

Evaluating Agent Interactions Through Episodic Knowledge Graphs

Selene Báez Santamaría and Piek Vossen and Thomas Baier

Vrije University Amsterdam

{s.baezsantamaria, p.t.j.m.vossen, t.baier}@vu.nl

Abstract

We present a new method based on episodic Knowledge Graphs (eKGs) for evaluating (multimodal) conversational agents in open domains. This graph is generated by interpreting raw signals during conversation and is able to capture the accumulation of knowledge over time. We apply structural and semantic analysis of the resulting graphs and translate the properties into qualitative measures. We compare these measures with existing automatic and manual evaluation metrics commonly used for conversational agents. Our results show that our Knowledge-Graph-based evaluation provides more qualitative insights into interaction and the agent’s behavior.

1 Introduction

In order to develop open-domain conversational agents, it is crucial to have automatic and reproducible ways of evaluating the interaction and the agent’s role. However, interaction with people is challenging to evaluate for several reasons: 1) people behave differently in each interaction, 2) people appreciate the interaction for different reasons and aspects, 3) different goals and sub-goals may play a role simultaneously, and 4) personal relationships and past experiences have an impact on every interaction. For these reasons, most evaluations of interactive systems use human judges and questionnaires in analogy to user-satisfaction methods.

In addition to these questionnaires, conversational systems are often evaluated by comparing system responses to human responses on a turn-by-turn basis, where the prompts and the gold responses are taken from human-human conversations. Standard measures such as BLUE (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), BERTscore (Zhang* et al., 2020) test the similarity between the system response and a gold response, whereas USR (Mehri and Eskenazi, 2020) tests the coherence of the system response to the previous

prompt and context. However, these measures do not truly assess the quality of the system’s interpretation and the relevance of the prompt, and they *punish systems for being creative and making responses personal*. By detaching prompt-response pairs from the whole conversation, these metrics evaluate the reactivity of an agent, not its ability to engage in a coherent interaction.

Deriu et al. (2021) mention five general requirements for evaluation: 1) automatic to reduce human labour and subjectivity, 2) repeatable when applied to the same dialogue, 3) correlate with human judgments, 4) differentiated for various strategies, and 5) explainable. None of the existing approaches satisfy all these criteria. In this paper, we demonstrate that graph properties can be used as an additional and independent evaluation of the **effectiveness** of the communication. This evaluation is an automatic measure of **semantic quality** that is also explainable and reproducible, meeting three of the five previous requirements.

We present a novel evaluation method that qualifies conversations using an episodic Knowledge Graph (eKG) (Baez Santamaria et al., 2021). We calibrate different groups of graph measures in relation to human evaluations and ground truth independent measures. To test our selected metrics, we compare the quality of types of conversations and show that the proposed evaluation framework holds, regardless of differences in system design. Our contributions are:

1. We provide an reference-free and explainable method for evaluating the interaction of conversational agents.
2. We compare our method to other standard evaluation methods and show its complementary value.
3. We demonstrate that our method can be applied across multiple conversations and different (types of) participants.

2 Related work

Dialogue systems have been studied for several decades and are further developed within Conversational AI systems. In their survey, [Deriu et al. \(2021\)](#) discuss different types of dialogue systems and how they are evaluated. They conclude that evaluating open conversational agents is an open problem due to the lack of a goal and variable structure. Therefore, evaluation approaches focus on appropriateness and human likeness of responses or specific linguistic properties such as variability, lexical complexity, coherence, correctness, and relevance of system responses. Evaluations can furthermore be done at a turn-level or conversation-level. There is a many-to-many problem in both cases: multiple responses can be correct, multiple dialogues can lead to the right/same result, and every interaction is unique.

Attempts to automate these notions often rely on metrics such as BLUE ([Papineni et al., 2002](#)), ROUGE ([Lin and Och, 2004](#)), METEOR ([Banerjee and Lavie, 2005](#)), and BERTscore ([Zhang* et al., 2020](#)) to measure the similarity of the agent response to one or more ground-truth responses; or they borrow from information retrieval measures when systems need to select the appropriate responses from a set of possible alternatives. In contrast, USR ([Mehri and Eskenazi, 2020](#)) is a new method created by fine-tuning RoBERTa-base ([Liu et al., 2019](#)) on the training set of Topical-Chat. Whereas METEOR and BERTscore compare a system response to a ground-truth response, USR gives a quality evaluation score without a ground-truth response (reference-free) by measuring the coherence of the system response with the human prompt and the previous context.

Evaluation regimes with ground-truth responses limit agents' "freedom and creativity" to generate other responses that may also fit the purpose. Therefore, it is unsurprising that evaluations often fall back on a posteriori evaluation by human judges. However, human evaluations suffer from several pitfalls: expensive, time-consuming, inconsistent across experiments, difficult to reproduce, and challenging to scale. Researchers tried to harmonize the evaluation criteria to address the inconsistency and lack of coherence in terminology and methodology for human evaluations of open-domain dialogue. [Howcroft et al. \(2020\)](#) survey of 165 papers with human evaluations reports more than 200 quality criteria (such as Fluency, Accuracy, or Readability.)

that have been used in Natural Language Generation. Independently, [Fitriani et al. \(2019\)](#) analyzed the proceedings of the conference of Intelligent Virtual Agents¹ between 2013 and 2018. They found 189 constructs from 89 questionnaires reported in 81 papers, which they reduced to 19 measurement instruments. Measurements range from how enjoyable, correct and useful to how fluent.

In an attempt to automate evaluations and make them more reproducible and scalable, the *Ninth Dialog System Technology Challenge, Track on Interactive Evaluation of Dialog* (DSTC9, Track 3)² carried out a variety of automatic and human evaluations on 33 systems submissions to the Topical-Chat challenge ([Gopalakrishnan et al., 2019](#)). Topical-Chat consists of conversations between two Amazon Mechanical Turk workers who were given prior knowledge or information to refer to during their conversation. Systems need to respond to turns from these conversations, replacing one worker. A human evaluation of system responses was done using the questionnaire from the FED dataset ([Mehri and Eskenazi](#)). An automatic evaluation was done using the measures METEOR, BERTscore, and USR. [Gunasekara et al. \(2020\)](#) report that the USR (0.3 Spearman) correlates better with human judgments than METEOR (0.23 Spearman) and BERTscore (Spearman 0.22), although they also admit that the correlation is not very high. It is still to be seen how easy USR can be transferred to other dialogues and contexts, as it was trained and tested on Topical-Chat.

We present a reference-free approach that is not based on coherence but measures the interpretability of (multimodal) situations and accumulates these over time. The basic idea is that effective interactions result in rich and high-quality representations that can be measured in a Knowledge Graph. Factors determining the communication effect are the agent's response quality and the collaboration between participants.

3 Problem formalization

We represent an interaction as a series of tuples $[t, s, g, f, p]$, such that:

$t \in \mathbf{T}$, a set of time points

$s \in \mathbf{S}$, a set of situations

$g \in \mathbf{G}$, a set of graphs

¹<https://dl.acm.org/conference/iva/proceedings>

²<http://dialog.speech.cs.cmu.edu:8003>

$\mathbf{f} \in \mathbf{F}$, a set of unknown features part of situation \mathbf{s}

$\mathbf{p} \in \mathbf{P}$, a set of defined properties that can make up graph \mathbf{g}

A graph \mathbf{g} represents the interpretation of a sequence of situations \mathbf{s} at time point \mathbf{t} . Each situation \mathbf{s} can be modeled as a bundle of unknown features \mathbf{f} and each graph \mathbf{g} as a set of properties \mathbf{p} that are defined a priori. To quality the conversations, we measure how many and which properties \mathbf{p} are extracted from each turn and the cumulative effect of adding these properties over time to the graph \mathbf{g} . The instrument’s effectiveness in measuring quality depends on the ability to detect the properties \mathbf{p} given the features \mathbf{f} . The effectiveness of the communication depends on the predefined properties \mathbf{p} that are chosen for the evaluation.

The properties used to define the quality of the conversation can be 1) mathematical, e.g. measuring the average degree and sparseness, 2) semantic, e.g. number and type of triples, 3) knowledge integrity, e.g. conflicts, outliers, analogies, completeness, 4) subjective values, e.g. sentiment and emotion, certainty, trust, and 5) dialogue properties, e.g. turn-property ratio’s, utterance type distribution and density, style and quality of expression. These measures depend on the capability to extract \mathbf{p} from the unknown features \mathbf{f} -implicit in image, audio and text signals- and the modelling of these properties in the knowledge graph \mathbf{g} . An interaction will result in a series of graphs over time. A cumulative graph can be seen as an episodic Knowledge Graph (eKG) (Baez Santamaria et al., 2021) for which the qualitative evaluation over time can provide valuable additional insight.

In this paper, we implement the above formal model and propose a set of properties \mathbf{p} , defined as RDF triples, that correlate with human judgements and can be used across different conversational setups and for different property detection systems.

3.1 Model

Figure 1 shows the interpretation of the statement "I like reading 1984" in an eKG. The core triple reflects that "Nicole" "reads" something labelled as "1984". The triple itself is a named graph representing a claim. The claim is mentioned (denotedBy) by the speaker Nicole (attributedTo) and perspective values are attributed to this mention, such as sentiment:positive, polarity:positive, certainty:certain. The model can represent multiple mentions of the same triple, with different perspective values and/or attributed to different speakers.

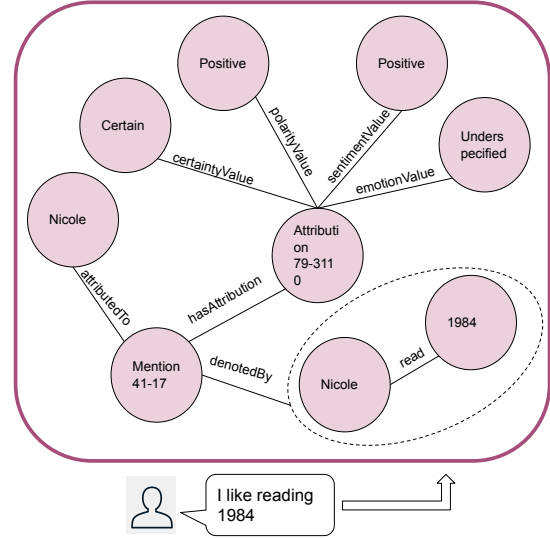


Figure 1: Example of how an utterance is converted into an episodic graph with source perspective values.

4 Capturing interaction and episodic knowledge

In order to apply graph metrics, we need to generate properties from the (multimodal) signals produced during the interaction. As we represent these properties as RDF triples, we rely on a text-to-triple instrument to detect these properties. This instrument’s effectiveness determines the evaluation’s depth, precision and richness; therefore, the most precise and standardised instrument is preferable. However, using the same instrument across interactions, we can compare the interactions and draw conclusions about its conversational variables.

Another factor for our evaluation is the graph properties modelled in the eKG. Some properties are generic and can be measured in any graph, whereas other properties depend on the semantics of the data model. In the following subsections, we discuss conversational variables when comparing interactions, the metrics that can be applied to the generated graphs and how they depend on specific property types.

4.1 Conversational variables

Our framework is agnostic to an interaction setup. We can thus have various combinations of prompt and response participants, and we can use different triples extractors, as shown in Figure 2. Schematically, the interlocutors (P1 and P2) both produce conversational signals to which we can apply triple-extraction, updating the eKG. Our evaluation framework is applied to this eKG, which

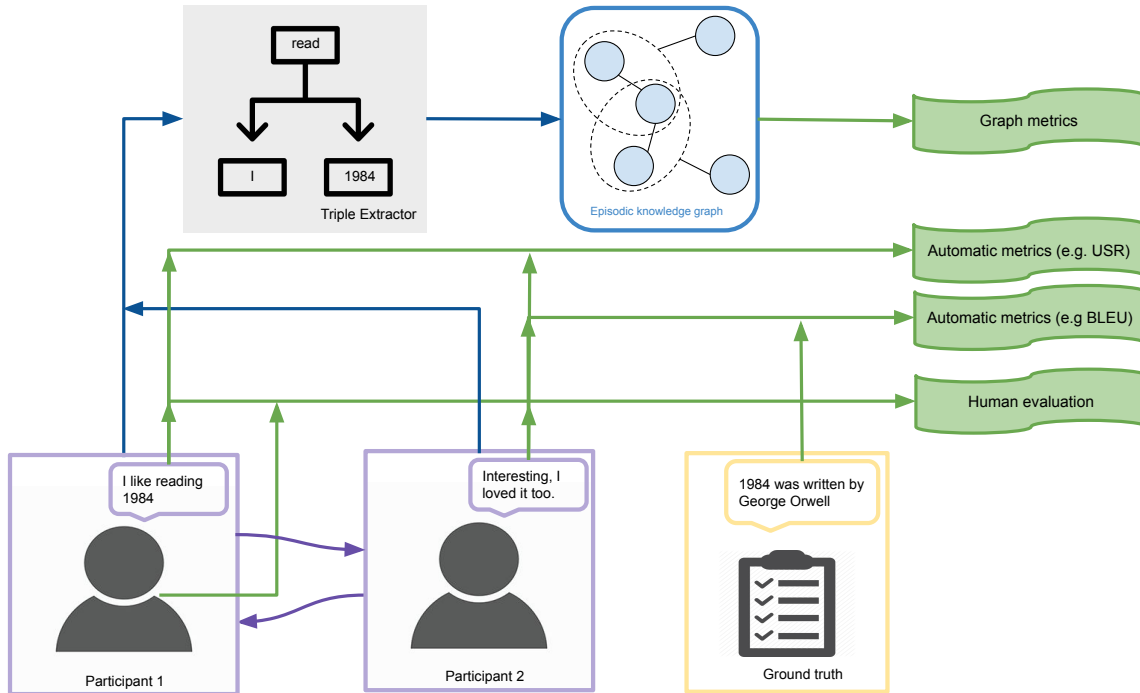


Figure 2: Schematic representation of interaction process (in purple), golden data (in yellow), EKG construction (in blue) and evaluation (in green). Arrows flowing into the metrics determine the components needed for their calculation.

is reference-free and considers the interaction between both interlocutors. This differs from other evaluation frameworks, such as USR, which only evaluates P2 as a coherent response to P1, and BLUE, METEOR and BERTscore, which evaluate P2 against a ground-truth response.

Triple extractor Regardless of the interlocutors, we can apply triple extraction to the utterances a-posteriori and derive an eKG from the communication. The extracted triples represent factoid information and possibly the speakers’ perspectives. Some approaches to extract triples are: StandforOpenIE (Angeli et al., 2015), spaCy’s Dependency parser (Vasilev, 2020), or tailored Context Free Grammars (CFG).

Agents As for the type of participants, we can take recorded human-human dialogues and apply triple extraction to each prompt from an actor. Similarly, machine-machine conversations can be generated where a chatbot mimics human input as a *prompt agent* and another chatbot functions as the *response agent*.

4.2 Graph metrics

Our formal model allows us to evaluate conversations under different frameworks: as a mathemati-

cal object (group A), as an RDF knowledge representation tool (group B), and as an eKG hosting the accumulation of interactions (group C). For groups A and B we used an implementation by Pernisch et al. (2020) while for groups C we implemented the metrics using *rdflib*.

After an exploratory analysis with 62 graph metrics (A-15, B-27, C-20), we select a subset of 24, as many of these metrics are compositional and therefore correlate highly and mostly reflect the same insights. Our selected Group A metrics include volume (number of nodes and edges), centrality (average node degree, degree centrality, and closeness), connectivity (average degree connectivity and assortativity), clique (number of strong connected components), entropy (centrality and closeness entropy) and density (sparseness). Group B metrics include volume (number of axioms) and density (average population). Finally, Group C metrics include density (ratios comparing claims to triples, perspectives to triples, conflicts to triples, perspectives to claims, mentions to claims, conflicts to claims), and interaction (average perspectives per claim, mentions per claim, turns per interaction, claims per source and perspectives per source).

The above measures can be applied to the evolving eKG during the interaction or a posteriori. Al-

Table 1: Statistics for 9 conversations with the number of turns, claims and perspectives. The conversation effect is measured by the average claim-triples (claim density) and perspective-triples (perspective density) per turn.

Category	P1	P2	Triple Extractor	Turns	Claims	Claim density	Perspective	Perspective density
Human - Machine	Student	Leolani	CFG	83	27	0.33	23	0.28
Human - Machine	Student	Leolani	CFG	57	18	0.32	14	0.25
Human - Machine	Student	Leolani	CFG	45	17	0.38	14	0.31
Human - Machine	Student	Leolani	CFG	55	14	0.25	11	0.20
Human - Machine	Student	Leolani	CFG	78	21	0.27	16	0.21
Human - Machine	Student	Leolani	CFG	80	22	0.28	17	0.21
Machine - Machine	Blenderbot	Leolani	CFG-spacy	298	6	0.02	4	0.01
Machine - Machine	Blenderbot	Eliza	CFG-spacy	207	24	0.12	22	0.11
Human - Human	Monica	Chandler	Stanford-OIE	243	109	0.44	0	0.00

though these measures were not originally intended as quality measures for interaction, we hypothesize that some of these measures can be used to characterize a conversation, and the resulting graphs can be compared independently. For example, the average claims per source may signal how much information the agent is getting per person it interacts with, the ratio of mentions to claims signals how much a factoid has been repeated in conversation, while the ratio of perspectives to claims signals how much diversity of opinions or sentiment has been expressed on claimed factoids.

5 Experimental setup

Table 1 shows an overview of the conversations analyzed. Details on the artificial agents and triple extractors can be found in the Appendix as well as an in-depth analysis of the triples extracted. We calculate the number of claims and the number of perspectives from each session. To measure the overall effectiveness in interpreting the conversation, we derive the density of claims and perspectives per turn³, which can be seen as the first crude measure of quality.

Human - Machine conversations Three groups of students had two conversations with Leolani (Vossen et al., 2018) during which they had to introduce themselves. Students were instructed to converse for 5 to 10 minutes or 30-40 turns per conversation. They were also instructed about the type of sentences and expressions from which the agent could extract triples to make the conversations more successful. Conversations will be more cumbersome in real open settings where users do not know what the agent can understand.

³Note that the number of turns includes Participant 2’s responses, while the extraction focuses on claims and perspectives from P1. About half of the total number of turns are utterances from the students, which makes densities of 0.25 and higher still effective.

Machine - Machine conversations We set up dialogues between Blenderbot (Roller et al., 2020) and Leolani, and Blenderbot with Eliza (Weizenbaum, 1966), where we extract triples using a tailored Context Free Grammar (CFG) and spaCy’s dependency parses.

Human - Human conversations We take all dyadic dialogues in the Friends dataset (Poria et al., 2018) between Monica and Chandler and extract triples using StanfordOIE (Angeli et al., 2015).

6 Comparing within conversational variables

To measure the quality of multiple interactions across different participants with the same agent, we collected human judgements of the human-machine conversations and compared these with the graph metrics groups.

6.1 Human evaluations

The students evaluated the agent’s responses in their conversation using the DSTC9 Track3 challenge evaluation metrics (Mehri et al., 2022), which form three submetric groups: 1) enjoyability (Interesting, Engaging, Specific, Relevant), 2) semantic correctness (Correct, Semantically Appropriate, Understandable), and 3) fluency. Each conversation got between two to four evaluations where students score each turn for all submetrics and overall score. Table 2 shows the aggregated results, showing the averaged scores for all six conversations. Although Leolani has limited communication skills, most ratings fall above mid-range.

Table 3 shows the average overall score (2.73) and the average over the submetrics (2.84), which hints that the submetrics comprehensively indicate the overall appreciation. The submetrics vary across but are close to the overall average. The enjoyable submetrics score lower than the semantic correctness and fluency ones.

Table 2: Human ratings of six conversations. Score range: 1 (very bad) to 5 (very good).

	Conversations					
	1.1	1.2	2.1	2.2	4.1	4.2
Interesting	2.65	2.97	1.56	1.60	2.98	1.91
Engaging	2.79	3.12	1.92	2.69	2.94	2.22
Specific	2.94	2.84	2.02	2.31	2.68	2.02
Relevant	3.25	3.86	2.83	2.76	2.97	2.25
Correct	3.15	3.93	2.63	2.55	3.01	2.20
Semantic appr.	3.06	3.91	2.44	2.48	2.98	2.27
Understandable	3.75	4.06	3.33	3.05	3.31	2.97
Fluent	3.64	4.13	2.67	2.45	3.16	3.08
Average submetrics	3.15	3.60	2.42	2.49	3.00	2.37
Overall hum.	3.16	3.41	2.35	2.57	2.74	2.12

6.2 Automatic evaluations

Following DSTC9-track3, we scored the agent responses using the USR model ("adamlin/usr-topicalchat-roberta_ft"). We implemented a likelihood score (USR LLH) that averages the masked-task prediction of the model for every token in the agent response given the preceding utterances as context (up to 300 characters). For every token, we get the top 20 predictions to get the token’s score, or a score of zero if it was not listed. We also average the likelihood of the highest-scoring token (USR MAX) as the perfect response according to the pre-trained model. We normalized the USR scores to a 5-point Likert scale to match it with the human ratings and averaged over all responses and conversations. The response of our agent scores significant lower (USR LLH=1.68) than the Overall Human Rating of the conversations (2.73). The maximal possible score by USR is close the to Overall Human Rating (USR MAX=2.78).

Table 3: Average human ratings. Score range: 1 (very bad) to 5 (very good). Mean Squared Error (MSE) for submetric and USR metric averages against the overall human rating.

	Avg	MSE against overall human score		
		Human	USR LLH	USR MAX
Interesting	2.28	0.05	0.15	0.16
Engaging	2.61	0.06	0.13	0.19
Specific	2.47	0.06	0.13	0.16
Relevant	2.99	0.04	0.13	0.16
Correct	2.91	0.03	0.14	0.13
Semantic appr.	2.86	0.03	0.14	0.15
Understandable	3.41	0.08	0.14	0.15
Fluent	3.19	0.09	0.19	0.18
Average submetrics	2.84	0.05	0.19	0.18
Overall hum.	2.73		0.15	0.16

We measured the Mean Squared Error (MSE) by comparing the USR scores and the human submetrics against the Overall Human Rating. The MSE scores for the human submetrics below 1 point. The USR LLH and USR MAX scores are 2 to 3 times

higher but remain below 2 points, which means they deviate less than average from the human norm. Finally, our agent response (USR LLH) in most cases correlates better than the most likely predicted tokens from the model itself (USR MAX).

6.3 Episodic knowledge graph evaluations

The previous evaluations (human questionnaires and automatic USR) do not evaluate the quality of the knowledge communicated. To that end, we apply the graph measures described in Section 4.2 to the eKGs of the student conversations. We compute the correlations between the graph metrics and the human and automatic metrics for each student conversation (Appendix Figure 4). Two patterns are visible: metric group A correlates more strongly with human evaluations, while metric group B correlates with automatic evaluation. Seven of the human metrics correlate the most with the average degree per node in the eKG. The other two human evaluations, Overall Human Rating and Relevance, correlate the most to sparseness.

7 Comparing across conversational variables

Three major factors determine the resulting graph: 1) the (human) participant, 2) the agent’s capability to understand the prompts, and 3) the agent’s capability to respond adequately. Our eKG-based evaluation, therefore, genuinely evaluates interaction from both ends. This makes it possible to evaluate interaction across different (types of) people with the same and/or different agents, resulting in different graphs due to the human input and/or the agents’ capabilities.

7.1 Correlation with human judgements

Table 4 shows the values between human evaluations, automatic evaluations, and the graph metric with highest correlation.⁴ For 6/9 human metrics, including Overall Human Rating, two graph metrics correlate more strongly than the USR metrics. We interpret this as evidence for these two graph evaluations to approximate human evaluations instead of USR evaluations. Recall that [Gunasekara et al. \(2020\)](#) reported a correlation of 0.3 for USR concerning the Topical-Chat evaluation, which is higher than the score obtained for our conversations. Since USR was fine-tuned on Topical-Chat

⁴Chatbot conversations have not been evaluated manually.

training data, it is expected to reflect stronger coherence relations on these conversations. Furthermore, Topical-Chat consists of human-human conversations replaced by system responses, whereas our conversations are naturally-born agent conversations that partly come from the inner drives of the agent. The task to generate an appropriate response is more challenging for our agent compared to fine-tuned language models that mimic human responses.

Table 4: Human evaluation, automatic scores and the most correlated graph metric.

	Average degree	Sparseness	USR LLH
Interesting	0.088	0.077	0.148
Engaging	0.158	0.145	0.076
Specific	0.124	0.067	0.072
Relevant	0.055	0.062	0.091
Correct	0.071	0.040	0.128
Semantically Appropriate	0.124	0.076	0.053
Understandable	0.119	0.050	-0.013
Fluent	0.184	-0.061	-0.039
Overall Human Rating	0.120	0.194	0.088

Average degree Average node degree reflects how many edges a node on the eKG has. Figure 3a exposes a relation between average degree and fluency. Conversations 1.2 and 1.1 scored the highest on fluency, while conversation 1.1 shows the lowest average degree curve. In contrast, conversation 2.2 scored the lowest on fluency and showed an incremental behaviour for this metric. Both Blenderbot conversations show a steep increase of the average degree as the conversation proceeds. Manual inspection of these dialogues reveals that Blenderbot becomes repetitive after several turns, resulting in an extreme increase of the average degree. Responses from Leolani are more repetitive than responses from Eliza, which is consistent as Blenderbot is trained with PersonaChat (Zhang et al., 2019), Empathetic Dialogues (Rashkin et al., 2018) and Wikipedia topic conversations (Dinan et al., 2018), making Blenderbot responsive for Eliza’s empathic prompts to talk about personal relations and emotions. The responses from Leolani, on the other hand, are based on its drives which can be more obscure and less "human", causing Blenderbot to fall back on standard responses rapidly. Finally, the conversation between Monica and Chandler shows a different pattern, where the degree drops linearly. The fast decreasing curve can be explained by high fluency between (scripted) human-human dialogues.

Sparseness Sparseness reflects how well connected a graph is. Figure 3b shows that, for all conversations, sparseness decreases as the conversation proceeds and the Overall Human Rating get higher as the eKG gets less sparse. Conversation 4.2 has the slowest decaying curve, while conversation 1.1 decays the fastest. Conversation 2.2 plateaus for a few turns, which is reflected by having the lowest Overall Human Rating. Since all eKGs have the same sparseness starting point, this suggests that conversations that fail to expand these initially dense graphs might not be successful. The Blenderbot-Leolani conversation is less steep, suggesting that it is less successful. The Blenderbot-Eliza conversation appears to be very similar in decreasing sparseness to the student-Leolani conversations, thus confirming that Blenderbot and Eliza are well aligned. The most effective conversation is shown by Monica and Chandler, having a curve that decays to the same level as conversation 1.1, but decays further.

7.2 Complementing human judgements

While some metrics correlate with human judgements, other graph metrics complement human evaluations. Evaluating conversations as eKGs allows to observe how much knowledge has been accumulated so far, how much diversity of opinions and conflicting information has been encountered, how often the same factoids are mentioned, how long conversations are, and how much knowledge has been acquired per source.

Ratio of mentions to claims Figure 3c shows the ratio of mentions to claims, which relates to how often the same topics are discussed in conversation. An increasing curve implies a repetitive conversation. The blender-Leolani conversation has the steepest curve, as the conversation stagnated with BlenderBot repeating the same factoid ("I have a dog"). In contrast, the Blenderbot-Eliza conversation goes well, almost like the student-Leolani conversations but still repetitive. Once again, the highest quality conversation is the human-human, with the lowest ratio.

Ratio of perspectives to claims Figure 3d shows the ratio of perspectives to claims, where higher means conversations contain more diverse views on the same topics. On the contrary, a lower ratio represents a series of broad conversations on their topics limited to the views of a few, if not only a single source. The lowest curve belongs to

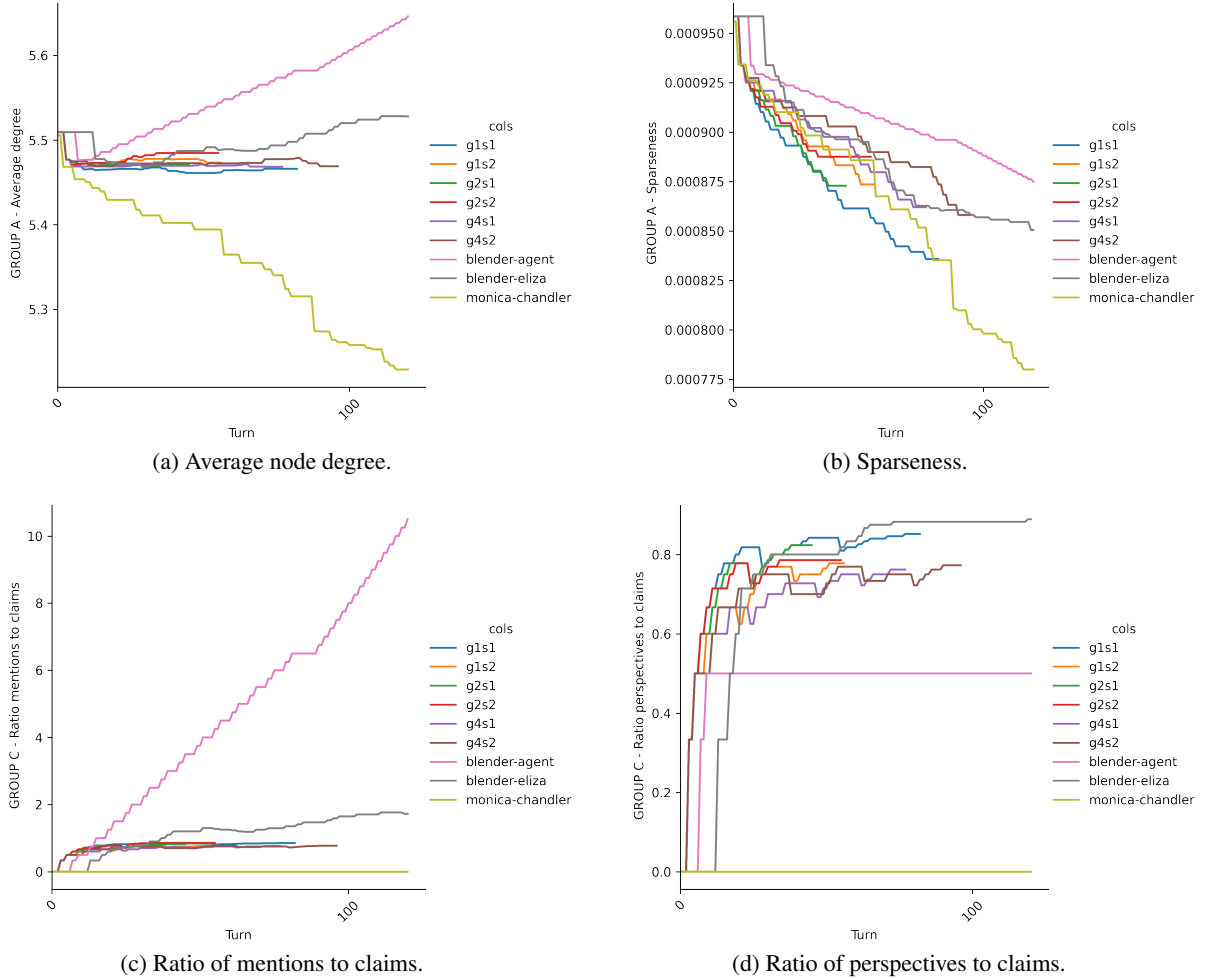


Figure 3: Selected graph metrics progression over turns for all student conversations

the Chandler-Monica conversation; due to using StanfordOIE, which does not extract perspectives.

8 Conclusion

We presented a new method for evaluating dyadic interactions that does not require a ground-truth interaction or a human judgment a posteriori. Our method analyses the episodic Knowledge Graph that results from interpreting prompts. Like the USR score, our method is automatic and reference-free. However, we provided evidence that our evaluation correlates better with human judgments and gives deeper insight into the knowledge built up from the conversation. We also conclude that the graph metrics provide nuanced information about the growth of knowledge resulting from the interaction. Although we cannot yet say anything about the absolute score of an interaction, we can compare different interactions based on the resulting graph and observe differences in graph properties

that may or may not be desirable.

Note that this line of work aims to approximate human evaluations in a cost-free manner. Yet, human evaluations are generally highly subjective and not reproducible to evaluate a conversation. Thus, even though Overall Human Ratings highly correlate to eKG sparseness, outliers arise due to the impact of individual judgments per student group.

In future work, we want to validate these metrics by evaluating benchmark datasets, similar to the methodology by Li et al. (2019). We also want to direct the conversations to aim for certain types of knowledge and perspectives to validate that our metrics can detect such different intentions. We want to demonstrate that we can associate graph structures with functionality, e.g., having broad or deep knowledge of subjects, being able to direct people to trustworthy informants, or directing people to shared interests.

Supplemental Material Statement: Code and data are publicly available from <https://github.com/selBaez/evaluating-conversations-as-ekg.git>.

Acknowledgements This research was funded by the Vrije Universiteit Amsterdam and the Netherlands Organisation for Scientific Research (NWO) via the *Spinoza* grant (SPI 63-260) awarded to Piek Vossen, the *Hybrid Intelligence Centre* via the *Zwaartekracht* grant (024.004.022) and the research project (9406.D1.19.005) *Communicating with and Relating to Social Robots: Alice Meets Leolani*.

References

- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354.
- Selene Baez Santamaria, Thomas Baier, Taewoon Kim, Lea Krause, Jaap Kruijt, and Piek Vossen. 2021. [EMISSOR: A platform for capturing multimodal interactions as episodic memories and interpretations with situated scenario-based ontological references](#). In *Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR)*, pages 56–77, Groningen, Netherlands (Online). Association for Computational Linguistics.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, 54(1):755–810.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. *arXiv preprint arXiv:1811.01241*.
- Siska Fitriani, Merijn Bruijnes, Deborah Richards, Amal Abdulrahman, and Willem-Paul Brinkman. 2019. What are we measuring anyway? -a literature survey of questionnaires used in studies reported in the intelligent virtual agent conferences. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, pages 159–161.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, Dilek Hakkani-Tür, and Amazon Alexa AI. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.
- Chulaka Gunasekara, Seokhwan Kim, Luis Fernando D’Haro, Abhinav Rastogi, Yun-Nung Chen, Mihail Eric, Behnam Hedayatnia, Karthik Gopalakrishnan, Yang Liu, Chao-Wei Huang, et al. 2020. Overview of the ninth dialog system technology challenge: Dstc9. *arXiv preprint arXiv:2011.06486*.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Margaret Li, Jason Weston, and Stephen Roller. 2019. Acute-eval: Improved dialogue evaluation with optimized questions and multi-turn comparisons. *arXiv preprint arXiv:1909.03087*.
- Chin-Yew Lin and FJ Och. 2004. Looking for a few good metrics: Rouge and its evaluation. In *Ntcir workshop*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Shikib Mehri and Maxine Eskenazi. [Unsupervised evaluation of interactive dialog with dialogpt](#).
- Shikib Mehri and Maxine Eskenazi. 2020. [USR: An unsupervised and reference free evaluation metric for dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Shikib Mehri, Yulan Feng, Carla Gordon, Seyed Hossein Alavi, David Traum, and Maxine Eskenazi. 2022. Interactive evaluation of dialog track at dstc9. pages 5731–5738.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

- Romana Pernisch, Daniele Dell’Aglia, Mirko Serbak, and Abraham Bernstein. 2020. Chimp: Visualizing ontology changes and their impact in protégé. In *Fifth International Workshop on Visualization and Interaction for Ontologies and Linked Data*, pages 47–60. CEUR Workshop Proceedings.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2018. Towards empathetic open-domain conversation models: A new benchmark and dataset. *arXiv preprint arXiv:1811.00207*.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M Smith, et al. 2020. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*.
- Yuli Vasiliev. 2020. *Natural Language Processing with Python and SpaCy: A Practical Introduction*. No Starch Press.
- Piek Vossen, Selene Baez, Lenka Bajčetić, Suzana Bašić, and Bram Kraaijeveld. 2019. *Leolani: A robot that communicates and learns about the shared world*. In *ISWC 2019 Satellites*, CEUR Workshop Proceedings, pages 181–184. CEUR-WS. 2019 ISWC Satellite Tracks (Posters and Demonstrations, Industry, and Outrageous Ideas), ISWC 2019-Satellites ; Conference date: 26-10-2019 Through 30-10-2019.
- Piek Vossen, Selene Baez, Lenka Bajčetić, and Bram Kraaijeveld. 2018. Leolani: a reference machine with a theory of mind for social communication. In *International conference on text, speech, and dialogue*, pages 15–25. Springer.
- Joseph Weizenbaum. 1966. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*.

A Appendix

A.1 Leolani, the agent

The Leolani agent creates an eKG on the fly and formulates a response to a prompt based on the changes in the graph. At the start, an empty eKG is initialized with the Nice-to-Meet-You (N2MU) ontology (Vossen et al., 2019) that defines basic things in the world such as people, locations, and some social properties, but without any knowledge about individual people. Triple extraction during the conversation adds such knowledge.

In addition to passive processing of human prompts, our agent Leolani reflects on the changing eKG during the conversation. After each change, a series of semantically meaningful SPARQL queries is applied to the eKG to detect specific graph properties that are deemed relevant to respond to. These handmade query patterns detect phenomena like conflicts, uncertainty, analogies, generalizations, knowledge gaps, trust in knowledge sources, the novelty of factoids, and the relevance of factoids to interlocutors. We consider these queries as "thoughts" that reflect on the current state of the eKG, given the latest incoming information. Each type of *thought* is associated to a broad higher level goal, which we call "*drives*", for example, to improve the correctness or completeness of the agent’s current knowledge, or to improve the engagement or social connection to the interlocutor. These *drives* result in speech acts such as statements or questions that the agent directs to the interlocutor, aiming of obtaining a relevant response. In the next short dialog, the agent Leolani learned a new concept "reading" and tries to generalise this to all people by asking a question to the human. The next statement by the human triggers the agent to ask another follow-up question:

- H: I like reading
- A: I never heard about reading before! can you tell me, is reading something people do?
- H: Yes, people read books
- A: I did not know that, what else do people read?

A.2 Context-Free-Grammar Triple Extractor

Triple extraction is performed by a dependency parser with a Context-Free Grammar (CFG) consisting of 12 rules and a lexicon with mainly closed-class words. It is adapted to resolve deictic references to the addressee or other observed objects in a situation. As the conversation is open and can be about any topic, part-speech-tagging and named entity-recognition is used to detect patterns of open-class words as subjects, predicates, and objects for triples.

A.3 Sample student-Leolani conversations

```
Leolani: "What's up? What is your name? Stranger?"
Student1: "Student1"
Leolani: "So your name is Student1?"
Student1: "yes"
Leolani: "Would you like to chat? I'll do my best to keep up"
Student1: "I like chatting"
Leolani: "I would like to know. What types of Instance like chatting do person or Instance usually like"
Student1: "Person likes convos"
Leolani: "I am curious. What types of noun or Instance like convos do tops or Instance usually like"
Student1: "Convos are not people"
Leolani: "Let me ask you something. Has convos work at institution?"
Student1: "I work at institution"
Leolani: "Let me ask you something. Has thomas experience touch?"
Student1: "I have experience touch"
Leolani: "I am curious. Has thomas own object?"
Student1: "I own object"
Leolani: "Interesting! I am excited to get to know about you!"
Student1: "What do you want to know?"
Leolani: "I know agent usually want to verb.cognition, but I do not know this case"
Student1: "I like sushi"
Leolani: "Exciting news! I did not know anything that Student1 like"
Student1: "I also like cats"
Leolani: "If you don't mind me asking. What types of animal or Instance like cats do person or Instance usually like"
```

Listing 1: Extract from conversation with Student 1

A.4 Analysis of student-Leolani conversations

In total, 81,631 triples were generated from the student conversations. Note that not all triples count as factual knowledge. 23 predicates were extracted of which "know" and "be" are most frequent (Table 5). The "know" predicates are mostly derived from introducing oneself to the agent. The predicates further show that the communication is open, although the semantics is not further defined beyond the predicate itself, e.g. there is no relation between "smell" and "can-smell". The triples further contained 31 unique subjects and 78 unique objects; most occurred only once in the communication (Table 6).

Table 5: Predicate labels and their frequency in student conversations.

know	43	can-fly	2	favourite-animal-is	1
be	39	like-to	2	favourite-cat-is	1
like	22	work-at	2	fly	1
live-in	9	be-in	1	hair-color-is	1
sense	7	be-to	1	own	1
have	6	can-learn	1	smell	1
be-from	3	can-smell	1	wear	1
love	3	could-help	1		

Table 6: Entity labels and their frequency in student conversations.

Leolani	23	a-dog	1	go	1
student1	11	a-girl	1	great-that-i	1
student2	10	a-man	1	his-uncle	1
student3	8	a-shame	1	institution	1
student4	8	a-student	1	student10	1
student5	8	a-wise-man	1	know	1
student6	7	a-woman	1	lasagne	1
a-flamingo	6	airplanes	1	phd2-student5	1
student7	6	amstelveen	1	phd2-name	1
student8	5	an-animal	1	Leolani-new-things	1
amsterdam	4	an-aunt	1	my-daughter	1
student9	4	an-emotion	1	my-parents	1
my	3	an-uncle	1	student5-friend	1
orange	3	brown	1	now	1
reading	3	bulgaria	1	object	1
alkmaar	2	business-class	1	other-things	1
cats	2	candy	1	parents	1
chatting	2	city	1	people	1
convos	2	cook	1	person	1
garfield	2	cook-by-myself	1	pink	1
japanese-food	2	deloitte	1	rotterdam	1
love	2	dogs	1	shrimp	1
None	2	every-agent	1	student	1
phd1	2	experience-touch	1	sushi	1
10-fingers	1	favorite-of-student6	1	tapas	1
a-bird	1	flamingo	1	the-south-of-holland	1
a-color	1	food	1	two-hands	1
a-company	1	front-camera	1	what	1
a-country	1	garfield-favourite-food	1	yes-candy	1
a-daughter	1	glasses	1		

A.5 Full list of metrics tested

The full list of 62 metrics used, sorted by group:

- GROUP A
 - Total nodes
 - Total edges
 - Average degree
 - Average degree centrality
 - Average closeness
 - Average betweenness
 - Average degree connectivity
 - Average assortativity
 - Average node connectivity
 - Number of components
 - Number of strong components
 - Shortest path
 - Centrality entropy
 - Closeness entropy
 - Sparseness
- GROUP B
 - Total classes
 - Total properties
 - Total instances
 - Total object properties
 - Total data properties
 - Total equivalent class properties
 - Total subclass properties
 - Total entities
 - Total inverse entities
 - Ratio of inverse relations
 - Property class ratio
 - Average population
 - Class property ratio
 - Attribute richness
 - Inheritance richness
 - Relationship richness
 - Object properties ratio
 - Datatype properties ratio
 - Total concept assertions
 - Total role assertions
 - Total general concept inclusions
 - Total domain axioms
 - Total range axioms
 - Total role inclusions
 - Total axioms
 - Total aBox axioms
 - Total tBox axioms
- GROUP C
 - Total triples
 - Total world instances
 - Total claims
 - Total perspectives
 - Total mentions
 - Total conflicts
 - Total sources
 - Total interactions
 - Total utterances

- Ratio claim to triples
- Ratio perspectives to triples
- Ratio conflict to triples
- Ratio perspectives to claims
- Ratio mentions to claims
- Ratio conflicts to claims
- Average perspectives per claim
- Average mentions per claim
- Average turns per interaction
- Average claims per source
- Average perspectives per source

A.6 Correlation matrix

Hereby we show the full correlation matrix between our proposed metrics and the human and automatic evaluations. Metrics related to conflicts are not informative since the short conversations did not produce conflicting information. Metrics related to perspectives are hindered because the simple triple extractor is limited in the range of perspectives it can extract.



Figure 4: Correlation matrix for evaluation metrics.