

Grounding Conversations with Improvised Dialogues

Hyundong Cho and Jonathan May

Information Sciences Institute
University of Southern California
{jcho, jonmay}@isi.edu

Abstract

Effective dialogue involves grounding, the process of establishing mutual knowledge that is essential for communication between people. Modern dialogue systems are not explicitly trained to build common ground, and therefore overlook this important aspect of communication. Improvisational theater (improv) intrinsically contains a high proportion of dialogue focused on building common ground, and makes use of the *yes-and* principle, a strong grounding speech act, to establish coherence and an actionable objective reality. We collect a corpus of more than 26,000 *yes-and* turns, transcribing them from improv dialogues and extracting them from larger, but more sparsely populated movie script dialogue corpora, via a bootstrapped classifier. We fine-tune chit-chat dialogue systems with our corpus to encourage more grounded, relevant conversation and confirm these findings with human evaluations.

1 Introduction

For humans, dialogue is fundamentally a collaborative, cooperative process by which *partners* coordinate via *turns* or *acts* to jointly construct a common *world state* (Bohm and Nichol, 2004). Without coordination, partners may establish different or conflicting world states, leading to solipsism in the best case and conflict in the worst. Clark and Schaefer (1989), describe five dimensions of *grounding*, by which partners cooperate to establish *common ground*, or a shared world state. The dimension of “initiation of next relevant contribution” is the most effective of these in expressing understanding of an ongoing dialogue, and yet is the least observed in dialogue systems.

Improvisational theater (improv) is a form of theater in which most or all of what is performed is unscripted, created spontaneously by the actors in real time. Because the performance is not scripted and there is typically little to no scenery or other es-

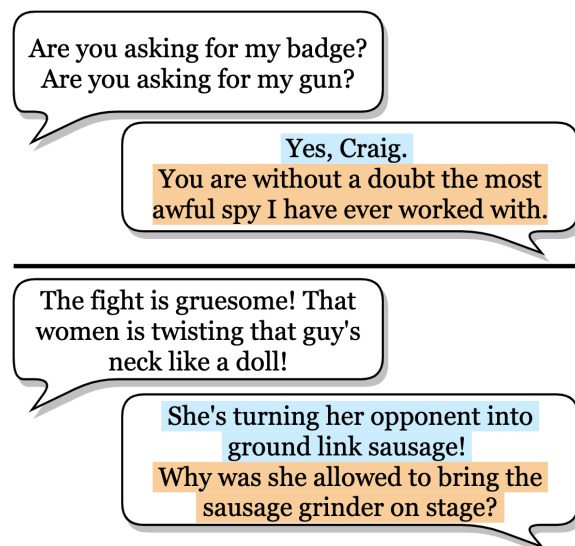


Figure 1: Explicit (top) and implicit (bottom) examples of *yes-and*s in the SPOLIN corpus. The text highlighted in light blue reflects acceptance of the context established in the prompt (“yes”) and the text highlighted in orange initiates a new relevant contribution to the dialogue (“and”).

tablished environment,¹ there is no objective reality that can naturally ground the scene. Hence, actors must mainly rely on dialogue in order to build a coherent scene and progressively establish a common world view. This necessitates accelerated use of the “initiation of next relevant contribution,” which in improv is known as the *yes-and* principle. The *yes-and* principle is a rule-of-thumb that suggests that a participant should accept the reality of what the other participant has said (“yes”) and expand or refine that reality with additional information (“and”). Since actors consciously abide by this principle during improv performances, there is a high proportion of these turns embedded in improv dialogue, which helps ensure scenes are coherent and interesting.

¹except for, on occasion, external stimulus such as a suggestion from the audience

Open-domain neural dialogue systems, by contrast, specifically lack coherence and interest- ingness. They commonly repeat previous utterances (Li et al., 2016c) or generate non-committal, generic statements such as *I don't know* that are log- ically coherent as a response but preempt further conversation (Sordoni et al., 2015; Serban et al., 2015; Li et al., 2016a). Either of these develop- ments leads to a conversational black hole and discourages participation in further dialogue turns. This is a critical shortcoming for open-domain dia- logue agents, which, unlike task-oriented dialogue systems, are not guided by specific objectives other than entertainment (Huang et al., 2020). It would behoove such systems to adopt the strategies im- provisers include by habit in their dialogues and, consequently, incorporating improv acts should be a key focus for the dialogue community.

Yet, to the best of our knowledge, this has not been previously done. There has been work in applying improv to build believable agents that in- teract with humans (Bruce et al., 2000; Winston and Magerko, 2017) or generate improvised stories (Martin et al., 2016), but development of improv- capable systems in the neural era is largely absent, stymied, we suspect, by the lack of substantial cor- pora. This is unsurprising; while improv speech acts such as *yes-and* are crucial in all dialogues, they are only highly concentrated in improv dia- logues. And improv dialogues are quite difficult to collect; research collections (Busso and Narayanan, 2008) have been far too small to be useful in the modern ML era. The art form has historically been mostly ephemeral, performed live in regional venues on shoestring budgets and rarely recorded.² Transcripts are all but absent and mainstream media products are rare.³ However, the liberalization of high quality audio *podcasts* since 2014 has enabled the availability of a long tail of niche products, improv included (McHugh, 2016).

²The art form has long roots, extending to the Italian *Com- media dell'arte* tradition from the 16th century and farces from the Roman era, but we constrain our focus to the post-20th century form developed and championed by e.g. Keith Johnstone (Johnstone, 2017), Del Close (Halpern et al., 1994), and our corpus' namesake, Viola Spolin (Spolin et al., 1986). Spolin was the originator of *Theater Games*, acting exercises that encourage the development of specific theatrical skills. As our corpus is similarly designed to elicit specific skills, we backronym it in recognition of her influence.

³One exception, the long-running TV show *Whose Line Is It Anyway*, has, despite a large number of episodes, sur- prisingly little continuous improvised dialogue, due to the rapid-fire nature of the program.

Therefore we set our objective as collecting *yes- and*-type dialogue pairs (*yes-ands*) to enable their modeling by corpus-driven dialogue systems. We mine podcasts and existing movie script corpora for dialogue that abides by the *yes-and* principle and extract dialogue pairs from these sources to build the *Selected Pairs Of Learnable Improvisa- tion* (**SPOLIN**) corpus. **SPOLIN** is a collection of more than 26,000 English dialogue turn pairs, each consisting of a *prompt* and subsequent *response*, which abide by the *yes-and* principle, though in diverse manners. Examples of *yes-and* type dia- logue pairs collected for **SPOLIN** are in Figure 1. The corpus is substantial enough to be usable for fine-tuning existing dialogue models to encourage more *yes-and* behavior, and beyond that may prove a valuable knowledge base for empirical sociolin- guistic studies on this dialogue act.

Our contributions are summarized as follows:

- We carefully curate *Selected Pairs Of Learn- able Improvisation* (**SPOLIN**), the first large- scale corpus of *yes-and* dialogue acts, sourced from improv and movie dialogues.
- We iteratively build a high-precision *yes-and* classifier, which we use to mine additional *yes- ands* from dialogue corpora with high volume but low *yes-and* density.
- We fine-tune existing open-domain conversa- tional models with our corpus and confirm via human evaluations that this approach im- proves creative grounding.
- We release our models and data for public use, including a 64,000 turn pair extension of the core **SPOLIN**, at [https://justin-cho. com/spolin](https://justin-cho.com/spolin).

2 Data Collection

Our data collection has five stages:

1. Manually extract *yes-ands* from a rich corpus of improv to obtain an initial set of *yes-ands*.
2. Construct a *yes-and classifier* from the corpus of collected *yes-and* data and negative exam- ples.
3. Use the classifier from step 2 to automatically extract *yes-and* candidates from a much larger but sparser dialogue corpus.

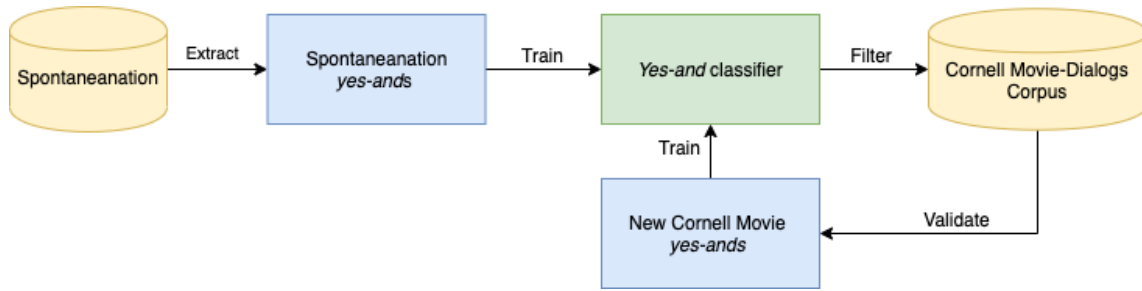


Figure 2: An illustration of the *yes-and* collection workflow. The core **SPOLIN** corpus comprises *Spontaneation yes-and*s and *Cornell yes-and*s (in blue boxes). However, **SPOLIN** can be augmented by including other general-purpose dialogue corpora in place of *Cornell* in this workflow, as described in Section 5.

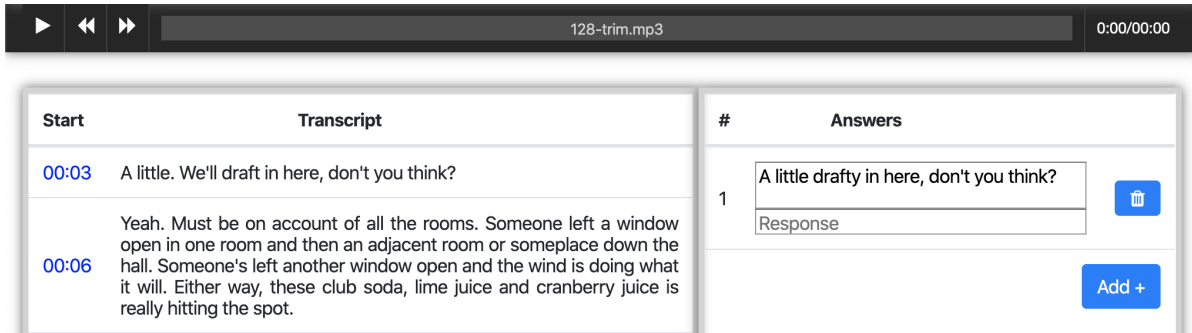


Figure 3: Amazon Mechanical Turk interface for transcribing *yes-and*s from *Spontaneation* episodes. Approximate transcriptions with speaker turns and time stamps generated from Amazon Transcribe are provided for additional guidance.

4. If necessary, manually validate candidates before adding them to the *yes-and* corpus.
5. Repeat from step 2 as needed.

An overview of this process is shown in Figure 2.

2.1 Core *yes-and* Collection from *Spontaneation*

We select the *Spontaneation*⁴ podcast as a source of concentrated *yes-and*s for its relatively noise-free recording quality and high-quality volume of broad domain improv dialogue. Each episode of this podcast includes an approximately 30 minute-long improv session performed by professional improvisers. Over its 201 episodes, we identified a total of 43K lines of useful spoken dialogue.

Given the confluence of a lack of objective reality, and uninterrupted multiturn dialogue, the improvisers mostly abide by the *yes-and* principle, and therefore *Spontaneation* is a rich resource for natural, high-quality *yes-and*s. As it exists only in audio form, and automatic transcription services are too noisy for high quality annotation use, we

⁴<https://www.earwolf.com/show/spontaneation-with-paul-f-tompkins/>

ask Amazon Mechanical Turk workers (Turkers) to listen to the improv sessions, view Amazon Transcribe preliminary transcriptions, and re-transcribe all of the *yes-and*s that they hear using our transcription interface, shown in Figure 3. The interface is based on oTranscribe, an open-source transcription service. Although the quality of transcriptions is poor, we find that including them assists the Turkers in identifying speaker turns and also understanding parts that are sometimes incomprehensible without helping context.

2.1.1 Recruiting Quality Crowdworkers for Difficult Annotation Tasks

One of the main challenges for the data collection process is to recruit competent Turkers who are able to develop a good understanding of the *yes-and* principle. We actively recruit potential annotators to our task by inviting denizens of the sub-Reddit TurkerNation, rather than simply inviting workers through Amazon’s native task posting interface based on HIT approval rate and total number of HITs approved. Our approach enables more human-level engagement, making it easier to determine Turkers’ English fluency and experience with improv. To ensure their competence,

Iteration	1	2	3	4
<i>Spontaneanation</i> +	10,459	10,459	10,459	10,459
<i>Spontaneanation</i> -	-	-	3,225	5,587
<i>Cornell</i> +	-	3,327	8,464	12,220
<i>Cornell</i> -	10,459	13,786	15,698	17,092
Total Training Samples	20,198	27,572	37,846	45,358
Dev Set Acc. (Spont)	80.9%	73.6%	71.6%	73.0%
Dev Set Acc. (<i>Cornell</i>)	52.2%	56.8%	62.1%	64.5%
Confidence Threshold	95%	70%	50%	50%
New Extraction Volume	12,360	12,802	5,150	3,515
New Proportion of <i>yes-and</i> s	26.9%	44.0%	72.9%	78.4%

Table 1: Iterative data collection results over *Cornell*. + indicates *yes-and*s and - indicates non-*yes-and*s. These counts exclude 500 turns collected from each of *Spontaneanation* and *Cornell* to form the validation set. The New Extraction Volume row indicates the new number of *yes-and* candidates identified with the given confidence threshold, and the New Proportion of *yes-and* row show as a percentage how many of these candidates were indeed evaluated as *yes-and*s by Turkers. The proportion of *yes-and*s increases after each iteration despite the lower confidence threshold used to filter the new predictions with the updated classifier.

Turkers first read *yes-and* guidelines (in the appendix) then demonstrate their level of understanding through qualification Human Intelligence Tasks (HITs), which test whether the candidates can identify if a *yes-and* exists in a 30 second audio segment and transcribe it if there is one. s

Even after inviting Turkers for the actual HIT of transcribing *yes-and*s, we frequently monitor the quality of the data they collect and provide feedback for incorrectly identified *yes-and*s. Apart from base pay for each episode they work on, we provide incentives for extracting more *yes-and*s. The pay for our HITs averages well above California minimum wage. From all of the episodes, we extract 10,959 *yes-and*s, indicating about 25% of the total number of dialogue turns in *Spontaneanation* are *yes-and*s.

2.2 Guided Extraction from the *Cornell* *Movie-Dialogs* Corpus

Although larger than any improv corpus, let alone *yes-and* corpus known to date, we seek to increase our corpus volume from 10,959 turn pairs. The *Cornell* *Movie-Dialogs* Corpus (Danescu-Niculescu-Mizil and Lee, 2011, *Cornell*) contains 304,713 turns, nearly an order of magnitude more than *Spontaneanation*, and it is one of the closest in domain to improv among existing dialogue datasets. However, a sample annotation of 300 randomly selected turn pairs by Turkers reveal only 11.1% of pairs are *yes-and*s. We thus use the already-collected

*yes-and*s to probe *Cornell* for likely candidates, to speed the search process. Recent developments of language models pre-trained on massive text data enable the training of high-accuracy models for down-stream tasks even with a small number of samples, by leveraging the contextualized embeddings that these models learn (Devlin et al., 2019; Radford et al., 2019). We thus fine-tune an initial BERT-based sequence classifier based on the implementation of Wolf et al. (2019a) with the *yes-and*s from the *Spontaneanation* episodes to determine if a given dialogue pair is a *yes-and*, using a high threshold (initially, a 95% probability of being *yes-and*) to bias for precision. We ask Turkers to validate the turn pairs identified by the classifier and add the validated pairs to our *yes-and* corpus. This procedure can be iterated.

For the first iteration, we train the classifier with a balanced number of non-*yes-and*s chosen by random sampling from *Cornell*, a reasonable assumption due to the relatively low concentration of *yes-and*s observed. The same Turkers that extracted *yes-and*s from *Spontaneanation* are invited to validate the *yes-and* candidates filtered out by the classifier using the interface shown in Figure 4. In order to ensure consistent annotation standards among Turkers, they are given a small number of overlapping HITs against which we validated. For 90 samples of unfiltered *yes-and* candidates from *Cornell*, the two workers yield a reasonably high Cohen’s κ value of 0.74. Turkers are paid at rates consistent with their rates on the extraction-from-*Spontaneanation* task.

After the set of *Cornell* *yes-and* candidates are validated, the *yes-and*s and non-*yes-and*s are added to the training set to train a new classifier, and the same process is repeated. We hold out 500 dialogue pairs from each subcategory (i.e. *Spontaneanation* *yes-and*s) as the development set for monitoring the classifier’s performance after each iteration. We incrementally lower the classification threshold for choosing a *yes-and* candidate as the classifier improved. We set this threshold on each iteration except for the first by retrospective evaluation of the classifier on the actual *yes-and* candidates’ labels from previous iterations. The threshold with the highest F1 score is chosen to filter new *yes-and* candidates to be validated.

We balance each progressively larger corpus with negative sample turn pairs, which are either randomly selected from *Cornell* (round 1), selected

Indicate whether each of these dialogue exchanges is a "Yes, and..." or not:

[Click for instructions](#)

Q1

Prompt
Great. Is that what you're going to say when I put you on the stand?

Response
No. When you put me on the stand, I'll say your client is catatonic and exhibits classic symptoms of a schizophrenic, sociopathic personality. And he doesn't sleep.

Yes No Typo/Fix

Figure 4: Amazon Mechanical Turk interface for validating *yes-and* candidates determined by the *yes-and* classifier. Turkers are asked to correct minor errors in grammar, spelling, and punctuation for qualifying *yes-and* candidates, which are then categorized as ‘Typo/Fix.’

from the rejected-but-extracted turn pairs from *Cornell* (round 2 and later), or sampled from non-*yes-and* turn pairs in *Spontaneation* formed by random coupling of prompts and responses of the *Spontaneation yes-ands* (round 3 and later). The latter forces the classifier to make decisions based on semantic features relevant to a *yes-and* instead of only stylometric features in *Spontaneation yes-ands*. We stop this iterative process after four rounds, when fewer than 5,000 new *yes-and* candidates are identified by the classifier, yielding a total corpus size of 26,435 *yes-ands* and 23,938 negative samples. An overview of this iterative process is summarized in Table 1. The negative sampling procedure, while somewhat ad-hoc, ultimately provides a mix of turn pairs from both corpora that is sufficient to allow extraction of *yes-ands* from new corpora at high precision rates, and is sufficient for our goals.

2.3 Additional Notes on *yes-and* Criteria

Although the concept of a *yes-and* is easy to define and understand, there are borderline cases between a *yes-and* and a non-*yes-and* that make the validation phase more difficult than originally expected. One of the cases that confused Turkers in the earlier stages of data collection is the case of *yes-buts*. A *yes-but* is a *yes-and* with a response that is coherent with the provided reality, but does not appear to provide an affirmative acceptance of a suggestion given in the prompt. These are different from *contradictions* that do not align with the previously established reality. In addition, there are instances where the response is a *yes-and*, but is accepted by a speaker other than the one to whom the prompt is directed. Some *yes-and* responses initiates a re-

pair of a problem encountered while accepting the prompt, due to a confusion or a possible inconsistency, by asking for clarification (Clark and Schaefer, 1989). While these responses may not strictly *establish* more detail, they provide information for ultimately establishing new information. We elide these edge cases under the umbrella category *yes-and* in **SPOLIN** as they further our top-level goal of providing relevant, actionable turn responses. Examples of some of these subtle differences are shown in Table 2.

3 Dataset Analysis

In order to provide a better understanding on the characteristics of our corpus, we annotate 200 *yes-ands* and 200 non-*yes-ands* in **SPOLIN**'s development set to categorize them into specific *yes-and* or non-*yes-and* types.

We classify *yes-ands* into explicit *yes-ands*, implicit *yes-ands*, or *yes-buts*. Only 15% of all *yes-ands* are explicit *yes-ands*, containing phrases such as “Yeah” or “Sure” that reflects agreement. Even with such phrases, identifying explicit *yes-ands* is not a trivial task because it requires semantic understanding of the relevance of the context established in the prompt and that introduced in the response. In fact, there are non-*yes-ands* that contain phrases affirming agreement but have no contributions or have new contributions that lack relevance. The majority (78%) of *yes-ands* are implicit *yes-ands*, meaning that the agreement is implied, often in a subtle manner. The remaining 7% are *yes-buts*.

Non-*yes-ands* are divided into *contradictions* and *others*. Most of the non-*yes-and* were *others*, as only 5% of candidates extracted from *Cornell* are *contradictions*, which are dialogue pairs with

Type	Example	%
Explicit	P: Does this map look homemade to you?	15%
	R: Yeah, it looks like someone without a grasp of English drew it.	
<i>yes-and</i>	P: Alright, pull up that plate so I can take a picture.	78%
	R: Sorry, the coleslaw is definitely giving off a lot of glare.	
<i>yes-but</i>	P: We all must say the chant that we say to the king. R: No, it's too erotic, please don't.	7%
Contra	P: Hey, hey, aren't you afraid you'll burn out a tonsil?	5%
	R: Tonsil? Me? No! Me burn a tonsil? My tonsils won't burn - As life's corners I...	
non- <i>yes-and</i>	Other P: I feel different right now. R: You wait and see. You're going to marry a big hero!	95%

Table 2: Examples and proportions of *yes-and* and non-*yes-and* types from annotations of 200 *yes-and*s and non-*yes-and*s in **SPOLIN**'s development set. Determining whether a given dialogue pair is a *yes-and* or not is a non-trivial task, as the agreement or contradiction of the previous dialogue turn's context is usually implicit.

	<i>yes-and</i> s	non- <i>yes-and</i> s
<i>Spontaneation</i>	10,959	6,087*
<i>Cornell</i>	15,476	18,351
Total	26,435	24,438

Table 3: Composition of **SPOLIN**, including the development set. *yes-and*s and non-*yes-and*s from *Cornell* are validated by Turkers. **Spontaneation* non-*yes-and*s are sampled from random combination of prompts and responses in *Spontaneation yes-and*s to balance the dataset for training the classifier in the final iteration, as shown in the last column of Table 1.

a response that actively negates the reality in the prompt. *Others* encompass any dialogue pairs with a response that lacks coherence to the prompt or adds no or minimal contributions. The distribution and examples of different types of *yes-and*s and non-*yes-and*s are shown in Table 2.

The main focus of our work is on *yes-and*s, but we provide non-*yes-and*s as part of **SPOLIN** for those interested in training their own classifiers. The negative samples are collected using the methods described in Section 2.2. The composition details of **SPOLIN** are shown in Table 3.

4 Experiments

To evaluate the effect of **SPOLIN** on generating *yes-and* responses and thus improving generated dialogue quality, we train a common architecture with a variety of fine-tuning data configurations, both with and without **SPOLIN**. Specifically, for each data configuration we fine-tune a doublehead GPT-2 model (117M-parameter version based on

the implementation by Wolf et al. (2019b)), which achieved state-of-the-art performance on *Persona-chat* for the ConvAI-2 dialogue competition (Zhang et al., 2018). We fine-tune the models using two learning objectives, which we weigh equally in calculating loss:

1. Predicting the next word.
2. Predicting the next correct candidate that best fits the dialogue given the dialogue history.

The language modeling component uses pre-trained weights from OpenAI, while the candidate classification head is trained from scratch. For evaluation, we use the language modeling component of the fine-tuned model to generate single-turn responses for the *yes-and* prompts in the development set. We use nucleus sampling (Holtzman et al., 2020) for the decoding step to keep only the top tokens with a cumulative probability that together exceed 0.9, from which the next token is chosen with multinomial sampling.

4.1 Data Configurations

For our experiments, we use several established dialogue datasets as baselines, namely *Persona-chat* (Zhang et al., 2018), *Cornell* (Danescu-Niculescu-Mizil and Lee, 2011) (the unfiltered corpus out of which we extract *yes-and*s, as described in Section 2.2), and *DailyDialog* (Li et al., 2017b). Each of these is an English-language open-domain casual conversation corpus with 100k–300k turns. For each of these datasets, we either simply fine-tune on that dataset, or fine-tune and then further

For each of the questions, rank the given options in order of which is the best "Yes, and" response to the given prompt.

[Click for instructions](#)

Q1	
Prompt Sorry we didn't do better, Jack. I feel like I let you down.	
4	Response1 what do you do for a living?
2	Response2 I feel like you're taking responsibility for everything, you know.
1	Response3 Naw, you didn't let me down. It was a long shot all the way. We gave 'em a good run at it.
3	Response4 You're right. I do feel like I have to make a choice.

Figure 5: Interface used by human evaluators to rank responses based on their quality as a *yes-and*, where a rank of 1 is most preferred. The correct ranking is shown for this example. The option ranked 1 is a *yes-but*: it does not reject a reality but rather rejects a suggestion and provides refining information that is most coherent to the prompt.

fine-tune with **SPOLIN**. In another configuration, we also fine-tune directly with **SPOLIN** on top of GPT-2. The original GPT-2 implementation prepends the personalities given in *Persona-chat* to the dialogue sequence input before tokenization. For fine-tuning to datasets apart from *Persona-chat*, we simply do not prepend any auxiliary information to the dialogue sequence input.

4.2 Human Evaluation

Automatic metrics that rely on n-gram overlap, such as BLEU, ROUGE, and METEOR, are often used for generative models when there is little variability in the target output (Papineni et al., 2002; Lin, 2004; Banerjee and Lavie, 2005). However, there can be a wide variety of responses that qualify as a good *yes-and*, a problem common to open-domain generation tasks. An adequate evaluation of our models requires assessing the main *yes-and* criteria: agreement with the context and the quality of the new relevant contribution, both of which are not feasible with the aforementioned metrics. Therefore, we ask human evaluators to compare the quality of the *yes-and*s generated by various models and the actual response to the prompt in **SPOLIN** that is used as the input.

We ask human evaluators to rank a set of four responses given a prompt, comparing the responses of a model trained only with **SPOLIN**, a model trained with an existing dialogue corpus, a model trained with both, and the actual response pair from

the development set, denoted as “Gold.” These four responses are randomly ordered for each question to prevent evaluators from developing a bias for responses that frequently have a good or poor response in a set order, as shown in Figure 5. The evaluators are permitted to provide the same rank for different responses if they are equal in quality. This evaluation set contains 100 such prompts, and each is evaluated twice by different evaluators. The results of the average ranking and some of the examples generated by the models are shown in Table 4.

Results show that models trained only with **SPOLIN** or with **SPOLIN** and another dialogue dataset are preferred to the models trained only with another dialogue dataset, although in the case of *DailyDialog*, the average ranking improves only by at most 0.06 after fine-tuning with **SPOLIN**. However, even the responses generated by models trained with **SPOLIN** are not ranked as well as the actual responses in the development set, indicating our models are still inferior to professional human improviser quality.

5 Extracting from Other Corpora

The approach to classifier-based mining we describe in Section 2.2 can naturally be applied to other dialogue corpora. We thus next consider mining the gigantic (441M sentence) OpenSubtitles (Lison and Tiedemann, 2016) collection. As OpenSubtitles contains undesirable material, such as subtitles for media with minimal dialogue, we instead mine from the (3.3M sentence) *SubTle* corpus (Ameixa et al., 2013), a preprocessed subset of OpenSubtitles that heuristically combines subtitle sequences into dialogue form.

By iterating through half of this corpus, we collect more than 40,000 *yes-and*s from it alone, which, when added to **SPOLIN**, yields what we call **SPOLIN-extended**, which contains about 68,000 *yes-and*s, more than 2.5 times the size of the core **SPOLIN**. Heuristics for finding alternations mean that *SubTle*’s utterances are shorter than those in *Spontaneanation* and *Cornell*, so once the proportion of utterances longer than the average length of in *Spontaneanation* and *Cornell* (18.5 words) is less than 40%, we stop further collection in the remainder of the dataset. **SPOLIN-extended** is available in the same public repository as **SPOLIN**. Details of the iterative process as applied to *SubTle* are in the appendix.

Dataset	Avg Rank ↓	Example Prompt	Example Responses
<i>Persona-chat</i>	3.67	I know alotta women and I'm sure she remembers me.	oh my goodness, i don't know her
SPOLIN	3.41		Yeah, she's a bit of a mystery.
<i>Persona-chat+SPOLIN</i>	3.17		So you remember her? I remember her in the shower.
Gold	1.92		She does. From when you were a boy.
<i>Cornell</i>	3.73	I made this beautiful salmon mousse that just looked like a puddle of spit.	What?
SPOLIN	3.39		And it's delicious!
<i>Cornell+SPOLIN</i>	3.34		That's the kind of thing you do when you're in love.
Gold	2.01		It was genius. It's making me hungry thinking about it.
<i>DailyDialog</i>	3.37	Excuse me. Is anybody in here? I'd like a bagel. Hello? Is anybody in here?	I'm sorry, sir. I'm not in the mood. I'm not in the mood.
SPOLIN	3.32		I'm in here. I'm just trying to make sure I can get a bagel.
<i>DailyDialog+SPOLIN</i>	3.31		Oh, yeah, the guy who left the bagel.
Gold	1.87		I can help you. The problem is that the bagels are burned.

Table 4: Average human ranking of responses to prompts from *Spontaneation* generated by models trained with **SPOLIN**, an existing dialog corpus, or both, based on the *yes-and* criteria. Rank is scaled from 1 to 4, lower is better.

Dataset	Source	Size*
DailyDialog (Li et al., 2017b)	Crowdsourced	104K
Cornell Movie-Dialogs Corpus (Danescu-Niculescu-Mizil and Lee, 2011)	Movie scripts	304K
Persona-chat (Zhang et al., 2018)	Crowdsourced	162K
The Ubuntu Dialogue Corpus (Lowe et al., 2015)	Ubuntu chat logs	7M
Twitter Triple Conversations (Sordoni et al., 2015)	Social media	6K
OpenSubtitles (Lison and Tiedemann, 2016)	Subtitles	441M sentences
SubTle (Eng) (Ameixa et al., 2013)	Subtitles	3.3M pairs
London-Lund Corpus (Greenbaum and Svartvik, 1990)	Various sources	500K words
London-Lund Corpus 2 (Pöldvere et al., 2017)	Various sources	500K words
SPOLIN (<i>yes-and</i> only)	Improv, Movie scripts	26K pairs
SPOLIN-extended (<i>yes-and</i> only)	Improv, Movie scripts, subtitles	68K pairs

Table 5: A survey of publicly available English language text-based corpora frequently used for open-domain dialogue systems. The last two rows are our contribution. *Size is measured as the number of total utterances (dialogue turns) unless otherwise specified.

6 Related Work

Many works have identified the same issues of repetitive or non-committal responses generated by neural conversational systems that are at least partially related to the lack of sufficiently high quality *yes-and*s we deal with in this work; approaches that mitigate these problems vary. The majority of recent works focus on diversifying the responses by modifying the training and decoding objectives (Li et al., 2016a,b, 2017a, 2016c; Xu et al., 2017; Shao et al., 2017). Other methods introduce latent variables to encourage diversity (Serban et al., 2017; Zhao et al., 2017). Some explore methods of re-weighting training instances that encourage diversity (Liu et al., 2018; Lison and Bibauw, 2017; Du and Black, 2019).

Our approach is complementary to all the model-

based approaches described here, as it simply deals with the production of a particularly useful *corpus*, that can be used to fine-tune on top of these methods.

We provide a survey of publicly available text-based datasets frequently used for open-domain dialogue systems and discuss their limitations for our purpose of generating grounded responses (see Table 5 for an overview).

DailyDialog is a collection of multi-turn dialogue with manually annotated emotion and intent labels (Li et al., 2017b). Danescu-Niculescu-Mizil and Lee (2011) created the *Cornell Movie-Dialogs Corpus*, a compilation of dialogue sequences paired with meta data about the movie and characters. *Persona-chat* provides dialogue sequence coupled with corresponding personas (Zhang et al., 2018).

The *Ubuntu Dialogue Corpus* contains 1 million dialogue turns extracted from Ubuntu chat logs, which discuss Ubuntu-related technical support (Lowe et al., 2015). The *Twitter Triple Corpus* is a dataset of 4K dialogue triples extracted from Twitter (Sordoni et al., 2015). *OpenSubtitles* is a huge collection of subtitles that span various genres, but the absence of speaker turn annotations make it difficult to modify into dialogue format (Lison and Tiedemann, 2016). Ameixa et al. (2013) use heuristics to reformat OpenSubtitles into dialogues with some limited success. Clark and Schaefer (1989) illustrate grounding in conversations with examples from the London-Lund Corpus (Greenbaum and Svartvik, 1990), a corpus of full conversations annotated with prosodic and paralinguistic features. A second version of the corpus was compiled with the same annotations standards as the first using more recent spoken and text data (Pöldvere et al., 2017).

These corpora were not collected with the criteria for *yes-and*s in mind. Even for datasets with dialogue taking place in a similar domain as improv, they naturally contain only a small proportion of *yes-and*s. However, the relatively large sizes of these datasets still make them useful for dialogue systems. They can be used effectively for grounded conversations if the *yes-and*s or other desirable dialogue acts can be filtered out or given higher weights in training to enforce their characteristics in the responses generated.

Our data collection approach is similar to the method of Yarowsky (1995), which formalizes the bootstrapping mechanism of iteratively improving a classifier and label unlabeled data. The main difference from the Yarowsky algorithm and our approach is that, rather than using a fully automated process for increasing training data, we use a probability threshold to regulate recall, followed by human judgment to ensure high precision.

Apart from Clark and Schaefer (1989) there have been other taxonomies of grounding. For example, Traum (1999) considers six categories; among these are *acknowledge* and *continue*, which, taken together, map nicely to *yes-and*. Magerko et al. (2009) and Fuller and Magerko (2010) note the importance of establishing common ground in improv.

7 Conclusion

Inspired by *yes-and*s in improv, we carefully construct **SPOLIN**, a collection of dialogue pairs with responses that are not only coherent with dialogue context but also initiate the next relevant contribution. We extract high-quality *yes-and*s from *Spontaneanation* and build a classifier with them, which is then used to mine additional *yes-and*s from the *Cornell Movie-Dialogs Corpus*. We further use our mining technique to elicit a corpus of more than 68,000 *yes-and* turn pairs, easily the largest collection of this dialogue act known to exist. From human evaluations of dialogue models trained with various data configurations we find that **SPOLIN** is useful—when including it we are able to build models that can generate *yes-and*s more consistently than when we leave it out. Nevertheless, our models are still inferior at producing good *yes-and*s when compared to professional improvisers. We plan to continue our data-driven approach for grounded conversations by expanding our dataset through our iterative data collection process with other larger text-based open-domain dialogue corpora and extend our work to model and collect longer conversations exhibiting more complex improv-backed turns.

Acknowledgments

Many thanks to Nanyun Peng and Xinyu Wang for key contributions in a preliminary study, to Paul F. Tompkins, Colin Anderson, and Earwolf for allowing us to include *yes-and*s extracted from *Spontaneanation* in **SPOLIN**, to Paul Elsberg, Risa Harms, P.T. McNiff, and Peter Schell for initial inspiration, and to Jordan Boyd-Graber for feedback on the final draft. This material is based on research sponsored by the AFRL and DARPA under agreement number FA8650-18-C-7878. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the AFRL, DARPA, or the U.S. Government.

References

- David Ameixa, Luísa Coheur, and Rua Alves Redol. 2013. From subtitles to human interactions: Introducing the SubTle corpus. Technical report, INESC-ID.
- Satanjeev Banerjee and Alon Lavie. 2005. **METEOR**:

- An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- David Bohm and Lee Nichol. 2004. *On Dialogue*. Routledge classics. Routledge.
- Allison Bruce, Jonathan Knight, Samuel Listopad, Brian Magerko, and Illah R. Nourbakhsh. 2000. Robot improv: using drama to create believable agents. *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*.
- Carlos Busso and Shrikanth S Narayanan. 2008. Scripted dialogs versus improvisation: Lessons learned about emotional elicitation techniques from the IEMOCAP database. In *Ninth annual conference of the international speech communication association*.
- Herbert H Clark and Edward F Schaefer. 1989. Contributing to discourse. *Cognitive science*, 13(2):259–294.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*, pages 76–87, Portland, Oregon, USA. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wenchao Du and Alan W Black. 2019. Boosting dialog response generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 38–43, Florence, Italy. Association for Computational Linguistics.
- Daniel Fuller and Brian Magerko. 2010. Shared mental models in improvisational performance. In *Proceedings of the Intelligent Narrative Technologies III Workshop, INT3 '10*, New York, NY, USA. Association for Computing Machinery.
- Sidney Greenbaum and Jan Svartvik. 1990. *The London-Lund corpus of spoken English*, volume 7. Lund University Press.
- Charna Halpern, Del Close, and Kim Johnson. 1994. *Truth in comedy: The manual of improvisation*. Meriwether Publishing.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2020. The curious case of neural text degeneration. In *Proceedings of the Eighth International Conference on Learning Representations*.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems*, 38(3):1–32.
- Keith Johnstone. 2017. *Impro: Improvisation and the Theatre*. Performance Books. Bloomsbury Publishing.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016c. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169, Copenhagen, Denmark. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. DailyDialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- 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.
- Pierre Lison and Serge Bibauw. 2017. Not all dialogues are created equal: Instance weighting for neural conversational models. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 384–394, Saarbrücken, Germany. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *LREC*.

- Yahui Liu, Wei Bi, Jun Gao, Xiaojiang Liu, Jian Yao, and Shuming Shi. 2018. [Towards less generic responses in neural conversation models: A statistical re-weighting method](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2769–2774, Brussels, Belgium. Association for Computational Linguistics.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. [The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- Brian Magerko, Waleed Manzoul, Mark Riedl, Allan Baumer, Daniel Fuller, Kurt Luther, and Celia Pearce. 2009. [An empirical study of cognition and theatrical improvisation](#). In *Proceedings of the Seventh ACM Conference on Creativity and Cognition*, page 117–126, New York, NY, USA. Association for Computing Machinery.
- Lara J. Martin, Brent Harrison, and Mark O. Riedl. 2016. [Improvisational computational storytelling in open worlds](#). In *Interactive Storytelling*, pages 73–84, Cham. Springer International Publishing.
- Siobhan McHugh. 2016. [How podcasting is changing the audio storytelling genre](#). *Radio Journal: International Studies in Broadcast & Audio Media*, 14(1):65–82.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Nele Pöldvere, V Johansson, and C Paradis. 2017. [The London-Lund corpus 2: A new corpus of spoken British English in the making](#). In *theICAME 38 Conference, Prague, Czech Republic*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). Technical report, OpenAI.
- Iulian Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2015. [Hierarchical neural network generative models for movie dialogues](#). *ArXiv*, abs/1507.04808.
- Iulian Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Louis Shao, Stephan Gouws, Denny Britz, Anna Goldie, Brian Strope, and Ray Kurzweil. 2017. [Generating long and diverse responses with neural conversation models](#). *arXiv preprint arXiv:1701.03185*.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205, Denver, Colorado. Association for Computational Linguistics.
- Viola Spolin, Arthur Morey, and Mary Ann Brandt. 1986. *Theater Games for the Classroom: A Teacher's Handbook*. Northwestern University Press.
- David R Traum. 1999. [Computational models of grounding in collaborative systems](#). In *Psychological Models of Communication in Collaborative Systems-Papers from the AAAI Fall Symposium*, pages 124–131.
- Lauren Winston and Brian Magerko. 2017. [Turn-taking with improvisational co-creative agents](#). In *Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019a. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019b. [Transfertransfo: A transfer learning approach for neural network based conversational agents](#). *arXiv preprint arXiv:1901.08149*.
- Zhen Xu, Bingquan Liu, Baoxun Wang, Chengjie Sun, Xiaolong Wang, Zhuoran Wang, and Chao Qi. 2017. [Neural response generation via GAN with an approximate embedding layer](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 617–626, Copenhagen, Denmark. Association for Computational Linguistics.
- David Yarowsky. 1995. [Unsupervised word sense disambiguation rivaling supervised methods](#). In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 189–196, Cambridge, Massachusetts, USA. 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, Melbourne, Australia. Association for Computational Linguistics.

Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.

A Appendix

Iteration	4	5	6	7
<i>Spontaneanation</i> +	10,459	10,459	10,459	10,459
<i>Spontaneanation</i> -	5,587	5,587	5,587	5,587
<i>Cornell</i> +	12,220	14,976	14,976	14,976
<i>Cornell</i> -	17,092	17,701	17,701	17,701
<i>SubTle</i> +	-	2,621	20,617	33,155
<i>SubTle</i> -	-	7,865	14,799	17,325
Total Training Samples	45,358	59,209	84,319	99,203
Dev Set Acc. (<i>Spont</i>)	73.0%	72.1%	68.4%	75.2%
Dev Set Acc. (<i>Cornell</i>)	64.5%	63.3%	63.3%	61.0%
Confidence Threshold	50% / 70%*	70%	70%	70%
New Extraction Volume	3,515 / 10,486*	36,608	15,424	14,979
New Proportion of <i>yes-ands</i>	78.4% / 25.0%*	58.4%	83.2%	76.0%

Table 6: Continuation of Table 1 with the extended version of **SPOLIN** that includes extracted *yes-ands* from *SubTle*. *SubTle* is collected from the fourth iteration onwards. *Statistics for *Cornell/SubTle* are shown separately. The same classifier is used for extracting candidates from *Cornell* and *SubTle*, but they are datasets with significantly different characteristics.

A.1 *yes-and* Guidelines for Turkers

We provide detailed annotation guidelines, shown in Figures 6–9, to the Turkers as a result of having continuous discussions with them and monitoring their submissions. Contrary to our expectations, it is difficult to make a binary decision on whether a dialogue turn is a *yes-and* or non-*yes-and*, and therefore these fine-grained details are crucial for collecting *yes-ands* in **SPOLIN**.

A.2 Iterative data collection results for *SubTle*

Due to *SubTle*’s relatively large size, we split *SubTle* into 20 equal blocks that each contains 10,486 dialogue turns. Note that every successive iteration of *SubTle* was not performed on the same block but on the next block. This is different from *Cornell*, for which every iteration is on the same set of dialogue turns. This difference is not due to any characteristics in the dataset but because of practical reasons arising from the size of the *SubTle* corpus.

The first extraction proportion for *SubTle* is low because of the prevalence of self-*yes-ands* in this corpus. Self-*yes-ands* are prompt and response pairs that evidently originate from the same speaker but otherwise meet the criteria of a *yes-and*. There are many incorrectly combined dialogue turns that actually come from the same speaker because of the heuristics employed for building *SubTle*. By providing labeled self-*yes-and* as negative samples, the classifier quickly learns to remove these self-*yes-ands*, leading to a significantly higher proportion of *yes-ands* in subsequent iterations. This is demonstrated in the specifics of the additional iterations, which are shown in Table 6.

Objective

The objective of this task is to collect a corpus consisting of "Yes, and..." type dialogue pairs (henceforth simply "Yes, and"), which will be used to train neural network models for creative language generation. For each of these HITs, you are expected to identify all the "Yes, and"s you hear from an improv session within a podcast episode.

What is a "Yes, and"?

Definition

"Yes, and..." is a rule-of-thumb in improvisational comedy that suggests that a participant in a dialogue should accept what another participant has stated ("Yes") and then expand on that line of thought or context ("and...").¹ In short, a "Yes, and" is a dialogue exchange in which a speaker responds by adding new information on top of the information/setting that was constructed by another speaker.

Note that a "Yes, and" **does not** require someone explicitly saying 'yes, and...' as part of a dialogue exchange, although it could be the case if it agrees with the description above. There are many ways in which a response could implicitly/explicitly agree to the prompt without specifically saying 'yes, and...!'

¹Reference: https://en.wikipedia.org/wiki/Yes,_and...

Basic Structure of a "Yes, and"

Prompt	"Yes, and" Response
Provides some information. (Either as a statement or a question)	Implicitly / explicitly accepts what has been said by the prompt and adds new relevant information.

If this is your first time doing this HIT or if you are unsure as to what qualifies as a "Yes, and", read through the ["Yes, and" Annotation Examples section](#) to develop a better understanding for this task.

Example

Given the following dialogue sequence, those selectively shown in the table below are examples of "Yes, and"s.

1. [A] Well, which one speaks to you?
2. [B] Well, the one over here in the northwest corner. The crow with the beady eyes that seem to follow you. Something about it. It reminds me of something.
3. [A] Yeah! I know what it reminds me of. It reminds me of that look that comes over Dean's face when he's done something terrible, when he says...
4. [B] When he says, "No, Mommy, I won't clean my room."
5. [A] That's right!
6. [B] It's a crow body, but with the face of our child. We'll give you fifty dollars for it.
7. [C] That's exactly what I am asking. There's a sign that says fifty.
8. [B] I didn't see the sign. So that was just a lucky guess. I'm a little bit psychic.
9. [A] My wife is being modest. She's very psychic.
10. [B] It is true. In fact, all the women in my family, we have abilities.
11. [C] Please tell me a psychic vision you have had that has come true.
12. [A] Oh, tell about, remember that time we were out with the kids and we were going to go to Six Flags, and then you said we shouldn't get on the roller coaster.

"Yes, and" Dialogue pair	Explanation
[B] Well, the one over here in the northwest corner. The crow with the beady eyes that seem to follow you. Something about it. It reminds me of something. [A] Yeah! I know what it reminds me of. It reminds me of that look that comes over Dean's face when he's done something terrible, when he says...	Speaker B establishes that she is interested in a painting with a crow in it and that it reminds her of something. Speaker A adds on to her remark and suggests details of what it reminds him of.
[A] Yeah! I know what it reminds me of. It reminds me of that look that comes over Dean's face when he's done something terrible, when he says... [B] When he says, "No, Mommy, I won't clean my room."	Speaker B establishes that the painting reminds him of Dean's face when he's done something terrible, and Speaker A provides a detailed situation in which Dean makes the face that Speaker B was mentioning.
[B] It's a crow body, but with the face of our child. We'll give you fifty dollars for it. [C] That's exactly what I am asking. There's a sign that says fifty.	Speaker B suggests that she will pay \$50 for the painting. Speaker C agrees with her price and adds on the fact that there is a sign with the price on it. Agreeing with the price is not necessary but rejecting a setting Speaker B has established should be refrained. For example, "I want \$100 for it. There's a sign with the price." is okay but "I know you don't have \$50" does not qualify as a "Yes, and".
[C] That's exactly what I am asking. There's a sign that says fifty. [B] I didn't see the sign. So that was just a lucky guess. I'm a little bit psychic.	Speaker B understands that there is a sign that she didn't see, and suggests that she has psychic abilities that allowed her to guess the price correctly.
[C] Please tell me a psychic vision you have had that has come true. [A] Oh, tell about, remember that time we were out with the kids and we were going to go to Six Flags, and then you said we shouldn't get on the roller coaster.	Speaker C suggests that Speaker B has psychic abilities. Speaker A, knowing of Speaker B's psychic powers, interrupts and suggests a specific instance in which she demonstrated her powers.

The other possible sequential pairs are not "Yes, and" dialogue pairs. The response ignores or contradicts the information present by the prompt, or it simply acknowledges the information without refining it or adding new information. If you had listened to an audio block containing any subset of this dialogue sequence, an ideal submission would have identified all the "Yes, and" dialogue pairs here and have them accurately transcribed.

Figure 6: Explanation of the objective and *yes-and* in the *yes-and* guideline provided to Turkers.

Label space

Here is a list of subcategories for a "Yes, and" and for those that are not. You **do not** need to specify subcategories for this task and the following information should only be used as an extra guideline to help you better identify "Yes, and"s. These labels are also used in the "Yes, and" [Annotated Example](#) section for succinct explanations.

"Yes, and" subcategory

- Yes-strong: specifics provided by the prompt are refined
- Yes-but: acknowledges and adds, but rejects a specific suggestion offered by the prompt, often with 'but' or 'however'; it's different from contradiction in that the reality itself isn't rejected but an offer within it isn't taken.

Non-"Yes, and" subcategory

- No (no-info): Not contradictory, but does not add more information to that added by the prompt
- No (ignore): The response ignores statements in the prompt; there is no plausible way the response and prompt could be connected
- No (contra): The response may add some information but contradicts the reality established by the prompt

Instructions

1. The task of identifying individual "Yes, and"s is the same as that of the qualification HITs. If a dialogue pair can be made into a "Yes, and" with easily identifiable corrections or by ignoring cross-talk, write the dialogue pair as a "Yes, and" in your answers as well. The main difference is that you are asked to identify all "Yes, and"s in a longer audio file.
2. Whenever you identify a "Yes, and" while listening to the improv session, pause and transcribe the "Yes, and" in the answer section to the right side of the window. Make sure to insert the prompt and response in separate boxes.
3. A rough transcript of the audio is shown on the left of the window, which is divided into different lines that are estimations of speaker turns. **The transcript and speaker turns are not always accurate.** They are for your guidance only and you may ignore them.
4. To add a new answer, click the "Add +" button to add a new prompt and response slot. To delete any answer you've written, click the trash icon on the right side of the answer.
5. The sequence of the "Yes, and"s you identify **does not** matter! Don't worry about placing your answers in the order they appeared in the audio.
6. Here are some keyboard shortcuts and tips you can use to complete this task:
 - Esc: Play / Pause
 - F2: Skip backward 1.5 seconds
 - F4: Skip forward 1.5 seconds
 - F5: Speed up audio by 0.25 (max: 2x)
 - F6: Speed down audio by 0.25 (min: 0.5)
 - You can click on the buttons on the left side of the audio control panel for the same functionalities, except for controlling the audio speed.
 - Click on the timestamp under the 'Start' column to jump to that time on the audio.
 - Click on the transcript to copy the transcript to your clipboard. You can paste the transcript to your answer with the regular paste command on your device.
 - You can also click on or slide along the audio bar (the middle section in dark grey of the control panel) to navigate yourself through different parts of the audio. The current time and total time of the audio is shown on the right side of the audio control panel.
7. Once you have reached the end of the audio and you are confident that you have identified most, if not all, the "Yes, and"s in the improv session, click the "Submit" button at the bottom. Before submission, please make sure that both prompts and responses are filled in for all answers and delete any answers that are empty.

Figure 7: Explanation of the label space for *yes-ands* and non-*yes-ands* and the detailed instructions for the transcription task.

Common Mistakes

Here are some common mistakes that we have seen in previous HIT submissions. Make sure to understand why

Common mistake 1: Write prompts as statements or questions that do not introduce information

Undesirable answer:

P: What do you think?

R: It looks nice, but it's like the eyes are following me.

Desirable answer:

P: Why do you want to do that? **If they are alive and they have been caged their entire lives, why do you want to have them suffer?**

R: Caged? They've been spoon-fed all kinds of luxuries in that castle. It's time for them to learn the tough lessons of life.

→ Although the prompt is a question, it establishes that 'they' were alive and caged their entire lives and that the following speaker wants them to suffer.

Common mistake 2: Response that simply agrees to the prompt but does not significantly add/refine information provided by the prompt

Undesirable answer:

P: They left me a note to read to you.

R: Oh did they? Let me see what that says.

Desirable answer:

P: They left me a note to read to you.

R: Oh did they? Let me see what that says. **Is it that red piece of paper in your hands?**

→ The added question in the response establishes additional information that the previous speaker is holding onto a red piece of paper in her hands. (may or may not be the note)

Common mistake 3: Response logically follows the prompt but does not add/refine information provided by the prompt

Undesirable answer:

P: God, I'm gonna punch you so hard in the stomach!

R: You've got to warn me first.

Desirable answer:

P: God, I'm gonna punch you so hard in the stomach!

R: You've got to warn me first. **I have a bad tummy right now and I don't want there to be an accident in your carpeted living room.**

→ The original answer could also be argued as a "Yes, and", but it is weak in saying that a warning is necessary is more an approval of the punch he's about to take than an addition of information. A better answer would provide a reason why a warning is necessary.

Figure 8: Common mistakes that Turkers made in the early stages of data collection were corrected and added to the guidelines to aid new Turkers.

"Yes, and" Annotated Examples

(Click to expand/hide)

X: I hate to quibble with you Rob, but she married her first boyfriend. Kevin Bannister.

Y: You gotta be kidding me.

Annotation: No (no-info). No contradiction, simply a mild expression of incredulity. A good answer would have included something like "Kevin Bannister the circus clown?" or "I can't believe Jill would do that".

X: Well, I wish you'd stop criticizing and picking on her.

Y: Forgive my crudity, darlin'. All I'm saying is that a girl who would wear clothes like that is going to get in trouble sooner or later.

Annotation: Yes (strong). The "forgive my crudity" acknowledges that the respondent is indeed 'criticizing and picking on her' and the "all I'm saying" part explains the behavior further.

X: With all she was doin'. With all the shit she kept doing! You stayed stuck to that bitch's ass and you wouldn't let go.

Y: I know about how she was like. But we was different. I's the only person she talked to about it. How she's abused. Terrible things, Gill, just terrible...

Annotation: Yes (strong) [not particularly funny, but that's not relevant]. Explains why Y "stayed stuck". Adds more to the character.

X: Yeah, man. Let's throw a bachelor party with drugs, booze and broads.

Y: Yeah. Right. All the things that make life worth living.

Annotation: No(no-info). Simply acknowledges X. A good answer could have included "Oh, I know where we can get some ketamine!" or "Let's have it in vegas!"

X: Instead of busting my chops you should do something about that girl. Fire her. Or something.

Y: Lisa's an extremely valuable member of the Skywire team. We've got our eyes on her. You keep yours on Milo.

Annotation: No (contra) Definitely adds info but is combative and contradictory; rejects the assertions offered by X. A better answer (a yes-but) would have preceded what is provided with, e.g. "I agree that Lisa has messed up, but". A true "yes, and" would have been "She's a valuable member of the team so we've been resisting, but you're right, she's gone too far. Milo is too, so you should also fire him"

X: Aw, come on Jerry. We've gotten all we're gonna get out of this place and its starting to rain.

Y: Shit, it is only sprinkling and it's worth the trouble. Hold on for two seconds.

Annotation: No (contra). Denies that it's not useful to stay around. Note that this does acknowledge the rain ("it is only sprinkling" supports "its starting to rain") but per the guidelines, there must be no contradictions for a 'yes'. Additionally, this doesn't add anything concrete. A good answer would be e.g. "yeah, well it is only sprinkling but it's true, staying here is too much trouble".

X: ...don't do it to me...

Y: She came to talk to me...

Annotation: No (ignore). Y is not responsive to X. A good answer might be "I'm just going to give you a little novocaine so you won't feel it when we remove your tooth"

Figure 9: Annotated examples provided to Turkers for understanding the label space of the *yes-and* transcription task.