

TIAGE: A Benchmark for Topic-Shift Aware Dialog Modeling

Huiyuan Xie¹ Zhenghao Liu² Chenyan Xiong³ Zhiyuan Liu⁴ Ann Copestake¹

¹Department of Computer Science and Technology, University of Cambridge, UK

²Department of Computer Science and Technology, Northeastern University, China

³Microsoft Research, United States

⁴Department of Computer Science and Technology, Tsinghua University, China

{hx255, aac10}@cl.cam.ac.uk, liuzhenghao@cse.neu.edu.cn

chenyan.xiong@microsoft.com, liuzy@tsinghua.edu.cn

Abstract

Human conversations naturally evolve around different topics and fluently move between them. In research on dialog systems, the ability to actively and smoothly transition to new topics is often ignored. In this paper we introduce TIAGE, a new topic-shift aware dialog benchmark constructed utilizing human annotations on topic shifts. Based on TIAGE, we introduce three tasks to investigate different scenarios of topic-shift modeling in dialog settings: topic-shift detection, topic-shift triggered response generation and topic-aware dialog generation. Experiments on these tasks show that the topic-shift signals in TIAGE are useful for topic-shift response generation. On the other hand, dialog systems still struggle to decide when to change topic. This indicates further research is needed in topic-shift aware dialog modeling.¹

1 Introduction

Existing dialog models (Ghandeharioun et al., 2019; Einolghozati et al., 2019; Liu et al., 2018) have been reported to perform well in generating on-topic utterances in dialog scenarios. However, those models still struggle to proactively generate appropriate topic-shift utterances in conversations (Holtzman et al., 2020; Zhang et al., 2020a).

It is beneficial for dialog systems to be able to shift topics fluently. As shown in Figure 1, topic-shift behaviors are commonly observed in human conversations (Brown and Yule, 1983). Fluent topic shifts therefore are crucial for dialog models to be able to model or mimic human conversational patterns. Proactively using topic shifts can help chatbots guide conversations to a pre-defined target (Tang et al., 2019). Furthermore, switching topics allows chatbots to maintain engaging conversations with users. Without the ability to actively

¹Code and data available at: <https://github.com/HuiyuanXie/tiage>.



Figure 1: An example of topic-shift behaviors in human conversations. Topic-shift utterances are highlighted in green and in *italic*. Changing the topic helps keep the conversation going on.

shift topics away from tired topics, chatbots risk generating dull responses or repeating themselves regarding a specific topic.

To facilitate research on topic-shift dialog modeling, we curate a Topic-shift Aware dialog dataset (TIAGE) by augmenting the PersonaChat dataset (Zhang et al., 2018) with topic-shift annotations. To the best of our knowledge, TIAGE is the first dataset that focuses on topic-shift behaviors in open-domain dialog data. TIAGE contains a human annotated dataset with 7,861 gold standard topic-shift annotations, and a weak supervision dataset to adapt pretrained NLG systems to PersonaChat-style data. The inter-annotator agreement for topic-shift annotations in TIAGE is 0.479.

With TIAGE, we propose three tasks to study topic-shift behaviors: topic-shift detection, topic-shift triggered response generation and topic-aware dialog generation. The topic-shift detection task asks models to detect whether the ongoing topic has shifted or should shift. The other two tasks focus on modeling topic-shift behaviors in response generation. Specifically, the topic-shift triggered response generator receives a fixed topic-shift signal to generate topic-shift responses, whilst the topic-aware dialog generation task requires dialog systems to

predict the topic-shift trigger by themselves.

Our experiments reveal that the topic-shift signals in TIAGE indeed improve dialog systems’ ability to generate topic-shift responses. However, it is difficult for dialog models to predict when it is appropriate to change topics. These observations highlight the need for better modeling of topic shifts in dialog generation. We hope our benchmark can motivate further research on topic-shift aware dialog modeling.

2 Related Work

Existing work in dialog systems falls into two broad categories. *Task-oriented* dialog systems (Budzianowski et al., 2018; Liu et al., 2018) help users complete tasks in specific domains. *Open-domain* dialog systems (Chen and Gao, 2017; Tang et al., 2019) allow agents to have open-ended conversations with users. Most existing dialog models (Fang et al., 2018; Zhang et al., 2020b; Ghandeharion et al., 2019) emphasize end-to-end response generation, and do not explicitly address the topic-shift problem in dialog generation.

Early work in topic detection and segmentation (Hirschberg and Litman, 1993; Passonneau and Litman, 1997) focused on identifying cue phrases (such as *on a different note*) or examining lexical cohesion to segment topical chunks. Other work (Fiscus and Doddington, 2002) investigated topic detection and tracking (TDT) in a stream of broadcast news stories. More recent work (Glavas and Somasundaran, 2020) has explored utilizing neural networks to address topic segmentation. Although some of the existing work (Galley et al., 2003; Arnold et al., 2019) has investigated topic detection in dialog-style data, the generation aspect of topic-shift modeling in dialog settings is still unclear.

3 Topic-Shift Aware Dialog Dataset

In this section we introduce the rationale for our choice of data source, the human annotation process of topic-shift labelling in TIAGE and its data statistics. We also analyze the linguistic patterns of topic-shift utterances in TIAGE.

Rationale for our choice of data source. We construct TIAGE by augmenting the PersonaChat dataset (Zhang et al., 2018) with topic-shift human annotations. We view PersonaChat as a suitable dataset for topic-shift annotation for the following reasons: (1) the Personachat data was collected online in a textual form by mimicking chit-chat sce-

	WEAKSUPO _{train}	WEAKSUPO _{dev}
#Dialogs	7,939	1,000
#Instances	108,711	13,788
#AvgTurns	14.7	14.8

(a) The weak supervision data split.

	ANNO _{train}	ANNO _{dev}	ANNO _{test}
#Dialogs	300	100	100
#Instances	4,767	1,546	1,548
#AvgTurns	15.6	15.5	15.6

(b) The human annotated data split.

Table 1: Data statistics. #AvgTurns denotes the average number of turns per dialog. Each *instance* is a (context, response) pair around a specific dialog turn. The average number of tokens per utterance is 11.8. In the human annotated data split, the average number of topic-shift turns per dialog is 3.5. The vocabulary size of the entire dataset is around 18K.

narios, where natural shifts of topics are more likely to happen; (2) dialogs in this dataset contain more than 10 dialog turns, and longer dialog contexts tend to exhibit a conversational flow with more topics; and (3) despite the fact that some participants in PersonaChat may have rushed into changing topics to quickly exchange their profile information, we observed that most of the participants still manage to change topics in a more natural and coherent way, making this dataset a favorable choice to study topic-shift behaviors.

Human annotation process. For the annotation pool, we have a total number of 25 human annotators. We randomly selected 500 dialogs from the original PersonaChat dev/test datasets, resulting in 7,861 dialog turns to label. Each dialog turn was randomly assigned to and independently labeled by 2 annotators. For each dialog turn, we asked annotators to indicate whether they think the conversational topic is changed at that turn. During the annotation process, all annotators were talked through the general aim of this annotation task and given the same annotation guidelines (see Appendix A.1 for details).

Since topic is co-constructed, it is rather limiting to analyze a turn for itself when trying to identify topic transitions. To facilitate the recognition of slowly transitioned topics, we encouraged the annotators to take into account both the previous two turns and the following two turns of the target dialog turn to make a decision. This helped decision making for cases where topics are slowly developed and transitioned.

After annotating, we obtained a dialog dataset

with gold standard topic-shift labels for 7,861 dialog turns. The Cohen’s Kappa score for all annotations is 0.479². Annotated examples of TIAGE dialogs are shown in Appendix A.2.

Dataset statistics. As shown in Table 1, TIAGE provides *weak supervision data* and *human annotated data* to train dialog models. *Weak supervision data* is selected from the original PersonaChat training set and helps adapt NLG models to PersonaChat-style data. The weak supervision data consists of 8,939 dialogs and is split into two sets: WEAKSUPO_{train} and WEAKSUPO_{dev}. *Human annotated data* consists of 500 annotated dialogs with topic-shift annotations at each dialog turn. We split them into 300 ANNO_{train}, 100 ANNO_{dev} and 100 ANNO_{test} dialogs respectively. As each dialog has multiple dialog turns, we extract (context, response) pairs as instances for all turns in each dialog.

Analysis of topic-shift patterns. We examine a number of topic-shift utterances labeled by human annotators. We find that many of the topic-shift responses demonstrate an interesting pattern of [comment; topic shift]. More specifically, the response that changes the conversational topic is typically a brief *comment* on the previous dialog context, tailed by a *topic-shift* sentence with a different conversational focus. The *comment* usually corresponds to the sentiment previously expressed in the dialog.

This pattern echoes some of the findings in pragmatics research (Brown and Levinson, 1987; Goldsmith, 2007). When speakers introduce a new topic, it is a common positive politeness strategy (Leech, 2014) to first respond to the content uttered by other speakers. This pattern is potentially useful for dialog systems seeking to generate topic-shift utterances in a natural and coherent manner. Before introducing a new topic, it is favorable for dialog systems to first generate a comment regarding the previous topic that expresses either *approbation* (e.g., “great”, “that’s cool”) or *sympathy* (e.g., “that’s too bad” or “I’m sorry to hear that”). This shows that they are attuned to the users’ interests and needs.

²The Cohen’s Kappa score ranges from 0.41 to 0.60 indicating moderate agreement, which confirms the quality of the human annotations of TIAGE.

4 Tasks of Topic-Shift Modeling

Along with dialog utterances, TIAGE also provides gold standard topic-shift labels for dialog turns. This enables us to model topic shifts in dialog scenarios. We first introduce two tasks: topic-shift detection and topic-shift triggered response generation. They can be considered as intermediate steps of the topic-aware dialog generation task.

4.1 Preliminary of Response Generation

When considering a specific turn in a dialog, we denote the current utterance and all its previous utterances as the *context* $\mathbf{X}_T = \{\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_N\}$ where \mathbf{x}_i is the i -th utterance in the dialog history, and N is the context length. Then we expect the *response* to be generated after the current utterance \mathbf{x}_N . We denote a *topic-shift response* as $\mathbf{s}_{TS} = \{s_1, \dots, s_i, \dots, s_T\}$ where s_i is the i -th token in the response and T is the sentence length. Similarly, an *on-topic response* is denoted as $\bar{\mathbf{s}}_{NTS} = \{\bar{s}_1, \dots, \bar{s}_i, \dots, \bar{s}_M\}$ where \bar{s}_i is the i -th token in the response and M is the sentence length.

4.2 Topic-Shift Detection

Topic-shift detection is a fundamental task that evaluates models’ ability to detect topic-shift occurrence at dialog turns.

Task definition. We introduce two settings for this task. In the *retrospective* setting, models have access to both the dialog context \mathbf{X}_T and the corresponding response (either \mathbf{s}_{TS} or $\bar{\mathbf{s}}_{NTS}$) to detect topic-shift occurrence, whilst in the *predictive* setting, models are asked to make topic-shift predictions based on the context \mathbf{X}_T only.

Topic shift classifiers. We first implement three retrospective classifiers. We employ GenEnc which uses the GEN Encoder (Zhang et al., 2019) to separately encode dialog context and response into embeddings to estimate the topic-shift intents. GenEnc uses a cosine similarity threshold of 0.25 to filter out (context, response) pairs, and classify them as topic-shift occurrences. Then we implement a BERT-Wiki727k model (Devlin et al., 2019) trained on the WIKI-727K dataset (Koshorek et al., 2018). We also employ a T5 model (Raffel et al., 2019) finetuned on the ANNO_{train} data with topic-shift labels as our retrospective T5 topic-shift classifier (denoted as RetroTS-T5).

For the predictive setting, we implement a T5-based topic-shift manager (denoted as T5Manager) and finetune it on the ANNO_{train}

data. The major difference between `Retros-T5` and `TManager` is that `Retros-T5` has access to both the dialog context and the response, while `TManager` makes topic-shift predictions based solely on the context.

4.3 Topic-Shift Triggered Response Generation

This task examines models’ ability to generate topic-shift utterances when a need to change topics is signaled.

Task definition. Given a dialog context \mathbf{X}_T , the topic-shift triggered response generation task requests models to directly generate a response s_{TS} that shifts the conversation to a different topic.

Topic-shift triggered generator. We build a `T5-NLGTS` response generator using the pre-trained T5 model. We first train the T5 model on the `WEAKSUPOtrain` data, and then further fine-tune it on the topic-shift instances (*i.e.*, where topic shifts occur) in the `ANNOtrain` data.

Compared approaches. We also try a number of topic-insensitive NLG models for comparison. We train a `T5-NLG` model on the `WEAKSUPOtrain` data without any topic-shift signals. We use the `DialogGPT` model (Zhang et al., 2020b) finetuned on the same data as another baseline.

4.4 Topic-Aware Dialog Generation

The third task we propose targets more difficult and realistic topic-shift modeling in dialog generation.

Task definition. More formally, given a dialog context \mathbf{X}_T , the goal of the topic-aware dialog generation task is to generate a topic-shift response s_{TS} if a change of topic is needed, or an on-topic response \bar{s}_{NTS} if otherwise. The topic-aware dialog generation task asks models to identify the need to change topics by themselves and generate topic-shift or on-topic responses according to the prediction.

Topic-aware dialog system. Our topic-aware dialog system (`TADial`) is a pipeline system. We separately train two T5-based response generators: `T5-NLGTS` and `T5-NLGNTS`. We switch between the two response generators to produce either a topic-shift or on-topic response, guided by the topic-shift signals from `TManager`. `T5-NLGTS` aims to generate topic-shift responses, while `T5-NLGNTS` is finetuned on non-topic-shift instances to generate on-topic responses.

Compared approaches. We use the `T5-NLG`

Approaches	Precision	Recall	F1-score
BERT-WIKI727K	0.412	0.020	0.038
GENENC	0.337	0.199	0.250
RETROS-T5	0.709	0.657	0.682
TSMANAGER	0.340	0.170	0.220
Human³	0.687	0.607	0.644

Table 2: Model performance on the topic-shift detection task.

Model	BLEU-2	METEOR	ROUGE_L	CIDEr
DIALOGPT	0.060	0.077	0.125	0.104
T5-NLG	0.079	0.086	0.161	0.170
T5-NLG _{TS}	0.092	0.092	0.177	0.175

Table 3: Evaluation results of topic-shift triggered response generation on topic-shift instances in `ANNOtest`.

and `DialogGPT` models finetuned on the `WEAKSUPOtrain` data as baselines for comparison.

5 Evaluation Results

We report here the evaluation results for baseline systems on the above three tasks.

Topic-shift detection. We test topic-shift classifiers on the annotated `ANNOtest` split. From Table 2 we observe that `Retros-T5` outperforms other approaches by a large margin and is on par with human performance. This indicates that topic shifts in `PersonaChat` dialogs exhibit certain patterns, which can be captured from our human-labeled topic-shift annotations by our retrospective T5 classifier. We also notice that there is a clear gap in classification performance between `Retros-T5` and `TManager`. The predictive setting of `TManager` is inherently harder than `Retros-T5`, as it is asked to predict topic-shift labels based solely on dialog context.

Topic-shift triggered response generation. In Table 3 we report evaluation results⁴ of our topic-shift triggered response generator (`T5-NLGTS`) and two topic-insensitive models (`DialogGPT` and `T5-NLG`). Models are tested on the *topic-shift* instances in `ANNOtest`. We observe that `T5-NLG` yields better performance than `DialogGPT`. Furthermore, `T5-NLGTS` achieves better evaluation results on topic-shift test instances, outperforming `T5-NLG` by 16.46% in BLEU-2 and 9.94% in ROUGE_L. The better performance of `T5-NLGTS`

²We use the annotations from one annotator as gold standard references, and calculate human performance on the annotations from the other annotator.

⁴We use the `nlg-eval` package for automatic evaluation. <https://github.com/Maluuba/nlg-eval>.

Model	BLEU-2	METEOR	ROUGE_L	CIDEr
DIALOGPT	0.063	0.077	0.134	0.125
T5-NLG	0.082	0.087	0.159	0.175
TADIAL	0.082	0.087	0.162	0.177

Table 4: Evaluation results of topic-aware dialog generation on all instances in ANNO_{test}.

validates the effectiveness of topic-shift signals in improving topic-shift response generation. It also proves that explicitly modeling topic-shift behaviors can potentially benefit dialog generation.

Topic-aware dialog generation. We test TADial and two topic-insensitive baselines on all instances in ANNO_{test}. From Table 4, we can see that TADial with a dedicated topic-shift management component does not yield better performance over the T5-NLG model which is simply trained on dialog instances with no topic-shift labels. This points out that due to the deficiency of TManager signals, hard-wiring a topic-shift management component into the generation pipeline falls short to improve generation results. It remains a challenging task to produce well-timed and good-quality topic-shift signals based on dialog context only, which hinders overall topic-aware dialog generation.

6 Conclusion and Future Work

We construct the TIAGE dataset with human annotated topic-shift labels on the basis of the PersonaChat dataset. Based on TIAGE, we introduce three tasks: topic-shift detection, topic-shift triggered response generation and topic-aware dialog generation. Empirical results show that topic-shift labels in TIAGE are useful for topic-shift response generation. However, it remains a challenging task for dialog models to predict good-quality topic-shift signals based on dialog context only. Further research is needed on selecting appropriate topics to shift to among multiple references. Natural topic shifts can be both a precaution against, and a remedy to, dull and repetitive response generation in real-world dialog applications. TIAGE with its topic-shift annotations can help direct future investigation on the incorporation of topic-shift tactics in dialog models, which allows more effective control over topic-shift aware dialog generation.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. We thank our annotators from

the University of Cambridge and Tsinghua University for their help in annotating the topic-shift labels used in this paper. Huiyuan Xie is grateful for being supported by the CSC Cambridge Scholarship. Zhenghao Liu is supported by National Natural Science Foundation of China (NSFC) under grant No. 61872074 and 61772122. This work is partly supported by the National Key Research and Development Program of China (No. 2020AAA0106501).

References

- Sebastian Arnold, Rudolf Schneider, Philippe Cudré-Mauroux, Felix A. Gers, and Alexander Löser. 2019. [SECTOR: A neural model for coherent topic segmentation and classification](#). *Transactions of the Association for Computational Linguistics*, pages 169–184.
- Gillian Brown and George Yule. 1983. *Discourse analysis*. Cambridge University Press.
- Penelope Brown and Stephen C Levinson. 1987. *Politeness: Some universals in language usage*, volume 4. Cambridge University Press.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of EMNLP*, pages 5016–5026.
- Yun-Nung Chen and Jianfeng Gao. 2017. [Open-domain neural dialogue systems](#). In *Proceedings of IJCNLP*, pages 6–10.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Arash Einolghozati, Sonal Gupta, Mrinal Mohit, and Rushin Shah. 2019. [Improving robustness of task oriented dialog systems](#). In *Proceedings of Conversational AI Workshop at NeurIPS*.
- Hao Fang, Hao Cheng, Maarten Sap, Elizabeth Clark, Ari Holtzman, Yejin Choi, Noah A. Smith, and Mari Ostendorf. 2018. [Sounding board: A user-centric and content-driven social chatbot](#). In *Proceedings of NAACL-HLT*, pages 96–100.
- Jonathan Fiscus and George Doddington. 2002. [Topic detection and tracking evaluation overview](#). *Topic Detection and Tracking: Event-based Information Organization*, pages 17–31.
- Michel Galley, Kathleen R. McKeown, Eric Fosler-Lussier, and Hongyan Jing. 2003. [Discourse segmentation of multi-party conversation](#). In *Proceedings of ACL*, pages 562–569.

- Asma Ghandeharioun, Judy Hanwen Shen, Natasha Jaques, Craig Ferguson, Noah Jones, Àgata Lapedriza, and Rosalind W. Picard. 2019. [Approximating interactive human evaluation with self-play for open-domain dialog systems](#). In *Proceedings of NeurIPS*, pages 13658–13669.
- Goran Glavas and Swapna Somasundaran. 2020. [Two-level transformer and auxiliary coherence modeling for improved text segmentation](#). In *Proceedings of AAAI*, pages 7797–7804.
- Daena J Goldsmith. 2007. Brown and levinson’s politeness theory. *Explaining communication: Contemporary theories and exemplars*, pages 219–236.
- Julia Hirschberg and Diane Litman. 1993. [Empirical studies on the disambiguation of cue phrases](#). *Computational Linguistics*, (3):501–530.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *Proceedings of ICLR*.
- Omri Koshorek, Adir Cohen, Noam Mor, Michael Rotman, and Jonathan Berant. 2018. [Text segmentation as a supervised learning task](#). In *Proceedings of NAACL-HLT*, pages 469–473.
- Geoffrey N Leech. 2014. *The pragmatics of politeness*. Oxford University Press.
- Bing Liu, Gokhan Tür, Dilek Hakkani-Tür, Pararth Shah, and Larry Heck. 2018. [Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems](#). In *Proceedings of NAACL-HLT*, pages 2060–2069.
- Rebecca J. Passonneau and Diane J. Litman. 1997. [Discourse segmentation by human and automated means](#). *Computational Linguistics*, (1):103–139.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv preprint arXiv:1910.10683*.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. [Target-guided open-domain conversation](#). In *Proceedings of ACL*, pages 5624–5634.
- 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 EMNLP*, pages 38–45.
- Hongfei Zhang, Xia Song, Chenyan Xiong, Corby Rosset, Paul N. Bennett, Nick Craswell, and Saurabh Tiwary. 2019. [Generic intent representation in web search](#). In *Proceedings of SIGIR*, pages 65–74.
- Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2020a. [Grounded conversation generation as guided traverses in commonsense knowledge graphs](#). In *Proceedings of ACL*, pages 2031–2043.
- 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 ACL*, pages 2204–2213.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. [DialogGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of ACL*, pages 270–278.

A Appendix

A.1 Human Annotation Guidelines

Here we present the annotation guidelines used for the human annotation process in this work.

Task description. Chitchat systems are expected to have the ability to proactively change conversational topics when necessary. For occasions when a chat agent runs out of things to say or the current discussion is starting to get boring, topic shifting is a common tactic to keep the conversation going on. In this work, we aim to model topic-shift phenomenon in open-domain dialog settings. To achieve this, we need to construct a new dialog dataset with topic-shift signals.

[Speaker1:]	My dad works for the New York Times.
[Speaker2:]	Oh wow! You know, I dabble in photography; maybe you can introduce us sometime.
[Speaker1:]	<i>Photography is the greatest art out there.</i> → not a topic shift
(a) Commenting on the previous context.	
[Speaker1:]	Do you teach cooking?
[Speaker2:]	<i>No, since I'm a native of Mexico, I teach Spanish.</i> → not a topic shift
(b) Question answering.	
[Speaker1:]	Pets are cute!
[Speaker2:]	<i>I heard that Huskies are difficult dogs to take care of.</i> → not a topic shift
(c) Developing the conversation to sub-topics.	
[Speaker1:]	You are an artist? What kind of art, I do American Indian stuff.
[Speaker2:]	Yes, I love to draw. <i>I love to eat too, sometimes too much.</i> → topic shift
(d) Introducing a relevant but different topic.	
[Speaker1:]	What do you do for fun?
[Speaker2:]	I drive trucks so me and my buds go truckin in the mud.
[Speaker1:]	Must be fun! My version of that's running around a library!
[Speaker2:]	<i>Do you have a favourite animal? Chickens are my favourite. I love them.</i> → topic shift
(e) Completely changing the topic.	

Table 5: Different scenarios of dialog response in conversations.

Data annotation. For each utterance in a dialog, annotators are asked to decide whether the topic of the conversation changes when transiting from the current utterance to the following response. If there is a topic shift, annotators should label the response with “1”, otherwise label it with “0”.

	Dialog	TS Label
[Speaker1:]	Hi! How are you this evening?	N/A
[Speaker2:]	Good. I spent all afternoon walking my dogs. I've three Labradors.	0
[Speaker1:]	Cool, that's a lot of dogs. Do you like music? I love it.	1
	Dialog	TS Label
[Speaker1:]	I think you are great. You are my best friend.	N/A
[Speaker2:]	My best friend is a bear, bears don't have friends, that's why they're my favourite.	0
[Speaker1:]	Webster's dictionary defines weddings as the fusing of two metals with a hot torch.	1

Table 6: Annotated dialog examples in TIAGE.

In conversations, the response of a speaker to the dialog context usually falls into one of the following cases (see examples in Table 5):

- Commenting on what the other participant just said (the most common scenario);
- Question answering;
- Developing the conversation to sub-topics;
- Introducing a relevant but different topic;
- Completely changing the topic.

Other tips for data labeling. A number of words and phrases are often used as indicators for topic shifts, including but not limited to: but, speaking of, talking about, anyway, by the way, that reminds me, before I forget, I want to mention, let's talk about, we need to discuss, funny you should mention that, not to change the subject but, changing the topic slightly, totally unrelated, on a different/relevant note.

A.2 Examples of Labeled Data in TIAGE

In Table 6 we showcase examples of labeled dialogs selected from TIAGE.

A.3 Implementation Details

For the topic-shift classifiers, we use the base version of BERT and T5 models, initialized from their pretrained weights. For the dialog response generation experiments, we use the small version of DialoGPT and the base version of T5. Our implementation is based on the HuggingFace Transformers library (Wolf et al., 2020). All models are optimized using Adam with a learning rate of 5e-5 and a batch size of 64. We set the maximum input sequence length to 512. The training is carried out using 1 Nvidia RTX 8000 GPU and takes around 15 hours.