

MultiTalk: A Highly-Branching Dialog Testbed for Diverse Conversations

Yao Dou,¹ Maxwell Forbes,^{1 2} Ari Holtzman,¹ Yejin Choi^{1 2}

¹ University of Washington,

² Allen Institute for AI

{douy, mbforbes, ahai, yejin}@cs.washington.edu

Abstract

We study conversational dialog in which there are many possible responses to a given history. We present the MultiTalk Dataset, a corpus of over 320,000 sentences of written conversational dialog that balances a high branching factor (10) with several conversation turns (6) through selective branch continuation. We make multiple contributions to study dialog generation in the highly branching setting. In order to evaluate a diverse set of generations, we propose a simple scoring algorithm, based on bipartite graph matching, to optimally incorporate a set of diverse references. We study multiple language generation tasks at different levels of predictive conversation depth, using textual attributes induced automatically from pretrained classifiers. Our culminating task is a challenging theory of mind problem, a controllable generation task which requires reasoning about the expected reaction of the listener.

Introduction

A range of work in artificial intelligence has studied natural language generation in conversational dialog (Adiwardana et al. 2020; Zhang et al. 2018; Serban et al. 2016). Dialogs provide rich scenarios for understanding social interaction and real-world knowledge. But within a dialog, the question often arises: What if a character had said something else? Such counterfactuals can be powerful tools to learn from, and children tend to naturally explore such possibilities naturally by asking “what if” questions (Bates 1976).

Inspired by these aims, we introduce the large-scale MultiTalk Dataset of highly-branching dialogs, where each character gives several (10) responses to the full conversation history at that point. The high branching factor leads to many different conversational paths, which we manage by continuing the dialog for only a subset of responses. In total, the dataset contains over 320,000 sentences across 120 conversation trees, where each tree contains a unique prompt and pair of characters that alternate responses.

To study model performance on this dataset, we introduce a new scoring technique that accounts for the diversity of responses. Naively applying an evaluation metric like BLEU (Papineni et al. 2002) would assign high points to a model that repeatedly outputs a single well-scoring response. To

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

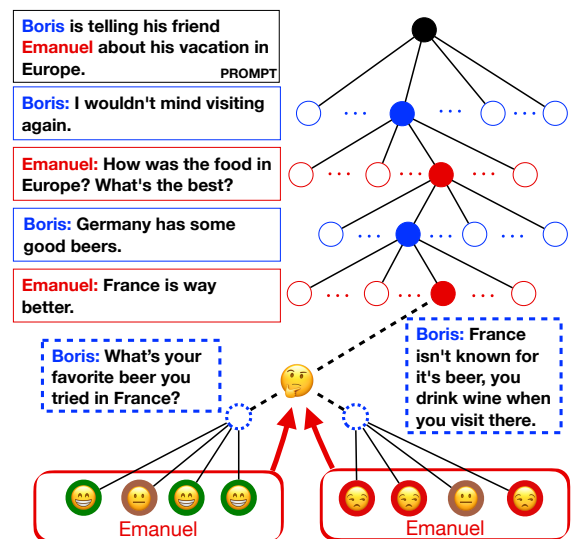


Figure 1: We introduce the highly branching *MultiTalk* dialog dataset, with over 320k sentences. We study several tasks in this new domain, including a *theory of mind* task, where a model estimates how listeners will react to its output.

account for this, we propose a simple scoring algorithm based on weighted bipartite graph matching. This technique dynamically matches candidates to references, assigning the maximum score to candidates while ensuring each reference is used only once. The technique requires no changes to models, can accept any individual scoring function to compare a candidate and reference (e.g., BLEU, BERT-Score (Zhang et al. 2020), etc.), and is solvable in polynomial time thanks to the Kuhn-Munkres algorithm (Kuhn 1955).

We also use the MultiTalk Dataset as an opportunity to explore models that attempt to optimize for particular dialog paths. First, we provide benchmark performances for a diverse response generation task, using MultiTalk Dataset as a testbed. Then, we propose two tasks to study controllable generation on a dialog tree using induced textual attributes. The second of these tasks emulates a *theory of mind* setting, where models must induce attribute-valued replies from a conversation partner who lacks explicit attribute knowledge.

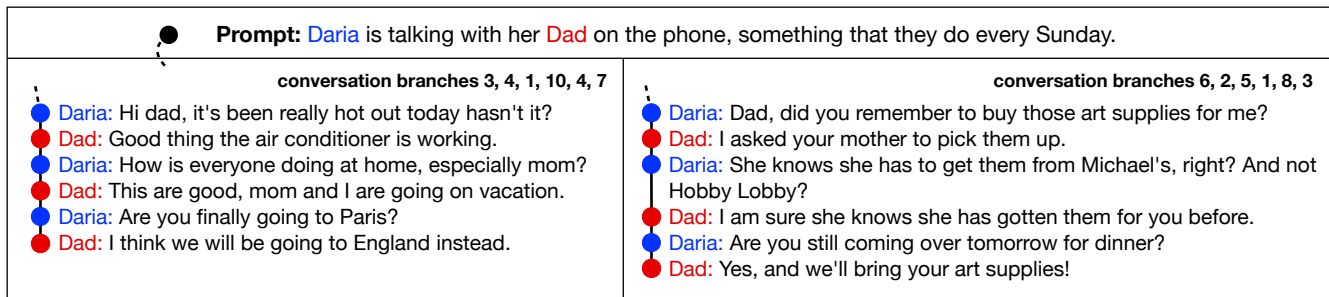


Figure 2: With a high (10) branching factor, a single prompt in the MultiTalk Dataset leads to a wide variety of conversations. Pictured here are two example paths from the training set that follow from a single prompt.

Branching factor (b)	10
Continuation factor (c)	3
Max response depth (d)	6
Total prompts (u_0)	120
Total sentences	320,804
Avg. sentences per prompt	2673
Avg. sentence length	8.5

Table 1: Statistics of our proposed MultiTalk Dataset.

MultiTalk Dataset

We collect and release¹ a large dataset of highly-branching written conversations. The dataset contains 320,804 individual responses in a conversation tree. Figure 2 shows two example conversation paths from the same prompt in the dataset. Summary statistics of the dataset are given in Table 1. Each dialog begins with a pre-specified scenario that establishes two characters and a setting. From this scenario, the conversation branches and deepens. Each character provides $b = 10$ potential responses or *branches* to the state of the conversation so far. From a randomly-chosen continuation subset $c = 3$ of those responses, the conversation continues with the next character. This continues to a maximum depth of $d = 6$ conversation turns. The following sections describe both the scenarios and conversation turns.

Scenarios: Prompts and Characters

We manually write $n = 120$ prompts. The prompts are motivated by everyday situations. We draw inspiration from a list of conversation questions for English language learners.² Each prompt is a single sentence.

The characters names are manually chosen using random name generators. We begin from a list of gender-balanced names, and assign a gender pronoun for each character. The list contains representations from names common among different cultural groups. The base prompt is combined with the two characters in order to form the scenario. We adjust the wording of scenarios to match the genders (e.g., *girl-friend* or *boyfriend*). Relationships are formed from all can-

¹<https://uwnlp.github.io/multitalk/>

²<http://iteslj.org/questions/>

Model	Diversity	Relevance	Fluency
GPT-2 XL	65.0%	98.8%	99.0%
Dialog Tree Model	85.6%	100.0%	99.0%

Table 2: How often the MultiTalk dataset (reference) annotation is rated as superior to top performing models by humans along three dimensions. Each comparison pits references against 10 model generations; 500 comparisons per model. *Takeaway:* the MultiTalk Dataset has significantly more diverse responses than top performing models, and overwhelmingly better relevance and fluency.

didate gender pairings. Some prompt/name pairings are purposefully left gender ambiguous.

Conversation Turns

At each turn, an annotator is provided with the full conversation context so far, and writes all $d = 10$ responses for the current speaker. A single annotator completes each step to minimize cognitive load; rather than read and characterise a partial set of existing responses, an annotator must only reason about the set of responses they will write. Annotators are free to mix both surface and semantic diversity (Tevet and Berant 2020). We perform manual quality control by checking a sample of work from each annotator and conversation tree during each round of annotation.

Dataset Quality and Utility

We briefly present two measurements of the dataset to assess its quality and utility. While these measurements use models and techniques described in later sections of the paper, we present them here as they primarily illustrate properties of the dataset.

Table 2 assesses dataset *quality*. It compares top performing models (§) against the dataset itself in a blind human evaluation. Models are tasked to perform *Diverse Generation* (§)—in other words, generate 10 responses to a context, just like the dataset. We observe the dataset is consistently more diverse, and resoundingly more relevant and fluent, than state-of-the-art neural language models.

Figure 3 demonstrates an aspect of the dataset’s *utility*. This graph controls the number of unique generations from




Figure 3: Varying the number of references in order to distinguish diverse generations. *Takeaway*: With only 1 reference, metrics have difficulty distinguishing between more and less diverse model outputs. With 10 references (our proposed dataset), models can be clearly stratified by diversity.

a model (different lines), and scores them against a varying number of references (x -axis) with a matching algorithm (§). We observe that having many references allows evaluations on the MultiTalk dataset to distinguish between more and less-diverse model outputs.

Scoring Algorithm

Evaluating Diverse References In dialogs with a high branching factor, multiple references are available for any candidate generation. How should we evaluate generations based on those references?

A naive approach is to simply run a metric like BLEU out-of-the-box. Figure 4 (top, left) illustrates an issue with this approach. Each generation can use the same reference to receive their score. This means that a model could achieve a high score simply by producing a single high-scoring response many times. This is not an issue in semantically-constrained generation tasks, as in machine translation, because all references attempt to say the same thing. However, for a dataset like ours where responses are intentionally diverse, this behavior is problematic.

Assignment Problem To address this issue, we can pose this task as an assignment problem. First, we add our set of references R and generations G to a graph as vertices. We can assign an edge (r, g) between each reference $r \in R$ and each generation $g \in G$ with a weight of $s(r, g)$ based on a scoring function of our choosing. For example, we could use the BLEU score between r and g edge weight between them. We now have a weighted bipartite graph, and we can find the assignment

$$f^* = \operatorname{argmax}_{f \in \mathcal{F}} \sum_{r \in R} s(r, f(r)) \quad (1)$$

where \mathcal{F} is the set of injective functions $f : R \rightarrow G$. The matching f^* produces the highest sum of edges that use each reference exactly once (Figure 4 top, center).

The Kuhn-Munkres algorithm, also known as The Hungarian Algorithm, is a polynomial time solution to that assignment problem (Kuhn 1955). While the algorithm is tra-

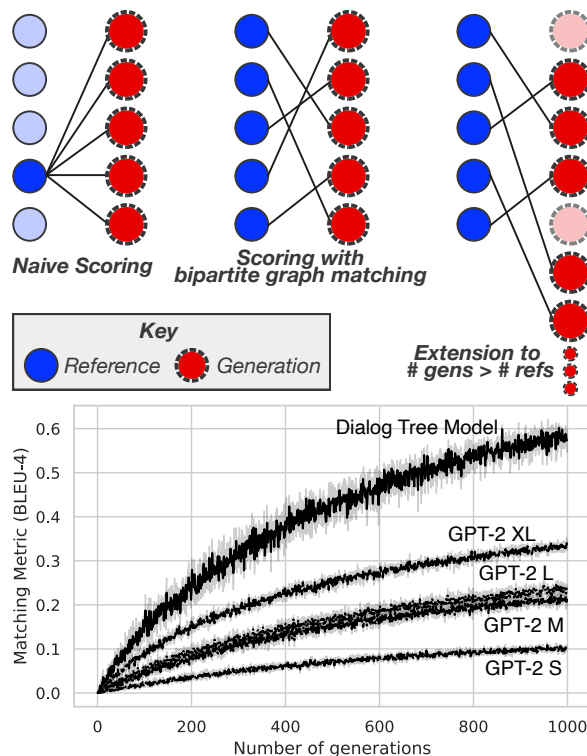


Figure 4: *Top*: By casting evaluation as a matching problem on a weighted bipartite graph, we can harness any pairwise scoring metric to evaluate a set of generations against a potentially diverse set of references. This technique extends naturally to a greater (or fewer) number of generations. *Bottom*: Match scoring against a diverse set of references with a tough metric (BLEU-4) smoothly scaling to significantly more generations.

ditionally used to assign a matching with minimum cost, we may simply invert scores beforehand and use it as usual.³

This technique has the advantage of being generous in its scoring of generations: each generation is allowed to match to any of the references, and this assignment changes optimally as new generations are introduced. Furthermore, the algorithm is fair, in that each reference is only allowed one usage. And lastly, the underlying models need not consider this process at all, nor specify any assignments themselves.

More Generations When $|G| = |R|$, an equal number of generations and references, f^* will be a bijection. But in practice, matching to even 10 references (e.g., in our dataset) can produce low scores depending on the metric used.⁴ Fortunately, because F is defined to be the set of *injective* functions, this approach extends naturally to scoring more generations $|G| > |R|$ (see Figure 4 top, right). The matching f^* will still use each reference $r \in R$ once, and will simply

³We use the *linear sum assignment* implementation in SciPy (<https://scipy.org/>).

⁴Intuitively, note that the space of potential reasonable responses in a conversational dialog is vast, certainly $\gg 10$.

pick the set of highest scoring $|R|$ generations from G .

Tasks

We use the MultiTalk Dataset to pose three tasks that we study in upcoming sections.

Diverse Generation The first task is a straightforward language generation task applied to the proposed dataset. A model is given a conversation history (u_0, \dots, u_{i-1}) and must produce n candidate generations for u_i . These generations are compared against the set of references for u_i using our proposed matching metric to account for diversity.

Emotion-Conditioned Generation For the second task, we first train an emotion classifier (details in §) to label each response in the dataset with an emotion $e \in E$. Then, the task is to generate $|E|$ responses for u_i given the conversation history and each emotion (u_0, \dots, u_{i-1}, e) . Models must succeed according to both the emotion classifier, which rates conformance to the provided emotions, as well as an *evaluative* language model fine-tuned on the dataset, which judges the coherence of the generation.

Theory of Mind Generation The final task we term a *theory of mind task*. In this task, a model is paired with a conversation partner: a fine-tuned language model that has no knowledge of emotions. A model must again produce $|E|$ responses for u_i given the conversation history and each emotion (u_0, \dots, u_{i-1}, e) . However, in this setting, it is not the model’s generation u_i , but its conversation partner’s reply u_{i+1} , that must match the provided emotion e .

Models

We present adaptations of pretrained language models to the three tasks we consider. In practice, we use the GPT2 architecture from (Radford et al. 2019), but our objectives are agnostic to the particular language model used.

Dialog Tree Model

For the diverse generation task, we perform a typical language model fine-tuning, with some changes to take into account the tree structure of the dialog data. We concatenate the prompt and full conversation history together for the full path leading to each utterance in the dataset $\mathbf{u} = [u_0; u_1; \dots; u_i]$. Then, we fine-tune a pretrained language model using the standard language modeling objective

$$p(\mathbf{u}) = p(w_1 \dots w_T) = \prod_{t=1}^T p(w_t | w_{<t}). \quad (2)$$

However, given the structure of our data, earlier utterances (e.g., u_1) are seen exponentially more frequently than later utterances (e.g., u_6). To prevent biasing the language model to utterances higher in the dialog tree, we compute loss for the model only for tokens in the final utterance:

$$\text{loss}(w_t) = \begin{cases} -\log(p(w_t | w_{<t})) & \text{if } w_t \in u_i \\ 0 & \text{otherwise} \end{cases}$$

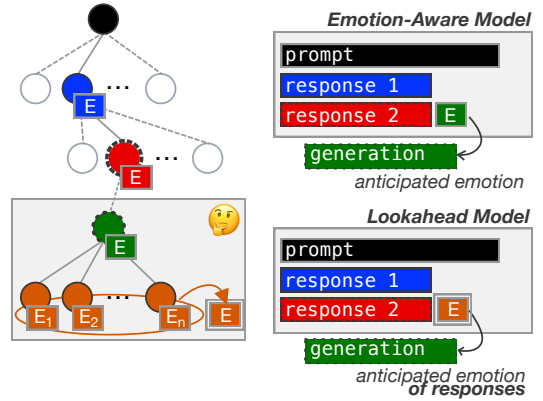


Figure 5: We study an Emotion-Aware Model, which anticipates the emotion of the next response, as well as a Lookahead Model, which anticipates the emotions of *responses* to its own generations.

Emotion-Aware Model

For the emotion-conditioned generation task, we develop a model that accounts for emotion state. The Emotion-Aware Model uses the same basic structure as the Tree Dialog Model, including only computing loss on tokens for the target utterance u_i and not the history $u_{<i}$. However, it is also given the emotion e of the latest response, which it conditions on when computing probabilities

$$p(w_t | w_{<t}, e). \quad (3)$$

A sketch of this model is given in Figure 5 (top right).

Lookahead Model

Finally, we present a lookahead model that addresses the theory of mind task. Its training objective models responses u_i that *result* in future responses u_{i+1} with the emotion e .

Intuitively, a model that would like to account for the emotion e of all responses below it needs to consider the emotions of all children. If we were interested in getting the emotional state of the complete tree below the utterance u_i , we would compute the emotion of all subtrees as a *depth-weighted estimate*, recursively defined as:

$$d(u_i) = \frac{1}{|c(u_i)|} \sum_{u_{i+1} \in c(u_i)} e(u_{i+1}) + \gamma d(u_{i+1}) \quad (4)$$

where $c(x)$ are the children of x , $e(x)$ is the (estimated) emotion of x , and γ is a discount factor either down or up-weighting utterances deeper in the conversation subtree. For practical purposes, we restrict our modeling to strictly optimize for the theory of mind task, and set $\gamma = 0$ to account only for the emotion of a response’s immediate children, and choose the emotion e that has the highest average in $d(u_i)$.

The per-token probability model for the Lookahead Model is then

$$p(w_t | w_{<t}, d(u_i)), \quad d(u_i) \in E \quad (5)$$

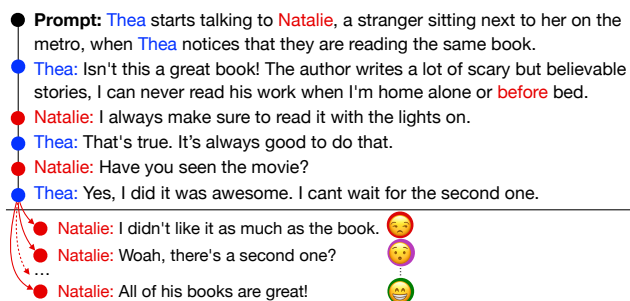


Figure 6: We induce emotion labels on the MultiTalk Dataset with a classifier trained on external emotion datasets. Here, we observe a diverse distribution of emotions for an example set of responses.

with $d(u_i)$ now the estimated emotion with the largest magnitude among the children of u_i .

Experiments and Results

Emotion Recognition

Classifier We train BERT-Large (Devlin et al. 2018) as an emotion classifier, fine-tuning it on multiple emotion-labeled corpora. We use two datasets from the EmotionX 2019 Challenge (Shmueli and Ku 2019) and part of the EMPATHETICDIALOGUES dataset (Rashkin et al. 2018). We use the seven emotions from the EmotionX datasets: joy, sadness, fear, anger, surprise, disgust, and neutral.

The accuracy of the trained emotion classifier is 82.6%. Figure 6 three classified responses from the dataset.

Label Distribution Because our dataset is collected without emotion labels, the classified emotion frequencies end up naturally imbalanced. While this has the advantage of potentially reflecting a more realistic distribution of conversational emotions (Poria et al. 2018), it creates an additional challenge during modeling. We thus also consider oversampling to smooth this distribution.

Oversampling For the purposes of the emotion-conditioned tasks, we create augmented versions of our dataset with two techniques. First, we add German and Russian back-translations of sentences with rare emotions (e.g., *disgust*) using (Ng et al. 2019; Ott et al. 2019). Second, we sample utterances from each emotion equally to create an emotion-balanced version of our dataset. We refer to this version of the dataset as *oversampled* when training models with it.

Training and Inference

We substitute all characters names with the special tokens `[speaker1]` and `[speaker2]`. We fine-tune GPT-2 M (345M params.) (Radford et al. 2019) on all training objectives (§).⁵ At inference time, we sample from all models using top-p sampling with $p = 0.9$ (Holtzman et al. 2019).

⁵GPT-2 Medium is the largest model we are able to train with our available resources.

Model	PPL	Matching Metric (BLEU-4)	Matching Metric (ROUGE-L f1)
GPT-2 S	51.63	0.06	3.62
GPT-2 M	41.95	0.14	3.87
GPT-2 L	35.67	0.16	4.03
GPT-2 XL	33.91	0.14	4.04
Dialog Tree Model	13.21	0.28	4.32

Table 3: Test performance for Task 1 (Diverse Generation).

Model	Diversity	Relevance	Fluency
GPT-2 S	3.91	2.05	2.81
GPT-2 M	3.78	2.67	3.42
GPT-2 L	3.68	2.86	3.58
GPT-2 XL	3.92	3.02	3.81
Dialog Tree Model	3.93	3.28	4.17
Gold Responses	4.06	4.49	4.61

Table 4: Human evaluation results (5-point Likert scale) for the Diverse Generation Task (Task 1). Bold numbers are a statistically significant improvement over the next highest result (t -test, $p < 0.001$). *Takeaway*: A fine-tuned model is able to improve upon Relevance and Fluency, while maintaining Diversity; still, there is significant room for improvement to human level, especially in Relevance.

Task 1: Diverse Generation

This task is a “standard” dialog response generation task, but where a model must produce a diverse set of responses to each provided context to be evaluated against 10 references. Our primary model of study is the Dialog Tree Model (§).

Automated Evaluation We compute the perplexity of all models on the test set. To evaluate model generations, we generate 200 responses from each model for 40 sampled contexts, and compare all model responses to the 10 gold generations for each context using our proposed matching algorithm (§), using BLEU-4 and ROUGE-L F1 (Lin 2004) as scoring functions.

Human Evaluation We have annotators evaluate the models’ generations. We give participants the prompt, conversation history, and a set of 10 model-generated responses, and ask them asked to rate three aspects of the responses (Diversity, Relevance, and Fluency) on a 5-point Likert scale. We collect 1000 ratings (of 10 generations each) per model.

Results Tables 3 and 4 show the results for the Diverse Generation Task, comparing the Dialog Tree Model (a finetuned GPT-2 M) against all out-of-the-box pretrained GPT-2 sizes. We observe an expected improvement on all metrics by carefully finetuning the model on the tree-structured corpus. In human evaluation (Table 4), the Dialog Tree Model significantly improves Relevance and Fluency, but ties at response Diversity. While it is initially surprising that a tree-trained model is not more diverse, the results are intuitive: the *standalone diversity* of generations is strongly influenced by the sampling procedure (Holtzman et al. 2019); and maximiz-

Model	Setup	Eval	Emotion classifier accuracy							Emotion Acc.	
		PPL	Average	Neutral	Joy	Sadness	Fear	Anger	Surprise	Disgust	Human Eval.
GPT-2 M		63.41	0.14	0.66	0.09	0.07	0.03	0.04	0.08	0.05	1.59
Dialog Tree Model		7.03	0.14	0.65	0.08	0.05	0.04	0.05	0.08	0.04	1.64
Retrieval	most likely	16.88	0.14	0.67	0.11	0.06	0.03	0.06	0.06	0.02	1.63
	w/ emotion e	17.64	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	<i>1.00</i>	2.13
Emotion-Aware Model	normal	7.53	0.51	0.84	0.60	0.38	0.40	0.39	0.56	0.42	2.18
	oversampled	8.33	0.54	0.76	0.64	0.42	0.48	0.48	0.60	0.40	2.11

Table 5: Test set performance for Emotion-Conditioned Generation task. For human evaluation (3-point Likert scale, rightmost column), bolded numbers are a statistically significant improvement from others (t -test, $p < 0.001$). *Takeaway*: The Emotion-Aware Model successfully balances generation coherency (as measured by *Eval PPL*) with adhering to the specified emotion.

ing diversity, *as well as* measures of quality, indicate good generations (Caccia et al. 2020). Promisingly, the matching metric (Table 3) captures the strongest (fluent and relevant) model, even when it exhibits the same diversity as others.

We present these results as baseline numbers for the diverse generation task, and leave efforts to explicitly target even greater generation diversity to future.

Task 2: Emotion-Conditioned Generation

In this task, a model is given the complete context and desired emotion (u_0, \dots, u_{i-1}, e), and should generate a response u_i with emotion e . We introduce the Emotion-Aware Model for this task (§).

Baselines In addition to GPT-2 Medium (both out-of-the-box and fine-tuned), we also use retrieval baselines. The first retrieval baseline finds the training set context with the highest similarity. We average GloVe embeddings (Pennington, Socher, and Manning 2014)⁶ of all words in the contexts and score with cosine similarity. The second retrieval baseline performs the same similarity calculation, but is restricted to choosing from utterances with the desired emotion e .

Automated Evaluation We sample 20 contexts from each prompt in the test set, and generate 7 responses per prompt, one for each emotion $e \in E$. Model generations are evaluated both on whether they are classified as emotion e (emotion acc.), and with a PPL score given to them by an emotion-agnostic LM fine-tuned on the dataset (coherence).

Human Evaluation We also measure human ratings of emotion accuracy on a 3-point Likert scale for 1000 generations per model. Whereas in Task 1, workers saw a set of 10 generations (to evaluate diversity), here workers rate each generation individually.

Results Table 5 shows the experimental results for the Emotion-Conditioned Generation task. While the retrieval baseline is trivially able to retrieve a response with the correct emotion, it suffers by having significantly worse perplexity as judged by the independent fine-tuned LM, and lower relevance and fluency (by humans). The Emotion-Aware model is able to balance coherent replies (low Eval. PPL) while accurately generating responses with the conditioned emotion (both automatic and human-evaluated).

⁶Specifically: 40B, 300d GloVe vectors.

Training the Emotion-Aware Model on an oversampled subset leads to a slight improvement in emotion adherence at the cost of higher perplexity.

Task 3: Theory of Mind Generation

In this task, models generate a response u_i that will lead to a conversation partner to a response u_{i+1} with emotion e . We propose the Lookahead Model for this task (§).

Baselines We run all of the baselines and models from Task 2 as baselines for this task. Though the goal is different, it is intuitively plausible that simply generating a response u_i with emotion e may cause a conversation partner to *mirror* that same emotion e in their reply u_{i+1} .

In addition, we create an emotion transition matrix T from the training data that contains which emotions are likely to lead to each other. We denote $T[\rightarrow e]$ the emotion of utterance u_i most likely to *lead* to emotion e for the reply u_{i+1} . We use $T[\rightarrow e]$ as additional variants both for the retrieval baseline, and for conditioning the Emotion-Aware Model.

Evaluation As with Task 2, we sample 20 contexts from each prompt in the test set, and generate 7 responses per prompt, one corresponding to each emotion $e \in E$. Then, for each response, we have a fine-tuned LM as a conversation partner continue the conversation by generating a u_{i+1} , which the emotion classifier then assesses for a match to e . The fine-tuned LM also scores the perplexity of the replies u_i . We report the average of 5 runs for all numbers to mitigate generation randomness.

Results Table 6 presents the experimental results for the Theory of Mind task. Unlike the previous two tasks, this task does not present a clear winner. The Lookahead Model, whose training scheme most closely mimics the task, achieves the lowest evaluated perplexity according to the external model (6.62). However, it struggles as much as emotion-agnostic baselines to influence emotion states of its conversation partner (0.14). The strongest models at this task (0.16) are *mirroring* approaches: the retrieval baseline and Emotion-Aware Model, both using the conversation partner’s target emotion e as their parameter.

Even the oracle—which uses responses from the validation set known to lead to the desired emotion—is not far ahead (0.17). This results hints at the significant challenge of this task, which we discuss in the next section.

Model	Setup	Eval									
		PPL	Average	No-neutral	avg.	Neutral	Joy	Sadness	Fear	Anger	Surprise
GPT-2 M		67.79	0.14	0.04	0.74	0.07	0.05	0.02	0.03	0.06	0.04
Dialog Tree Model		7.11	0.15	0.05	0.77	0.09	0.05	0.02	0.03	0.05	0.03
Retrieval	most likely	16.72	0.14	0.04	0.74	0.06	0.05	0.02	0.04	0.06	0.03
	w/ e	18.33	0.16	0.06	0.73	0.13	0.06	0.04	0.05	0.04	0.07
	w/ T[$\rightarrow e$]	17.90	0.15	0.05	0.77	0.08	0.04	0.02	0.04	0.07	0.04
Emotion-Aware Model	w/ e	7.68	0.16	0.06	0.74	0.11	0.05	0.02	0.06	0.07	0.03
	w/ T[$\rightarrow e$]	7.21	0.14	0.04	0.73	0.09	0.04	0.03	0.03	0.05	0.03
	w/ e + oversampled	8.47	0.15	0.06	0.75	0.10	0.07	0.03	0.04	0.04	0.04
	w/ T[$\rightarrow e$] + oversampled	8.50	0.15	0.04	0.78	0.07	0.05	0.03	0.03	0.05	0.03
Lookahead Model	normal	6.62	0.15	0.05	0.76	0.08	0.05	0.03	0.03	0.05	0.03
	oversampled	7.30	0.14	0.04	0.75	0.06	0.05	0.03	0.03	0.06	0.02
Oracle		13.90	0.17	0.07	0.78	0.13	0.09	0.04	0.04	0.09	0.03

Table 6: Test set performance for Theory of Mind Generation task. All numbers are averaged over 5 runs. *Takeaway:* While the Lookahead Model achieves the lowest perplexity, all models struggle to reliably influence the emotion of the next prediction. Approaches that mirror the desired emotion e reliably perform best, indicating this may be a hard-to-beat strategy.

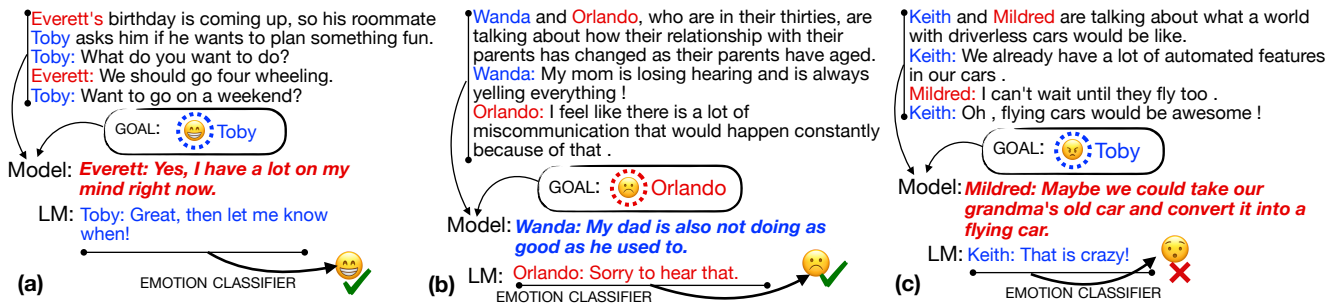


Figure 7: Three generation examples from the Lookahead Model for the Theory of Mind Task.

Qualitative Analysis

Figure 7 gives three examples from the Lookahead Model’s output on the dev set, which illustrate some of the main response categories. In (a), the model performs what appears to be relatively complex maneuver: by outputting a relatively melancholy response (*Yes, I have a lot on my mind right now*), the conversation partner appears to “cheer them up” by being excited about their plans. Example (b) shows *mirroring*: the model says something sad, and the LM replies with something that is also sad. The approach served other models well at directing conversations. Finally, example (c) illustrates a common problem: it can be difficult to consistently direct an LM conversation partner. While a language model fine-tuned on our dataset is exposed to a variety of responses for many contexts, this same diversity can lead to challenges reliably influencing the behavior of the model. Investigating the role of cooperative agents (Lewis et al. 2017) in a diverse dialog setting may prove an interesting avenue of future work.

Related Work

Dialog Datasets and Personas Previous dialog datasets have been created from scraped conversations of social media, such as Twitter (Ritter, Cherry, and Dolan 2010) and

Weibo (Shang, Lu, and Li 2015). Other recent work has proposed that Chit-chat models become more engaging when endowed with a consistent “persona” (Li et al. 2016; Zhang et al. 2018; Shuster et al. 2019). Several datasets have taken advantage of this idea and included personas for the participants. PERSONA-CHAT (Zhang et al. 2018) dialogues revolve around fictional speakers’ personas. Our dataset contains implicit personas from the conversation prompts, rather than explicit attributes assigned to each character.

Dialogue Generation Zhou and Wang (2017); Zhou et al. (2017); Wang and Wan (2018); Li et al. (2019) generated responses with corresponding emotions using Conditional Variational Autoencoders (CVAEs) (Sohn, Lee, and Yan 2015), GANs (Goodfellow et al. 2014), and Deep RL. The core of our approach, like other recent work, is strategically finetuning large pretrained LMs (Radford et al. 2019).

Conclusion

We present the large-scale, highly branching conversational MultiTalk Dataset. In studying this dataset, we propose a multi-reference scoring algorithm, as well as models that look ahead to optimize for future conversation states. We hope the introduction of our data and methods will spur more exciting work in multi-path conversational dialog.

References

- Adiwardana, D.; Luong, M.-T.; So, D. R.; Hall, J.; Fiedel, N.; Thoppilan, R.; Yang, Z.; Kulshreshtha, A.; Nemade, G.; Lu, Y.; et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977* .
- Bates, E. 1976. *Language and context: The acquisition of pragmatics*. Academic Press.
- Caccia, M.; Caccia, L.; Fedus, W.; Larochelle, H.; Pineau, J.; and Charlin, L. 2020. Language GANs Falling Short. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=BJgza6VtPB>.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Holtzman, A.; Buys, J.; Du, L.; Forbes, M.; and Choi, Y. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751* .
- Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2(1-2): 83–97.
- Lewis, M.; Yarats, D.; Dauphin, Y. N.; Parikh, D.; and Batra, D. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125* .
- Li, J.; Galley, M.; Brockett, C.; Spithourakis, G. P.; Gao, J.; and Dolan, B. 2016. A persona-based neural conversation model. *arXiv preprint arXiv:1603.06155* .
- Li, J.; Sun, X.; Wei, X.; Li, C.; and Tao, J. 2019. Reinforcement Learning Based Emotional Editing Constraint Conversation Generation. *arXiv preprint arXiv:1904.08061* .
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Ng, N.; Yee, K.; Baevski, A.; Ott, M.; Auli, M.; and Edunov, S. 2019. Facebook FAIR’s WMT19 News Translation Task Submission. *arXiv preprint arXiv:1907.06616* .
- Ott, M.; Edunov, S.; Baevski, A.; Fan, A.; Gross, S.; Ng, N.; Grangier, D.; and Auli, M. 2019. fairseq: A Fast, Extensible Toolkit for Sequence Modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. URL <http://www.aclweb.org/anthology/D14-1162>.
- Poria, S.; Hazarika, D.; Majumder, N.; Naik, G.; Cambria, E.; and Mihalcea, R. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508* .
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8): 9.
- Rashkin, H.; Smith, E. M.; Li, M.; and Boureau, Y.-L. 2018. Towards empathetic open-domain conversation models: A new benchmark and dataset. *arXiv preprint arXiv:1811.00207* .
- Ritter, A.; Cherry, C.; and Dolan, B. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 172–180.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. *arXiv preprint arXiv:1503.02364* .
- Shmueli, B.; and Ku, L.-W. 2019. SocialNLP EmotionX 2019 Challenge Overview: Predicting Emotions in Spoken Dialogues and Chats. *arXiv preprint arXiv:1909.07734* .
- Shuster, K.; Humeau, S.; Hu, H.; Bordes, A.; and Weston, J. 2019. Engaging image captioning via personality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12516–12526.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, 3483–3491.
- Tevet, G.; and Berant, J. 2020. Evaluating the Evaluation of Diversity in Natural Language Generation. *arXiv preprint arXiv:2004.02990* .
- Wang, K.; and Wan, X. 2018. SentiGAN: Generating Sentimental Texts via Mixture Adversarial Networks. In *IJCAI*, 4446–4452.
- Zhang, S.; Dinan, E.; Urbanek, J.; Szlam, A.; Kiela, D.; and Weston, J. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? *arXiv preprint arXiv:1801.07243* .
- Zhang, T.; Kishore*, V.; Wu*, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=SkeHuCVFDr>.
- Zhou, H.; Huang, M.; Zhang, T.; Zhu, X.; and Liu, B. 2017. Emotional chatting machine: Emotional conversation generation with internal and external memory. *arXiv preprint arXiv:1704.01074* .
- Zhou, X.; and Wang, W. Y. 2017. Mojtalk: Generating emotional responses at scale. *arXiv preprint arXiv:1711.04090* .