

Neural Path Hunter: Reducing Hallucination in Dialogue Systems via Path Grounding

Nouha Dziri*, Andrea Madotto†, Osmar Zaiane*§, Avishek Joey Bose‡

*University of Alberta, †Mila, McGill University, §Canada CIFAR AI Chair

‡The Hong Kong University of Science and Technology

dziri@cs.ualberta.ca

Abstract

Dialogue systems powered by large pre-trained language models exhibit an innate ability to deliver fluent and natural-sounding responses. Despite their impressive performance, these models are fitful and can often generate factually incorrect statements impeding their widespread adoption. In this paper, we focus on the task of improving faithfulness and reducing hallucination of neural dialogue systems to known facts supplied by a Knowledge Graph (KG). We propose NEURAL PATH HUNTER which follows a generate-then-refine strategy whereby a generated response is amended using the KG. NEURAL PATH HUNTER leverages a separate token-level fact critic to identify plausible sources of hallucination followed by a refinement stage that retrieves correct entities by crafting a query signal that is propagated over a k -hop subgraph. We empirically validate our proposed approach on the OpenDialKG dataset (Moon et al., 2019) against a suite of metrics and report a relative improvement of faithfulness over dialogue responses by 20.35% based on FeQA (Durmus et al., 2020). The code is available at <https://github.com/nouhadziri/Neural-Path-Hunter>.

1 Introduction

Conversation within a dialogue can be thought of as an exchange of utterances between two speakers. Each utterance is not independent of one another but is instead grounded within a larger dialogue context known to both parties (Jurafsky and Martin, 2018; Sordoni et al., 2015; Serban et al., 2016; Dziri et al., 2019). Indeed, if a response to an utterance fails to be faithful to some given knowledge—i.e. by producing false information—it is uninformative and runs the risk of jeopardizing the entire enterprise of conversation. More precisely, this means that in addition to being fluent, topical, and

* Corresponding author.

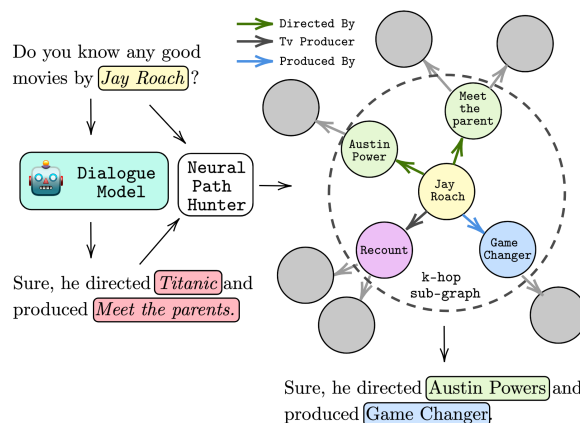


Figure 1: NEURAL PATH HUNTER overview.

grammatical, utterances within a dialogue must also be *factually correct*.

The faithfulness of responses is of principal importance when designing dialogue systems that are grounded using auxiliary knowledge such as Knowledge Graphs (KG). Despite maintaining plausible general linguistic capabilities, dialogue models are still unable to fully discern facts and may instead hallucinate factually invalid information. Moreover, empirical evidence for hallucination in Language Models (LM) runs contrary to known studies that these large models are capable of recalling factual knowledge, e.g. entities and relations in a KG, (Roberts et al., 2020; Petroni et al., 2019). This suggests that this inherent lack of controllability may be remedied by leveraging external oracle knowledge. However, existing approaches to knowledge grounding often suffer from a source-reference divergence problem whereby the reference contains additional factual information and simply training on the reference is insufficient to guarantee faithfulness (Wiseman et al., 2017; Parikh et al., 2020; Tian et al., 2019). Consequently, ensuring the faithfulness of knowledge grounded dialogue systems—via precise alignment of the source and reference—remains an open challenge.

Present Work. In this work, we focus on address-

ing the open problem of hallucination of factually invalid statements in knowledge grounded dialogue systems where the source of knowledge is a KG. We first identify prominent modes of hallucination by conducting a systematic human study on generated responses which reveals one major source of hallucination as the (mis)-use of wrong entities to describe factual content (Kryscinski et al., 2020), a problem that persists when naively applying language models in dialogue systems.

To enforce faithfulness to the misattribution of entities in grounded dialogue systems, we introduce NEURAL PATH HUNTER (NPH), a module that operates on hallucinated responses. NPH follows a generate-then-refine approach by augmenting conventional dialogue generation with an additional refinement stage enabling the dialogue system to correct potential hallucinations by querying the KG. NPH grounds dialogue generation by constraining the flow of conversation to be supported by a valid path on the KG. To do so, the module combines a token-level hallucination critic that masks out entities of concern in an utterance, followed by a pre-trained non-autoregressive LM which prescribes contextual representations for each masked entity. This is then fed sequentially to an autoregressive LM to obtain output representations. These output representations can then be used to efficiently launch a query on the KG—effectively modelling dialogue as a signal being propagated on a local k -hop subgraph whereby locality is enforced through the conversation history—returning factually correct entities. Our proposed approach is applicable to any generated response whenever an available KG is provided and works without further fine-tuning. The high-level overview of our proposed approach is outlined in Fig. 1 and exemplar machine-generated responses post-refinement are presented in Table 8 in §H. Our main contributions are summarized as follows:

- We conduct a comprehensive human study on hallucinations generated by state-of-the-art dialogue systems which reveals that the main mode of hallucinations is through the injection of erroneous entities in generated responses.
- We propose NEURAL PATH HUNTER, which leverages facts supplied by a KG to reduce hallucination in any machine-generated response.
- We empirically demonstrate that NEURAL PATH HUNTER substantially reduces hallucinations in

KG-grounded dialogue systems with a relative improvement of 20.35% in FeQA, a QA-based faithfulness metric (Durmus et al., 2020), and an improvement of 39.98% in human evaluation.

2 Hallucination in KG-grounded Dialogue Systems

We consider the task of generating factual and grounded dialogue when presented with auxiliary structured knowledge. In particular, we focus on factoids taken from multi-relational graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$, termed Knowledge Graphs (KG). Each KG consists of a set of directed edge triples $t = \langle [\text{SBJ}], [\text{PRE}], [\text{OBJ}] \rangle$, where $[\text{SBJ}], [\text{OBJ}] \in \mathcal{V}$ are nodes denoting subject and object entities and $[\text{PRE}] \in \mathcal{R}$ is a predicate that can be understood as a relation type. Broadly speaking, we say that a neural dialogue system is guilty of hallucinating whenever it generates a factual sentence that is not supported by a valid path in a k -hop subgraph $\mathcal{G}_c^k \subset \mathcal{G}$ of the original KG anchored around a context entity c .

As a starting point for our investigation, we study the various types of hallucinations a model may inject into an otherwise satisfactory response. Specifically, we explore the circumstances under which LMs are likely to exhibit unfaithful behaviour through misappropriation of entities (e.g. Barrack Obama was the President of Canada). Inspired by (Maynez et al., 2020) for KG-grounded dialogue systems we hypothesize—among other possible mechanisms—hallucination can take form as either *intrinsic* or *extrinsic* to the provided KG.

Definition 2.1 (Extrinsic Hallucination). *An extrinsic hallucination corresponds to an utterance that brings a new span of text that does not correspond to a valid triple in \mathcal{G}_c^k .*

From the perspective of definition 2.1, an utterance that might be partially faithful is still guilty of hallucination if there exists any injection of knowledge not authentically captured in \mathcal{G}_c^k . Despite this, external hallucinations can often be easier to identify due to their egregious nature. For example, the dialogue sample in Fig. 1 contains an external hallucination as the entity in question “Jay Roach” did not direct the movie “Titanic” and it is not supported within the 1-hop subgraph. On the other hand, the generated response may identify the correct set of entities but make false claims about their relationship which leads to the following definition.

Definition 2.2 (Intrinsic Hallucination). *An intrinsic hallucination corresponds to an utterance that*

History	A: Do you know the book The Witches? B: The Witches is written by Roald Dahl. He also wrote The Champion of the World.
GPT2-KG	A_{gen} Yes he did. He also wrote The Time Machine and The Invisible Man .
Gold knowledge	T_1 : [Charlie and the Chocolate Factory, written by, Roald Dahl] T_2 : [Charlie and the Chocolate Factory, has genre, Fantasy]
Top-5 Paths	T'_1 : [The BFG, written by, Roald Dahl] T'_2 : [Charlie and the Chocolate Factory, written by, Roald Dahl] T'_3 : [You Only Live Twice, written by, Roald Dahl] T'_4 : [James and the Giant Peach, written by, Roald Dahl] T'_5 : [Tales of the Unexpected, TV regular appearance, Roald Dahl]
NPH response	A_{fix} Yes he did. He also wrote The BFG and Charlie and the Chocolate Factory .

Table 1: A selected response based on a GPT2-KG test response before and after applying NEURAL PATH HUNTER. The span of texts highlighted in red indicate the hallucinated entity mentions whereas the ones highlighted in green indicate the retrieved correct entity mentions.

misuses either $[SBJ]$ or $[OBJ]$ in \mathcal{G}_c^k such that there is no direct path between the two entities.

Intrinsic hallucinations inject false information by condensing information from the KG in a wrong way. For instance, claiming that “Jay Roach” produced “Meet the Parents” is an incorrect association of the true relationship between these entities.

To ascertain the degree to which KG-grounded dialogue systems hallucinate and the nature of these hallucinations, we conduct a systematic evaluation by soliciting human judgement. We first fine-tune a LM on the OpenDialKG dataset (Moon et al., 2019) which contains a turn-based dialogue between two speakers on extracted triples from a known KG. The sequential nature of such turn-based dialogues grounded via extracted KG-triples effectively renders the entire conversation as a path traversed on the KG (see §A for dataset details).

2.1 Modes of Hallucination

Experimental Protocol. As a demonstrative example, we use a pre-trained GPT-2 model (Radford et al., 2019) as the backbone of a neural dialogue system. To fine-tune GPT2, we concatenate the dialogue history, the KG-triples $\langle [SBJ], [PRE], [OBJ] \rangle$ and the ground truth response and then train the model to predict the next word in the response. To explore the effect of different decoding strategies and their impact in injecting hallucinations, we sample 300 responses from each decoding approach. We investigate greedy search, beam search, nucleus sampling (Holtzman et al.,

GPT2-KG	Hallucination			Faith.	Gen.
	Ex	In	B		
Greedy	17.66	2.00	1.66	69.00	9.66
Beam Search	18.33	3.33	4.00	68.00	6.33
Nucleus 0.9	25.33	4.00	2.33	64.66	3.66
Nucleus 0.5	23.33	5.33	4.33	59.90	7.00
Top20	28.33	7.00	5.00	55.00	4.66

Table 2: Human assessment of random 1500 GPT2 dialogue responses generated using OpenDialkg. “Ex”, “In” and “B” mean extrinsic, intrinsic, and both hallucinations respectively. Each cell shows the mean percentage of responses with a specific dialogue property (see §B for confidence intervals).

2020) and top- k sampling (Radford et al., 2019) as representative decoding strategies.

For each dialogue sample, we crowd-source human judgement by soliciting evaluations from 3 different annotators from Appen¹, a high-quality annotation platform. Each annotator is tasked to first identify the presence of hallucination in the generated response when provided the dialogue history and KG triples. For samples where hallucination is present, we further ask the human annotators to identify whether the hallucination is extrinsic, intrinsic or both. If the response is not hallucinated, we ask them whether the response is faithful (i.e., supported by the triples) or generic (e.g., “I don’t know about that”). The results of the human assessment are shown in Table 2. Overall, we report the average Krippendorff’s alpha coefficient to be 0.72 on the annotator responses to the different questions which indicates high agreement. Using

¹<https://appen.com/>

Table 2, we make the following key observations:

Observation 1. Humans notice most hallucinations in KG-grounded dialogue systems are extrinsic.

Observation 2. A hallucination occurs the least in dialogue responses generated using a greedy decoding scheme. Conversely, top- k sampling results in the highest hallucination percentage (40.33%).

Observation 3. Increased diversity in response generation —i.e.(less generic), is positively correlated with an increase in hallucination e.g. Nucleus=0.9.

Observation 1 indicates that the dominant mode of hallucination for all decoding strategies in KG-grounded dialogue systems is extrinsic rather than intrinsic. In fact, we find that in the OpenDialKG dataset, 54.80% of the responses contain extra entity mentions that are not supported by either \mathcal{D} or \mathcal{G}_c^1 which may partially explain empirical observations. Observation 2 suggests that the model—when conditioned on factual knowledge—often assigns the highest probability mass to the correct response and sampling based on other distributions (e.g. top- k) invites hallucination in the generation process—a fact also observed in language modelling (Keskar et al., 2019). Observation 3 suggests an implicit trade-off between the different goals of response generation whereby improving the diversity of response can negatively impact its faithfulness. This reveals that in certain cases responses might be originally faithful to \mathcal{G}_c^k but increasing diversity encourages the model to hallucinate. In light of these important observations, the main goal of this paper is not necessarily to advance state-of-the-art decoding methods but instead to instrument an efficient technique to identify hallucinations as well as retrieve the correct entities from the KG.

3 Neural Path Hunter

We seek to design a dialogue refinement system capable of fixing generated utterances such that they are semantically relevant given the conversation history *and* supported within a provided KG. To do so, we introduce NEURAL PATH HUNTER (NPH) a refinement strategy that can be easily applied to any generated response without retraining the model. NPH is composed of two modules: A token-level hallucination critic and an entity mention retriever. The first module flags and masks out hallucinated entities in an existing response and can be trained offline. The second module accepts masked representations identified by the critic and builds contextual representation of these problem-

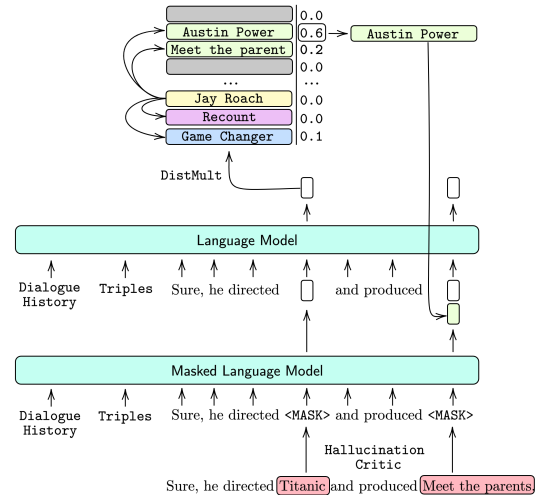


Figure 2: Entity Mention Retriever architecture.

atic tokens which are then used to retrieve more faithful entities by running a query over \mathcal{G}_c^k . We assume the local k -hop subgraph is either provided or extracted based on the dialogue history. The following sections describe the data preparation, training, and inference procedures for these submodules.

3.1 Problem Formulation

Each instance in the dataset is composed of a dialogue history $\mathcal{D} = (x_1, \dots, x_n)$, a set of j triples at turn n , $\mathcal{K}_n = (t_1, t_2, \dots, t_j)$ which together with \mathcal{D} must be used towards generating the response \bar{x}_{n+1} . Here, each individual triple $t_i = \langle [SBJ], [PRE], [OBJ] \rangle$ is extracted from a provided KG. Thus, the task is to generate a response \bar{x}_{n+1} that is faithful to a non-empty subset $M_n \subset \mathcal{K}_n$ —i.e., it can optionally talk about a few triples but not none. Specifically, the response \bar{x}_{n+1} may contain entity mentions $m_i \in \mathcal{V}$ which indicates a factual response that potentially needs to be refined using NPH. For our purposes, it is most convenient to represent each mention as a tuple of three elements that indicates the beginning of the mention at position m_i^b and the end at position m_i^e . In other words, we represent an entity mention m_i as $m_i = (m_i, m_i^b, m_i^e)$. These entity mentions may not be faithful at all if they do not belong to either a [SBJ] or [OBJ] in M_n (extrinsic hallucination) or they could inject false relationships between mentions via an unsupported path in \mathcal{G}_c^k by incorrectly utilizing a [PRE] (intrinsic hallucination). We target and correct these unfaithful entities through retrieval over \mathcal{G}_c^k in §3.3.

3.2 Token-level hallucination critic

To enforce faithfulness via refinement, we first identify the exact sources of hallucination in a given response. Based on the findings of human judgement in Tab.2 and §2.1, we find hallucination errors in a dataset like OpenDialKG are often associated with entity mentions such as names of people, movies titles, locations, etc. To flag entities of concern, we design a token-level hallucination critic C that consumes $\mathcal{D}, \mathcal{K}_n, \bar{x}_{n+1}$ and outputs the set of hallucinated entity mentions M_c . To train C , we choose to cast the problem as a sequence labelling task where a binary label is predicted at each word position. As there is no labelled training data available for this task, we create a synthetic dataset consisting of ground truth dialogue samples and corrupted negative samples. We explore two corruption processes that convert a regular clean ground-truth response x_{n+1} to its corresponding hallucinated one \hat{x}_{n+1} based on the type of hallucination we might expect to encounter —i.e. extrinsic and intrinsic.

1. **Extrinsic Negatives.** We replace each m_i in x_{n+1} with entities of the same type (e.g., person, location, etc...) but crucially not within \mathcal{G}_c^k and the dialogue history \mathcal{D} .
2. **Intrinsic Negatives.** We simply swap every pair [SBJ] and [OBJ] in x_{n+1} . For example, the response “Crescendo was written by Becca Fitzpatrick” \rightarrow “Becca Fitzpatrick was written by Crescendo” results in an intrinsic hallucination as in this case [PRE] is not bidirectional.

Overall, we apply a 60%/40% split of extrinsic versus intrinsic corruption strategies to the original train OpenDialKG to obtain a synthetic dataset to train C which is taken to be a pre-trained LM that is then fine-tuned on this binary classification task.

3.3 Entity Mention Retriever

An overview of the Entity Mention Retriever is depicted in Fig. 2. Having identified entities of concern in \bar{x}_{n+1} , we now wish to craft a query that can be efficiently run over \mathcal{G}_c^k . To do so, we model the generated response \bar{x}_{n+1} as a signal being propagated over \mathcal{G}_c^k which serves to capture the highest probability paths starting from the context node c the conversation may take if it was faithful. The context node c is extracted from ground truth triples available in the dataset and or \mathcal{D} . In order to run an effective query over \mathcal{G}_c^k , it is critical that the representation of all flagged $m_i \in M_c$ and edge triples

$\mathcal{E} \in \mathcal{G}_c^k$ are in the same representation space. Inspired by the Cloze task (Taylor, 1953), we obtain contextual representations of all m_i ’s identified by the critic by first masking them out before using a Masked Language Model (MLM). Operationally, we feed $\mathcal{D}, \mathcal{K}_n$, as well as the flagged set of entities to obtain contextual hidden state representations:

$$H = \text{MLM}(\mathcal{D}, \mathcal{K}_n, M_c) \quad (1)$$

As the MLM may return multiple hidden d -dimensional state representation for each $m_i \in M_c$, we simply apply a pooling operation to obtain a single representation for each entity —i.e. $h_i = \text{MaxPool}(h_b, h_e)$. To obtain the actual query q_i , we use an autoregressive LM which iteratively consumes an order dependent representation of h_i given by applying a learnable projection map $W : \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ to a concatenation of the current hidden state and the retrieved entity embedding e_{i-1} using previous query q_{i-1} as shown in Fig. 2,

$$q_i = \text{LM}(W(\text{concat}[e_{i-1}, h_i])),$$

KG-Entity Memory. Viewed another way, each q_i can be interpreted as a relation embedding for the masked position in \bar{x}_{n+1} . To effectively query \mathcal{G}_c^k , we must also represent all nodes in the same embedding space as q_i and in doing so effectively build a representation of \mathcal{G}_c^k which we call KG-Entity Memory. We explore two approaches towards this goal. The first uses the final hidden layer of a pre-trained GPT2 to obtain initial embeddings for each node in \mathcal{G}_c^k ². Our second approach uses CompGCN (Vashishth et al., 2020), which is a Graph Convolutional Network (Kipf and Welling, 2017) purposely built for multi-relational data. We initialize the CompGCN network offline with GPT2 embeddings for all entities and relations in the full graph \mathcal{G} before running a few rounds of message passing by optimizing for a standard relation prediction objective. Both approaches to KG-Entity memory embeddings can be further updated during training. Finally, to retrieve the correct entity for query q_i , we simply use a scoring function s to score every KG-Entity memory triple in \mathcal{G}_c^k —i.e. $t_i = \langle c, q_i, [\text{OBJ}] \rangle$. The retrieved entity is the [SUB] or [OBJ] that achieves the highest score.

3.4 Training the Entity Mention Retriever

To train the Entity Mention Retriever, we augment the conventional maximum likelihood objective

²Actually, GPT2 returns word piece representations and we use a pooling operation to get a single representation.

with an additional contrastive loss \mathcal{L}_{NCE} that encourages faithful retrieval. In particular, we use Noise Contrastive Estimation (NCE) (Gutmann and Hyvärinen, 2010) which forces the Entity Mention Retriever to learn a scoring rule such that $s(t_i) > s(t'_i), \forall t_i \in \mathcal{E}, t'_i \in \bar{\mathcal{E}}$ where $t_i = \langle c, q_i, [\text{OBJ}] \rangle$ is the edge-triple based on KG-entity memory and $t'_i = \langle c, q_i, [\text{OBJ}]^- \rangle$ is a negative sample where $[\text{OBJ}]^-$ ³ is sampled from a corruption distribution over edge triples $\bar{\mathcal{E}}$ not in \mathcal{G}_c^k . To compute \mathcal{L}_{NCE} , we draw n negative samples uniformly over all entities for each query q_i .

$$\mathcal{L}_{\text{NCE}} = -\log(s(t)) - \log\left(s(t) + \sum_{j=1}^n s(t'_j)\right).$$

At training time, we use teacher forcing (Williams and Zipser, 1989); first, we mask out all entity mentions within the gold response x_{n+1} , get their representations through a MLM and provide the ground truth entity mention concatenated with h_i at each time step in the LM. For the scoring function, we use DistMult (Wang et al., 2014) due to its simplicity in the absence of known structure over the modified triples e.g. translation, rotation, which are exploited in other popular scoring functions for KGs. By optimizing \mathcal{L}_{NCE} , we encourage the model to leverage the dialogue history, the position of the masked entity in x_{n+1} , and the k -hop subgraph to identify more faithful entities that are relevant to the conversation history. To train the Entity Mention Retriever, we thus jointly optimize \mathcal{L}_{NCE} and \mathcal{L}_{MLE} for the main language modelling task,

$$\mathcal{L} = \mathcal{L}_{\text{MLE}} + \lambda \mathcal{L}_{\text{NCE}}. \quad (2)$$

4 Experiments

We evaluate the ability of NEURAL PATH HUNTER towards reducing hallucinations in KG-grounded dialogue systems on the OpenDialKG dataset (Moon et al., 2019). At present, OpenDialKG is the only publicly available dataset that provides open-ended dialogue responses grounded on paths from a given KG, this is why we limit our experiments on this dataset. As there are no established metrics for this task, we consider a suite of task-specific and automated metrics to assess the different components of NPH and the degree of hallucination present. We use standard classification metrics such as F1-score, precision and recall to evaluate C and PPL

³or $[\text{SUB}]^-$ if c is an object

Model	FeQA	Critic	BLEU
GPT2-KG	26.54	19.04	11.79*
+ NPH	28.98*	11.72*	11.29
+ NPH-w/o NCE	26.02	17.91	10.98
+ NPH-w. COMPGCN	26.89	15.41	11.10
+ NPH-w/o MLM	27.01	15.02	10.88
+ NPH-w/o CRITIC	18.23	19.65	6.49
AdapterBot	23.11	26.68	10.56
+ NPH	27.21*	18.51*	10.74*
+ NPH-w/o NCE	24.02	25.02	9.98
+ NPH-w. COMPGCN	25.83	20.23	10.11
+ NPH-w/o MLM	26.02	21.04	10.06
+ NPH-w/o CRITIC	16.21	27.22	5.64
GPT2-KE	19.54	28.87	6.24*
+ NPH	26.21*	20.34*	6.06
+ NPH-w/o NCE	20.34	24.32	5.89
+ NPH-w. COMPGCN	23.23	21.21	6.01
+ NPH-w/o MLM	24.01	22.40	5.99
+ NPH-w/o CRITIC	15.89	30.71	3.49
Gold response	33.34	5.2	-

Table 3: Measuring the degree of hallucination of different models pre and post-refinement on generated samples based on the OpenDialkg test data. A higher FeQA score indicates an increase in faithfulness. The hallucination Critic (Critic) measures the percentage of hallucinated responses in the dataset. (* p -value < 0.001). NPH uses GPT2 emb. for the KG-Entity Memory.

to measure the quality of the LM. Similarly, we use retrieval metrics like Hits@ k , Mean Rank (MR), and Mean Reciprocal Rank (MRR) to evaluate the Entity Mention Retriever. Precise implementation details can be found in §D.

Hallucination Metrics. We consider 3 different hallucination metrics **M1-M3** that provide a multi-faceted measure of performance. Appendix §E outlines these metrics in detail. Succinctly, **M1**. BLEU (Papineni et al., 2002) **M2**. Hallucination Critic which we reapply to the refined response. For **M3**. we repurpose the FeQA measure (Durmus et al., 2020)—known to be an effective faithfulness measure in text summarization—by considering our document as the concatenation of \mathcal{D} and all \mathcal{G}_c^1 triples while the summary is the response \bar{x}_{n+1} .

Negative Candidates. We consider two different negative sampling strategies in order to compute \mathcal{L}_{NCE} : SANS (Ahrabian et al., 2020) and In-batch-negatives. SANS selects hard negatives by leveraging the graph structure and selecting negative samples from a context entity’s k -hop subgraph (e.g. \mathcal{G}_c^1). Meanwhile, In-batch-negatives considers the ground truth triple of each sample within a batch as a negative candidate for the other samples in the same batch. Using this approach, the number of candidates is equal to the batch size.

4.1 Main Experimental Questions

Our experiments answer the following questions:

- Q1) Identifying Hallucinations.** Can C identify both extrinsic and intrinsic hallucinations?
- Q2) Reducing Hallucinations.** Is NPH effective in reducing hallucinations?
- Q3) Query Generation.** Can NPH retrieve the correct entities and is \mathcal{L}_{NCE} important to learn query representations q_i ?
- Q4) Impact of MLM and Critic.** Is MLM essential to our training strategy or can we only use an autoregressive LM? Analogously, can we simply bypass the critic during refinement?
- Q5) Impact of global graph structure.** Is the global graph structure important for learning KG-Entity memory representations?

4.2 Results

Throughout our experiments, we rely on three representative baselines for response generation: GPT2-KG, AdapterBot (Lin et al., 2020), and GPT2-KE (Madotto et al., 2020). GPT2-KG is a small pre-trained GPT2 model (Radford et al., 2019) fine-tuned on the dialogue corpus. AdapterBot uses a fixed backbone conversational model such as DialGPT (Zhang et al., 2020) and encodes multiple dialogue skills via different adapters (Houlsby et al., 2019). Both GPT2-KG and AdapterBot process inputs by concatenating \mathcal{D} , \mathcal{K}_n and the generated response. GPT2-KE on the other hand uses a GPT2 model trained on a knowledge-augmented training set.

Q1: Identifying Hallucinations

Analogous to the study conducted in §2.1, we ask humans to identify the span of text that is hallucinated w.r.t. to the given triples in 500 responses generated greedily from GPT2-KG. We report the average Krippendorff’s alpha coefficient to be 0.73 on the annotator responses. Table 5 outlines our results. To explore the robustness of our corruption strategies as discussed in §3.2, we fine-tune a large RoBERTa model (Liu et al., 2019a) on three different synthetic datasets: (i) RoBERTa-Extrin corresponds to the negative examples crafted using an extrinsic hallucinations, where entity mentions are first extracted using the SpaCy NER tagger (Honnibal and Montani, 2017). (ii) RoBERTa-Intrin consists of negative examples that contain intrinsic hallucinations. (iii) Finally, RoBERTa-Intrin-Extrin

corresponds to examples that were either corrupted using an extrinsic or intrinsic strategy but not both simultaneously. For (i) and (ii), the examples are obtained by corrupting the full train OpenDialKG data. We observe that RoBERTa-Intrin-Extrin achieves the highest F1 (70.35%), compared to the classifiers trained on the first two synthetic datasets. Such a result highlights that our RoBERTa-Intrin-Extrin classifier can indeed detect both kinds of hallucinations and also that our corruption strategies are effective. In the rest of the experiments, we take RoBERTa-Intrin-Extrin as the hallucination classifier C .

Q2: Reducing Hallucinations

We evaluate the ability of NPH in fixing hallucination in generated responses in the three response generation baselines. We also perform ablation for each model using the different components of NPH. We present the results in Table 3 which show the degree of hallucination prior to and after applying NPH on each response generation method. We find that NPH consistently performs favourably in reducing hallucination across FeQA and the hallucination Critic. In particular, we observe that the strongest iteration of each baseline model is the original model paired with the full NPH module. For example, in AdapterBot, NPH decreases the Critic score by 8.17 points and increases faithfulness by 6.67 points on FeQA. With respect to BLEU scores, we observe inconsistent performance across the different baselines with AdapterBot+NPH incurring a marginally higher score. While we use BLEU as a proxy for faithfulness, it is still an imperfect measure as it is computed solely between the n-gram overlap between a reference and generated text which neglects the important fact that there is a multitude of different ways to generate a faithful response w.r.t. a KG.

Q3: Query Generation

We now investigate NPH’s ability to retrieve the correct entity using the crafted query. We present the results in Table 4 along with different ablation studies. We find that key metrics such as Hits@3 and Hits@10 are nearly saturated when using the complete NPH module with GPT2 embeddings for the KG-Entity memory. Furthermore, we notice that all retrieval metrics drop dramatically (e.g. \downarrow 70 Hits@1) when \mathcal{L}_{NCE} is omitted. Finally, we observe that SANS negatives lead to lower perplexity

	Model	Neg. candidates	PPL	Hits@1	Hits@3	Hits@10	MR	MRR
GPT2-Emb	NPH	SANS	8.56	0.73	0.92	0.99	1.76	0.83
		In-Batch Negatives	8.67	0.42	0.75	0.94	3.08	0.68
	NPH-w/o NCE	-	9.64	0.02	0.05	0.1	35.49	0.07
	NPH-w/o MLM	SANS	9.73	0.47	0.76	0.96	2.83	0.64
In-Batch Negatives		9.70	0.20	0.43	0.75	9.22	0.36	
CompGCN-Emb	NPH	SANS	8.99	0.13	0.26	0.52	14.27	0.25
		In-Batch Negatives	10.04	0.08	0.17	0.43	15.75	0.16
	NPH-w/o NCE	-	10.61	0.04	0.12	0.27	26.50	0.12
	NPH-w/o MLM	SANS	9.63	0.08	0.21	0.47	15.52	0.20
In-Batch Negatives		9.64	0.02	0.05	0.16	80.52	0.07	

Table 4: Ablation studies on NEURAL PATH HUNTER on the gold responses from the OpenDialKG test data.

Model	Precision	Recall	F1
RoBERTa-Intrin	44.9	32.54	37.73
RoBERTa-Extrin	68.65	46.94	55.76
RoBERTa-Intrin-Extrin	83.05*	61.02*	70.35*

Table 5: Performance of the hallucination critic on the 500 human-annotated data (* p -value < 0.001)

and better retrieval performance across the board. This is unsurprising since local negative samples are known to be harder and thus provides a richer learning signal (Ahrabian et al., 2020).

Q4: Impact of MLM and Critic

We now gauge the importance of using MLM and Critic within NPH. To assess the MLM component, we replace each contextual representation $m_i \in M_c$ with randomly initialized values. We highlight our findings in Table 3 where NPH-w/o MLM performs worse than NPH across all models. Investigating further in Table 4, we observe that performance without MLM degrades substantially (e.g. \downarrow 26 Hits@1) when using pre-trained GPT2 embeddings as entity memory and similarly for CompGCN embeddings. These findings suggest that MLM facilitates the learning of rich masked representations that are useful in downstream applications, a fact which is in line with other works that leverage MLM (Roberts et al., 2020; Devlin et al., 2019; Joshi et al., 2020). To judge the impact of the critic, we mask out all entity mentions as opposed to only masking out potential hallucinated ones during refinement. In Table 3, we find that NPH-w/o CRITIC performs the worst in every metric compared to all baselines which underlines that simply masking all entities—hallucinated or otherwise—in a response is not a productive strategy for effective refinement.

Q5: Impact of global graph structure

We now investigate the representation of entities in our KG-Entity Memory. We explore two vari-

Model	Hallucination	Fluency
GPT2-KG	97.5 \pm 0.6	92.5 \pm 1.6
GPT2-KG (+ NPH)	56.5 \pm 1.2	88.5 \pm 0.7
AdapterBot	95.5 \pm 0.8	90.5 \pm 0.4
AdapterBot (+ NPH)	59.0 \pm 0.5	87.5 \pm 1.2
GPT2+KE	97.0 \pm 0.2	91.5 \pm 0.7
GPT2+KE (+ NPH)	58.5 \pm 0.6	86.0 \pm 0.9

Table 6: Human Evaluation on 1200 responses (200 \times 6) from different response generation baselines.

ants: 1) Initializing embeddings as the output of a pre-trained GPT-2 model. 2) Utilizing node embeddings learned by a CompGCN network trained on a standard relation prediction task over the entire graph \mathcal{G} . In both these approaches, the embeddings are updated throughout training using Eq. 2. As per Table 4, we notice a dramatic difference in both perplexity and retrieval performance in favour of using simply the output of a pre-trained GPT-2 model. Such a result may be reconciled by noticing that any specific turn in dialogue local information (e.g. previous turn)—as conversation topics may drift—is significantly more important to generate a faithful response. Thus, enriching entity embeddings with global structure in \mathcal{G} is less beneficial than aligning \mathcal{G}_c^k with the representation space of the autoregressive LM, which for us is also GPT2.

4.3 Human Evaluation

In addition to the automated hallucination metrics, we conduct human evaluation to assess NPH’s ability to reduce hallucination. We provide human annotators with 200 hallucinated responses per baseline (§4.2) as identified by our hallucination critic §3.2. The faithfulness of each response is evaluated by 3 humans who are provided \mathcal{D} , \mathcal{K}_n , and the retrieved path from \mathcal{G}_c^k . We further request annotators to evaluate the fluency of the responses before and after refinement. Results are depicted in Table 6. We see that the hallucination critic achieves a precision of 97.5% for GPT2-KB responses, 95.5% for AdapterBot and 97.0% for

GPT2-KE. In contrast, generation methods when paired with NPH reduce hallucinations by a large margin 42.05% for GPT2-KB responses with a marginal drop in fluency (4.32%). We also observe similar performance gains for responses generated from AdapterBot and GPT2-KE.

5 Related Work

Knowledge Graphs. Building large-scale repositories of knowledge has been one of the principle directions of research in artificial intelligence since the inception of the field (Newell and Simon, 1956; Newell et al., 1959). Often represented as large scale multi-relational graphs, KGs have seen wide application in a variety of domains, such as question answering (Yao and Van Durme, 2014; Hao et al., 2017), and natural language processing (Berant et al., 2013; Yu and Dredze, 2014) to name a few. Beyond academic research, public KG’s like FreeBase (Bollacker et al., 2008) have been invaluable in industrial applications forming symbolic backbones of most important products (Singhal, 2012). KG’s have also risen in prominence in the context of dialogue models that propose to explicitly embed symbolic knowledge representations into a neural embedding space (Liu et al., 2019a; Zhu et al., 2017; Moon et al., 2019; Zhou et al., 2018; Xu et al., 2020). Liu et al. (2019b) use a knowledge retriever component that conditions the response by retrieving relevant facts from the KG based on the current utterance. Similarly, Young et al. (2018) and Zhou et al. (2018) use a commonsense KG to inject commonsense knowledge into the response of the conversational model. Tuan et al. (2019) explore the effects of using a dynamic KG in the dialogue model. On the other hand, Moon et al. (2019) propose a conversational reasoning model that traverses a large scale KG to retrieve a relevant path given a starting node and a classifier to predict the next node a response show follow. Unlike the KG path traversal problem, this work focuses on removing hallucinations in generated responses using a KG.

Hallucination. The injection of false information is a well-known phenomena in data-to-text generation (Tian et al., 2019; Dhingra et al., 2019; Parikh et al., 2020), machine translation (Koehn and Knowles, 2017; Lee et al., 2019), image captioning (Rohrbach et al., 2018), machine summarization (Maynez et al., 2020; Durmus et al., 2020) and question answering (Feng et al., 2018). In the

context of dialogue systems, Dušek et al. (2018, 2020) demonstrate that state-of-the-art natural language generation (NLG) models can hallucinate by missing important entities. Few NLG models have been proposed to cope with the issue, but are often custom-made for task-oriented dialogue (Balakrishnan et al., 2019). Recently, little progress has been made for studying hallucination in open-domain dialog systems. Dziri et al. (2021) study hallucination in knowledge-grounded dialogue systems and introduce a the BEGIN benchmark for measuring groundedness in dialogue systems. Finally, Rashkin et al. (2021) propose a dialogue system that is more faithful to the source knowledge by adding control tokens at training time that guide the model towards generating more objective sentences which have higher overlap with the source.

6 Conclusion

In this work, we investigate the open problem of hallucination in KG-grounded dialogue systems and demonstrate that these models are more susceptible to extrinsic hallucinations which predominantly manifest as the injection of erroneous entities. To tackle this challenging problem, we propose a new module NEURAL PATH HUNTER that aims to enforce faithfulness in KG-grounded dialogue systems by identifying and refining hallucinations via queries over a k-hop subgraph. We empirically observe that NPH is capable of reducing hallucination when paired with a number of base dialogue models with relative improvements of 20.35% over vanilla GPT2 on FeQA. Our findings also reveal the crucial role the representation of the local subgraph plays as external memory compared to the full global graph. In this work, we considered a paired KG aligned with dialogue but in many other applications, such dialogue to KG alignment may be difficult to easily obtain necessitating the usage of the full graph which is interesting direction for future work.

Acknowledgements

We are grateful to the anonymous reviewers for helpful comments. This research is supported by Natural Sciences and Engineering research Council of Canada, the Alberta Machine Intelligence Institute Fellow Program and the Canadian Institute for Advanced Research AI Chair Program. This research is also supported in part by Compute Canada.

References

- Kian Ahrabian, Aarash Feizi, Yasmin Salehi, William L. Hamilton, and Avishek Joey Bose. 2020. [Structure aware negative sampling in knowledge graphs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6093–6101, Online. Association for Computational Linguistics.
- Anusha Balakrishnan, Jinfeng Rao, Kartikeya Upasani, Michael White, and Rajen Subba. 2019. [Constrained decoding for neural NLG from compositional representations in task-oriented dialogue](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 831–844, Florence, Italy. Association for Computational Linguistics.
- Hannah Bast, Florian Baurle, Björn Buchhold, and Elmar Haußmann. 2014. [Easy access to the freebase dataset](#). In *Proceedings of the 23rd International Conference on World Wide Web*, pages 95–98.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuvan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics.
- Esin Durmus, He He, and Mona Diab. 2020. [FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5055–5070, Online. Association for Computational Linguistics.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2018. [Findings of the E2E NLG challenge](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 322–328, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2020. [Evaluating the state-of-the-art of end-to-end natural language generation: The e2e nlg challenge](#). *Computer Speech & Language*, 59:123–156.
- Nouha Dziri, Ehsan Kamaloo, Kory Mathewson, and Osmar Zaiane. 2019. [Augmenting neural response generation with context-aware topical attention](#). In *Proceedings of the First Workshop on NLP for Conversational AI*, pages 18–31, Florence, Italy. Association for Computational Linguistics.
- Nouha Dziri, Hannah Rashkin, Tal Linzen, and David Reitter. 2021. [Evaluating groundedness in dialogue systems: The begin benchmark](#). *ArXiv preprint*, abs/2105.00071.
- et al. Falcon, WA. 2019. [Pytorch lightning](#). *GitHub*. Note: <https://github.com/PyTorchLightning/pytorch-lightning>, 3.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. [Pathologies of neural models make interpretations difficult](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3719–3728, Brussels, Belgium. Association for Computational Linguistics.
- Michael Gutmann and Aapo Hyvärinen. 2010. [Noise-contrastive estimation: A new estimation principle for unnormalized statistical models](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304.
- Yanchao Hao, Yuanzhe Zhang, Kang Liu, Shizhu He, Zhanyi Liu, Hua Wu, and Jun Zhao. 2017. [An end-to-end model for question answering over knowledge base with cross-attention combining global knowledge](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 221–231, Vancouver, Canada. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Matthew Honnibal and Ines Montani. 2017. [spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing](#). *To appear*, 7(1).
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly.

2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Daniel Jurafsky and James H Martin. 2018. Speech and language processing (draft). *Chapter A: Hidden Markov Models (Draft of September 11, 2018)*. Retrieved March, 19:2019.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. [Ctrl: A conditional transformer language model for controllable generation](#). *ArXiv preprint*, abs/1909.05858.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. 2019. Hallucinations in neural machine translation. In *International Conference on Learning Representations*.
- Zhaojiang Lin, Andrea Madotto, Yejin Bang, and Pascale Fung. 2020. [The adapter-bot: All-in-one controllable conversational model](#). *ArXiv preprint*, abs/2008.12579.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv preprint*, abs/1907.11692.
- Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019b. [Knowledge aware conversation generation with explainable reasoning over augmented graphs](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Andrea Madotto, Samuel Cahyawijaya, Genta Indra Winata, Yan Xu, Zihan Liu, Zhaojiang Lin, and Pascale Fung. 2020. [Learning knowledge bases with parameters for task-oriented dialogue systems](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2372–2394, Online. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. [OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy. Association for Computational Linguistics.
- Allen Newell, John C Shaw, and Herbert A Simon. 1959. Report on a general problem solving program. In *IFIP congress*, volume 256, page 64. Pittsburgh, PA.
- Allen Newell and Herbert Simon. 1956. The logic theory machine—a complex information processing system. *IRE Transactions on information theory*, 2(3):61–79.
- 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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ankur Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. [ToTTo: A controlled table-to-text generation dataset](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1173–1186, Online. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*

- Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*.
- Hannah Rashkin, David Reitter, Gaurav Singh Tomar, and Dipanjan Das. 2021. [Increasing faithfulness in knowledge-grounded dialogue with controllable features](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 704–718, Online. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. [Object hallucination in image captioning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4035–4045, Brussels, Belgium. Association for Computational Linguistics.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784. AAAI Press.
- Amit Singhal. 2012. Introducing the knowledge graph: things, not strings. *Official google blog*, 5.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205, Denver, Colorado. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Wilson L Taylor. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism quarterly*, 30(4):415–433.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P Parikh. 2019. [Sticking to the facts: Confident decoding for faithful data-to-text generation](#). *ArXiv preprint*, abs/1910.08684.
- Yi-Lin Tuan, Yun-Nung Chen, and Hung-yi Lee. 2019. [DyKgChat: Benchmarking dialogue generation grounding on dynamic knowledge graphs](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1855–1865, Hong Kong, China. Association for Computational Linguistics.
- Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. 2020. [Composition-based multi-relational graph convolutional networks](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. [Asking and answering questions to evaluate the factual consistency of summaries](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. [Knowledge graph embedding by translating on hyperplanes](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 1112–1119. AAAI Press.
- Ronald J Williams and David Zipser. 1989. Experimental analysis of the real-time recurrent learning algorithm. *Connection science*, 1(1):87–111.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2263, Copenhagen, Denmark. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Hu Xu, Seungwhan Moon, Honglei Liu, Bing Liu, Pararth Shah, Bing Liu, and Philip Yu. 2020. [User memory reasoning for conversational recommendation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5288–

5308, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Xuchen Yao and Benjamin Van Durme. 2014. [Information extraction over structured data: Question answering with Freebase](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 956–966, Baltimore, Maryland. Association for Computational Linguistics.

Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. [Augmenting end-to-end dialogue systems with commonsense knowledge](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4970–4977. AAAI Press.

Mo Yu and Mark Dredze. 2014. [Improving lexical embeddings with semantic knowledge](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 545–550, Baltimore, Maryland. Association for Computational Linguistics.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.

Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. [Commonsense knowledge aware conversation generation with graph attention](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4623–4629. ijcai.org.

Wenya Zhu, Kaixiang Mo, Yu Zhang, Zhangbin Zhu, Xuezheng Peng, and Qiang Yang. 2017. Flexible end-to-end dialogue system for knowledge grounded conversation. *arXiv*, pages arXiv–1709.

A OpenDialKG

We use OpenDialKG (Moon et al., 2019), a crowd-sourced English dialogue dataset where two workers are paired together to chat about a certain topic. The first speaker is asked to initiate the conversation about a given entity and the second speaker is tasked to form a factual response based on a set of facts extracted from an existing KG, Freebase (Bast et al., 2014). Those facts represent paths in the KG that are either 1-hop or 2-hop from the initial entity. Once the second speaker sends a response, the first speaker continues discussing the topic engagingly and new multi-hop facts from the KG are presented to the second speaker. The conversation can be regarded as traversing multiple paths in the KG. However, not all utterances within the same dialogue are grounded on facts from the KG. The second speaker can choose not to select a path from the KG to form an answer and instead forms a “chit-chat” response. Overall, the dataset consists of four domains: movie, music, sport and book where each second speaker’s utterance is annotated with paths from the KG. The KG corresponds to a large subgraph extracted from Freebase with $\sim 1.2\text{M}$ triples (subject, predicate, object), $\sim 101\text{k}$ distinct entities and 1357 distinct relations. No official split is provided in the original dataset, and thus we randomly split the dataset in 80/10/10 for the train/valid/test, respectively. The data consists of 61778 train, 7933 valid and 7719 test. Some utterances in the dataset are chit-chat and thus are not annotated with a path from the KG. Thus, we filter the dataset by keeping only the dialogue examples that are annotated with a path from the KG. We ended up with 23314 training examples, 2954 valid examples and 2954 test examples.

B Human Evaluation for Modes of Hallucination

Workers, fluent in English, were trained for the task before starting the evaluation process. If the workers fail to achieve at least 80% accuracy in answering the different test questions, they would not be allowed to start the evaluation process. These workers are hired from Appen⁴. Each worker was presented with a dialogue history, knowledge triples including the gold triples and the 1-hop paths from the centre node in \mathcal{G}_c^k . Each example was evaluated by 3 workers and majority vote was considered.

⁴<https://appen.com/>

Workers were asked the following questions:

1. Is this response hallucinated with respect to the gold knowledge triples? (Most definitely, Not at all)
 - (a) If the response is hallucinated, does it represent extrinsic hallucination, intrinsic hallucination or both? (Extrinsic, Intrinsic, Both)
2. If the response is not hallucinated, is it faithful to the source or generic? (Faithful, Generic)
3. Is this a coherent response with respect to the dialogue history even if it was identified as hallucinated? (Most definitely, Not at all)
4. Is this response grammatically correct? (Most definitely, Not at all)

C KG-Entity Memory

GPT2 embeddings OpenDialKG contains a textual description, called “render”, for triples extracted from the KG. Note that not all triples in the dataset are associated with “render”. To get a contextual representation for each entity mention, we feed “render” to GPT2 and then extract hidden states representations for each entity’s word piece and finally obtain a final representation by applying a MaxPool over the hidden representations. For entity mentions that are not described in “render”, we get their representations directly from the last hidden states in GPT2.

D Implementation Details

NPH: NPH is implemented using the Pytorch Huggingface Transformers library (Wolf et al., 2020) and the Pytorch-lightning library (Falcon, 2019). Concretely, we use a small RoBERTa model (Liu et al., 2019a) as the MLM and the base GPT2 model (Radford et al., 2019) as our autoregressive LM. During training, we use the Adam optimizer (Kingma and Ba, 2015) with Dropout (Srivastava et al., 2014) on a batch size of 16 with a learning rate of 6.25×10^{-5} that is linearly decayed. The maximum dialogue history length is set to 3 utterances. The coefficient λ in Eq. 2 is set to 0.5. We varied the factor from 0.1 to 1 and 0.5 was chosen based on the best results on the validation set. The number of negative examples is set to 50 for SANS. The model early-stops at epoch 10 and we save the best model based on the validation set. Our

GPT2	Hallucination			Faithfulness	Generic	Coherence	Fluency
	Ex	In	B				
Greedy	17.66 ± 2.6	2.00 ± 3.5	1.66 ± 0.5	69.00 ± 3.2	9.66 ± 2.7	81.66 ± 3.2	95.67 ± 1.6
Beam Search	18.33 ± 2.8	3.33 ± 3.8	4.00 ± 1.8	68.00 ± 3.9	6.33 ± 2.7	83.33 ± 1.6	97.00 ± 1.9
Nucleus 0.9	25.33 ± 2.1	4.00 ± 3.6	2.33 ± 3.6	64.66 ± 2.3	3.66 ± 3.2	83.66 ± 2.4	99.10 ± 0.6
Nucleus 0.5	23.33 ± 2.2	5.33 ± 3.1	4.33 ± 0.8	59.90 ± 2.5	7.00 ± 2.6	87.66 ± 2.1	98.34 ± 0.4
Top20	28.33 ± 1.5	7.00 ± 2.6	5.00 ± 1.5	55.00 ± 0.6	4.66 ± 1.8	80.33 ± 1.6	97.34 ± 0.5

Table 7: Human assessment of random 1500 GPT2 dialogue responses (300×5) generated based on the test OpenDialogk data (Moon et al., 2019)(mean preferences $\pm 90\%$ confidence intervals).

hyperparameters search is done via greed search. The average runtime of this model is 4 hours.

GPT2-KG: Similarly, we implement this baseline using the Pytorch Huggingface Transformers library (Wolf et al., 2020) and the Pytorch-lightning library (Falcon, 2019). During training, we use the Adam optimizer (Kingma and Ba, 2015) with Dropout (Srivastava et al., 2014) on a batch size of 32 with a learning rate of 6.25×10^{-5} that is linearly decayed. The maximum dialogue history length is set to 3 utterances. The model early-stops at epoch 6. The average runtime of this model is 2 hours.

AdapterBot and GPT2-KE: We use the code that’s publicly available by the authors at <https://github.com/HLTCHKUST/adapterbot> and <https://github.com/HLTCHKUST/ke-dialogue> and we follow closely their training procedure described in (Lin et al., 2020) and (Madotto et al., 2020). We use the GPT2-KE with 9K iterations. The average runtime of these models is 3 hours.

Training for all models, including baselines, is done on an Nvidia V100 GPU 32GB and for inference, we use greedy search.

Hallucination Critic: We use a pre-trained RoBERTa-large classifier (Liu et al., 2019a) provided by the Huggingface Transformers library (Wolf et al., 2020). The model was trained using the Adam optimizer with a learning rate of 2×10^{-5} for 5 epochs on one Nvidia V100 GPU 32GB. The average runtime of this model is 2 hours.

E Hallucination Metrics

Although BLEU measures the extent to which the generated response is similar to the reference faithful response, it can be misleading in the case where the generated response is very distant from the

ground-truth response but faithful to the knowledge triples. We consider 2 other metrics that focus on measuring the degree of hallucination in the generated responses:

Hallucination Critic We use our trained token-level hallucination critic as a sentence-level hallucination detector. We consider the utterance as hallucinated if at least one token was identified as hallucinated. As input, the critic receives the dialogue history, the gold triples and the generated response and the output is a binary label indicating hallucination or not. To use this critic for the output of NPH, we augment the gold triples with the path extracted based on the Entity Mention Retriever.

FeQA Durmus et al. (2020) has been shown successful in measuring faithfulness in the text summarization task. It generates questions from the candidate summaries and then answers them against the input documents. It measures the average F1 score against the gold answers from the document. Through asking and answering questions, FeQA measures the semantic correctness of the generated responses. To adapt FeQA to our dialogue task, we flatten each path into a pseudo sentence by joining the $\langle [SBJ], [PRE], [OBJ] \rangle$ with a simple space, e.g., [Crescendo, written by, Becca fitzpatrick] \rightarrow “Crescendo written by Becca Fitzpatrick”. We consider our document as the concatenation of \mathcal{D} and all \mathcal{G}_c^1 triples and the candidate summary as the generated/refined response. FeQA takes a given generated grounded response as input, and generates questions. It then employs a QA system to answer the generated questions based on the knowledge the response was grounded in.

We use the code made publicly available by the authors⁵. A similar work to FeQA is QAGS (Wang et al., 2020) which corresponds to asking and answering questions to evaluate the factual consist-

⁵<https://github.com/esdurmus/feqa>

tency of summaries.

F Human Evaluation of NPH responses

Analogous to evaluating modes of hallucination, we solicit human evaluation from Appen⁶ where we train English-speaking annotators for the task before starting the evaluation process. To evaluate the responses generated by our response generation baselines, annotators were presented with \mathcal{D} , \mathcal{K}_n and the generated response. And, to evaluate NPH’s responses, annotators were presented with \mathcal{D} , \mathcal{K}_n , the retrieved path from \mathcal{G}_c^k and the refined response. Humans were asked to answer the following questions:

1. Is this response hallucinated with respect to \mathcal{K}_n ? (Most definitely, Not at all)
2. Is this a fluent response, i.e., a response that’s grammatically correct? (Most definitely, Not at all)

In total, humans evaluated 1200 responses: 600 responses (200 from each response generation baseline before refinement) and 600 responses after refinement.

G Error Analysis

To gain insight into the potential shortcomings of NEURAL PATH HUNTER, we conduct an error analysis on refined responses that still contain undesirable hallucinations. Examples of failed refinements using NPH are listed below. Recall that for effective retrieval NPH requires oracle access to \mathcal{G}_c^k which pre-supposes the existence of the correct entity in the subgraph. However, based on the examples below, we observe that many of the failed retrievals correspond to entities that might exist in \mathcal{G} but are critically not supported within \mathcal{G}_c^k . To highlight this point, let us consider the following example:

Previous turn: Could you recommend a book similar to *Thirteen Reasons Why*?

Gold triple:[[‘Thirteen Reasons Why’, ‘has genre’, ‘Young-adult fiction’]]

GPT2-KB Response: Sure, there is a book called **The Sea of Monsters** by **John Green**.

Critic: John Green.

Context nodes: Thirteen Reasons Why, Young-adult fiction

Top-1 Path:[[Thirteen Reasons Why, written by, Jay Asher]]

NPH response: Sure, there is a book called *The Sea of Monsters* written by **Jay Asher**.

In this example, we notice that the original response contains hallucination as “John Green” did not write the book titled “The Sea of Monsters”. Moreover, to fix this response NPH has access to \mathcal{G}_c^1 which contain the context nodes “Thirteen Reasons Why” and “Young-adult fiction” but not the entity of interest “The Sea of Monster”. As a result the proposed fix using NPH is the wrong entity “Jay Asher”. It is important to note that the perceived mistakes by NPH is in fact a limitation of our assumption that the correct entities not only exist in \mathcal{G} but can be extracted by running a query on the extracted k -hop subgraph. If this fails to be the case in practice, then we cannot hope for NPH to retrieve the correct entity. We hypothesize that extending NPH to reason over multiple subgraphs or even \mathcal{G} in an efficient manner will serve to alleviate this failure mode.

In addition to using the incorrect subgraph for retrieval, another source of errors is caused by the failure to identify the hallucinated entities to fix in the first place. This can occur in practice as the Critic C is not perfect and may occasionally miss hallucinated entities and thus propagate the error forward. The example below illustrates this failure mode:

Previous turn: Yes, I watched them. Could you tell me more about *The Heart Is Deceitful Above All Things*, please.

Gold triple:[[The Heart Is Deceitful Above All Things, starred actors, Jimmy Bennett]]

GPT2-KB Response: Sure, it was written by **Tony Kaye** and stars **Scarlett Johansson**.

Critic: Scarlett Johansson.

Context nodes: The Heart Is Deceitful Above All Things, Jimmy Bennett

Top-1 Path:[[The Heart Is Deceitful Above All Things, starred actors, Jimmy Bennett]]

NPH response: Sure, it was written by **Tony Kaye** and stars **Jimmy Bennett**.

We believe a stronger Critic or other mechanisms to identify token-level hallucinations can improve the downstream performance of NPH.

The final source for error—while not frequent—is the fluency of the refined response. We find that

⁶<https://appen.com/>

in an effort to remove hallucinations, other undesirable properties such as a lack of grammaticality may be introduced. The example below demonstrates such a failure mode. This error can be mitigated if we associate each node in the graph by an entity type.

Previous turn: Sure. Isn't he married to Coleen Rooney

Gold triple:[[United Kingdom, Country of nationality, Wayne Rooney]

GPT2-KB Response: No, Yes, he is married to **Caroline Wozniacki**.

Critic: Scarlett Johansson.

Context nodes: United Kingdom, Wayne Rooney

Top-1 Path:[[United Kingdom, Country of nationality, Wayne Rooney]

NPH response: Yes, he is married to **United Kingdom**.

H Sample Refined Responses

Some examples of the refined responses are presented in Table 8.

History	<i>A:</i> I love Shakespeare. Romeo and Juliet is my favorite. <i>B:</i> I like that one as well. The play was written by Kylie Scott. <i>A:</i> She also wrote Deep, did she not?
GPT2-KG	B_{gen} Yes she did. She also wrote The Tragedy of King Richard the Third .
Gold knowledge	T_1 [Deep, written by, Kylie Scot]
Top-5 Paths	T'_1 : [Play, written by, Kylie Scott] T'_2 : [Dirty, written by, Kylie Scott] T'_3 : [Lick, written by, Kylie Scott] T'_4 : [Deep, written by, Kylie Scott] T'_5 : [Trust, written by, Kylie Scott]
NPH response	B_{fix} Yes she did. She also wrote Play .
History	<i>A:</i> Do you know of a movie directed by Jean Pierre Jeunet?
GPT2-KG	B_{gen} Jean-Pierre jeunet directed The Bourne Ultimatum .
Gold knowledge	T_1 : [Micmacs, written by, Jean Pierre Jeunet] T_2 : [Micmacs, has genre, French]
Top-5 Paths	T'_1 : [Alien: Resurrection, directed by, Jean Pierre Jeunet] T'_2 : [A Very Long Engagement, directed by, Jean Pierre Jeunet] T'_3 : [Amélie, directed by, Jean Pierre Jeunet] T'_4 : [The Extraordinary Voyage, starred actors, Jean Pierre Jeunet] T'_5 : [Micmacs, written by, Jean Pierre Jeunet]
NPH response	B_{fix} Jean-Pierre jeunet directed Alien: Resurrection .

Table 8: Selected responses based on GPT2-KG test responses before and after applying NEURAL PATH HUNTER. The span of texts highlighted in red indicate the hallucinated entity mentions whereas the ones highlighted in green indicate the retrieved correct entity mentions.