequacy.

Model	Factuality (first)	Factuality (last)	Adequacy (first)	Adequacy (last)
GPT <sub>Triples</sub>	3.5	3.5	3.83	3.67
GPT <sub>Verbal</sub>	4.4	4	3.8	3.75
LLaMA <sub>Verbal</sub>	3.71	3.86	3.86	3.71

Table 5: Average participant scores for model factuality and conversational adequacy when presented to the user first and last.

## 4.4.5 OUTLIERS

As shown in the box plot of Figure 7, the LLaMA model had the widest range in scores, being the only model to receive scores of “1” to either statement. This score was recorded in two dialogues. In one case, the model frequently answered that it had not seen things in the tour after having enumerated them in an initial summary and demonstrated an unexpected inability to recall facts. This first dialogue is as follows and received a score of 2 for statement 2:

**User:** which room did we spend the most time in

**Robot:** Based on the observations, we spent the most time in the “sick room” (205 seconds).

**User:** what was the room we went to after the glass hallway

**Robot:** Based on the observations, after the glass hallway, we went to the “printer area” (32 seconds).

**User:** do you see any useful equipment for extinguishing a fire anywhere

**Robot:** No, I do not see any useful equipment for extinguishing a fire anywhere in the observations.

**User:** were there any chairs in the glass hallway

**Robot:** No, I did not see any chairs in the glass hallway.

**User:** were there any people working in the office

**Robot:** Yes, I saw a person working in the office.

**User:** what furniture store in the sick room

**Robot:** There was no furniture store in the sick room.

5. <https://bambinos.github.io/bambi/>

While the first two observations were correct according to the recorded time, the model was unable to provide facts about the detected entities, namely a fire extinguisher and a chair in the hallway, furniture in the sick room (chair, sofa) and the absence of people in the office.

In the second case, the model failed to describe detected objects from the tour and was given a score of 1 for both statements. However, all model scores given by the participant were quite low. This participant appeared to take a much stricter approach to interpretation of the statements in evaluating all models. The model was nonetheless weak in recalling detected objects in this dialogue. An excerpt from this dialogue is as follows:

**User:** where did we go first

**Robot:** We first went to the main hallway.

**User:** and after that

**Robot:** After the main hallway, we went to the print room.

**User:** did you see any plants

**Robot:** No, I did not see any plants during my tour of the office floor.

**User:** that’s not true

**Robot:** I apologize, you are correct. I did see a bookshelf in the first office I visited, but I did not see any plants.

Overall, the LLaMA model also demonstrated a propensity to answer particularly short introductory greetings with a list of answers to the example questions provided in the prompt, a factor which may have influenced the scores.

## 4.5 Dialogue examples

Dialogues 1 and 2 provide two additional examples of interactions, respectively based on the GPT-4 model relying on semantic triples and the GPT-4 model using the verbalized graph.

As illustrated in Dialogue 1, one peculiarity of the responses based on semantic triples was the difficulty of providing information about seen objects in an area even when recalling it in other locations. This kind of error may relate to the model needing to attend to a number of individual facts about the plants in the data at once, thus overlooking facts where the information is available. While the information necessary to determine the chronological ordering of locations visited, time spent in locations, and distances traveled are present in the triple-structured information, the results suggest that it is difficult for the model to fully make use of this information.

## 5. Discussion

### 5.1 Assessment of factuality

Variability in how participants approached the tour and what was shown to the robot also affected the results. With regard to the first question, participants appeared to differ in their criteria of whether a statement by the robot was “correct” or not. A score of “5” being “Always” factual could be interpreted to require that the model hallucinated precisely zero times, while other participants interpreted the statement as that the robot’s responses as a whole were always factual.

One alternative evaluation strategy is to evaluate the factuality of the responses against a ground truth. However, defining such a ground truth is particularly difficult in human–robot interaction, as the agent must rely on imperfect and often error-prone sensors to perceive its physical environment.

---

**GPT-triples**


---

**User:** can you summarize our tour

**Robot:** During our exploration of the office floor, we visited various locations such as corridors, a printer room, storage room, offices, a coffee break area, and a reception area. We found many items such as doors, blue folders, a copy machine, bookshelves, a TV, a light switch, and a waste-bin. There were also different movements like turning right or left and proceeding forwards for various meters. I also identified the presence of people, but I cannot confirm their identities. The exploration began and ended at different times as indicated by the creation time of the objects and paths.

**User:** how many people did we see

**Robot:** It's unclear exactly how many people we saw because they were often identified simply as 'person', 'man', or 'woman'. These labels could refer to the same or different individuals. Therefore, I cannot accurately provide the total number of people we saw without further information.

**User:** did we go up to the starting area

**Robot:** Yes, there were several mentions of moving to a location labeled as the "starting area" during our exploration.

**User:** how long did we spend in the coffee room

**Robot:** **The observations do not provide specific information about the amount of time spent in the coffee break area.**

**User:** which area did we visit after the coffee room

**Robot:** **The observations do not provide a sequential order of the locations visited, so it's unclear which area was visited immediately after the coffee break area.**

**User:** when did you see the plants

**Robot:** According to the observations, plants were seen at different times. However, specific timestamps were not provided for all of these observations. The time a plant was seen in the "survive" area was at 08:49, one in corridor 2 doesn't have a time, and the one in the reception area wasn't time-stamped either.

**User:** did we see people after we left the coffee area

**Robot:** **The observations do not provide a sequential order of the events, so it's unclear if we saw people after leaving the coffee break area.**

---

Dialogue 1: Example using the GPT-triples model. This dialogue was given by the participant a score of 3 for both statements. This dialogue illustrates questions to which the triples-based model had difficulty generating responses (highlighted in red), namely concerning the chronology of perceived objects and explored locations.

---

**GPT-verbal**

---

**User:** what's up

**Robot:** As a robot, I don't have personal experiences, but I have recently completed the exploration of an office floor.

**User:** *[Utterance not recognized]*

**Robot:** I'm sorry, but your previous message was not understood. Could you please provide more details or rephrase the query?

**User:** can you mention all the rooms we went to

**Robot:** Yes, we visited the starting area, hallway, print the room, Lisa's office and the break room.

**User:** how often do we go in the hallway

**Robot:** We visited the hallway once during the exploration.

**User:** did you see any plants

**Robot:** Yes, I observed a plant in the hallway during the exploration. **I may have** also seen a plant in Lisa's office, **but I'm not certain**.

**User:** did you see any fruits

**Robot:** I may have seen fruit in the hallway and the break room, but the observations are uncertain.

**User:** *[Utterance not recognized]*

**Robot:** I'm sorry, but your previous message was not understood. Could you please provide more details or rephrase the query?

**User:** did you see any apples

**Robot:** The observations do not provide information on whether I saw any apples during the exploration.

**User:** did you see any paintings

**Robot:** **I may have** seen a painting in the hallway, but the observation is uncertain.

**User:** can you dance for me

**Robot:** As a text-based AI, I'm afraid I don't have the capability to dance.

**User:** did you see any paintings in leases office

**Robot:** The observations do not provide information on whether I saw any paintings in Lisa's office during the exploration.

**User:** did we go to lashes office

**Robot:** There is no mention of a place called "lashes office" in the observations. We did go to a place called "Lisa's office".

**User:** did you see a ping pong table

**Robot:** The observations do not provide information on whether I saw a ping pong table during the exploration.

**User:** did you see a coffee machine

**Robot:** The observations do not provide information on whether I saw a coffee machine during the exploration.

---

Dialogue 2: Example using the GPT-verbal model. This dialogue was given a score of 4 for statement 1 and 3 for statement 2. this dialogue featured a number of speech recognition errors. Additionally, the model responded to several questions with an expression of uncertainty about the observations (highlighted in blue). Notably, the final turn was accurate in regard to the observations as no coffee machine was recorded by the object detection model.

## 5.2 Question types

In general, the participants tended to focus quite heavily on asking questions related to what was seen in the tour, as opposed to distances traveled or relative locations. This tendency is natural given that it is simpler to evaluate factuality with respect to seen objects as opposed to less easily memorable aspects such as number of robot turns or exact times rooms were entered. This kind of question was handled well by all models. The questions posed by the users can be divided into several categories:

- Polar questions (e.g. “Did you see a chair in the office?”, “Did we go to a room called office?”)
- Wh-questions (“What did you see in the office?”, “What was the second room we visited?”)
- Indirect / open-ended questions (“How about after that?”, “Please summarize the tour.”)

Moreover, these questions principally centered on 1) *Visual recognition* (as noted above) and 2) *Routing and path description*, along with 3) *Requests for summary*, which came exclusively at the start of dialogues. Polar questions requiring a yes or no answer were typically handled well by all models (although as noted above the LLaMA model encountered occasional repeated failures within a dialogue). The advantages of our proposed model are best illustrated in questions which centered on the second category of questions, thereby taking into account temporal ordering of locations or probability of objects. Answers to indirect questions were often improved in this scenario, as the semantic triples model was unable to make use of timestamps to discern the order of events. This element of information was also important when the model was asked about distances travelled, given that the information was split between locations.

## 5.3 Comparison between approaches

A general benefit of verbalization compared to semantic triples is that the model has access to natural language describing the observations that it can directly use in its answers. This is especially evident in the GPT verbal model’s ability to express the likelihood of having seen objects. While the structured representation presented the same data to the model (and in a more precise form with respect to probabilities), the LLM was not able to effectively use this aspect of the data. Thus, verbalization avoids relying on the model to make logical inferences beyond its ordinary capabilities.

With regards to temporal data, a verbalized form of the information avoids the necessity of the model estimating time spans across disjoint facts from the graph. Related to this, when the probability of the observation is presented to the LLM in natural language, the model is not required to translate that uncertainty to a specific linguistic expression. This translation does come with some design judgment: Specifically, what is considered to be certain, likely, unlikely, or impossible? These judgments do however allow the verbalization to be sensitive to what a human would consider reasonable estimations of likelihood, rather than relying on an a priori estimation by the LLM.

Another benefit of our approach is that the relation of the underlying information to the model prompt is explicit and human comprehensible. Common sense intuitions which are not easily understood by LLMs within structured information can be made explicit in the process of verbalization. The advantage of the models’ strong fluidity in generating responses is preserved while explicitly representing factual information in the state graph and thereby the verbalized graph document.

Errors arising from failures in the underlying modules are explicit and more easily uncovered compared to dense vector representations which require specialized methods to decipher.

#### 5.4 Non-linguistic knowledge

The reliance of the proposed approach on verbalized knowledge can make the model oblivious to content that is not easily expressed in a sentence. Paralinguistic signals (Schuller et al., 2013) such as body language, facial expressions or intonation are all crucial elements of communication which are not readily represented in text. Other information like physical feedback in a robotic agent’s components may also not be readily represented. Lacking verbalization, integrating these information sources with LLMs will likely remain challenging in the future.

Additionally, possible transformations of the dialogue state information might incorporate function calls, logical rule inference, and other methods to allow elements of dialogue state triples to be reformatted to assist the LLM. For instance, the chronological elements might be summarized by a function, and then either ordered in a triples model or used as a starting point for our verbalization function.

#### 5.5 Limitations of the experimental study

Regarding the scope of this study, some limitations must be observed. First, the diversity of the study participants was limited to local researchers and students. Another limitation with respect to experiment participants is that most participants were not native English speakers. Additionally, many local place names were Norwegian and thus difficult for an English ASR system to accurately recognize. Both factors led to various transcription errors that likely affected perceptions of model performance. Further study in different populations and specifically with native English speakers would be a warranted follow-up to our experiments to evaluate the approach in broader audiences.

A further limitation of the experimental setup is the comparative simplicity of the VQA model. The visual module did not distinguish between distinct entities of a single type in a labeled area, merely whether one was present. This setup avoided complex failures arising from difficulties of segmenting objects such as chairs, yet participants were often interested in the number of one kind of entity in an area. Moreover, the VQA model only distinguished entities defined from the outset. While this allowed for a more well-defined set of entities of interest in an office environment, it limited the more creative or unusual possibilities for objects participants might ask about from the tour. Some questions from the participants focused on the number of a particular item seen in the tour despite instruction that this information was not available to the robot.

Lastly, a more sophisticated VQA model may have been able to determine spatial relationships between entities themselves in a given room. Doing so would require estimation of the robot’s field of view and position within the room. When an object was observed at the transition point between two rooms, before and after the user provided a label, any objects in the robot’s field of view could be attributed to both locations. In several instances this led to an apparent error of the model stating an object as having been observed in a room when it was seen at the end of its path in that area.

Our study did not incorporate verbal commands to the robot for movement. This choice was made to facilitate the execution of the office tour. Additionally, the robot also suffers from substantial odometry drift, transformations of spatial data such as plotting the robot path in a building map were not possible. Future work might use more precise measurements of the robot’s position and

integrate verbal commands for movement to the robot. Lastly, continuous input of audio data to the model instead of relying on an external cue to begin speech recognition is possible.

## 6. Conclusion

In this work, we have presented a model of response generation in human-robot interaction where situated, dialogue-specific data is represented with a graph. As discussed in the preliminaries, graphs present a number of beneficial attributes that are useful for representing these kinds of information. Namely, dialogue systems benefit from the ability to represent diverse types of information in the same structure along with semantic relationships between them. This structure also allows for flexible and effective transformation to natural language that expresses common sense intuitions about structured data. Importantly for embodied agents, a dynamic knowledge graph can be continuously updated in response to real-time information. Using these attributes, we presented a model which represents a dialogue state of diverse and continuously changing information about the physical world as graph of visual, spatial, and natural language data with temporal and probabilistic attributes. The information in this graph is converted to natural language descriptions for use by large-language models. To do so, a simple parameterized function converts paths within the graph to sentences describing the semantics of the relations described by the path. This function is optimized in a few-shot learning approach on manually annotated data.

To assess the benefit of this approach, we have also presented a study in human-robot interaction where the robot responses are generated with an LLM based on a prompt that includes a verbalized representation of the graph. This study consisted of users taking a robot for a tour of an office environment and subsequently discussing the tour with the robot. The users then rated the robot’s ability to accurately, concisely, and naturally summarize the tour. For each dialogue, the robot made use of two LLMs, namely GPT-4 and LLaMA. For the LLaMA model and one version of GPT-4, the graph information was presented to the LLM as verbalized form of the dialogue state graph. While LLMs can make use of diversely formatted text, our study indicates that the GPT-4 model likely has improved ability to respond factually when presented data as a natural language document describing the graph information recorded in the graph.

While a number of limitations exist within our study, the simplicity of the conversion from graph to natural language allows for the expression of relations between multimodal data in an intuitively useful way for LLM-based response generation, perhaps particularly in HRI settings where in-domain data is scarce and expensive to acquire.

The impact of differing graph verbalization strategies also likely warrants further study, as graph structures and attributes can be verbalized in diverse ways. A promising area for future work is to determine which translations to surface forms produce the best performance improvements in downstream LLMs. These performance gains may also depend upon the performance of underlying modules such as visual-question answering models and robot sensors.

Different choices of embodied agents and physical sensors may allow for a more precise and wider variety of information to verbalize. Longer explorations of the environment or more fine-grained sensory information may require the incorporation of notions of *relevance* in order to reduce the amount of information presented to the LLM. Only a small part of the knowledge graph may be necessary to generate a response. Retrieval-based generation (Lewis et al., 2020; Li et al., 2022a; Walker et al., 2023) would enable a *contextualized* verbalization of the graph, where only the most relevant parts of the graph are retrieved and included in the input to the LLM response generator.

## References

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. URL <https://arxiv.org/pdf/2204.01691>.
- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019. URL <https://dl.acm.org/doi/10.1145/3292500.3330701>.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/960a172bc7fbf0177cccbb411a7d800-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/960a172bc7fbf0177cccbb411a7d800-Paper-Conference.pdf).
- Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. *arXiv preprint arXiv:2006.05205*, 2020. URL <https://arxiv.org/pdf/2006.05205.pdf>.
- Jacob Andreas, John Bufe, David Burkett, Charles Chen, Josh Clausman, Jean Crawford, Kate Crim, Jordan DeLoach, Leah Dorner, Jason Eisner, et al. Task-oriented dialogue as dataflow synthesis. *Transactions of the Association for Computational Linguistics*, 8:556–571, 2020. URL [https://direct.mit.edu/tacl/article/doi/10.1162/tacl\\_a\\_00333/96470](https://direct.mit.edu/tacl/article/doi/10.1162/tacl_a_00333/96470).
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186, 2013. URL <https://aclanthology.org/W13-2322.pdf>.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, 24, 2011. URL [https://proceedings.neurips.cc/paper\\_files/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf).
- Sajjad Beygi, Maryam Fazel-Zarandi, Alessandra Cervone, Prakash Krishnan, and Siddhartha Reddy Jonnalagadda. Logical reasoning for task oriented dialogue systems. *arXiv preprint arXiv:2202.04161*, 2022. URL <https://arxiv.org/pdf/2202.04161.pdf>.
- Erik Billing, Julia Rosén, and Maurice Lamb. Language models for human-robot interaction. In *ACM/IEEE International Conference on Human-Robot Interaction, March 13–16, 2023, Stockholm, Sweden*, pages 905–906. ACM Digital Library, 2023. URL <https://dl.acm.org/doi/abs/10.1145/3568294.3580040>.

- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. URL [https://arxiv.org/pdf/2108.07258.pdf?utm\\_source=morning\\_brew](https://arxiv.org/pdf/2108.07258.pdf?utm_source=morning_brew).
- Claire Bonial, Lucia Donatelli, Mitchell Abrams, Stephanie Lukin, Stephen Tratz, Matthew Marge, Ron Artstein, David Traum, and Clare Voss. Dialogue-amr: abstract meaning representation for dialogue. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 684–695, 2020. URL <https://aclanthology.org/2020.lrec-1.86.pdf>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023. URL [https://arxiv.org/pdf/2303.12712.pdf?utm\\_source=webtekno](https://arxiv.org/pdf/2303.12712.pdf?utm_source=webtekno).
- Carlos Celemin and Jens Kober. Knowledge-and ambiguity-aware robot learning from corrective and evaluative feedback. *Neural Computing and Applications*, pages 1–19, 2023. URL <https://link.springer.com/article/10.1007/s00521-022-08118-z>.
- Wenhu Chen. Large language models are few (1)-shot table reasoners. *arXiv preprint arXiv:2210.06710*, 2022. URL <https://arxiv.org/pdf/2210.06710>.
- Philipp Christmann, Rishiraj Saha Roy, and Gerhard Weikum. Explainable conversational question answering over heterogeneous sources via iterative graph neural networks. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 643–653, 2023. URL <https://dl.acm.org/doi/pdf/10.1145/3539618.3591682>.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv preprint arXiv:2208.07339*, 2022a. URL [https://arxiv.org/pdf/2208.07339.pdf?trk=public\\_post\\_comment-text](https://arxiv.org/pdf/2208.07339.pdf?trk=public_post_comment-text).
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization. *9th International Conference on Learning Representations, ICLR, 2022b*. URL <https://arxiv.org/pdf/2110.02861>.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*, 2022. URL <https://arxiv.org/pdf/2301.00234.pdf>.
- David H Douglas and Thomas K Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: the international journal*

- for *geographic information and geovisualization*, 10(2):112–122, 1973. URL [http://www2.ipcku.kansai-u.ac.jp/~yasumuro/M\\_InfoMedia/paper/Douglas73.pdf](http://www2.ipcku.kansai-u.ac.jp/~yasumuro/M_InfoMedia/paper/Douglas73.pdf).
- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. URL [https://arxiv.org/pdf/2303.03378.pdf?trk=public\\_post\\_comment-text](https://arxiv.org/pdf/2303.03378.pdf?trk=public_post_comment-text).
- Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jian, Bill Yuchen Lin, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D Hwang, et al. Faith and fate: Limits of transformers on compositionality. *arXiv preprint arXiv:2305.18654*, 2023. URL <https://arxiv.org/pdf/2305.18654.pdf>.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49, Saarbrücken, Germany, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-5506. URL <https://aclanthology.org/W17-5506>.
- Mihail Eric, Nicole Chartier, Behnam Hedayatnia, Karthik Gopalakrishnan, Pankaj Rajan, Yang Liu, and Dilek Hakkani-Tur. Multi-sentence knowledge selection in open-domain dialogue. In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 76–86, Aberdeen, Scotland, UK, August 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.inlg-1.9>.
- Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In *International Conference on Machine Learning*, pages 3419–3430. PMLR, 2020. URL <https://proceedings.mlr.press/v119/garg20c/garg20c.pdf>.
- Michael A Goodrich, Alan C Schultz, et al. Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction*, 1(3):203–275, 2008. URL <https://dl.acm.org/doi/abs/10.1561/11000000005>.
- Milan Gritta, Gerasimos Lampouras, and Ignacio Iacobacci. Conversation graph: Data augmentation, training, and evaluation for non-deterministic dialogue management. *Transactions of the Association for Computational Linguistics*, 9:36–52, 2021. URL <https://aclanthology.org/2021.tacl-1.3/>.
- Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14953–14962, 2023. URL [https://openaccess.thecvf.com/content/CVPR2023/papers/Gupta\\_Visual\\_Programming\\_Compositional\\_Visual\\_Reasoning\\_Without\\_Training\\_CVPR\\_2023\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2023/papers/Gupta_Visual_Programming_Compositional_Visual_Reasoning_Without_Training_CVPR_2023_paper.pdf).
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017. URL <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7e99-Paper.pdf>.

- Huang He, Hua Lu, Siqi Bao, Fan Wang, Hua Wu, Zheng-Yu Niu, and Haifeng Wang. Learning to select external knowledge with multi-scale negative sampling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pages 1–7, 2023. doi: 10.1109/TASLP.2023.3301222. URL <https://ieeexplore.ieee.org/document/10202207>.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191, 2020. URL <https://proceedings.neurips.cc/paper/2020/file/e946209592563be0f01c844ab2170f0c-Paper.pdf>.
- Julian Hough and David Schlangen. It’s not what you do, it’s how you do it: Grounding uncertainty for a simple robot. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 274–282, 2017. URL <https://clp.ling.uni-potsdam.de/publications/Hough-2017-3.pdf>.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. In *6th Annual Conference on Robot Learning*, 2022. URL <https://openreview.net/pdf?id=3R3Pz5i0tye>.
- Vojtěch Hudeček and Ondřej Dušek. Are llms all you need for task-oriented dialogue? *arXiv preprint arXiv:2304.06556*, 2023. URL <https://arxiv.org/pdf/2304.06556>.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023. URL <https://dl.acm.org/doi/10.1145/3571730>.
- Jaehun Jung, Bokyung Son, and Sungwon Lyu. AttnIO: Knowledge Graph Exploration with In-and-Out Attention Flow for Knowledge-Grounded Dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3484–3497, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.280. URL <https://aclanthology.org/2020.emnlp-main.280>.
- Jack Lanchantin, Sainbayar Sukhbaatar, Gabriel Synnaeve, Yuxuan Sun, Kavya Srinet, and Arthur Szlam. A data source for reasoning embodied agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 8438–8446, 2023. URL <https://ojs.aaai.org/index.php/AAAI/article/view/26017/25789>.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lema Liu. A survey on retrieval-augmented text generation. *arXiv preprint arXiv:2202.01110*, 2022a.
- Xinghang Li, Di Guo, Huaping Liu, and Fuchun Sun. Embodied semantic scene graph generation. In *Conference on Robot Learning*, pages 1585–1594. PMLR, 2022b. URL <https://proceedings.mlr.press/v164/li22e/li22e.pdf>.

- Qika Lin, Rui Mao, Jun Liu, Fangzhi Xu, and Erik Cambria. Fusing topology contexts and logical rules in language models for knowledge graph completion. *Information Fusion*, 90:253–264, 2023. URL <http://ww.sentic.net/knowledge-graph-completion.pdf>.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023. URL [https://dl.acm.org/doi/pdf/10.1145/3560815?trk=public\\_post\\_comment-text](https://dl.acm.org/doi/pdf/10.1145/3560815?trk=public_post_comment-text).
- Corey Lynch, Ayzaan Wahid, Jonathan Tompson, Tianli Ding, James Betker, Robert Baruch, Travis Armstrong, and Pete Florence. Interactive language: Talking to robots in real time. *IEEE Robotics and Automation Letters*, 2023. URL <https://ieeexplore.ieee.org/iel7/7083369/7339444/10182264.pdf>.
- Blanca Miller and David Feil-Seifer. Embodiment, situatedness, and morphology for humanoid robots interacting with people. *Humanoid Robotics: A Reference*, pages 1–23, 2016. URL [https://rrl.cse.unr.edu/media/documents/2017/10.10072F978-94-007-7194-9\\_130-1.pdf](https://rrl.cse.unr.edu/media/documents/2017/10.10072F978-94-007-7194-9_130-1.pdf).
- Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. *ECCV*, 2022. URL [https://www.ecva.net/papers/eccv\\_2022/papers\\_ECCV/papers/136700714.pdf](https://www.ecva.net/papers/eccv_2022/papers_ECCV/papers/136700714.pdf).
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1081. URL <https://aclanthology.org/P19-1081>.
- Tomáš Nekvinda and Ondřej Dušek. Aargh! end-to-end retrieval-generation for task-oriented dialog. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 283–297, 2022. URL <https://aclanthology.org/2022.sigdial-1.29.pdf>.
- Preksha Nema, Shreyas Shetty, Parag Jain, Anirban Laha, Karthik Sankaranarayanan, and Mitesh M. Khapra. Generating descriptions from structured data using a bifocal attention mechanism and gated orthogonalization. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1539–1550, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1139. URL <https://aclanthology.org/N18-1139>.
- OpenAI. Gpt-4 technical report, 2023. URL <https://arxiv.org/abs/2303.08774>.
- Alexandros Papangelis and Stefan Ultes. Towards meaningful, grounded conversations with intelligent agents. *arXiv preprint arXiv:2006.15768*, 2020. URL <https://arxiv.org/pdf/2006.15768>.

- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. Few-shot natural language generation for task-oriented dialog. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 172–182, 2020. URL <https://aclanthology.org/2020.findings-emnlp.17.pdf>.
- Martin J Pickering and Simon Garrod. Prediction and embodiment in dialogue. *European Journal of Social Psychology*, 39(7):1162–1168, 2009. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/ejsp.663>.
- Plamen Prodanov and Andrzej Drygajlo. Bayesian networks for spoken dialogue management in multimodal systems of tour-guide robots. In *Proceedings of the 8th European Conference on Speech Communication and Technology (Eurospeech)*. ETH Zürich, 2003. URL <https://www.research-collection.ethz.ch/bitstream/handle/20.500.11850/82551/1/eth-8214-01.pdf>.
- Ian C Rankin, Seth McCammon, and Geoffrey A Hollinger. Robotic information gathering using semantic language instructions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4882–4888. IEEE, 2021. URL <https://par.nsf.gov/servlets/purl/10315084>.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. Parallel context windows for large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6383–6402, 2023. URL <https://aclanthology.org/2023.acl-long.352.pdf>.
- Vipula Rawte, Prachi Priya, SM Tonmoy, SM Zaman, Amit Sheth, and Amitava Das. Exploring the relationship between llm hallucinations and prompt linguistic nuances: Readability, formality, and concreteness. *arXiv preprint arXiv:2309.11064*, 2023. URL <https://arxiv.org/pdf/2309.11064>.
- Liliang Ren, Kaige Xie, Lu Chen, and Kai Yu. Towards universal dialogue state tracking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2786, 2018. URL <https://aclanthology.org/D18-1299.pdf>.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, et al. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, 2021. URL <https://aclanthology.org/2021.eacl-main.24.pdf>.
- Björn Schuller, Stefan Steidl, Anton Batliner, Felix Burkhardt, Laurence Devillers, Christian Mülller, and Shrikanth Narayanan. Paralinguistics in speech and language—state-of-the-art and the challenge. *Computer Speech & Language*, 27(1):4–39, 2013.
- Dhruv Shah, Błażej Osiański, Sergey Levine, et al. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*, pages 492–504. PMLR, 2023. URL <https://proceedings.mlr.press/v205/shah23b/shah23b.pdf>.

- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR, 2023. URL <https://proceedings.mlr.press/v202/shi23a/shi23a.pdf>.
- Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. *arXiv preprint arXiv:2212.04088*, 2022. URL <https://embodied-ai.org/papers/2023/7.pdf>.
- Yan Tai, Weichen Fan, Zhao Zhang, Feng Zhu, Rui Zhao, and Ziwei Liu. Link-context learning for multimodal llms. *arXiv preprint arXiv:2308.07891*, 2023. URL <https://arxiv.org/pdf/2308.07891>.
- Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:25–55, 2020. URL <https://h2r.cs.brown.edu/wp-content/uploads/tellex20.pdf>.
- Blaise Thomson, Jost Schatzmann, and Steve Young. Bayesian update of dialogue state for robust dialogue systems. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4937–4940. IEEE, 2008. URL <http://svr-www.eng.cam.ac.uk/~sjy/papers/thsy08.pdf>.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022. URL [https://arxiv.org/pdf/2201.08239.pdf?trk=public\\_post\\_comment-text](https://arxiv.org/pdf/2201.08239.pdf?trk=public_post_comment-text).
- David Thulke, Nico Daheim, Christian Dugast, and Hermann Ney. Efficient retrieval augmented generation from unstructured knowledge for task-oriented dialog. *arXiv preprint arXiv:2102.04643*, 2021. URL <https://arxiv.org/pdf/2102.04643>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. URL <https://arxiv.org/pdf/2302.13971>.
- Kai-Wen Tuan, Yi-Jyun Chen, Yi-Chien Lin, Chun-Ho Kwok, Hai-Lun Tu, and Jason S. Chang. Learning to find translation of grammar patterns in parallel corpus. In *Proceedings of the 33rd Conference on Computational Linguistics and Speech Processing (ROCLING 2021)*, pages 301–309, Taoyuan, Taiwan, October 2021. The Association for Computational Linguistics and Chinese Language Processing (ACLCLP). URL <https://aclanthology.org/2021.rocling-1.39>.
- Yi-Lin Tuan, Yun-Nung Chen, and Hung-yi Lee. DyKgChat: Benchmarking dialogue generation grounding on dynamic knowledge graphs. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1855–1865, Hong Kong, China,

- November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1194. URL <https://aclanthology.org/D19-1194>.
- Tomer Ullman. Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv preprint arXiv:2302.08399*, 2023. URL <https://arxiv.org/pdf/2302.08399.pdf>.
- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Large language models still can’t plan (a benchmark for llms on planning and reasoning about change). *arXiv preprint arXiv:2206.10498*, 2022. URL <https://arxiv.org/pdf/2206.10498>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Joshua Wainer, David J Feil-Seifer, Dylan A Shell, and Maja J Mataric. The role of physical embodiment in human-robot interaction. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pages 117–122. IEEE, 2006. URL <https://rrl.cse.unr.edu/media/documents/2016/498.pdf>.
- Nicholas Thomas Walker, Torbjørn Dahl, and Pierre Lison. Dialogue management as graph transformations. In *Conversational AI for Natural Human-Centric Interaction*, pages 219–227, Singapore, 2022. Springer Nature Singapore. ISBN 978-981-19-5538-9. URL [https://nr.brage.unit.no/nr-xmlui/bitstream/handle/11250/3055827/IWSDS\\_2021.pdf?sequence=1](https://nr.brage.unit.no/nr-xmlui/bitstream/handle/11250/3055827/IWSDS_2021.pdf?sequence=1).
- Nicholas Thomas Walker, Stefan Ultes, and Pierre Lison. Retrieval-augmented neural response generation using logical reasoning and relevance scoring. In *Proceedings of the 27th Workshop on the Semantics and Pragmatics of Dialogue, Full Papers*, 2023. URL <https://arxiv.org/pdf/2310.13566.pdf>.
- Yanmeng Wang, Ye Wang, Xingyu Lou, Wenge Rong, Zhenghong Hao, and Shaojun Wang. Improving dialogue response generation via knowledge graph filter. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7423–7427. IEEE, 2021. URL <https://ieeexplore.ieee.org/document/9414324>.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022. URL <https://arxiv.org/pdf/2206.07682.pdf?trk=cndc-detail>.
- Graham Wilcock and Kristiina Jokinen. Conversational ai and knowledge graphs for social robot interaction. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 1090–1094. IEEE, 2022. URL <https://helda.helsinki.fi/server/api/core/bitstreams/60e2a300-e51d-4b86-8558-06e7326e7f3b/content>.
- Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.

- Jason D Williams, Antoine Raux, and Matthew Henderson. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 7(3):4–33, 2016. URL <https://core.ac.uk/download/pdf/230775555.pdf>.
- Peng Wu, Bowei Zou, Ridong Jiang, and AiTi Aw. GCDST: A graph-based and copy-augmented multi-domain dialogue state tracking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1063–1073, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.95. URL <https://aclanthology.org/2020.findings-emnlp.95>.
- Bowen Xing and Ivor W Tsang. Relational temporal graph reasoning for dual-task dialogue language understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. URL <https://arxiv.org/pdf/2306.09114.pdf>.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, and Vadim Sheinin. Sql-to-text generation with graph-to-sequence model. *arXiv preprint arXiv:1809.05255*, 2018. URL <https://arxiv.org/pdf/1809.05255.pdf>.
- Pengfei Xu, Xiaojun Chang, Ling Guo, Po-Yao Huang, Xiaojiang Chen, and Alexander G Hauptmann. A survey of scene graph: Generation and application. *IEEE Trans. Neural Netw. Learn. Syst.*, 1:1, 2020. URL [https://www.researchgate.net/profile/Xiaojun-Chang-6/publication/340528049\\_A\\_Survey\\_of\\_Scene\\_Graph\\_Generation\\_and\\_Application/links/5e8eedb9a6fdcca789020994/A-Survey-of-Scene-Graph-Generation-and-Application.pdf](https://www.researchgate.net/profile/Xiaojun-Chang-6/publication/340528049_A_Survey_of_Scene_Graph_Generation_and_Application/links/5e8eedb9a6fdcca789020994/A-Survey-of-Scene-Graph-Generation-and-Application.pdf).
- Xiaoyu Yang, Feng Nie, Yufei Feng, Quan Liu, Zhigang Chen, and Xiaodan Zhu. Program enhanced fact verification with verbalization and graph attention network. *arXiv preprint arXiv:2010.03084*, 2020. URL <https://arxiv.org/pdf/2010.03084.pdf>.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179, 2013. URL <https://cs.brown.edu/courses/csci2951-k/papers/young13.pdf>.
- Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, et al. Socratic models: Composing zero-shot multimodal reasoning with language. *arXiv preprint arXiv:2204.00598*, 2022. URL <https://arxiv.org/pdf/2204.00598.pdf>.
- Min Zhang, Jingxiang Chen, Pengfei Li, Ming Jiang, and Zhe Zhou. Topic scene graphs for image captioning. *IET Computer Vision*, 16(4):364–375, 2022. URL <https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/cvi2.12093>.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, 2020. URL <https://aclanthology.org/2020.acl-demos.30.pdf>.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*, 2023. URL <https://arxiv.org/pdf/2309.01219.pdf>.

Xufeng Zhao, Mengdi Li, Cornelius Weber, Muhammad Burhan Hafez, and Stefan Wermter. Chat with the environment: Interactive multimodal perception using large language models. *arXiv preprint arXiv:2303.08268*, 2023. URL <https://arxiv.org/pdf/2303.08268.pdf>.

## Appendix A. VQA Objects

“a person”, “a woman”, “a man”, “a computer”, “a desk”, “a table”, “a coffee mug”, “a chair”, “a whiteboard”, “a garbage can”, “a door”, “a window”, “a plant”, “a fire extinguisher”, “a pillar”, “fruit”, “apples”, “bananas”, “flowers”, “magazines”, “books”, “a computer mouse”, “a couch”, “a TV”, “a wastebin”, “a pen”, “a pencil”, “scissors”, “folders”, “a light switch”, “cables”, “notebooks”, “paper”, “a printer”, “a bookshelf”, “bookshelves”, “a painting”, “a camera”, “food”, “cake”, “a light”, “a lamp”, “a stapler”, “a red folder”, “a blue folder”, “a yellow folder”, “a green folder”, “a broom”, “a blue curtain”, “a wooden pallet”, “a marker”, “a copy machine”, “cardboard boxes”

## Appendix B. Prompt Format

*“You are a robot that has just explored an office floor and recorded observations. Answer each question in a short, concise sentence, from the robot perspective. Example responses are:*

*Yes, I saw a woman in the main hallway.*

*I may have seen a fire extinguisher, but I’m not sure.*

*Sorry, I didn’t see any plants in the office.*

*We traveled 10 meters in the hallway. I saw a bookshelf in the first office, but not in the second office.*

*We first went to an office, then to the reception area.*

*Only use the provided observations to arrive at your answer.*

*If you don’t know the answer, just say that you don’t know. Don’t try to make up an answer.*

*Observations:*

[triples or verbalized facts]”