



Heriot-Watt University
Research Gateway

OTTERS

Citation for published version:

Sevegnani, K, Howcroft, DM, Konstas, I & Rieser, V 2021, OTTERS: One-turn topic transitions for open-domain dialogue. in *59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference*. vol. 1, Association for Computational Linguistics, pp. 2492-2504, Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Virtual, Online, 1/08/21. <https://doi.org/10.18653/v1/2021.acl-long.194>

Digital Object Identifier (DOI):

[10.18653/v1/2021.acl-long.194](https://doi.org/10.18653/v1/2021.acl-long.194)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference

Publisher Rights Statement:

©2021 Association for Computational Linguistics

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

OTTERS: One-turn Topic Transitions for Open-Domain Dialogue

Karin Sevegnani, David M. Howcroft, Ioannis Konstas, Verena Rieser

The Interaction Lab, MACS Heriot-Watt University

Edinburgh, Scotland, UK

{karin.sevegnani, i.konstas, v.t.rieser}@hw.ac.uk

Abstract

Mixed initiative in open-domain dialogue requires a system to pro-actively introduce new topics. The one-turn topic transition task explores how a system connects two topics in a cooperative and coherent manner. The goal of the task is to generate a “bridging” utterance connecting the new topic to the topic of the previous conversation turn. We are especially interested in commonsense explanations of how a new topic relates to what has been mentioned before. We first collect a new dataset of human one-turn topic transitions, which we call OTTERS¹. We then explore different strategies used by humans when asked to complete such a task, and notice that the use of a bridging utterance to connect the two topics is the approach used the most. We finally show how existing state-of-the-art text generation models can be adapted to this task and examine the performance of these baselines on different splits of the OTTERS data.

1 Introduction

For a conversation to be truly engaging, we typically assume that both participants take initiative, e.g. by introducing a new topic. We call this a *mixed-initiative* dialogue. Open-domain systems trained on vast amounts of data (Jiang et al., 2020; Zhang et al., 2020; Gao et al., 2018; Li et al., 2017, 2016; Vinyals and Le, 2015), however, are often purely responsive, make abrupt transitions, or fail to take initiative (see examples in Table 1). In this paper, we consider the case where the system pro-actively introduces a new topic in a conversation by providing a *commonsense link* of how this new topic relates to what was mentioned previously (see Fig. 1). We call this transition strategy “bridging”. Humans deploy a range of strategies

User A	Source Topic: I spend a lot of time outside .
User B	Transition: I like the outdoors as well, especially <u>gardening</u> . It destresses me.
	Target Topic: I enjoy relaxing and getting flowers .
	Entity Path: <u>outside</u> - <u>garden</u> - <u>flower</u>
User A	Source Topic: I like seafood a lot.
User B	Transition: Since you like <u>seafood</u> , is <u>Swedish fish</u> a candy that you might enjoy?
	Target Topic: I have no self control when it comes to candy .
	Entity Path: <u>seafood</u> - <u>Swedish fish</u> - <u>candy</u>
User A	Source Topic: I think I am getting engaged soon.
User B	Transition: I have two children from a previous marriage
	Target Topic: My children are my life.
	Entity Path: <u>engagement</u> - <u>marriage</u> - <u>child</u>

Figure 1: Example topic transitions from OTTERS. User A introduces a topic with a short sentence (main concept in **bold**). Then User B responds with a (optionally multi-sentence) “bridging” transition before introducing the new topic (the main concepts for the transition and target topic are denoted with underline and **bold**, respectively). Each example is accompanied by an entity path, comprising Knowledge Graph entities (denoted with `teletype`) instantiating the main concepts of the dialogue turn.

in addition to bridging, including disengagement, discourse markers or silence (Riou, 2015). We hypothesise that introducing a new topic by making a connection with the previous dialogue turn can be perceived as a less abrupt transition.

More specifically, we investigate bridging transitions between two user utterances in the form of one or more sentences that contain at least one main *linking* concept. These inherently can allow for better grounding to external resources such as entities

¹<https://github.com/karinseve/OTTERS>

in large Knowledge Graphs (KG) (e.g., Wikidata), or named entities mentioned in documents (e.g., Wikipedia, or news articles), ultimately leading to more controlled and interpretable outputs.

To this end, we crowdsource a corpus of human-written topic transitions focused on these “bridging” strategies, where humans introduce a “missing link” concept, given a source and target topic in the form of two short user utterances (Fig. 1). By grounding the topics on a KG using automatically recognised entities associated with each topic, we can then identify “commonsense” connections which are similar to these missing links.

By modelling such topic transitions in the form of Cause-Effect relationships in a KG, we can then perform abductive inference on commonsense knowledge for which we provide a language generation baseline. In particular, we fine-tune a multi-hop reasoning model (Ji et al., 2020) which was trained on a similar task called Abductive NLG (α NLG) to generate an explanatory hypothesis given two observations. We find that combining a reasoning module over a KG (ConceptNet) with a language model achieves the best performance on our “topic transition” task for both the predicted entity path as well as the generated utterance. In addition, we show that existing multi-topic dialogue datasets, such as PersonaChat (Zhang et al., 2018) and TopicalChat (Gopalakrishnan et al., 2019), cannot be easily adapted to this task, due to the different nature of the tasks they were designed for. Our contributions are as follows:

- We propose a new Natural Language Generation task based on one-turn topic transitions for open-domain dialogue based on a “bridging” strategy, which promotes grounding on KG entities.
- We collect a crowdsourced dataset, OTTers, and present a rigorous analysis in terms of transition strategies, linguistic properties and entity linking to a KG.
- We show that our KG-grounded dataset can effectively leverage the reasoning component of an existing Transformer-based model (Ji et al., 2020) to generate better output compared to a vanilla GPT-2 (Radford et al., 2019) decoder, both in in-domain and out-of-domain data splits.

2 Related Work

Topic Transitions in the Linguistic Literature.

There is no common definition for the term *topic* (Goutsos, 1997; Purver et al., 2011); however, there are a number of definitions which are helpful for our purposes. Goutsos (1997) divide a “topic” into two main components: 1) what constitutes a topic (the “what”) and 2) how participants perceive and manage a topic (the “how”). An early work from Brown and Yule (1983) declares that “topics should be described as the most frequently used, unexplained term in the analysis of discourse”. In general, “discourse topics” can be explained as what a portion of the interaction is about, therefore the “aboutness” (Berthoud and Mondada, 1995; Porhiel, 2005). More specifically Chafe (1994) defines the notion of topic as “the totality of information that is semiactive at one time”.

Prior work has shown that the introduction of a new topic usually co-occurs with cues such as wrapping things up about the current topic (Maynard, 1980), preceding silence, or the use of discourse markers (Riou, 2015). Also, backchannel signals, e.g., *yeah, right, you know*, indicate that both agents are involved in the interaction and show consent for the topic development (James, 1995). Beyond these overt cues, James (1995) and Geluykens (1993) describe semantic topic transitions: “each topic has a tendency to lead to the next; to provide the opening for another” (James, 1995), and topics are typically “co-constructed”, requiring each speaker to contribute to the conversation for further progression and development (Geluykens, 1993). The identification of topic transition is indeed not an easy task. It is not only about linguistic cues such as discourse markers and prosodic cues, as sometimes a topic switch can be identified with the introduction of a new entity (James, 1995). Additionally, in a conversation topics are created and introduced by participants themselves in real time, making topics participant- and interaction-specific (Mondada, 2001, 2003). Moreover, “the entities in focus at a given point in the discourse will be that partially-ordered subset of activated entities which are likely to be continued as topics of subsequent utterances” (Gundel et al., 1993). These cooperative elements emphasise the importance of mixed-initiative topic management for open-domain dialogue systems.

PersonaChat
User A: I do not like carrots. I throw them away.
User B: Really. But, I can sing pitch perfect.
User A: I also cook, and I ride my bike to work.
User B: Great! I had won an award for spelling bee.
TopicalChat
User A: Yeah and saltwater fish are lucky because they can do that and drink through their mouths.
User B: Seems like fresh water fish got the short end of the stick with that one. Have you ever been to a cat cafe?

Table 1: Examples of abrupt topic transitions from the PersonaChat and TopicalChat datasets.

Current Multi-topic Open-domain Systems.

Previous work in open-domain dialogue systems has largely avoided explicitly modelling topic transitions and instead focused on grounding system behaviour in a “persona” (a set of statements about hobbies, demographics, or preferences) (Zhang et al., 2018; Li et al., 2016) or by conditioning conversations on knowledge sources such as newspaper articles, fun facts or Wikipedia articles (Gopalakrishnan et al., 2019; Dinan et al., 2019) to generate engaging responses while avoiding generic replies, improving coherence, and raising new and interesting topics. These approaches often lead to poor topic transitions, as illustrated in Table 1. The PersonaChat example shows neither initiative nor common sense while transitioning to a new topic; it only displays passive acknowledgement from User B. Whereas the TopicalChat example presents a very abrupt topic shift by User B. Our dataset is the first corpus focused specifically on one-turn topic transitions; however, there are several human-to-human dialogue corpora wherein participants discuss assigned topics. Two prominent such corpora are TopicalChat (Gopalakrishnan et al., 2019) and PersonaChat (Zhang et al., 2018).

In TopicalChat both participants used source documents from Wikipedia to discuss a shared topic. The dialogues in this corpus tend to flow less naturally than those in PersonaChat with participants generally focusing on expressing the main facts, often by copy and pasting from their source documents rather than having a natural conversation. Therefore we focus on PersonaChat as a point of comparison.

PersonaChat dialogues consist of chit-chat conversations based on a set of “persona traits” assigned to each participant. Because participants seek to express their persona to each other, the con-

versations require mentioning various topics (i.e. their persona traits) in a natural way. Indeed, Zhang et al. (2018, Sec. 3.3) adjusted their design to encourage users to engage with each other’s topics and not simply state their own topics as quickly as possible to end the dialogue. PersonaChat does not contain annotations for the topic of each turn and participants had the freedom to mention their topics (i.e. persona traits) in any order.

We use PersonaChat in two different ways: 1) using their persona traits as starting and goal topics for our own data collection, and 2) as a point of comparison for our dataset.

Commonsense-Aware Neural Text Generation.

Large Language Models still suffer in cases where reasoning over underlying commonsense knowledge is required during generation, including dialogue generation (Zhou et al., 2018), story ending generation (Guan et al., 2019), and topic-to-essay generation (Yang et al., 2019). Recently, Guan et al. (2019); Bhagavatula et al. (2020) attempted to integrate external commonsense knowledge into generative pretrained language models, which we will also attempt in Section 4 using the Abductive NLG (α NLG) dataset (Bhagavatula et al., 2020). Our setup is similar in spirit to α NLG, which is a conditional generation task for explanations given observations in natural language. In particular, the model has to generate an explanatory hypothesis given two observations: the cause (e.g. *The Smith family went on a cruise for their summer vacation*) and the consequence (e.g. *From then on, the Smiths went to the beach each summer instead*). Here, a possible explanation might be: *The Smith family got seasick on the cruise*. The α NLG dataset contains 20k pairs observations and 200k explicative hypotheses, which we will later use for fine-tuning our models (see Section 4).

3 One-turn Topic Transitions

3.1 Task Design and Data Collection

Task Description. We assume there are topics t_a and t_b for utterances u_a and u_b (with $u. = t.$ for this paper). The goal of the task is to generate a *one-turn transition* utterance u_t to serve as a smooth link between t_a and t_b so that its concatenation with utterance u_b is a sensible response to u_a . A *bridging* transition occurs when one or more of the entities $e_t \in \mathbf{e}_t$ mentioned in u_t lies on a path in the knowledge graph between entities $e_a \in \mathbf{e}_a$ and

$e_b \in \mathbf{e}_b$ mentioned in u_a and u_b , respectively.

Knowledge Graph Construction. We use PersonaChat persona traits as the starting point for our data collection. In order to model commonsense connections, we built a knowledge graph (KG) using the entities found in each persona trait through the Yahoo Entity Linker (Blanco et al., 2015; Pappu et al., 2017). Each entity is linked to its correspondent Wikidata identifier, while a SPARQL query retrieved the entity’s super-classes and sub-classes, which were added to the KG. Furthermore, the KG has been augmented by retrieving the commonsense connections for each entity from ConceptNet (Speer et al., 2017) and by parsing Wikipedia abstracts mentions.

To select which traits to use for the data collection, we first selected all pairs of entities connected with k -hops ($1 < k < 20$) in the KG. Then, we recovered the entities mentions in the persona traits and saved every pair (nearly 30k) as potential pairs for our data collection.

Data Collection. We crowdsourced the data collection for OTTers on Amazon Mechanical Turk (AMT). Each user was provided with two topics **A**, **B** from the PersonaChat persona traits, along with instructions explaining the task. The instructions ask the user to imagine they are having a conversation where the first topic **A** from the pair represents the last turn of the other person, and the second topic **B** contains the final topic the user wants to talk about. The user then has to write a short utterance to transition to the new topic **B** in the least abrupt way possible. Additionally, in order to encourage crowd-workers to ground their utterances in actual topics, we asked them to report the “topics” mentioned in their sentence (see Figure 2).

For each topic pair in the study we collected three different transition utterances to provide more insight into the different strategies users adopt when transitioning to a new topic.

3.2 Corpus Properties

Basic Statistics. Table 2 provides summary statistics describing OTTers. Our corpus consists of 4,316 utterances for 1,421 unique topic pairs, with an average utterance length of 1.3 sentences and 16.4 words. The KG path statistics for OTTers are based on all of the paths found by the Yahoo Entity Linker between the 1421 unique topic pairs in the corpus, a total of just over $12k$ paths.

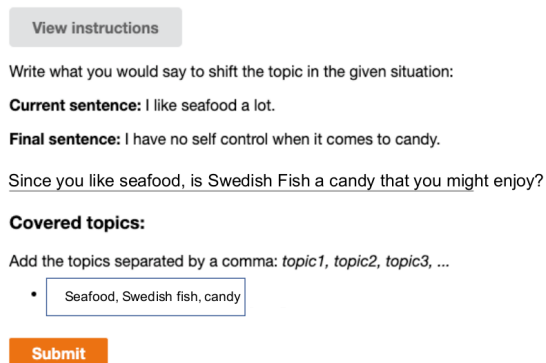


Figure 2: Screenshot of part of the interface users are presented with when accepting the study on AMT.

Property	OTTers
TTR	5085 / 70935
ent. TTR	3984 / 17054
turn length	2-68
—: mean	16.4
—: mode	9
sent.s/turn	1-7
—: mean	1.3
—: mode	1
KG path	1-14
—: mean	6.1
—: mode	5

Table 2: Properties of OTTers. *Type-token ratio* (TTR) for the different splits of the dataset. *Entity TTR* refers to the number of (unique) entities appearing in that portion of the dataset. *Turn length* is the length of the transition utterance in tokens, including punctuation. *Number of sentences per turn* measured by splitting on sentence-final punctuation.

KG coverage. We calculated the distance between each pair of topics in the knowledge graph described in Sec. 3.1 to facilitate analyses of the role of topic distance in transition strategy and transition quality. To extract entities from the utterances in our corpus, we extended the tagger built-in to the Yahoo Entity Linker with the spaCy Named Entity Recognizer to include all nouns and adjectives as potential entities.²

Using these extracted entities we analyse the overlap between entities mentioned in the given topics **A**, **B** and those mentioned in the crowdsourced transition utterances. The Jaccard distance between these two sets is 1 for nearly a quarter of the topic-pairs and utterances in our dataset, with a mean of 0.842, meaning that the overlap between entities mentioned in the utterances and entities mentioned in the topics is fairly low. This indicates that users

²This modified version allowed us to identify a wider range of topic-related entities.

overlap	Jaccard Dist.	
	Mean	Mode (freq)
utt.-topics	0.842	1.0 (1274)
utt.-KG path	0.751	0.667 (451)

Table 3: Overlap in entities between transition utterances and (1) the topic sentences (i.e. persona traits) and (2) the path between those topics in the KG.

	Spearman ρ	Pearson r
Cambridge	0.039*	0.046*
PDTB3	0.003	0.001
turn length	0.139	-0.001

Table 4: Correlations between the KG distance between topics **A**, **B** and the number of discourse markers used, as defined by *Cambridge Dictionary* and *Penn Discourse Treebank*, as well as correlation with turn length. * indicates statistical significance $p < 0.01$ with Bonferroni correction for multiple comparisons.

transition from Topic A to Topic B mentioning new unseen entities, following a “path” that can be grounded on a knowledge graph.

In contrast, the overlap between the entities in the KG path between the topics and the entities mentioned in the transition utterances is higher: both the mean and the mode Jaccard distances drop to below 0.8, suggesting that crowdworkers make similar connections to the ones we can find in our knowledge graph a substantial portion of the times. This suggests that our KG-grounded approach can find plausible entities to be mentioned to bridge between topics, similar to the commonsense connections made by humans shifting between topics.

3.3 Transition Strategies in OTTers

To examine the strategies humans applied while completing the OTTers task, we adapted the categories of Riou (2015) for a manual analysis of our data. Riou (2015) distinguishes between *disjunctive* and *stepwise* transitions between topics. Disjunctive transitions make no attempt to relate the new topic to the previous topic, switching abruptly to the new topic without acknowledging the previous topic, whereas stepwise transitions are akin to the previously described transition strategies.

We distinguish between *bridging* and *acknowledge & continue* strategies: in the former, the speaker aims to produce an utterance which connects the previous and new topics directly; in the latter, the speaker acknowledges the previous topic before introducing their own topic, without explicitly relating the two to one another. In addition to these categories, we also annotated utterances

as *off-task* (e.g. replying to or continuing the first topic without any attention paid to the second topic) or *off-topic* when the utterance had nothing to do with either of the two topics (e.g. random greetings or generic questions).

Two of the authors annotated 10 utterances from 10 different users, resulting in 200 total annotations. The initial inter-annotator agreement was 71%, classified as substantial (Krippendorff’s $\alpha = 0.34$), after which the annotators collaborated to reach a consensus annotation for each of the examples that presented a disagreement. Table 5 contains a prototypical example for each of the annotated classes.

More than 80% of the data contains some form of transition to the second topic, with 79% containing a bridging utterance, 5% applying an acknowledge and continue strategy, and only 2% using the disjoint transition strategy. 12% of the data is connected to one or more of the topics in some way but does not serve as a transition, and 2% of the data is completely off-topic. This analysis suggests that our corpus indeed represents the kind of knowledge-based transitions we are interested in.

KG distance and discourse markers. We hypothesize that speakers are less likely to use explicit topic management strategies (e.g. topic wrap-ups, discourse markers) when topics are more closely related to each other, e.g. as measured by graph distance in a large knowledge graph. This would be in line with findings about the use of explicit discourse markers versus leaving discourse relations implicit. Torabi Asr and Demberg (2012, 2013) found that explicit markers are more likely to be omitted when the discourse relation is highly predictable based on the content of the arguments.

Based on Riou (2015) we examined the frequency of discourse markers in utterances to test our hypothesis, examining both general conversational discourse markers and those associated with specific discourse relations. For conversational discourse markers we use the Cambridge Dictionary, which provides a list of spoken and written markers, including “well”, “you know”, etc., while for markers signalling particular discourse relations we use the list from the Penn Discourse Treebank (Webber et al., 2019; Prasad et al., 2008, PDTB); these include markers like “because” indicating a causal relationship or “in addition” for an additive relationship. We find a small but significant correlation (≈ 0.04) between conversational discourse

Acknowledge and continue	
A:	i like to eat the same thing as ninja turtles.
T:	I love pizza. I eat it while I skateboard.
B:	i enjoy riding around on a plank with wheels.
Bridging: Missing Link	
A:	i prefer things to be authentic.
T:	I think children are the truest form of authenticity because they say things unfiltered.
B:	i am not a fan of children.
Disjunctive	
A:	i like american made cars.
T:	I like liver cooked in butter – just throwing that in!
B:	i avoid eating broccoli.
Off-Task	
A:	i prefer things to be authentic.
T:	my bro just made some authentic thai chicken.
B:	i am not a fan of children.
Off-Topic	
A:	i learnt to drive.
T:	I had a rough night sleeping in my new bed last night.
B:	i like making a salmon entree.

Table 5: Prototypical examples of each annotation category for transition strategy or lack of transition in OT-Ters. (A) is the preceding topic or utterance; (T) is the collected utterance; and (B) is the goal topic and potential next utterance for Speaker B.

markers and no significant correlations between the use of PDTB3 discourse markers or the turn length and KG distance. This suggests that users are somewhat more likely to use conversational discourse markers as the distance between topics in the knowledge graph increases, in line with our hypothesis.

3.4 Validating the Corpus

We evaluate whether the transition strategies in OT-Ters are less abrupt than those found in PersonaChat by constructing a comparable subset of PersonaChat and performing a human evaluation.

Comparable Corpus Construction. We first extract a subset of PersonaChat where two consecutive turns contain different topics. In other words: turns where one speaker changed the topic from what the previous speaker has just said. Since PersonaChat turns do not incorporate topic annotations, we use a heuristic based on BERTScore to assign a topic to each turn. Given topics \mathbf{t} and turns \mathbf{u} for a dialogue in PersonaChat, we calculate the BERTScore similarity between each $u \in \mathbf{u}$ and each $t \in \mathbf{t}$. For each turn u we then assign $t = \operatorname{argmax}_t(\text{BERTScore}(u, t))$, if and only if

$$\frac{\text{BERTScore}(u, t)}{\text{BERTScore}(u, t')} > d \quad (1)$$

Two people are having a conversation.

Speaker B wants to change the topic of the conversation to Topic B. Choose the snippet (1 or 2) that makes the transition **less abrupt**.

Speaker A says: "I am trying to create my own startup company in the gaming industry."

Speaker B can say (1) "I am wonderful, my wife is laid off and she cooks all the time out of frustration."

or (2) "Well if you get to the point where you need someone in the human resources department of your new company, my wif recently got laid off from her job and could use a new one."

To reach **Topic B:** "my wife recently got laid off from her human resources job."

Which snippet was **less abrupt** as a transition between what Speaker A said and Topic B?

- Snippet (1) Snippet (2) Both are equally good
- Neither of them is good

Figure 3: Interface for crowdsourced validation.

where t' is the topic achieving the second highest BERTScore relative to u , and d is a threshold to ensure that we only assign a topic to a turn if it is a substantially better fit than the other topics.³ While this means that not every turn is assigned a topic, this is necessary to ensure that we do not assign topics to, e.g., greetings like "hi, how are you?".

This way of assigning topics yields a subset consisting of 22,010 utterances which have a different topic from the preceding utterance. Most of these topic-pairs (20,491) are only expressed through one utterance in the dataset, while 1,188 are expressed by two utterances, 248 by three, and 83 by more than 3 utterances. Moreover, there are 445 topic-pairs which also occur in our corpus.

Crowdsourced Validation. Using the comparable sub-corpus of PersonaChat, we asked crowdworkers to vote which of two potential transition utterances was "less abrupt" (Fig. 3) for 49 topic-pairs occurring in both datasets. We collected 3 votes for each utterance and only counted instances where 2/3 workers agreed on the same choice.

The results confirm that OT-Ters has less abrupt transitions: the utterances in OT-Ters were judged as less abrupt in 44/49 cases, with the comparable PersonaChat utterance judged less abrupt in one case, and both utterances rated "bad" in another. Only 3 cases did not present a majority class.

4 Experiments

Having confirmed the quality of our corpus, we now adapt two existing text generation models as baselines for this task. We also explore different train-dev-test splits and conduct an error analysis.

³We set $d = 1.27$ for our dataset construction, since this is the 50th percentile of BERTScore values observed in our data.

	Train	Dev	Test
id	693/1,929	404/1,160	303/1,158
ood	677/2,034	372/1,152	372/1,130

Table 6: Num. unique/total topic-pairs in each split.

4.1 Baselines

The first baseline we consider is a vanilla GPT-2 language model (Radford et al., 2019) fine-tuned on OTTers (vGPT2). Next, we test the recent MultiGen (Ji et al., 2020) on this task, which extends GPT-2 with multi-hop reasoning on commonsense knowledge graphs. In particular, this model combines the vocabulary distribution generated by GPT-2 with a concept distribution in order to produce knowledge grounded responses. The concept distribution is given by reasoning performed on the commonsense knowledge graph ConceptnetIO, using the context modeled through GPT-2.

4.2 Train-Dev-Test Splits

The first split is an out-of-domain split (ood), which ensures that none of the topics in the test-set are present in any of the topic-pairs in the train-set. For the second split, this restriction is relaxed to create an in-domain split (id), allowing one of the topics in each pair in the test-set to appear in the train-set, although with a different second topic.

The ood split resembles a zero-shot scenario, where the model has to generate a shift between two topics it has never been fine-tuned on. Hence, we expect results to be lower than the ones from id. The number of unique and total topic pairs for each split is illustrated in Table 6.

4.3 Evaluation

We evaluate two aspects of the transition task: 1) whether the model can find a sensible path through intermediate topics and 2) whether the model can generate a natural utterance which mentions such intermediate topics.

To evaluate the former, we assess the entities mentioned in the transition utterance to determine how well they bridge the gap between Topic A and Topic B. We use $\text{hits}@k$ ratio as an automatic approximation, which measures the number of relevant entities correctly predicted by the model, out of the k most important entities identified in the target references. This metric shows how well the models ground the concepts introduced in the two dialogue turns and how the reasoning compares to

the human standard presented in OTTers.

For (2) we adopt the same automated metrics used for evaluating MultiGen on the α NLG dataset for comparability: ROUGE-L (Lin, 2004), METEOR (Banerjee and Lavie, 2005), and CIDEr (Vedantam et al., 2015). However, we report the full BLEU score (Papineni et al., 2002)⁴ that accounts for the overlap across 1-4 ngrams instead of just 4-grams (BLEU-4). As word-overlap based metrics have been widely criticised due to their lack of correlation with human judgements (Novikova et al., 2017; Reiter, 2018), we also provide an example-based error analysis in Section 4.4.

4.4 Results

For each aforementioned split we evaluated three different models to compare performance: the pre-trained vGPT2 fine-tuned on each split for OTTers, the MultiGen model fine-tuned only on α NLG, and the same model additionally fine-tuned on OTTers (called α NLGft).

Overview of Results. Table 7 shows the results of these experiments. vGPT2 performs poorly on the one-turn transition task, regardless of the train-dev-test split, which we attribute to the small size of OTTers: with only a few thousand utterances, vGPT2 is unable to learn the task. We notice, however, that the system tends to repeat the main entity in Topic A, therefore scoring surprisingly well on the $\text{hits}@k$ metric, despite the fact that the utterances themselves are of low quality (see Table 8).

The reasoning component added by MultiGen leads to substantial improvements in most of the evaluation metrics but not $\text{hits}@k$ (α NLG in the table). Therefore, the improvements in text quality metrics appear to be due primarily to the similarity between the structure of the abductive NLG task and the increased amount of data for fine-tuning ($\approx 688k$ tokens) compared to fine-tuning vGPT2 on our $\approx 71k$ tokens alone.

Further fine-tuning MultiGen on OTTers leads to substantial improvements on all metrics for both in-domain & out-of-domain splits. The performance improvement is considerable especially because of the relatively small size of the training set (693 unique topic pairs on in-domain, see Table 6), further justifying the compatibility between the original task MultiGen was trained on and OTTers.

⁴we use SacreBLEU (Post, 2018).

	split	BLEU	METEOR	ROUGE-L	CIDEr	hits@1	hits@3
vGPT2	ood	1.26	8.37	12.4	4.65	22.94	12.54
vGPT2	id	1.58	10.26	14.67	3.75	57.14	30.79
α NLG	ood	1.52	16.35	23.26	12.12	21.85	12.03
α NLG	id	1.6	18.9	25.52	11.33	22.89	12.72
α NLGft	ood	4.0	22.8	29.27	21.7	35.33	18.67
α NLGft	id	4.27	25.03	31.54	18.48	38.11	21.51

Table 7: System performance on in-domain and out-of-domain data splits.

Ex. 1	Topic A:	I like babies.
	OTTers	I make babies laugh.
	Topic B:	I enjoy making other people laugh.
	vGPT2:	I like other people making babies.
	α NLG:	I decided to give birth to a baby.
	α NLGft:	I like to make babies laugh when I'm not working.
Ex. 2	Topic A:	I love those sleek Italian cars.
	OTTers	I wish I was as sleek as those cars.
	Topic B:	I am in good shape.
	vGPT2:	I love it. I love it.
	α NLG:	I bought a new car.
	α NLGft:	I like the look of Italian cars.
Ex. 3	Topic A:	I have prior agriculture experience.
	OTTers	My parents were farmers.
	Topic B:	I am a mom.
	vGPT2:	a lot of cooking. I am a lot of cooking.
	α NLG:	My mom taught me how to grow vegetables.
	α NLGft:	I have been a farmer for many years.
Ex. 4	Topic A:	I make blogs.
	OTTers	Do you know of any blogs that are on the subject of large families with babies?
	Topic B:	I have a large family with babies.
	vGPT2:	family.
	α NLG:	I have a lot of babies.
	α NLGft:	I have a large family.

Table 8: Representative outputs from each model.

Nonetheless, the BLEU scores from Table 7 indicate there is still space for improvement. We hypothesise METEOR are higher than BLEU scores, because they also consider paraphrases.

These results confirm that our newly introduced one-turn topic transition task needs a reliable language model combined with an advanced reasoning component.

Detailed Discussion and Model Limitations.

We further analyse the results to understand model limitations. First, we observe that Multigen’s hits@k ratio is quite low, especially when compared to vGPT2. This is surprising considering vGPT2’s generated sentences are mostly very short and repetitive, and the predicted concepts mostly match the ones contained in the ‘Topic A’ sentence. One possible explanation is that Multigen’s

reasoning module uses a gate loss, which determines whether to select a concept from the provided knowledge graph or a word from the GPT2 dictionary. We observed that the majority of the times the model will use a word from the GPT2 dictionary rather than selecting a concept from the knowledge graph.

Moreover, we observe that only 65% of the concepts found in the target sentences are actually nodes in Multigen’s subgraphs. One possible explanation is that Multigen’s reasoning model has a limited input capacity of up to 100 nodes that are at most 2 hops away in order to prune the very large knowledge graph from ConceptNet. The English vocabulary from ConceptNet contains approximately 1,500,000 nodes, which makes the process of determining the concept distributions very computationally expensive and time inefficient. Therefore, the pruning strategy adapted by Ji et al. (2020) overcomes these problems but cannot be applied to the OTTers task, as the selection of the concepts is just as important as the output sentence being fluent. Contrary to our expectations, expanding the size of the knowledge graphs from 100 nodes to 200 and 300 did not improve the hits@k ratio. Most likely because the concepts added to the graphs are either not relevant or misleading for the model. This suggests that improving concept selection is a promising future direction to improve the performance of the reasoning module, leading to overall better topic transitions.

Error Analysis. In addition, we preform an example-based error analysis to further understand the strengths and weaknesses of the individual models. Table 8 shows representative system outputs for each of the models on the in-domain data split. First, we observe that vGPT2 often generates very simple sentences (e.g., ‘family.’, in Ex. 4), repeated non-content bearing tokens (e.g., ‘I love it.’, in Ex. 2), or incoherent and often not specific enough output to form a successful bridging transition (e.g., ‘a lot of cooking.’, in Ex. 3, is not a well-formed sentence, and only loosely connected to Topic A

about ‘*agricultural experience*’), contributing to low BLEU scores. However, this also reinforces the idea that the `hits@k` scores are artificially inflated simply due to vGPT2 choosing to include one of the entities from the first topic.

The outputs from MultiGen tested on OTTers show a better performance than vGPT2, given that the topic selection for the model is grounded on ConceptNet. However, since the Abductive NLG task is different than the ‘Topic Transition’ task addressed in OTTers, there is a discrepancy in the use of the language. The model often outputs coherent sentences that use generic commonsense facts which may not be related to Topic B (e.g., ‘*I decided to give birth to a baby*’, in Ex. 1).

The texts generated from MultiGen fine-tuned on OTTers on the other hand, introduce interesting connections between Topic A and Topic B (e.g., ‘*I like to make babies laugh when I’m not working*’, in Ex. 1) and leverage commonsense (e.g., ‘*I like the look of Italian cars*’, in Ex. 2, where ‘*the look*’ creates a connection with ‘*being in good shape*’ from Topic B).

5 Discussion & Conclusion

Ethical Considerations. We recognise that any mixed-initiative dialogue system carries risks related to dual-use: in addition to helpful systems which serve to help users explore a new topic or discover more about the world, a system which can effectively change the topic of conversation could also be used to manipulate user behaviour. For example, bridging strategies for topic transitions could be used by virtual assistants to encourage users to make a purchase or to express their opinion or preference regarding sensitive subjects.

Conclusion. We have defined a new NLG task exploring one-turn topic transitions for mixed-initiative in open-domain systems. Our OTTers corpus provides training data for modelling topic transitions based on ‘missing link’ topics which connect the previous conversation subject to a new topic. Baseline models based on state-of-the-art approaches to text generation illustrate possible approaches to the task and show that there is room for improvement. In particular, we show that commonsense knowledge grounding is necessary for this task, outperforming fine-tuned large language models. In future work, we will explore model architectures specifically designed for topic transitions, as well as fine-tuning strategies to deal with

small datasets. We also plan to evaluate the impact of bridging transitions on user (dis)engagement in an open-domain dialogue system.

Acknowledgements

This research received funding from the EPSRC project MaDrIgAL (EP/N017536/1), as well as Google Research Grant to support NLU and dialog research at Heriot-Watt University.

References

- 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.
- A-C Berthoud and Lorenza Mondada. 1995. Traitement du topic, processus énonciatifs et séquences conversationnelles. *Cahiers de linguistique française*, (17):205–228.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen-tau Yih, and Yejin Choi. 2020. Abductive commonsense reasoning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Roi Blanco, Giuseppe Ottaviano, and Edgar Meij. 2015. Fast and space-efficient entity linking in queries. In *Proceedings of the Eight ACM International Conference on Web Search and Data Mining, WSDM 15, New York, NY, USA. ACM*.
- Gillian Brown and George Yule. 1983. *Discourse analysis*. Cambridge university press.
- Wallace Chafe. 1994. *Discourse, consciousness, and time: The flow and displacement of conscious experience in speaking and writing*. University of Chicago Press.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered conversational agents. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.
- Ronald Geluykens. 1993. Topic introduction in english conversation. *Transactions of the Philological Society*, 91(2):181–214.

- Karthik Gopalakrishnan, Behnam Hedayatnia, Qiang 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. *Proc. Interspeech 2019*, pages 1891–1895.
- Dionysis Goutsos. 1997. *Modeling Discourse Topic: sequential relations and strategies in expository text*, volume 59. Greenwood Publishing Group.
- Jian Guan, Yansen Wang, and Minlie Huang. 2019. Story ending generation with incremental encoding and commonsense knowledge. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6473–6480.
- Jeanette K Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language*, pages 274–307.
- Angela James. 1995. Topic shift in casual conversation. *The University of Western Ontario Journal of Anthropology*, 2(1).
- Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, Xiaoyan Zhu, and Minlie Huang. 2020. Language generation with multi-hop reasoning on commonsense knowledge graph. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Bin Jiang, Wanyue Zhou, Jingxu Yang, Chao Yang, Shihan Wang, and Liang Pang. 2020. Pednet: A persona enhanced dual alternating learning network for conversational response generation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4089–4099.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Douglas W Maynard. 1980. Placement of topic changes in conversation. *Semiotica*, 30(3-4):263–290.
- Lorenza Mondada. 2001. Gestion du topic et organisation de la conversation. *Cadernos de estudos linguísticos*, 41:7–36.
- Lorenza Mondada. 2003. Parler topical et organisation séquentielle: l’apport de l’analyse conversationnelle. *Verbum (Nancy)*, (2):193–219.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for nlg. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252.
- 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.
- Aasish Pappu, Roi Blanco, Yashar Mehdad, Amanda Stent, and Kapil Thadani. 2017. Lightweight multilingual entity extraction and linking. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 17*, New York, NY, USA. ACM.
- Sylvie Porhiel. 2005. Les marqueurs de thématisation: des thèmes phrastiques et textuels. *Travaux de linguistique*, (2):55–84.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191. Association for Computational Linguistics.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The Penn Discourse TreeBank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Matthew Purver, Arash Eshghi, and Julian Hough. 2011. Incremental semantic construction in a dialogue system. In *Proceedings of the Ninth International Conference on Computational Semantics (IWCS 2011)*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Ehud Reiter. 2018. A structured review of the validity of BLEU. *Computational Linguistics*, 44(3):393–401.
- Marine Riou. 2015. A methodology for the identification of topic transitions in interaction. *Discours. Revue de linguistique, psycholinguistique et informatique. A journal of linguistics, psycholinguistics and computational linguistics*, (16).
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge.

Fatemeh Torabi Asr and Vera Demberg. 2012. Implicitness of discourse relations. In *Proceedings of COLING 2012*, pages 2669–2684, Mumbai, India. The COLING 2012 Organizing Committee.

Fatemeh Torabi Asr and Vera Demberg. 2013. On the information conveyed by discourse markers. In *Proceedings of the Fourth Annual Workshop on Cognitive Modeling and Computational Linguistics (CMCL)*, pages 84–93, Sofia, Bulgaria. Association for Computational Linguistics.

Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.

Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.

Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The Penn Discourse Treebank 3.0 Annotation Manual. Technical report.

Pengcheng Yang, Lei Li, Fuli Luo, Tianyu Liu, and Xu Sun. 2019. Enhancing topic-to-essay generation with external commonsense knowledge. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2002–2012, Florence, Italy. Association for Computational Linguistics.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213.

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 *ACL, system demonstration*.

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-18*, pages 4623–4629.

A OTTers Data Collection

In order to avoid collecting noisy or out-of-task data, we established some worker requirements for turkers participating in our data collection. Workers needed to:

- be Masters (label assigned by Mechanical Turk to workers who achieve excellence across a variety of tasks),
- have a number of HITs approved greater than 500,
- have a HIT approval rate (%) greater than 80,
- being located in an English speaking country, namely Australia, Canada, New Zealand, United Kingdom, and United States.

A Worker received a reward of \$0.3 for each completed assignment. The reward was calculated based on an estimate of the time it would take a Worker to read the instructions and complete the task. The time for completing the task has been estimated at 1.5 minutes, and the reward was calculated accordingly to a \$12 hourly payment. Each task had been assigned to 3 unique workers.

Figure 4 shows the instructions that Workers were presented after opening the OTTers data collection task. The instructions explain that the context is a conversation with a newly-met person. After writing the sentence for transitioning the current topic to the ‘final’ one, workers are asked to list the topics they covered for the transition. Additionally, we provided an example:

Current sentence: ‘I have a love of reptiles.’

Final sentence: ‘I want to travel to NYC.’

Topic shifting sentence: ‘I know there is a cool snake species in the New York zoo. This is why I want to travel to NYC.’

Covered topics:

- Reptiles
- Zoo
- NYC

You've just met this new person, who is talking about something and you want to change the topic!

In this scenario, imagine that you are having a conversation with a person that you just met. You are currently talking about a given topic but you want to change the subject in a natural way. That is your challenge! Start by reading the two sentences we have provided.

The first sentence is the last thing your conversation partner said, and your goal is to switch the topic conveyed in the final sentence.

Write one or two sentences you would use to shift the topic in the most natural and coherent way, so that your topic introduction does not seem abrupt.

After you are done switching subject, list the topics that you covered in order to make this transition. Remember to first list the topic from the 'current' sentence, and to finalize your list with the subject from the 'final' sentence.

Please do NOT just copy & paste the final sentence. If you do so your HIT will be REJECTED.

Please remember you can withdraw from the task at any time.

Please do NOT include any personal information.

Figure 4: Instructions shown to Workers for the OTTers data collection.

Current sentence	I like the pool.
Topic transition	I reward myself after going to the pool by eating a hearty meal.
Final sentence	I like vegetables.
Covered topics	<ul style="list-style-type: none"> • Pool • Hearty meal • Vegetables

Table 9: Example of topic transition collected for OTTers.

Current sentence	I am a parent.
Topic transition	My kids each have a pet they take care of.
Final sentence	I like animals.
Covered topics	<ul style="list-style-type: none"> • Parent • Kids • Pet • Animals

Table 10: Example of topic transition collected for OTTers.

Current sentence	I have gone across the ocean.
Topic transition	My mom was in a band when she lived in France.
Final sentence	My mom is famous.
Covered topics	<ul style="list-style-type: none"> • Ocean • Mom • Band • France • Famous

Table 11: Example of topic transition collected for OTTers.

Current sentence	I love cuddling with my babies.
Topic transition	I wish I could still do that with my children, but after my back injury I have had to listen to my physician and really change my ways.
Final sentence	The Dr said no sitting up for me.
Covered topics	<ul style="list-style-type: none"> • Children • Medical care • Disability

Table 12: Example of topic transition collected for OTTers.

Current sentence	I like going to concerts.
Topic transition	I like going to concerts which means I normally have to take a break during a workweek.
Final sentence	I do not go a full week of employment without a break.
Covered topics	<ul style="list-style-type: none"> • Concerts • Break • Work

Table 13: Example of topic transition collected for OTTers.

Current sentence	I play basketball and football.
Topic transition	I bet my fiancee while at the park playing basketball.
Final sentence	My significant other and I will be having a wedding.
Covered topics	<ul style="list-style-type: none"> • Basketball • Football • Fiancee • Park • Wedding

Table 14: Example of topic transition collected for OTTers.