

Secondary Publication



Walker, Nicholas Thomas; Ultes, Stefan; Lison, Pierre

GraphWOZ: Dialogue Management with Conversational Knowledge Graphs

Date of secondary publication: 05.06.2023

Version of Record (Published Version), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-596685

Primary publication

Walker, Nicholas Thomas; Ultes, Stefan; Lison, Pierre: GraphWOZ: Dialogue Management with Conversational Knowledge Graphs. In: The 13th International Workshop on Spoken Dialogue Systems Technology, Los Angeles: arXiv:2211.12852, 2022, pp. 1-17.
DOI: 10.48550/arXiv.2211.12852

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holder(s).

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

GraphWOZ: Dialogue Management with Conversational Knowledge Graphs

Nicholas Thomas Walker, Stefan Ultes, and Pierre Lison

Abstract We present a new approach to dialogue management using conversational knowledge graphs as core representation of the dialogue state. To this end, we introduce a new dataset, GraphWOZ, which comprises Wizard-of-Oz dialogues in which human participants interact with a robot acting as a receptionist. In contrast to most existing work on dialogue management, GraphWOZ relies on a dialogue state explicitly represented as a dynamic knowledge graph instead of a fixed set of slots. This graph is composed of a varying number of entities (such as individuals, places, events, utterances and mentions) and relations between them (such as persons being part of a group or attending an event). The graph is then regularly updated on the basis of new observations and system actions. GraphWOZ is released along with detailed manual annotations related to the user intents, system responses, and reference relations occurring in both user and system turns. Based on GraphWOZ, we present experimental results for two dialogue management tasks, namely conversational entity linking and response ranking. For conversational entity linking, we show how to connect utterance mentions to their corresponding entity in the knowledge graph with a neural model relying on a combination of both string and graph-based features. Response ranking is then performed by summarizing the relevant content of the graph into a text, which is concatenated with the dialogue history and employed as input to score possible responses to a given dialogue state.

Nicholas Walker
Norwegian Computing Center, Postboks 114 Blindern, e-mail: walker@nr.no

Stefan Ultes
University of Bamberg, 96045 Bamberg, Germany, e-mail: stefan.ultes@uni-bamberg.de

Pierre Lison
Norwegian Computing Center, Postboks 114 Blindern, e-mail: plison@nr.no

1 Introduction

Task-oriented dialogue systems often include a dialogue management module in charge of tracking relevant information and selecting the next system response or action to perform [20]. This module often revolves around a representation of the current dialogue state comprising domain-specific and predefined “slots” that are necessary to complete a task within a domain.

While most previous work on dialogue management renders the dialogue state with a fixed list of predefined slots, this paper follows an alternative approach based on a dialogue state expressed as a graph of conversational entities [46, 47]. Graphs are increasingly employed in dialogue research, albeit typically as static graphs of background knowledge [35, 50, 53] rather than to capture the entirety of the dialogue state. Such knowledge graphs are frequently used with open domain dialogue systems, but have also been integrated with neural dialogue models for task-oriented dialogue systems [16]

Graph representations of dialogue state are particularly relevant for domains that must capture a rich conversational context, as in virtual assistants or human–robot interaction. Graphs makes it possible to directly express the connections between abstract entities and their lower-level groundings in the external environment, along with relations between those entities and the tasks to perform [4].

Several datasets for dialogue-oriented semantic parsing have been released in recent years [2, 4]. However, there is a lack of datasets that specifically target the problem of dialogue management with graph-structured dialogue states. This paper seeks to address this shortcoming, and presents GraphWOZ, a new Wizard-of-Oz dataset in which participants interact with a robot receptionist. The dialogue state underlying this domain includes a varying number of entities – such as persons, events or locations – connected with one another through relations (such as a person attending an event or being at a particular location).

Along with this new dataset, we present novel approaches to perform two dialogue management tasks based on this graph-structured dialogue state:

1. A neural feed-forward model for conversational entity linking based on a combination of string and graph distance features ;
2. A simple but effective method for response selection, which operates by converting a relevant part of the graph into text and running a fine-tuned language model to rank candidate responses.

The rest of this paper is as follows. We first review in Section 2 related work on the use of graphs for dialogue management. Section 3 describes the creation and annotation of the GraphWOZ dataset. In Section 4, we present the representation of the dialogue state as a graph, and the use of this representation for conversational entity linking and response selection. Finally, Section 5 provides experimental results on those two tasks, and Section 6 concludes.

2 Related Work

Graphs have been used in numerous contexts in dialogue systems research. For instance, the Microsoft Concept Graph has been used to connect background information to entities in the dialogue, along with a state graph to manage overall dialogue flow [38]. Similarly, graphs have been used to graphs represent multi-turn utterances and intents [49], while other work has used dataflow graphs representing executable programs [2]. Key-Value Retrieval Networks made use of task-specific information structured as a Knowledge graph (the KVRET* dataset) [15]. Dynamic knowledge graphs in particular have been used in a Question Answering (QA) task to represent domain-slot pairs [53]. The use of knowledge graphs has been prominent in QA, with uses including enhancement of pre-trained language models [43]. The Agent-Graph framework [10] used graph neural networks to improve sample efficiency in learning a dialogue management policy, while ConvGraph [17] used a graph-based approach to dialogue data augmentation. Recently, graphs have also been used to represent both dialogue context and external knowledge in a property graph [47] or to jointly encode dialogue context and a knowledge base for a graph neural network [33]. Graphs have also been used to encode relationships between visual and language data [22] and mental states of dialogue agents [39].

Entity linking is a core task in many NLP applications [36], and can be defined in the following way [44]:

Given a knowledge base containing a set of entities E and a text collection in which a set of named entity mentions M are identified in advance, the goal of entity linking is to map each textual entity mention $m \in M$ to its corresponding entity $e \in E$ in the knowledge base

Entity linking has been explored in various tasks in NLP, particularly in the rich literature on knowledge base construction [23], conversational QA, and knowledge base QA [30]. This often takes the form of retrieving information about entities from large scale knowledge bases [42]. Other work in conversational QA has explored the use of Graph Attention Networks to model relationships between entities [26]. Most existing work in entity linking has been focused on text, albeit few studies have also investigated entity linking on noisy ASR transcriptions [3].

Entity linking has also been explored in conversational settings, particularly for open domain dialogues [12]. The representation of these entities and their links to each other has taken various forms [24]. So-called entity centric models focus on learning representations of dialogue entities [1], without necessarily establishing links to a particular knowledge graph. An example is the Recurrent Entity, which Network [19] processes text data using a dynamic memory component to learn representations of dialogue. In such cases, it has been noted that there exists a close relationship between entity linking and coreference resolution [9], and this relationship has been explored by modelling coreference resolution and entity linking as a joint task [51]. Similarly, Slugbot [5] explored Named-Entity Recognition (NER) and entity linking with the Google Knowledge Graph. Other dialogue systems have explored the use of large knowledge graphs [13]. Recent work has also explored

the problem of plural mentions [52], and entity linking within multi-party dialogue, where it can be referred to as character identification [11].

Previous work has additionally explored the use of document representations of semantic graphs [32], an approach which stands closely to our response ranking experiments. Hand crafted features as well as neural methods have been employed for response ranking [37, 41]. Task-oriented dialogue systems specifically have made use of ranking modules with document-grounded large pre-trained language models [28, 45]. The PolyResponse model relied on a non-generative response ranking component to present both language and visual data to the user [21]. Other work has demonstrated improvements to dialogue ranking models by using synthesized negative training examples [18].

The GraphWOZ dataset was collected with a "Wizard-of-Oz" (WOZ) setup, which is a cost-effective method to collect dialogue data [27, 48]. MultiWOZ [6] and Taskmaster [8] are other datasets of task-oriented dialogues collected through WOZ experiments.

3 Data

3.1 Dialogue collection

GraphWOZ consists of interactions between a human participant that engages with a Pepper robot¹ acting as a receptionist. The dialogue state contains both physical and abstract entities (such as persons, groups, events, locations and time references) that have relevance to the task, and which may change in the course of the dialogue. GraphWOZ is released along with annotations which reflect both user intents, system actions, and the operations performed on the entity graph at each turn. GraphWOZ is made publicly available on Github².

At the onset of each dialogue, we generate a fictive organization composed of randomly generated entities. Each organization consists of between 40 and 60 fictive persons, and all individuals in the organization are assigned to a group and randomly generated office number along with a phone and email address. Subsequently, 30 to 50 events are generated with names combining randomly sampled jargon and a set of pre-defined meeting types (e.g. "web-readiness" + "status update"). Each event is randomly assigned to an organizer, a location and a start and end time. The event locations are randomly selected from a set of fictive conference rooms and the office of the person organizing the event.

Each dialogue begins with a task shown to the user on the robot tablet. The tasks are instantiated using predefined templates, which contain typed placeholders for elements such as people, events, places, and times. These slots are filled in by randomly sampling the generated organization. Tasks consist of combinations of

¹ <https://www.softbankrobotics.com/emea/en/pepper>

² <https://github.com/ntwalker/GraphWOZ>

1 VISITOR: Hello! My name is Wendy Parker and I am trying to find out who is organizing the users workshop.
 2 ROBOT: Hi Wendy! The users workshop is organized by Mark Suarez.
 3 VISITOR: Could you give me the room number for his office?
 4 ROBOT: Sure, his office is room 270.
 5 VISITOR: Thank you, that's all.
 6 ROBOT: Okay! Have a nice day!

Fig. 1 Short dialogue example with a fragment of the dialogue state graph. The entities under discussion are associated to nodes in the dialogue state graph. Orange represents **location** entities, green represents **people** entities, and purple is **events**.

topics such as scheduling, moving, or cancelling meetings, or finding attributes of entities in the organization. After the user finishes a task, a new task is created and displayed. Participants were given flexibility to interpret the task as desired and mark it as completed or failed to the best of their knowledge. The participants were instructed to work with the robot to negotiate a reasonable solution matching the intention of the task to the best of their ability. When the task is marked by the participant as completed or aborted, a new task is generated. Each user utterance is transcribed using Google Speech and the top ASR transcription is shown to the wizard. The wizard then composes a free-text response to the user's utterance, which is uttered by the robot using text-to-speech synthesis as a response to the user.

An example of a user task is as follows:

You are Richard Schroeder, and you want to arrange a meeting called methodologies discussion. Make a half hour long meeting in an available conference room, and invite Daniel Patton and at least one member of the Mathematics group.

A total of 15 distinct speakers (university students or researchers in data science) participated in the data collection, and interacted with the robot receptionist for about an hour each. The interactions took place in English (most participants were non-native speakers of English but had a high command of the language). The participants were paid for their participation, and provided their consent to the inclusion of their dialogues in a publicly released dataset.

3.2 Data Encoding

In total, GraphWOZ contains 119 dialogues from 15 distinct speakers. The dialogues are between 11 and 12 turns long on average. Figures 1 and 2 provide two concrete examples of dialogue in GraphWOZ.

Each dialogue is paired with the associated fictive organization generated for that dialogue, along with the task presented to the human user. The fictive organizations contain on average 90 entities of the types $\{people, event, room, group\}$, along with

1 VISITOR: Hello I am Andrew Fletcher
 2 ROBOT: Hello Andrew! How can I help you?
 3 VISITOR: Do you know if Stephanie Jules is attending a meeting called infrastructures discussion
 4 ROBOT: I can't find anyone named Stephanie Jules, did you mean Stephanie Shields?
 5 VISITOR: [UNK]
 6 ROBOT: Sorry, I didn't understand that. Could you repeat?
 7 VISITOR: Yes I meant Stephanie Shields
 8 ROBOT: Stephanie Shields is not attending infrastructures discussion.
 9 VISITOR: Could you add her to infrastructures discussion
 10 ROBOT: Ok! I've sent her an invitation to the event.
 11 VISITOR: Thank you very much
 12 ROBOT: You're welcome! Is there anything else I can help you with?
 13 VISITOR: No that will be all
 14 ROBOT: Okay! Have a nice day!

Fig. 2 Dialogue example including a misunderstanding due to ASR errors. The graph is updated on line 10 to add the individual "Stephanie Shields" as a participant to the event (creating an edge between her and the event). Green entities are *people*, Brown are *spurious entities* not found in the dialogue state, and purple are *events*.

start and end times for events and additional attributes such as emails and phone numbers which can in principle also be treated as distinct entities.

The raw recordings of each user turn are provided as WAV files along with Automatic Speech Recognition and "gold standard" transcriptions of the utterance. Each user turn was annotated with an abstract representation of the user intent, along with entity mentions in both ASR and gold transcriptions. The wizard responses were similarly annotated with a logical representation of the system actions along with entity mentions. Both user intents and agent actions are expressed with a flat predicate taking entities as arguments. The annotation of entity mentions describe entities that exist in the graph and can be linked from the utterance as well as spurious entities which cannot. Along with these links, we also annotate links between mentions referring to the speaker or a new event with a special token indicating that the link refers to a newly created entity in the task. As part of the calendar scheduling task, nodes are added to represent a newly created event for the user. Mentions from both user and system utterances are provided with annotations of which text span is the entity mention and which entity is the correct target. In the case of ambiguous names (e.g. two events with the same name on different days), the target entity is annotated by constraints describing the attributes of the target entity.

The annotations in GraphWOZ are provided in JSON format. The data is split into approximately 50 % for training, 25 % for development, and 25 % for the held-out test set.

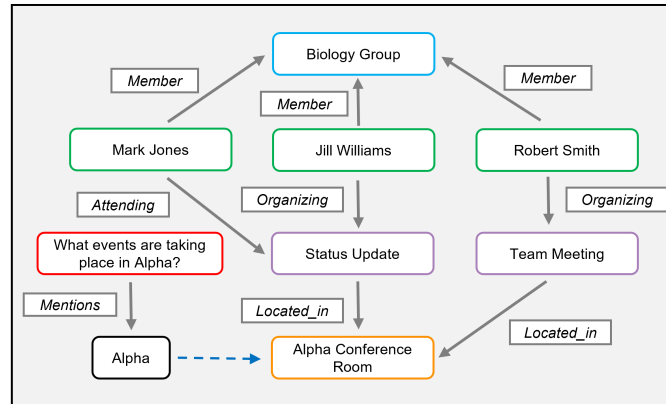


Fig. 3 Example (directed) dialogue graph with entity linking. Nodes represent people, places, or dialogue elements, and edges represent semantic relations between them. Green nodes are **people** entities, purple nodes are **events**, orange is a **location**, light blue is a **group**, and red is a **user utterance**. Edges are labeled as semantic relations. The dark blue dashed arrow represents a link created by **entity linking**, recognizing that "Alpha" refers to the Alpha Conference Room.

4 Approach

4.1 Graph representation

Our approach to dialogue management centers around a dynamic entity knowledge graph. The graph represents both general background knowledge for the interaction (in this case events, persons, groups and contact information known to the robot receptionist) but also nodes and edges representing the dialogue itself (such as utterances, intents, responses, and entity mentions). The relations contained in the dataset are expressed through directed edges, and each node is associated with a semantic type. A small example of state graph is provided in Figure 3.

In this setting, dialogue management centers on graph transformations which alter the graph of the dialogue state. Each new utterance is represented as a new node in the graph, with attributes including the ASR transcription itself as well as start and end timestamps. These utterances are added to the graph along with a linked node for each entity mention they contain, along with a link to the preceding utterance node (if any). The link between utterances represents a sequential relation between them. At this stage, the dialogue management system must determine which (if any) entities from the existing knowledge graph refer to the same entity as the entity mention. For instance, if a person is mentioned, they may be identified by name, a pronoun, or some other construction (e.g. "that person"). This identification step can readily be framed as an entity linking task.

To construct the graph from the annotated data, we consider each entity of type $\{people, event, room, group\}$ as an individual node in the graph. For this purpose, a person's office is represented as a room as well, although it could equally be viewed

as an attribute of the person entity. All nodes can in the simplest case, and for our experiments, be assigned the entity name as their attribute. Alternatively, these strings or a combination of attributes such as a person’s email or phone number can in principle also be encoded with a language model to assign a vector representation to each node. Similarly, the edges of the graph are assigned labels to represent the semantic relationship between them. For example, each event in the data has a set of attendees and an organizer, thus each person in these attributes is assigned a link to the relevant link node. For our experiments, the task of entity linking can be viewed as predicting the occurrence of the *refers_to* relationship between an entity mention (text span present in an utterance) and the actual entity it refers to in the background knowledge of the robot receptionist. Other semantic relations such as *attending* or *is_located* capture relationships that are valuable for spatio-temporal reasoning.

The knowledge graph also represents the set of user and agent utterances as nodes. Each node contains the raw string of the utterance (the ASR transcription for the user utterance) and is connected to nodes representing entity mentions in the utterance. An entity *mention* is thereby treated separately from whichever known entities it corresponds to in the graph, where an unlinked entity mention may be a spurious entity that cannot be connected to the known set of entities. Additionally, an underspecified entity links to multiple entities in the graph, where the exact reference should be clarified by the system at a later round.

4.2 Entity Linking

Along with the release of the GraphWOZ dataset, we present simple but effective approaches for two specific tasks: entity-linking and response selection.

Upon recognizing a new user utterance, a graph-based dialogue system must link entities recognized in the user utterance with any matching entities in the graph. For each mention $m \in M$ and a set of known entities E , we wish to link m to a set of entities $E' \subset E$. For our experiments, we consider all dialogue turns containing at least one mention where $E' \neq \emptyset$. Note that, in contrast to traditional entity-linking models, we allow a mention to connect to more than one entity, as is the case for e.g. plural pronouns.

Input Features

The following features were employed for the entity-linking model:

String distance: Many entity mentions in a dialogue refer to the entity directly by name. In these cases, it is appropriate to simply search for the the recognized entity string to the entities in the graph. For this purpose, we use string equality as a simple binary feature to the model. However, due to ASR error the transcription of the entity may be corrupted and require edit distance metrics to recognize the link. To this end, our baseline entity linking model employs several such metrics, namely:

Levenshtein Distance, *Jaro-Winkler Distance*, and *Longest-Common Substring*. We additionally add features which are also calculated from the edit-distance features divided by the length of the mention string (to account for pronouns). Additionally, we compare similarity on the token level by computing the set difference between the tokens in the mention and candidate target. Similarly, we also add the cardinality of the union of tokens between the pair as a feature.

Graph distance: At a given turn in the dialogue, there may exist previous entity mention nodes which have been linked to another entity in the graph. Upon addition of a new mention node to the graph, it is necessary to also consider whether the new node refers to a previously mentioned entity. The graph structure is useful to this purpose, as the distance to previously mentioned nodes corresponds directly to how recently they were mentioned. The graph distance can then be used as an indication of relevance and recency of possible targets for mentions at the current turn.

A new entity mention may in many cases refer back to a node which is more distant than other intervening mention-entity pairs, so it is desirable to evaluate whether the two mentions are likely part of the same *coreference chain*. In this way, entity linking in this graph can be viewed as a task of coreference resolution where we can also make use of linguistic structure to assist the model predictions. From this intuition, we use two pretrained coreference resolution models to extract coreference chains from the dialogue utterances:

1. The first model is based on token-level coreference resolution as implemented in SpaCy Coreferee³, and relies on contextualized word representations extracted via RoBERTa [34].
2. The second model employs the end-to-end coreference resolution method described by [31], as implemented in Allen NLP⁴, and using SpanBERT embeddings [25] to derive the coreference clusters.

Based on the coreference chains extracted with those models, we add a link between all entity mentions which exist in the same coreference chain. This procedure greatly shortens the graph distance from a new mention node to the correct target when it refers to a previously mentioned entity.

Classification model

The string and graph-based distance features described above are then used as inputs to train binary classification models for the task of predicting whether a (mention, entity) pair should be linked together. Given the modest size of the training set, those classification models need to be limited to a small set of parameters. In Section 5, we present results obtained with a simple logistic regression model and a feedforward neural model with two hidden layers.

³ <https://github.com/msg-systems/coreferee>

⁴ http://docs.allennlp.org/v0.9.0/api/allennlp.models.coreference_resolution.html

4.3 Response Ranking

Once entity mentions are linked to their corresponding entity in the graph-structured dialogue state, we proceed with the task of selecting the system response. The approach rests upon determining which part of the graph is relevant for selecting the next action and then converting this subgraph into a short text. This text is then prepended to the dialogue history and used to rank possible responses using a fine-tuned language model.

The response ranking approach relies on the following steps:

Extraction of relevant subgraph

We first extract a relevant subgraph of the dialogue state to filter out entities and relations that are unlikely to influence the response selection. To this end, we first determine the entities explicitly mentioned through the dialogue history, along with the nodes adjacent to those mentioned entities, and the set of all relations between those nodes. These entities are determined using the output of the MLP-based entity linking system described in the previous section.

Conversion of subgraph into NL sentences

In order to take advantage of large pretrained language models, the subgraph is then converted from a graph representation into text using predefined templates. Once generated, the text is then prepended to the dialogue history. As an example, the relevant subgraph containing entities mentioned in the dialogue in Figure 1 are converted into two sentences, as illustrated in Figure 4.

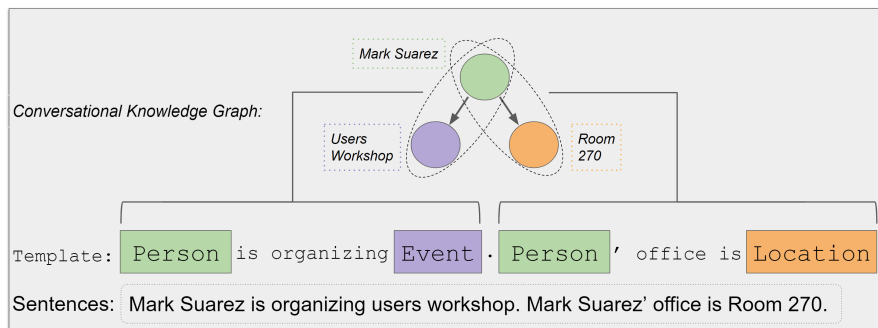


Fig. 4 Example of two sentences filled in by a template. These sentences along with others filled by templates describing entity attributes are combined into a document describing a person. The document created with these templates is combined with the dialogue history to form the context for the language model.

Fine-tuning of ranking model

Finally, we fine-tune a neural language model. Two alternative approaches are evaluated:

- A pointwise approach taking a (context, response) pair as input and returning a normalized score denoting the appropriateness of the response given the conversational context.
- A pairwise approach where the model is simultaneously provided with a positive pair (context, actual response) and a negative pair (context, random response). The model is then optimized to minimize the cross-entropy loss for the difference between the positive and negative examples [7].

To mitigate the scarcity of training data, we expand the initial set of (dialogue context, wizard response) pairs using a simple mechanism to swap entity names for persons, events and groups with alternative values from the knowledge graph. Those modified pairs are then added to the training set.

Implementation details: The neural ranking model is fine-tuned using a BERT language model [14] on 10 epochs (batch size of 5) with a cross-entropy loss and the Adam optimizer [29] with a learning rate set to 1×10^{-6} . For the negative examples, we experimented both with a random selection of responses and with a more informed approach where candidates are samples according to their cosine similarity, using dense embeddings from [40]. Those two negative sampling methods led to similar experimental results on GraphWOZ.

5 Results

5.1 Entity-linking

Two entity-linking models were tested, respectively based on logistic regression and a multilayer perceptron (MLP). For each (mention, entity) pair, we extract both string and graph features as described in section 4.2 and provide them as input to the classifier. For the MLP model, the best performing parameters were two layers of size 20 and 12, trained for 600 iterations optimized with Limited-memory BFGS. Those hyper-parameters were selected empirically based on the development set. For comparison, we also provide two simple baselines:

- *String equality*, which connects an entity mention to an entity represented by the exact same string.
- *String equality + Recency*, which first checks the existence of an entity of same name as above. In case no exact string match can be found, the baseline then connects the entity mention to the most recently mentioned entity.

Table 1 shows the evaluation results of the entity linking models with precision, recall, and F_1 -score. We observe that both the logistic regression and feedforward

neural model benefit from the inclusion of graph-distance features. It can also be noted that the string equality baseline achieves less than perfect precision due to ambiguity between different entities with the same name.

Model	Dev			Test		
	Precision	Recall	F1	Precision	Recall	F1
String Equality	0.92	0.22	0.35	0.96	0.32	0.48
String equality + Recency	0.60	0.28	0.38	0.59	0.38	0.46
Logistic Regression, String	0.85	0.33	0.48	0.89	0.46	0.61
Logistic Regression, String + Graph	0.87	0.46	0.61	0.95	0.59	0.73
MLP (best), String	0.87	0.53	0.65	0.89	0.59	0.71
MLP (best), String + Graph	0.89	0.64	0.74	0.88	0.67	0.76

Table 1 Entity Linking with string equality and recency heuristics, Logistic Regression, and MLP. Graph distance is computed as $\frac{1}{2^l}$ where l is the length of the path from the mention to the target node. If there is no path, the $\frac{1}{2^{10}}$ is used (an intermediate but finite value l , as many correct utterance links will have no existing path between the nodes).

5.2 Response ranking

Table 2 provides the evaluation results on the task of ranking possible system responses using standard IR metrics, namely the recall $R_{10}@1$, which measures how often the correct response is ranked highest among a set of 10 candidate responses, the recall $R_{10}@2$, which measures how often the correct response is placed in the top-2 among 10 candidates, and the Mean Reciprocal Rank (MRR), which is the harmonic mean of the rank of the correct response. Those results are obtained using the best entity-linking model from the previous section (MLP with both string and graph distance features) to link entity mentions to their corresponding entities.

As we can observe from Table 2, the performance of the response ranking models (both pointwise and pairwise) improves when the content of the relevant subgraph (converted to a text format through templates) is prepended to the dialogue history. The pairwise model also seems to perform slightly better than the pointwise variant.

6 Conclusion

This paper presented GraphWOZ, a new collection of Wizard-of-Oz dialogues where the dialogue state is structured as a graph including a varying number of entities (persons, events, groups, etc.) and relations between them. The dialogues were gathered as part of a human-robot interaction scenario where the robot was expected to operate as a receptionist and answer various requests regarding the agenda of a

Input	Ranking model	$R_{10}@1$		$R_{10}@2$		MRR	
		Dev	Test	Dev	Test	Dev	Test
Only dialogue history	Pointwise	0.76	0.72	0.91	0.91	0.85	0.83
	Pairwise	0.80	0.83	0.89	0.94	0.87	0.90
Relevant subgraph (in text form) + dialogue history	Pointwise	0.80	0.85	0.91	0.93	0.88	0.91
	Pairwise	0.85	0.84	0.91	0.95	0.95	0.91

Table 2 Evaluation results on the development and test sets of the GraphWOZ dataset for the task of selecting the next agent response. The results are computed based from a 10 candidate responses for each dialogue turn. In addition to the type of ranking model (pointwise or pairwise), we also show results either with or without the natural language sentences summarizing the content of the relevant subgraph for the dialogue history.

fictional organization. Each dialogue is coupled with a randomly generated knowledge graph, and is provided with a range of annotations, notably the intents (expressed as logical forms) associated with each user utterance and system response as well as the entities mentioned through the dialogue. The GraphWOZ dataset comprises a small but richly annotated set of dialogues in a novel domain which explicitly represents the dialogue state as a single, unified graph.

In addition to this dataset, we also present simple but effective approaches to two dialogue management tasks (conversational entity linking and response ranking) based on a graph-structured dialogue state. For entity-linking, experimental results on GraphWOZ demonstrate that the use of both string and graph distance features allows the classification model to perform entity linking for both direct references as well as pronouns and other indirect references.

We also provide empirical results on the task of response selection based on the dialogue state graph and dialogue history. The presented approach uses the results of the entity linking model to select a relevant subgraph based on the entities linked from user utterances. This subgraph is then converted into natural language sentences and prepended to the dialogue context. The resulting text is finally fed to a neural language model and employed to rank candidate responses.

The response selection model presented in this paper depends on the availability of templates that can express the content of the dialogue state graph into English sentences. Although we found those templates to be relatively easy to create for this particular domain, this requirement may be difficult to satisfy for other types of graph. Future work will focus on investigating the use of graph neural networks to perform dialogue state tracking and response selection on the GraphWOZ dataset without the need to convert the graph into text. Additionally, the ability of graphs to represent uncertainty by means of e.g. weighting of proposed entity links is a potential topic for exploration.

References

1. Aina, L., Silberer, C., Sorodoc, I., Westera, M., Boleda, G.: What do entity-centric models learn? insights from entity linking in multi-party dialogue. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 3772–3783 (2019)
2. Andreas, J., Bufe, J., Burkett, D., Chen, C., Clausman, J., Crawford, J., Crim, K., DeLoach, J., Dorner, L., Eisner, J., et al.: Task-oriented dialogue as dataflow synthesis. *Transactions of the Association for Computational Linguistics* **8**, 556–571 (2020)
3. Benton, A., Dredze, M.: Entity linking for spoken language. In: Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 225–230 (2015)
4. Bonial, C., Donatelli, L., Abrams, M., Lukin, S., Tratz, S., Marge, M., Artstein, R., Traum, D., Voss, C.: Dialogue-amr: abstract meaning representation for dialogue. In: Proceedings of the 12th Language Resources and Evaluation Conference, pp. 684–695 (2020)
5. Bowden, K.K., Wu, J., Oraby, S., Misra, A., Walker, M.: Slugnerds: A named entity recognition tool for open domain dialogue systems. arXiv preprint arXiv:1805.03784 (2018)
6. Budzianowski, P., Wen, T.H., Tseng, B.H., Casanueva, I., Ultes, S., Ramadan, O., Gasic, M.: Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 5016–5026 (2018)
7. Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., Hullender, G.: Learning to rank using gradient descent. In: Proceedings of the 22nd international conference on Machine learning, pp. 89–96 (2005)
8. Byrne, B., Krishnamoorthi, K., Sankar, C., Neelakantan, A., Goodrich, B., Duckworth, D., Yavuz, S., Dubey, A., Kim, K.Y., Cedilnik, A.: Taskmaster-1: Toward a realistic and diverse dialog dataset. 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), pp. 4516–4525 (2019)
9. Chen, H.Y., Zhou, E., Choi, J.D.: Robust coreference resolution and entity linking on dialogues: Character identification on tv show transcripts. In: Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017), pp. 216–225 (2017)
10. Chen, L., Chen, Z., Tan, B., Long, S., Gašić, M., Yu, K.: Agentgraph: Toward universal dialogue management with structured deep reinforcement learning. *IEEE/ACM Transactions on audio, speech, and language processing* **27**(9) (2019)
11. Chen, Y.H., Choi, J.D.: Character identification on multiparty conversation: Identifying mentions of characters in TV shows. In: Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, pp. 90–100 (2016)
12. Cui, W., Rolston, L., Walker, M., Hockey, B.A.: OpenEL: An annotated corpus for entity linking and discourse in open domain dialogue. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference. European Language Resources Association (2022)
13. Curry, A.C., Papaioannou, I., Suglia, A., Agarwal, S., Shalyminov, I., Xu, X., Dušek, O., Eshghi, A., Konstas, I., Rieser, V., et al.: Alana v2: Entertaining and informative open-domain social dialogue using ontologies and entity linking. *Alexa Prize Proceedings* (2018)
14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
15. Eric, M., Krishnan, L., Charette, F., Manning, C.D.: Key-value retrieval networks for task-oriented dialogue. In: Proceedings of the SIGDIAL 2017 Conference, p. 37–49 (2017)
16. Ghazvininejad, M., Brockett, C., Chang, M.W., Dolan, B., Gao, J., Yih, W.t., Galley, M.: A knowledge-grounded neural conversation model. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018)
17. Gritta, M., Lampouras, G., Iacobacci, I.: Conversation graph: Data augmentation, training, and evaluation for non-deterministic dialogue management. *Transactions of the Association for Computational Linguistics* **9**, 36–52 (2021)

18. Gupta, P., Tsvetkov, Y., Bigham, J.: Synthesizing adversarial negative responses for robust response ranking and evaluation. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 3867–3883. Association for Computational Linguistics (2021)
19. Henaff, M., Weston, J., Szlam, A., Bordes, A., LeCun, Y.: Tracking the world state with recurrent entity networks. In: International Conference on Learning Representations (2017)
20. Henderson, M.: Machine learning for dialog state tracking: A review. In: Machine Learning in Spoken Language Processing Workshop (2015)
21. Henderson, M., Vulić, I., Casanueva, I., Budzianowski, P., Gerz, D., Coope, S., Spithourakis, G., Wen, T.H., Mrkšić, N., Su, P.H.: PolyResponse: A rank-based approach to task-oriented dialogue with application in restaurant search and booking. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing: System Demonstrations (2019)
22. Hong, Y., Rodriguez, C., Qi, Y., Wu, Q., Gould, S.: Language and visual entity relationship graph for agent navigation. *Advances in Neural Information Processing Systems* **33**, 7685–7696 (2020)
23. Ji, H., Grishman, R.: Knowledge base population: Successful approaches and challenges. In: Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies, pp. 1148–1158 (2011)
24. Joko, H., Hasibi, F., Balog, K., de Vries, A.P.: Conversational entity linking: Problem definition and datasets. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2390–2397 (2021)
25. Joshi, M., Chen, D., Liu, Y., Weld, D.S., Zettlemoyer, L., Levy, O.: SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics* **8**, 64–77 (2020)
26. Kacupaj, E., Plepi, J., Singh, K., Thakkar, H., Lehmann, J., Maleshkova, M.: Conversational question answering over knowledge graphs with transformer and graph attention networks. In: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (2021)
27. Kelley, J.F.: An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Information Systems (TOIS)* **2**(1), 26–41 (1984)
28. Kim, B., Lee, D., Kim, S., Lee, Y., Huang, J.X., Kwon, O.W., Kim, H.: Document-grounded goal-oriented dialogue systems on pre-trained language model with diverse input representation. In: Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021), pp. 98–102 (2021)
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
30. Lan, Y., He, G., Jiang, J., Jiang, J., Zhao, W.X., Wen, J.R.: A survey on complex knowledge base question answering: Methods, challenges and solutions. *IJCAI* (2021)
31. Lee, K., He, L., Lewis, M., Zettlemoyer, L.: End-to-end neural coreference resolution. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 188–197. Association for Computational Linguistics, Copenhagen, Denmark (2017). DOI 10.18653/v1/D17-1018. URL <https://aclanthology.org/D17-1018>
32. Li, S., Namazifar, M., Jin, D., Bansal, M., Ji, H., Liu, Y., Hakkani-Tur, D.: Enhanced knowledge selection for grounded dialogues via document semantic graphs. arXiv preprint arXiv:2206.07296 (2022)
33. Liu, Q., Bai, G., He, S., Liu, C., Liu, K., Zhao, J.: Heterogeneous relational graph neural networks with adaptive objective for end-to-end task-oriented dialogue. *Knowledge-Based Systems* **227**, 107,186 (2021)
34. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A robustly optimized BERT pretraining approach. *CoRR* (2019). URL <http://arxiv.org/abs/1907.11692>
35. Madotto, A., Cahyawijaya, S., Winata, G.I., Xu, Y., Liu, Z., Lin, Z., Fung, P.: Learning knowledge bases with parameters for task-oriented dialogue systems. In: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 2372–2394 (2020)

36. Mudgal, S., Li, H., Rekatsinas, T., Doan, A., Park, Y., Krishnan, G., Deep, R., Arcaute, E., Raghavendra, V.: Deep learning for entity matching: A design space exploration. In: Proceedings of the 2018 International Conference on Management of Data, pp. 19–34 (2018)
37. Papaioannou, I., Curry, A.C., Part, J.L., Shalymov, I., Xu, X., Yu, Y., Dušek, O., Rieser, V., Lemon, O.: An ensemble model with ranking for social dialogue. NIPS Workshop on Conversational AI (2017)
38. Pichl, J., Marek, P., Konrád, J., Matulík, M., Nguyen, H.L., Šedivý, J.: Alquist: The alexa prize socialbot. arXiv preprint arXiv:1804.06705 (2018)
39. Qiu, L., Zhao, Y., Liang, Y., Lu, P., Shi, W., Yu, Z., Zhu, S.c.: Towards socially intelligent agents with mental state transition and human value. In: Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue, pp. 146–158 (2022)
40. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics (2019). URL <https://arxiv.org/abs/1908.10084>
41. Shalymov, I., Dušek, O., Lemon, O.: Neural response ranking for social conversation: A data-efficient approach. In: Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI, pp. 1–8. Association for Computational Linguistics, Brussels, Belgium (2018)
42. Shen, T., Geng, X., Qin, T., Guo, D., Tang, D., Duan, N., Long, G., Jiang, D.: Multi-task learning for conversational question answering over a large-scale knowledge base. 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), pp. 2442–2451 (2019)
43. Shen, T., Mao, Y., He, P., Long, G., Trischler, A., Chen, W.: Exploiting structured knowledge in text via graph-guided representation learning. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 8980–8994 (2020)
44. Shen, W., Wang, J., Han, J.: Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* **27**(2), 443–460 (2014)
45. Tan, C.H., Yang, X., Zheng, Z., Li, T., Feng, Y., Gu, J.C., Liu, Q., Liu, D., Ling, Z.H., Zhu, X.: Learning to retrieve entity-aware knowledge and generate responses with copy mechanism for task-oriented dialogue systems. arXiv preprint arXiv:2012.11937 (2020)
46. Ultes, S., Budzianowski, P., Casanueva, I., Barahona, L.M.R., Tseng, B.H., Wu, Y.C., Young, S., Gasic, M.: Addressing objects and their relations: The conversational entity dialogue model. In: Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, pp. 273–283 (2018)
47. Walker, N.T., Dahl, T., Lison, P.: Dialogue management as graph transformations. In: S. Stoyanchev, S. Ultes, H. Li (eds.) *Conversational AI for Natural Human-Centric Interaction*, pp. 219–227. Springer Nature Singapore, Singapore (2022)
48. Wen, T.H., Vandyke, D., Mrkšić, N., Gasic, M., Barahona, L.M.R., Su, P.H., Ultes, S., Young, S.: A network-based end-to-end trainable task-oriented dialogue system. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers, pp. 438–449 (2017)
49. Yang, K., Kong, X., Wang, Y., Zhang, J., De Melo, G.: Reinforcement learning over knowledge graphs for explainable dialogue intent mining. *IEEE Access* **8**, 85,348–85,358 (2020)
50. Young, T., Cambria, E., Chaturvedi, I., Zhou, H., Biswas, S., Huang, M.: Augmenting end-to-end dialogue systems with commonsense knowledge. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018)
51. Zheng, J., Vilnis, L., Singh, S., Choi, J.D., McCallum, A.: Dynamic knowledge-base alignment for coreference resolution. In: Proceedings of the Seventeenth Conference on Computational Natural Language Learning, pp. 153–162 (2013)
52. Zhou, E., Choi, J.D.: They exist! introducing plural mentions to coreference resolution and entity linking. In: Proceedings of the 27th International Conference on Computational Linguistics, pp. 24–34 (2018)

53. Zhou, H., Young, T., Huang, M., Zhao, H., Xu, J., Zhu, X.: Commonsense knowledge aware conversation generation with graph attention. In: IJCAI, pp. 4623–4629 (2018)