

Did they answer? Subjective acts and intents in conversational discourse

Elisa Ferracane¹, Greg Durrett², Junyi Jessy Li¹, and Katrin Erk¹

¹Department of Linguistics

²Department of Computer Science

The University of Texas at Austin

elisa@ferracane.com, gdurrett@cs.utexas.edu

jessy@austin.utexas.edu, katrin.erk@mail.utexas.edu

Abstract

Discourse signals are often implicit, leaving it up to the interpreter to draw the required inferences. At the same time, discourse is embedded in a social context, meaning that interpreters apply their own assumptions and beliefs when resolving these inferences, leading to *multiple*, valid interpretations. However, current discourse data and frameworks ignore the social aspect, expecting only a single ground truth. We present the first discourse dataset with multiple *and* subjective interpretations of English conversation in the form of perceived conversation acts and intents. We carefully analyze our dataset and create computational models to (1) confirm our hypothesis that taking into account the bias of the interpreters leads to better predictions of the interpretations, (2) and show disagreements are nuanced and require a deeper understanding of the different contextual factors. We share our dataset and code at http://github.com/elisaF/subjective_discourse.

1 Introduction

Discourse, like many uses of language, has inherent ambiguity, meaning it can have *multiple*, valid interpretations. Much work has focused on characterizing these “genuine disagreements” (Asher and Lascarides, 2003; Das et al., 2017; Poesio et al., 2019; Webber et al., 2019) and incorporating their uncertainty through concurrent labels (Rohde et al., 2018) and underspecified structures (Hanneforth et al., 2003). However, prior work does not examine the *subjectivity* of discourse: how you *resolve* an ambiguity by applying your personal beliefs and preferences.

Our work focuses on subjectivity in question-answer conversations, in particular how ambiguities of responses are resolved into subjective assessments of the **conversation act**, a speech act in conversation (Traum and Hinkelman, 1992), and the **communicative intent**, the intention underlying

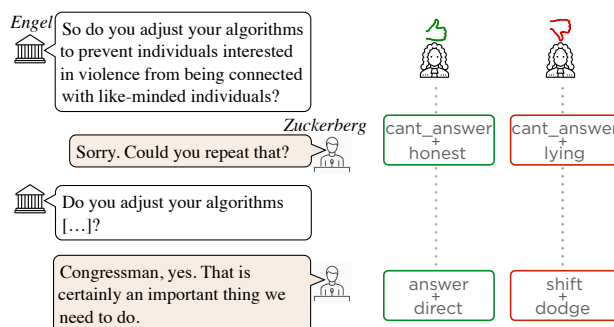


Figure 1: Conflicting interpretations of conversation acts + intents for witness responses in a U.S. congressional testimony.

ing the act (Cohen and Perrault, 1979). We choose conversation acts (or more broadly, dialogue acts) as a challenge to the view that dialog act classification may be an “easy” task that has never been approached from a subjective perspective. Moreover, they are a good fit for our question-answering setting and are intuitive for naive annotators to understand. Our data consists of witness testimonials in U.S. congressional hearings. In Figure 1, annotators give conflicting assessments of responses given by the witness Mark Zuckerberg (CEO of Facebook) who is being questioned by Congressman Eliot Engel.

To make sense of our setting that has speakers (witness, politicians) and observers (annotators), we are inspired by the game-theoretic view of conversation in Asher and Paul (2018). The players (witness, politicians) make certain discourse moves in order to influence a third party, who is the judge of the game (the annotator). Importantly, the judge makes biased evaluations about the type of the player (e.g., *sincere* vs. *deceptive*), which leads to differing interpretations of the same response.

In our example, the two annotators are the biased judges with differing judgments on what type of player Zuckerberg is: the first assumes *sincere* and the second *deceptive*. For Zuckerberg’s first response, the conversation act is interpreted un-

ambiguously: both annotators agree he is signaling he can't answer the question. The intent, however, is ambiguous, where the cynical annotator interprets the clarification question as lying in order to stall vs. being honest. The second response yields both diverging conversation acts and intents: the first judge interprets the conversation act as an answer with the intent to provide a direct response, whereas the second judge perceives the conversation act as a shift to answer a different question with the intent to dodge the original, unfavorable question. We detail our full label set in Section 3.2.

We create the first discourse dataset with multiple, valid labels that are subjective. They do not hold concurrently and vary depending on the annotator; we collect annotator sentiments towards the conversants as a rough proxy for annotator bias. We further elicit annotator explanations for a window into their rationalization. A careful annotation protocol and qualification process ensure high quality crowd-sourced annotators with a strong understanding of the task. Our dataset contains 6k judgments over 1k question-response pairs, with disagreements in 53.5% of the data. However, unlike our prior example, disagreements are not often trivially attributable to differing sentiments. Uncooperative moves are sometimes warranted, regardless of annotator sentiment. Interpretation of a response is further influenced by its question. A qualitative analysis of annotator explanations reveals strikingly different uses of subjective language across diverging interpretations.

Identifying all the possible interpretations of a response is a useful way of analyzing discourse in a realistic setting with multiple observers, and could aid in uncovering sociolinguistic aspects relevant to variations in discourse comprehension. With these goals in mind, we propose the task of predicting the complete set of annotator labels for a given response. We find a transformer-based model outperforms other neural and linear models. We confirm our assumption that incorporating the context of the judge helps the model make better predictions, but still leaves room for improvement.

In summary, the task together with the dataset present a valuable opportunity to understand perceptions of discourse in a non-cooperative environment. More broadly, we show the need and value for considering the subjectivity of NLP tasks. Our work introduces a framework for identifying, elic-

iting, and analyzing these subjective elements, to enable application for other tasks.

2 Background and Related Work

Asher and Paul (2016) apply their game-theoretic view of non-cooperative conversations to discourse moves in Segmented Discourse Representation Theory (Asher and Lascarides, 2003). Our work is applied instead to conversation acts and their communicative intents, which are more amenable to untrained annotators. Conversation acts are speech acts specific to conversation that can encompass entire turns in a conversation (Traum and Hinkelman, 1992). Speech act theory describes performative actions, i.e., how we can do things with words (Austin, 1962; Searle, 1969), but fails to account for how the act is perceived by an observer (the annotator in our scenario). Subsequent work in planning extends the theory to incorporate the cognitive context of an observer that includes the perceived *communicative intent* underlying a speech act (Cohen and Perrault, 1979; Pollack, 1986).

Speech act theory originally did not consider insincere speakers, but later work recognized that even in non-cooperative settings, conversants adhere to the conventions of dialogue, or *discourse obligations*, such as responding to a question (Traum and Allen, 1994; Potts, 2008). For this reason, we explicitly separate judgments on conversation acts (that usually fulfill a specific obligation) from communicative intents, which can be perceived as deceptive (or sincere).

Prior work examines how writer intentions are often misaligned with reader perceptions (Chang et al., 2020), which further motivates our focus on the reader (our annotator). While our work focuses on subjectivity, ambiguity is studied in many NLP tasks, including Natural Language Inference (Pavlick and Kwiatkowski, 2019; Nie et al., 2020), evaluation of NLG (Schoch et al., 2020), a recent SemEval 2021 shared task,¹ as well as several discourse tasks (Asher and Lascarides, 2003; Versley, 2011; Webber and Joshi, 2012; Das et al., 2017; Poesio et al., 2019; Webber et al., 2019). Only one study strives to understand how these ambiguities are resolved: Scholman (2019) shows different interpretations of ambiguous coherence relations can be attributable to different cognitive biases. However, our work focuses more generally on subject-

¹<https://sites.google.com/view/semEval2021-task12/home>

tivity rather than cognitive processes.

Related NLP tasks include dialog act classification, intent detection, deception detection and argumentation, though we importantly note these predict only a single interpretation. Dialog acts are similar to conversation acts that apply at the utterance level. Classification models typically combine representations of linguistic units (word, utterance, conversation-level) (Chen et al., 2018). In our work, we employ a hierarchical model to account for the levels in our label taxonomy. Intent detection is traditionally applied to human-computer scenarios for task-specific goals such as booking a flight. Our conversation data is not task-oriented, and we thus define our intents more closely aligned with beliefs in the sincerity of the speaker. Detection of deception is, unlike many other NLP tasks, challenging even for humans (Ott et al., 2011). Most datasets consist of instructed lies (where participants are told to lie). Our work contains naturally-occurring deception where we include not just lying but other more covert mechanisms such as being deliberately vague or evasive (Clementson, 2018), both frequent in political discourse (Bull, 2008).

Argumentation mining analyzes non-cooperative conversations, but typically requires expert annotators. Recent work decomposes the task into intuitive questions for crowdsourcing (Miller et al., 2019), inspiring our annotation schemes that assume little to no training. Closer to our setting is argument persuasiveness, where Durmus and Cardie (2018) find prior beliefs of the audience play a strong role in their ability to be persuaded, which further motivates our focus on the annotator’s bias.

3 Dataset

We create the first dataset with multiple, *subjective* interpretations of discourse (summarized in Table 1). Recalling our example in Figure 1, we focus on responses to questions: the *conversation act*, how the response is perceived to address the question (such as Zuckerberg saying he `cant_answer`); and the *communicative intent*, the sincere or deceptive intent behind choosing that form of response (such as one annotator believing the intent was `honest`). As our source of data, we choose the question-answer portions of U.S. congressional hearings (all in English) for several reasons: they contain political and societal controversy identifiable by crowdsourced workers, they have a strong signal of ambiguity as to the form and intent of

item	#sents/ turn	#toks/ turn	total sents	total toks	total spkrs
question	4.1	81.5	4096	82582	91
response	2.6	47.0	2634	48831	20

Table 1: Statistics of our 20 U.S. congressional hearings.

the response, and the data is plentiful.² A dataset statement is in Appendix D.

3.1 Dataset creation

Congressional hearings are held by committees to gather information about specific issues before legislating policies. Hearings usually include testimonies and interviews of witnesses. We focus on hearings that interview a single witness and that exceed a certain length (>100 turns) as a signal of argumentative discourse. To ensure a variety of topics and political leanings are included, we sample a roughly equal number of hearings from 4 Congresses (113th-116th) that span the years 2013-2019, for a total of 20 hearings. For each hearing, we identify a *question* as a turn in conversation containing a question posed by a politician that is immediately followed by a turn in conversation from the witness, which is the *response*. We thus extract the first 50 question-response pairs from each hearing. Each data point consists of a question followed by a response. Table 1 summarizes the dataset statistics.

3.2 Dataset annotation

We collect labels through the Amazon Mechanical Turk crowdsourcing platform. In the task, we ask a series of nested questions feasible for untrained annotators (from which we derive question response labels), then elicit annotator sentiment. Each HIT consists of five question-response pairs in sequential order from the same hearing; we group them to preserve continuity of the conversation while not overloading the annotator. We collect 7 judgments for each HIT.³ Screenshots of the task and the introductory example with all annotations are in Appendix A.

3.2.1 Annotations

For each question-response pair we collect three pieces of information: the question label, the re-

²Transcripts lack intonation and gestures, and thus a certain amount of information is lost from the original discourse.

³During our pilot, we experimented with increasing the number of judgments (up to 11) but found the number of chosen labels remains stable. We thus scaled back to 7.

(1) Q: How much of the financing was the Export-Import Bank responsible for?	R: We financed about \$3 billion.
(2) Q: If you were properly backing up information <i>required under the Federal Records Act, which would include the information she deleted from the server</i> , you'd have had all of those emails in your backup, <i>wouldn't you?</i>	R: All emails are not official records under any official records act.
(3) Q: So you're not willing to say that it doesn't meet due process requirements at this point?	R: Well, what I'd like to do is look at the procedures in place.

Table 2: Different question forms lead to different responses: (1) an information-seeking question leads to a direct answer. (2) A loaded question with a *presupposition* and *tag question* leads to an indirect answer because the responder rejects the presupposition. (3) A declarative question where the questioner commits to an unfavorable view of the responder leads to an indirect answer.

sponse label, and an explanation. At the end of each HIT, we collect two pieces of information: the annotator's sentiment towards the questioners, and sentiment towards the witness.⁴

Question We collect judgments on the question as it can influence the response. For example, an objective, information-seeking question lends itself to a direct answer (Table 2 example (1)). A *loaded* question with presuppositions can instead result in an indirect answer when rejecting these presuppositions (Walton, 2003; Groenendijk and Stokhof, 1984), as in example (2) of Table 2. *Leading* questions, often asked as declarative or tag questions, are conducive to a particular answer (Bolinger, 1957) and signal the questioner is making a commitment to that underlying proposition. A pragmatic listener, such as our annotator, is inclined to believe the questioner has reliable knowledge to make this commitment (Gunlogson, 2008). Challenging the commitment leads to indirect answers as in example (3) of Table 2.

To elicit the question intent without requiring familiarity with the described linguistic concepts, we ask the annotator a series of intuitive questions to decide if the question is an *attack* on the witness, *favoring* the witness, or is *neutral*. We use a rule-based classifier to determine the question type (*wh*, *polar*, *disjunctive*, *tag*, *declarative*).

Response For judging the response, we combine conversation acts with communicative intents as in Figure 2, in the spirit of the compositional semantic framework of Govindarajan et al. (2019). The taxonomy is a result of a combination of expert involvement, data observation and user feedback.⁵

⁴We elicit sentiments at the *end* because we do not expect annotators to be familiar with the hearing or conversants. Future annotations could elicit sentiments at the beginning to capture strong a priori biases in high-profile hearings.

⁵We consulted with existing taxonomies (SWBD-DAMSL Jurafsky et al. (1997), MRDA Shriberg et al. (2004), DialogBank Bunt et al. (2018), evasive rhetorical strategies in Gabrielsen et al. (2017), dialogue acts paired with content

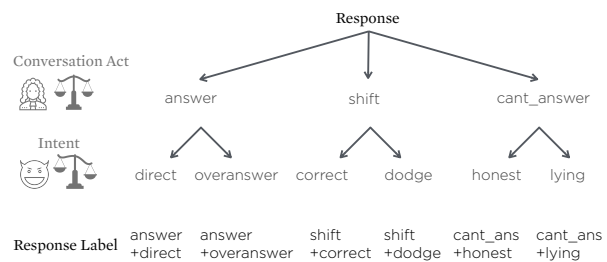


Figure 2: Hierarchical taxonomy of the perceived conversation act and intent for a response, forming the 6 response labels.

We next describe the taxonomy and its theoretical motivations.

In accordance with the discourse obligations of a conversation, a witness must respond in some form to a question (Traum and Allen, 1994). The function of the response is captured by the perceived *conversation act*, and is meant to be a more objective judgment (e.g., recognizing that Zuckerberg is using the ‘*can’t answer*’ form of a response, regardless of whether you believe him). This conversation act constitutes the top layer of the taxonomy. The conversation acts include the standard *answer* and *cant_answer*. Inspired by work on answerhood (Ginzburg et al., 2019; de Marnaffe et al., 2009; Groenendijk and Stokhof, 1984) and evasion in political discourse (Gabrielsen et al., 2017), we also include a more nuanced view of answering the question where giving a partial answer or answering a different question is labeled as *shift*.

The bottom layer of the taxonomy is the perceived *intent* underlying that conversation act, and is meant to be subjective. The intents hinge on whether the annotator believes the witness’s conversation act is sincere or not. For *answer*, the annotator may believe the intent is to give a *direct answer*, or instead an *overanswer* with the in-

features in Plüss and Piwek (2016)), and researchers in the dialogue field to construct the initial taxonomy, then conducted internal pilots with linguists and non-linguists, and finally conducted several iterations of an external pilot with crowdworkers to further refine the taxonomy.

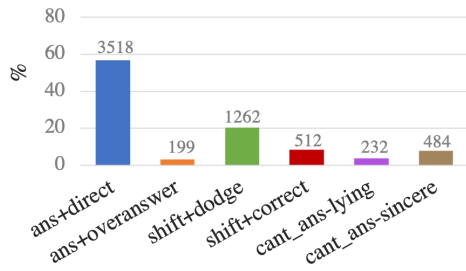


Figure 3: Distribution of response labels.

ment to sway the questioner (or even the public audience).⁶ If *shifting* the question, the annotator may believe the responder is *correcting* the question (e.g., to reject a false presupposition) or is *attempting to dodge* the question. If the witness says they *cant_answer*, the annotator may believe the witness is *honest* or is *lying*.

The annotation task implements a series of nested questions that mimic the hierarchy of the label taxonomy, which we map to conversation act and intent labels. That is, we first ask how the witness responds to the question (conversation act), then what is the intent and combine these into a single response label.

Explanation We ask annotators for a free-form *explanation* of their choices in order to elicit higher quality labels (McDonnell et al., 2016) and for use in the qualifying task as explained later.

Sentiment At the end of the HIT, we ask the annotator to rate their sentiment towards the politicians and towards the witness on a 7-point scale (we later collapse these into 3 levels: negative, neutral, positive). These ratings provide a rough proxy for annotator bias.

3.2.2 Worker qualification

Because the task requires significant time and cognitive effort, we establish a qualification process.⁷ In the qualifying task, we include question-response pairs already explained in the instructions, and unambiguous cases as far as the conversation act (e.g., a response of ‘Yes’ can only be construed as an answer). The criteria for qualification are: correctly labeling the conversation act for the instruction examples and unambiguous cases, providing explanations coherent with the intent label, and response times not shorter than the reading time.

⁶Overanswering with the intent to be *helpful* was included in our original taxonomy but then eliminated due to sparsity.

⁷This in addition to the requirements of >95% approval rating, >500 approved HITs, and living in the US for greater familiarity with the political issues.




Figure 4: Response label agreement (blue) and disagreement (between 2 or 3 labels in orange, green) for all hearings (a), a particular hearing with more (b) and less (c) disagreement.

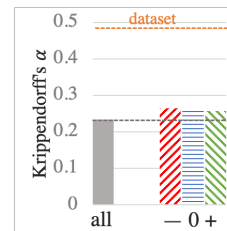


Figure 5: Response label IAA for all annotators (left) increases when grouped by sentiment (right); dashed orange line is response IAA for entire dataset.

This rigorous process yielded high quality data from 68 annotators who were genuinely engaged with the task. On average, an annotator labeled 91 question-response pairs, with 4 superannotators who provided labels for half of the data. During post-processing, we consider a label valid if it receives more than one annotator vote. The annotated dataset consists of 1000 question-response pairs with 6,207 annotations (3-7 annotations per item) on the first 50 question-responses from each of 20 congressional hearings.

3.3 Annotated Dataset Analysis

Here, we explore the annotated dataset to confirm its validity, focusing on the response labels (Figure 3) and sentiment towards the witness. We then conduct a word association analysis that finds meaningful lexical cues for the conversation act, but not for the intent label.

Is there disagreement? One initial question with collecting data on multiple interpretations is whether crowdworkers have sufficiently different viewpoints. However, we do find there is sufficient disagreement: Figure 4 (a) shows annotators disagree about the response label (the combined conversation act + intent) on roughly half the data (53.5%), though this trend can vary considerably from one hearing to the next as shown in (b) and (c).

Is disagreement real or noise? To understand whether disagreements are genuine or noise, we

Question	Response	Resp. Label	Anno.Sentiments
We used to be on the Small Business Committee, so we had definitions for what was a small business. What was yours?	What was my—say that again, sir?	cant_answer+ honest	[-1,-1,0,0,1,1]
And you spoke to their CEO immediately?	We shut down the app. We demanded—	shift+dodge	[-1,-1,0,0,1,1]

Table 3: All 6 annotators with different sentiments (negative: -1, neutral: 0, positive: 1) agree on the response label, even when it conflicts with their sentiment towards the witness (blue depicts sincere or positive, while red is deceptive or negative).

Response Label disagreement		%
ans+direct	vs. shift+dodge	15.9
shift+correct	vs. shift+dodge	8.2
ans+direct	vs. ans+overans	5.6
cant_ans+honest	vs. cant_ans+lying	4.9
ans+direct	vs. shift+correct	3.9
ans+direct	vs. shift+correct	3.9
	vs. shift+dodge	3.6

Table 4: Distribution of most frequent response disagreements; sincere intents in blue and deceptive in red.

examine the response label’s inter-annotator agreement (IAA) and which labels are disagreed upon.

We do not expect high IAA for the response label as we are eliciting disagreement. Overall, IAA is 0.494 in Krippendorff’s α (considered ‘moderate’; Artstein and Poesio (2008)), but importantly, we find higher agreement on the conversation act (0.652) compared to the intent (0.376). This finding confirms annotator understanding that the top-level label is more objective than the bottom-level one. We next group annotators with the same sentiments, expecting that when there is a disagreement, the same-sentiment groups will agree more with each other than with others. We partly confirm this intuition in Figure 5: grouping annotators by their sentiment increases agreement, but not by much. Sentiment is actually a more complicated signal, as we show in the following section.

Exploring annotator disagreements on the response label, we list the most frequent in Table 4. We find the disagreements often have opposing intents, but agree on the conversation act (e.g., shift+correct vs. shift+dodge). This result is encouraging, showing annotators have a shared understanding of the label definitions and further motivating our label taxonomy (Figure 2).

Is sentiment predictive of intent? We have pointed out how the annotator’s sentiment towards the witness can help explain the label they choose. Is annotator sentiment then an easy predictor of the intent label or is it a more complicated signal? A correlation study shows they are in fact only weakly correlated (correlation ratio $\eta = 0.34$ for coarse-

grained sentiment). There are two reasons for this result: (1) responses may have an unambiguous interpretation regardless of annotator sentiment, and (2) annotator sentiment towards the witness typically fluctuates throughout the hearing.

The most common unambiguous response is answer+direct (58%). Direct answers often leave little room for interpretation (e.g., ‘Yes, that is correct.’). More interestingly, annotators sometimes choose an intent that conflicts with their sentiment towards the witness (in 10% of unambiguous items). We illustrate the two cases in Table 3. In the first case, even the annotators with a negative view of the witness choose a sincere intent label. Conversely, in the second case, even the annotators with a positive view of the witness choose a deceptive intent label. While these are small phenomena, they illustrate the nuances of signaling sincerity and how they interact with the annotator’s sentiment towards the witness.

For the annotator’s sentiment across a hearing, a simplifying assumption is that it remains constant (recall the sentiment is reported at the end of each HIT, and HITs are presented to annotators in almost the same order as the original hearing). In practice it does not: 59% of annotators that label more than one HIT change their sentiment. As one annotator explained, “When he [the witness] said that, I got a different attitude towards him.”

Influence of question Earlier, we posited the question influences the response (Table 2). We find the question intent and type are weakly correlated with the response label. On a per-hearing basis, though, we observe stronger correlation for declarative question types in some hearings, partly confirming our hypothesis. We find qualitative evidence in explanations that annotators consider the question (“it was a terrible question to begin with”).

Lexical cues for labels To understand whether the response labels have lexical cues, we follow Schuster et al. (2019) to analyze the local mu-

	Annotator with positive sentiment	Annotator with negative sentiment
R: Congresswoman, it might be useful to <u>clarify</u> what actually happened. A developer who is a researcher–	Resp. Label: shift+correct Expl: Witness <i>wants</i> to <u>clarify</u> what happened.	Resp. Label: shift+dodge Expl: Mr. Zuckerberg <i>goes off on a tangent</i> to “clarify” the situation.
R: We are working through the process. <u>We have never said we would not provide those.</u>	Resp. Label: ans+direct Expl: Mr. Koskinen <i>answers</i> and does <i>say factually</i> that <u>they never said they would not provide the emails.</u>	Resp. Label: shift+dodge Expl: Koskinen <i>evades</i> the question, by saying that <u>he never said he wouldn’t provide the emails.</u>

Table 5: Explanations from annotators with opposing interpretations quoting the same response text (underlined) with subjective language in **blue** (neutral, positive) and **red** (negative).

Response label	top LMI n-grams
ans+direct	‘yes’, ‘be correct’
ans+overanswer	‘think’, ‘think that’
shift+dodge	‘we’, ‘-’
shift+correct	‘-’, ‘be something’
cant_ans+lying	‘I’, ‘?’, ‘not’
cant_ans+honest	‘I’, ‘not’, ‘?’

Table 6: n-grams with highest LMI scores in each label.

tual information (LMI) between labels and the response text n-grams (n=1,2,3). Unlike PMI, LMI highlights *high frequency* words co-occurring with the label. The top-scoring n-grams in Table 6 show most labels have a meaningful cue (the lower scoring words are not informative as they tend to be hearing-specific with much lower frequencies). The *ans+direct* cues signal straight answers. Dashes for both *shift* indicate the witness was interrupted (recall these include partial answers). Both *cant_answer* labels have the same cues, which include negation (to indicate not being able to answer) and question mark for clarification questions. We thus expect these cues may help identify conversation acts, but not the intents.

In summary, our analysis of the dataset shows there is ample and genuine disagreement. Interestingly, these disagreements are only partly attributable to differences in annotator sentiment. Furthermore, sentiment often fluctuates across a hearing, and can be influenced by what is said during the hearing. The question labels are not a straightforward signal for the response labels, but can vary by hearing. Finally, we find evidence of lexical cues for the conversation act label, but not for the intent.

3.4 Qualitative Analysis of Explanations

The explanations are a rich source of data for understanding annotator interpretations, with evidence they are applying personal beliefs (‘Bankers are generally evil’) and experiences (‘I have watched

hearings in congress’). We conduct a qualitative analysis to gain insight into the differing interpretations. Explanations are free-form, but annotators sometimes quote parts of the response. Interestingly, multiple annotators can quote the *same* text, yet arrive at *opposite* labels, as in Table 5. Studying these cases offers a window into what part of a discourse may trigger a subjective view, and how this view is expressed.

To this end, we examine the discourse and argumentative relations of the quoted text, and the linguistic devices used by the annotator to present the quote. We find the quoted text is often part of the response’s supporting argument, serving as the background or motivation that underpins the main claim. The annotator’s presentation of the quote differs drastically depending on their slant. Sincere labels use neutral or positive language (‘state’, ‘say factually’), whereas deceptive labels use negative words and framing (‘evades’, ‘goes off on a tangent’). Quotation marks in positive explanations become scare quotes in a negative one (first example in Table 5). On the negative side, we also find hedging (‘claim’) and metaphors (‘skirting the meaning’, ‘dances around’).

Our qualitative analysis shows annotators consider the side arguments underpinning the main claims, and employ rich linguistic devices to reflect their judgments.

4 Experiment

We propose the task of predicting all possible interpretations of a response (i.e., all perceived conversation act+intent labels) with the goals of analyzing discourse in a realistic setting and understanding sociolinguistic factors contributing to variations in discourse perception. We frame this task as a multi-label classification setting where 6 binary classifiers predict the presence of each of the 6 labels.⁸

⁸We experimented with a set-valued classifier that predicts the label set from all observed combinations (27-way multi-

We evaluate with macro-averaged F1 which gives equal weight to all classes, unlike micro-averaging which in our imbalanced data scenario (Figure 3) would primarily reflect the performance of the large classes.

4.1 Models

We experiment with pretrained language models with the intuition that a general language understanding module can pick up on patterns in the response to distinguish between the classes.

Training We split the data into 5 cross-validation folds, stratified by congressional hearing (to preserve the differing response distributions as seen in Figure 3). We reserve one fold for hyperparameter tuning and use the remaining 4 folds for cross-validation at test time.⁹

Baselines The ALL POSITIVE baseline predicts 1 for all labels. This baseline easily outperforms a majority baseline that predicts the most frequent label (`answer+direct`). LOG REGRESSION performs logistic regression with bag-of-words representations. CNN is a convolutional neural network as implemented in Adhikari et al. (2020). Other baselines performing lower than CNN are in Appendix C.

Pretrained We experiment with several pretrained language models, and find ROBERTA (Liu et al., 2019) performs the best on the held-out development fold. We use the implementation from Hugging Face.¹⁰ We feed in the tokenized response text and truncate input to 512 word pieces (additional inputs used in the model variants we describe next are separated by the `[SEP]` token).

Hierarchical We use two classifiers to mimic the hierarchy of our taxonomy: the first classifier predicts the conversation act while the second predicts the complete label (conversation act+intent). We train the classifiers independently, and condition the second classifier on the ground truth of the first classifier during training, only placing a distribution over intents consistent with that conversation act. At test time, we use predictions from the first classifier instead of ground truth.

+Question Building on top of the hierarchical model, this model incorporates the context of the class classification), but found this didn't work well.

⁹See Appendix B for training details and hyperparameters.

¹⁰<https://huggingface.co/transformers/>

question by including all interrogative sentences.¹¹

+Annotator This model incorporates annotators' coarse-grained sentiment towards the witness (fed in as a space-separated sequence of numbers, where each number is mapped from {negative, neutral, positive} sentiment to {-1, 0, 1}).

4.2 Results

The pretrained models easily outperform the baselines as seen in Table 7, where ROBERTA performs best. We next report results on incorporating hierarchy and context. Macro-F1 is calculated over the pooled results of the 4 folds; statistical significance is measured with the paired bootstrap test (Efron and Tibshirani, 1994) and $\alpha < 0.05$.

Model	macro-F1 (var)
ALL POSITIVE	35.0
LOG REGRESSION	40.5 (6.9)
CNN	46.0 (7.2)
BERT (uncased)	55.6 (3.8)
ROBERTA	58.5 (3.7)
HIERARCHICAL	58.2 (3.9)

Table 7: Results on the held-out fold's dev set, averaged across three random restarts.

Adding hierarchy As seen in Table 8, incorporating an additional classifier to predict the top-level conversation helps, but not significantly.¹² The per-class performance shows it mainly helps the less-represented conversation acts `shift` and `cant_answer`, with a better false negative rate for these classes. While the HIERARCHICAL model makes fewer errors of the kind intended to be corrected by the hierarchy as illustrated in Table 9 (by not predicting labels incompatible with the conversation act), the difference is very small. Jointly training these two classifiers with an adaptive learning curriculum may yield better results, which we leave for future work.

Adding context As shown in Table 8, adding the question in +QUESTION actually hurts performance, in particular by overpredicting the smaller classes `ans+overans` and `cant_ans+honest`. The lack of a benefit

¹¹We employ this truncation method because questions can be very lengthy (Table 1) We obtain poorer results using other forms of question context, including the entire question text, or only the last question.

¹²We nevertheless choose to build on this model as the subsequent models incorporating context exhibit more stable and significant differences.

Model	macro-F1	top-level class F1			bottom-level class F1					
		answer	shift	cant_ans	answer+direct	answer+overans	shift+dodge	shift+correct	cant_ans+lying	cant_ans+honest
ROBERTA	56.9	88.3	69.9	72.9	87.8	26.1	63.2	42.3	51.3	70.4
HIERARCHICAL	57.6 [†]	87.9	74.0	77.3	87.3	26.7	66.3	47.2	42.6	75.3
+QUESTION	56.4	87.9	76.0	75.3	87.6	23.8	65.5	45.6	43.8	72.0
+ANNOTATOR	60.5*	87.5	75.4	78.2	87.3	33.6	67.6	46.3	54.4	73.6

Table 8: Macro and class-level F1 on the test sets for the top-level (conversation act) and bottom-level (conversation act+intent) classes. † indicates not stat. sig. better vs. ROBERTA. * indicates stat. sig. better vs. HIERARCHICAL.

Context	Model Predictions		
	ROBERTA	HIERARCH.	+ANNOTATOR
Q: So my first question to you is, was the FBI aware of this Reddit post prior to offering Mr. Combetta immunity in May? R: I'm not sure. I know our team looked at it. I don't know whether they knew about it before then or not. Sentiments: 0, 0, -1, 0	<i>answer+direct</i> <i>cant_ans+lying</i> cant_ans+honest	<i>cant_ans+lying</i> cant_ans+honest	cant_ans+honest

Table 9: Example of model predictions (incorrect ones in *red*). Taking into account the hierarchy correctly eliminates labels for the absent conversation act of 'answer'. (Not shown: adding the question makes no corrections to this prediction). Adding the mostly neutral sentiments corrects the false positive for the lying intent, and is able to predict the entire label set correctly.

contradicts our expectations of the importance of the question and qualitative evidence, but is consistent with the weak correlation results. We hypothesize a different representation of the question is needed for the model to exploit its signal, which we leave for future work.

Incorporating the annotator sentiments in +ANNOTATOR provides a statistically significant benefit that helps both the false positive and false negative rate of the smaller classes *ans+overans* and *cant_ans+lying*. In the example of Table 9 which has mostly neutral sentiments, the model corrects the false positive made by the HIERARCHICAL model for *cant_ans+lying*.

From these results, we conclude that our task is heavily contextual with complex labels. On the one hand, taking into account the sentiments of the annotator leads to better predictions. On the other hand, we've shown annotator sentiment is not a simple reflection of intent. Furthermore, questions qualitatively influence the response labels, but linguistic features and labels of the question are not strongly correlated with the response and our model is not able to make effective use of it. The disagreements appear to reflect other axes, and this work begins to scratch the surface of understanding the subjective conversation acts and intents in conversational discourse.

5 Conclusion

In this paper, we tackle the *subjectivity* of discourse; that is, how ambiguities are resolved. We present a

novel English dataset containing *multiple* ground truths in the form of subjective judgments on the conversation acts and intents of a response in a question-response setting. We show the dataset contains genuine disagreements which turn out to be complex and not easily attributable to a single feature, such as annotator sentiment. The annotator rationales provide a window into understanding these complexities, and offer a rich source of linguistic devices. We propose a task to predict all possible interpretations of a response, whose results are consistent with our data analysis: incorporating the annotator bias helps the model significantly improve. We publicly release the dataset in hopes to spur further research by exploring the sequential nature of the hearings to employ CRF-type losses and other forms of aggregating annotator judgments.

6 Ethical Considerations

We provide a detailed dataset statement in Appendix D. The data collected in this dataset is produced by the U.S. government and is freely available to the public. The ids of the crowdsourced workers that contributed to the annotation are anonymized. Workers were compensated an average of \$1.20 per HIT (approximately \$8/hour), using the U.S. federal minimum wage as a minimum bar.

We recognize that crowdsourced workers, and thus the collected judgments in our dataset, are not representative of the U.S. population (Difallah et al., 2018).

Acknowledgments

We thank the annotators that contributed to this dataset. We thank reviewers and the first author's thesis committee for insightful feedback. We acknowledge the Texas Advanced Computing Center for grid resources. The first author was supported by the NSF Graduate Research Fellowship Program under Grant No. 2017247409.

References

- Ashutosh Adhikari, Achyudh Ram, Raphael Tang, William L. Hamilton, and Jimmy Lin. 2020. [Exploring the Limits of Simple Learners in Knowledge Distillation for Document Classification with DocBERT](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 72–77, Online. Association for Computational Linguistics.
- Ron Artstein and Massimo Poesio. 2008. [Survey Article: Inter-Coder Agreement for Computational Linguistics](#). *Computational Linguistics*, 34(4):555–596.
- Nicholas Asher and Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.
- Nicholas Asher and Soumya Paul. 2016. [Evaluating conversational success: Weighted message exchange games](#). In *Proceedings of the 20th Workshop on the Semantics and Pragmatics of Dialogue - Full Papers*, New Brunswick, NJ. SEMDIAL.
- Nicholas Asher and Soumya Paul. 2018. Strategic Conversations Under Imperfect Information: Epistemic Message Exchange Games. *Journal of Logic, Language and Information*, 27(4):343–385.
- John L Austin. 1962. How to Do Things with Words. *The William James Lectures delivered at Harvard University in 1955*.
- Dwight Bolinger. 1957. *Interrogative structures of American English: The direct question*, volume 28. Publication of the American Dialect Society.—University of Alabama Press.
- Peter Bull. 2008. “Slipperiness, Evasion, and Ambiguity”: Equivocation and Facework in Noncommittal Political Discourse. *Journal of language and social psychology*, 27(4):333–344.
- Harry Bunt, Volha Petukhova, Andrei Malchanau, Alex Fang, and Kars Wijnhoven. 2018. The DialogBank: dialogues with interoperable annotations. *Language Resources and Evaluation*, pages 1–37.
- Jonathan P Chang, Justin Cheng, and Cristian Danescu-Niculescu-Mizil. 2020. Don't let me be misunderstood: Comparing intentions and perceptions in online discussions. In *Proceedings of The Web Conference 2020*, pages 2066–2077.
- Zheqian Chen, Rongqin Yang, Zhou Zhao, Deng Cai, and Xiaofei He. 2018. Dialogue Act Recognition via CRF-Attentive Structured Network. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 225–234.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In *International Conference on Learning Representations*.
- David E Clementson. 2018. Deceptively dodging questions: A theoretical note on issues of perception and detection. *Discourse & Communication*, 12(5):478–496.
- Philip R. Cohen and C. Raymond Perrault. 1979. [Elements of a plan-based theory of speech acts](#). *Cognitive Science*, 3(3):177 – 212.
- Debopam Das, Manfred Stede, and Maite Taboada. 2017. [The Good, the Bad, and the Disagreement: Complex ground truth in rhetorical structure analysis](#). In *Proceedings of the 6th Workshop on Recent Advances in RST and Related Formalisms*, pages 11–19. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Scott Grimm, and Christopher Potts. 2009. [Not a simple yes or no: Uncertainty in indirect answers](#). In *Proceedings of the SIGDIAL 2009 Conference*, pages 136–143, London, UK. Association for Computational Linguistics.
- Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. [Demographics and dynamics of mechanical turk workers](#). In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM '18*, page 135–143, New York, NY, USA. Association for Computing Machinery.
- Esin Durmus and Claire Cardie. 2018. [Exploring the role of prior beliefs for argument persuasion](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1035–1045, New Orleans, Louisiana. Association for Computational Linguistics.
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. CRC Press.
- Jonas Gabrielsen, Heidi Jønch-Clausen, and Christina Pontoppidan. 2017. Answering without answering: Shifting as an evasive rhetorical strategy. *Journalism*, pages 1–16.
- Jonathan Ginzburg, Zulipiye Yusupujiang, Chuyuan Li, Kexin Ren, and Paweł Łupkowski. 2019. [Characterizing the response space of questions: a corpus study for English and polish](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 320–330, Stockholm, Sweden. Association for Computational Linguistics.

- Venkata Govindarajan, Benjamin Van Durme, and Aaron Steven White. 2019. Decomposing generalization: Models of generic, habitual, and episodic statements. *Transactions of the Association for Computational Linguistics*, 7:501–517.
- Jeroen Antonius Gerardus Groenendijk and Martin Johan Bastiaan Stokhof. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, Univ. Amsterdam.
- Christine Gunlogson. 2008. A question of commitment. *Belgian Journal of Linguistics*, 22(1):101–136.
- Thomas Hanneforth, Silvan Heintze, and Manfred Stede. 2003. Rhetorical Parsing with Underspecification and Forests. In *Companion Volume of the Proceedings of HLT-NAACL 2003 - Short Papers*, pages 31–33. Association for Computational Linguistics.
- Dan Jurafsky, Elizabeth Shriberg, and Debra Bisca. 1997. Switchboard SWBD-DAMSL labeling project coder’s manual. *Draft 13. Technical Report 97-02*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *International Conference on Learning Representations*.
- Yang Liu, Ivan Titov, and Mirella Lapata. 2019. Single document summarization as tree induction. 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 1745–1755, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. 2016. Why Is That Relevant? Collecting Annotator Rationales for Relevance Judgments. In *Fourth AAAI Conference on Human Computation and Crowdsourcing*.
- Tristan Miller, Maria Sukhareva, and Iryna Gurevych. 2019. A streamlined method for sourcing discourse-level argumentation annotations from the crowd. 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 1790–1796, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. What can we learn from collective human opinions on natural language inference data? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9131–9143, Online. Association for Computational Linguistics.
- Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T. Hancock. 2011. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 309–319, Portland, Oregon, USA. Association for Computational Linguistics.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent Disagreements in Human Textual Inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Brian Plüss and Paul Piwek. 2016. Measuring non-cooperation in dialogue. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1925–1936, Osaka, Japan. The COLING 2016 Organizing Committee.
- Massimo Poesio, Jon Chamberlain, Silviu Paun, Juntao Yu, Alexandra Uma, and Udo Kruschwitz. 2019. A Crowdsourced Corpus of Multiple Judgments and Disagreement on Anaphoric Interpretation. 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 1778–1789, Minneapolis, Minnesota. Association for Computational Linguistics.
- Martha E. Pollack. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. In *24th Annual Meeting of the Association for Computational Linguistics*, pages 207–214, New York, New York, USA. Association for Computational Linguistics.
- Christopher Potts. 2008. Indirect answers and cooperation: On Asher and Lascarides’s ‘Making the right commitments in dialogue’. *University of Michigan Linguistics & Philosophy Workshop on Implicatures*, 21:23.
- Hannah Rohde, Alexander Johnson, Nathan Schneider, and Bonnie Webber. 2018. Discourse coherence: Concurrent explicit and implicit relations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2257–2267, Melbourne, Australia. Association for Computational Linguistics.
- Stephanie Schoch, Diyi Yang, and Yangfeng Ji. 2020. “this is a problem, don’t you agree?” framing and bias in human evaluation for natural language generation. In *Proceedings of the 1st Workshop on Evaluating NLG Evaluation*, pages 10–16, Online (Dublin, Ireland). Association for Computational Linguistics.
- Merel Cléo Johanna Scholman. 2019. *Coherence relations in discourse and cognition: comparing approaches, annotations and interpretations*. Ph.D. thesis, Universität des Saarlandes.

- Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay. 2019. [Towards debiasing fact verification models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3419–3425, Hong Kong, China. Association for Computational Linguistics.
- John Rogers Searle. 1969. *Speech acts: An essay in the philosophy of language*, volume 626. Cambridge University Press.
- Elizabeth Shriberg, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The ICSI meeting recorder dialog act (MRDA) corpus. Technical report, INTERNATIONAL COMPUTER SCIENCE INST BERKELEY CA.
- David R. Traum and James F. Allen. 1994. [Discourse Obligations in Dialogue Processing](#). In *32nd Annual Meeting of the Association for Computational Linguistics*, pages 1–8, Las Cruces, New Mexico, USA. Association for Computational Linguistics.
- David R Traum and Elizabeth A Hinkelman. 1992. [Conversation acts in task-oriented spoken dialogue](#). *Computational Intelligence*, 8(3):575–599.
- Yannick Versley. 2011. [Multilabel Tagging of Discourse Relations in Ambiguous Temporal Connectives](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 154–161, Hissar, Bulgaria. Association for Computational Linguistics.
- Douglas Walton. 2003. The interrogation as a type of dialogue. *Journal of Pragmatics*, 35(12):1771–1802.
- Bonnie Webber and Aravind Joshi. 2012. [Discourse Structure and Computation: Past, Present and Future](#). In *Proceedings of the ACL-2012 Special Workshop on Rediscovering 50 Years of Discoveries*, pages 42–54, Jeju Island, Korea. Association for Computational Linguistics.
- Bonnie Webber, Rashmi Prasad, and Alan Lee. 2019. ["Ambiguity in Explicit Discourse Connectives"](#). In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 134–141, Gothenburg, Sweden. Association for Computational Linguistics.

A Annotation Task

Screenshots of the task are in Figures 6 and 7. For each HIT, we provide the hearing title, date and summary, along with titles of the politicians and witness. If there are intervening turns in the conversation that are not part of a question-response pair, we give the annotator the option to view the immediately preceding or following turn (e.g., see ‘Show/Hide Next Turn’ in Figure 6). Each HIT takes an average of 15 minutes to complete. To minimize context switching for annotators and roughly preserve the original conversation order, we publish only a small batch from one hearing at a time, waiting until it completes before publishing the sequentially next one or starting a new hearing.

Annotations that were collected are summarized in Table 10 for the question and the response, and in Table 11 for the HIT. The introductory example (Figure 1) is further labeled with *all* the annotations for illustrative purposes in Figure 8.

B Training Details

All models are trained with binary cross-entropy loss on an NVIDIA Quadro RTX 6000 GPU. For hyperparameter tuning, we search over the learning rates of [1e-5, 2e-5, 3e-5], warmup proportions of [0, 0.001, 0.01, 0.1] and weight decays of [0, 0.001, 0.01, 0.1]. We use early stopping based on development macro-F1 with a patience of 5 epochs and average results across 3 runs with different random initializations. For test, we train for 30 epochs and then evaluate on the test fold. If training does not improve by 40% in the first 10 epochs, then training is restarted.

For ROBERTA and all models that build on top of it, we use a learning rate of 3e-5, a warmup proportion of 0.1, a weight decay of 0.1 and batch size of 8, max sequence length of 512. The 4-fold cross-validation takes approximately 65 minutes.

C Model variants

Table 12 includes results with additional baselines, pretrained language models and adding other forms of context.

HAN is a Hierarchical Attention Network as implemented in Adhikari et al. (2020). LSTM is a regularized LSTM as implemented in Adhikari et al. (2020). For the pretrained models ALBERT (Lan et al., 2020) and ELECTRA (Clark et al., 2020), we use the implementations from Hugging Face.

For adding context, the models with * are the ones described in the main paper, and their cross-validation results are in Table 8. +ENTIRE QUESTION includes the *entire* question (not just the interrogative sentences as does +QUESTION). The +QUESTION INTENTS includes the annotators’ perceived intents of the question (fed in as a space-separated sequence of numbers, where each number is mapped from [attack, neutral, favor] to [-1, 0, 1]). The +FINE-GRAINED WITNESS SENTIMENT and +FINE-GRAINED QUESTIONER SENTIMENT include the 7-valued sentiment of the annotator towards witness (or questioner) (fed in the same style where the numbers are mapped from [very negative, negative, somewhat negative, neutral, somewhat positive, positive, very positive] to [-3, -2, -1, 0, 1, 2, 3]). The +COARSE-GRAINED QUESTIONER SENTIMENT includes the 3-valued annotator sentiment towards the questioner (mapping from [negative, neutral, positive] to [-1, 0, 1]).

Question-Response 5 of 5

Mr. Engel : So do you adjust your algorithms to prevent individuals interested in violence or nefarious activities from being connected with other like-minded individuals?

Mr. Zuckerberg : Sorry. Could you repeat that?

Show/Hide Next Turn

Please briefly explain all your choices for *Witness Intent*:

How does the politician pose their question(s)?
includes leading question, controversial assumption or attack? yes no

How does the witness respond?
 answers all reasonable questions
 answers partially or not all questions
 answers a different question
 says they don't understand/don't know/can't answer or asks a clarification question

Sentiment towards questioners and witness

In this congressional hearing:

- What is your overall sentiment towards the *questioners* (the congressional members asking the questions)?
 Very Negative Negative Somewhat Negative Neutral Somewhat Positive Positive Very Positive
- What is your overall sentiment towards the *witness*?
 Very Negative Negative Somewhat Negative Neutral Somewhat Positive Positive Very Positive

Figure 6: Overview of the crowdsourced task (only last of 5 question-response pairs shown for space).

How does the politician pose their question(s)?
includes leading question, controversial assumption or attack? yes no
includes question or statement favoring the witness?
 yes no

(a) question

How does the witness respond?
 answers all reasonable questions
 answers partially or not all questions
 answers a different question
 says they don't understand/don't know/can't answer or asks a clarification question

Witness provides extra info?
 yes
 to sway opinion in their favor
 to be helpful or add critical details
 no

(b) answer, overanswer

How does the witness respond?
 answers all reasonable questions
 answers partially or not all questions
 answers a different question
 says they don't understand/don't know/can't answer or asks a clarification question

Witness intent for giving a partial answer?
 to dodge an unfavorable question or question part
 for a justified reason

(c) partial answer

How does the witness respond?
 answers all reasonable questions
 answers partially or not all questions
 answers a different question
 says they don't understand/don't know/can't answer or asks a clarification question

Witness intent for answering a different question?
 to reformulate or evade an unfavorable question
 to correct the questioner or other justified reason

(d) shifted answer

How does the witness respond?
 answers all reasonable questions
 answers partially or not all questions
 answers a different question
 says they don't understand/don't know/can't answer or asks a clarification question

Witness intent for saying they can't answer?
 you believe they're lying
 you believe they're sincere

(e) can't answer

Figure 7: Nested questions that appear for the question (a) and for the response (b-e).

Annotation	Label	Example
Question	attack	Q: Could you take that \$250 million and ensure that every man, woman, and child in America has a CFPB tee shirt, ball cap, and koozie?
	favor	Q: So yes, you—exactly right. We weren't actually—we were expecting you to try to run out the clock like the last guy. But the—I want—I do want to congratulate you on your staff reduction of 0.0614 percent of your staff. So yes, that would be a sarcastic note to those that believe that you are gutting it all. I do want to give you an opportunity, though, to address a couple of things that were brought up. How many enforcement actions were taken under the former Bureau chief, Director Cordray in his first 6 months?
	neutral	Q: Thank you, Mr. Chairman. As I mentioned in my opening statement there are many reviews currently going on at the EPA, in the Inspector General's Office, Government Accountability Office, and other congressional committees, about some of these concerns you are hearing about today, Mr. Administrator, and that have been raised in the media. So my question is pretty easy. Will you commit the EPA will provide this committee with all the documents and information EPA produces for those inquiries?
Response	ans+direct	Q: So is what you are trying to do is make more information available or less information available? R: Yes, absolutely more information available.
	ans+overanswer	Q: OK. You have been attacked for flying first class. Is that illegal? R: Congressman, that was approved by the travel office and the security team at the EPA. I have since made changes to that. But that was—
	shift+correct	Q: To the public. So you are going to require that every one of these decisions or whatever they are based on, the data and the methodology as well as the conclusions are transparent and available to the public. Is that going to be on your website? How are we going to know this? R: Well, it is actually a proposed rule, Congressman. It is actually something that we are taking comment on, and I am sure there will be a wide array of comment on that very proposal. But the objective, once again, is to ensure transparency, reproducibility, with respect to the science that we rely upon in making our decisions in rulemaking.
	shift+dodge	Q: Well, you say that but that is not accurate. Do you know that manufacturers of methylene chloride paint strippers have been aware of deaths linked to this use for more than 28 years but continue to produce it? Yes or no. R: That is actually a solvent that we are considering under the—
	cant_ans+honest	Q: This is to get a little bit to the budget we are actually here to discuss, there is a program in your Agency called Leaking Underground Storage Tanks, short in acronym is LUST. The money that goes into that fund is supposed to be used to clean up or prevent leaks from underground storage tanks. To your knowledge, is there anything under current law that prevents a State from using it for other purposes? In other words, the money is supposed to be used to clean up these underground storage tanks, but my understanding is very few States use it for that purpose? R: You know, Congressman, I am not aware of that happening but it is something that we would investigate and look into if you have some information about that happening in your State and elsewhere.
	cant_ans+lying	Q: No. You answer to me whether it is, "yes" or "no." Your response? R: But I didn't quite catch the beginning of the question. I'm sorry.
	Explanation (free-form)	Q: Will you commit to working with Congress, and not against us, to make sure section 702 is reauthorized, either the way you want it or the way we want it? R: Congress gets to dispose; we get to give our opinion. Does not say if they are willing to commit to working with congress or not.

Table 10: Annotations and examples labeled for the question and the response.

Annotation	Labels
Sentiment towards questioners	very negative, negative, somewhat negative, neutral, somewhat positive, positive, very positive
Sentiment towards witness	very negative, negative, somewhat negative, neutral, somewhat positive, positive, very positive

Table 11: Annotations and labels for the HIT.

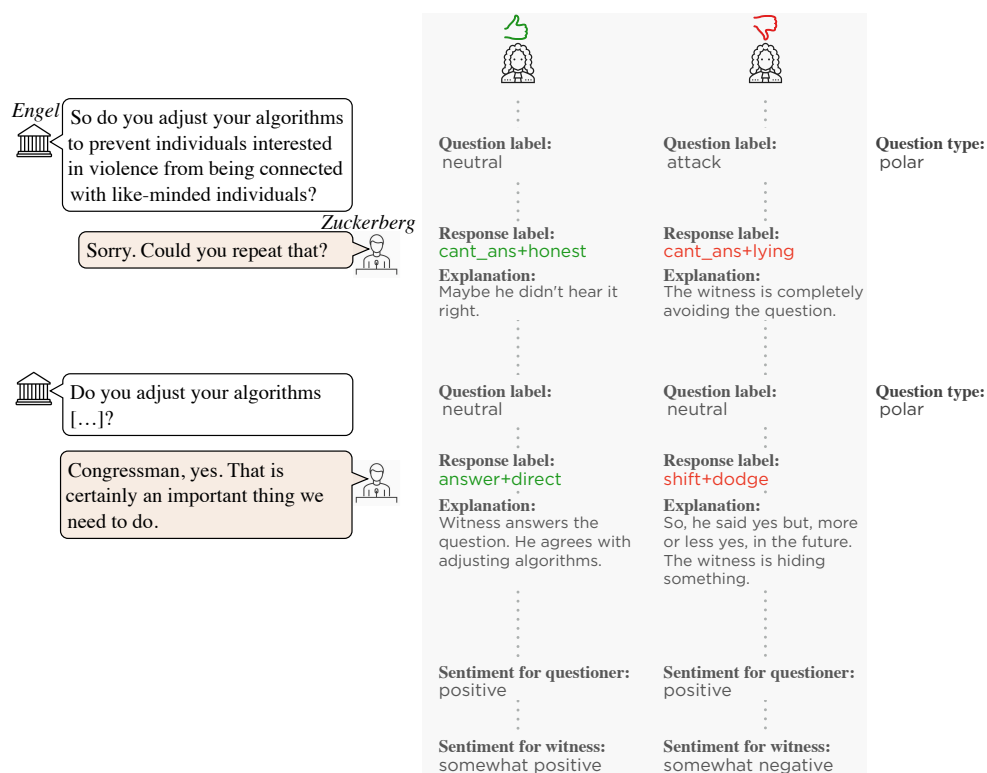


Figure 8: The introductory example with all annotations from two annotators with conflicting interpretations of the responses (the question type is determined by a rule-based classifier).

Model	macro-F1
LSTM	42.3 (6.6)
HAN	43.1 (11.9)
ALBERT	55.9 (8.0)
ELECTRA	55.8 (4.7)
+QUESTION* (All interrogatives)	57.6 (4.0)
+ENTIRE QUESTION	53.4 (22.0)
+QUESTION INTENTS	57.5 (18.0)
+ANNOTATOR* (Coarse-grained Witness)	62.0 (2.0)
+FINE-GRAINED WITNESS SENTIMENT	60.0 (2.0)
+FINE-GRAINED QUESTIONER SENTIMENT	57.7 (13.0)
+COARSE-GRAINED QUESTIONER SENTIMENT	58.9 (5.0)

Table 12: Results on the held-out fold’s dev set for additional baselines (top), pretrained language models (middle), and incorporating other contexts (bottom). The models with * indicate the contextual models described in the main paper (cross-validation results for these are in Table 8.)

D Data Statement

The latest version of the data statement is maintained at https://github.com/elisaF/subjective_discourse/blob/master/data/data_statement.md.

Data Statement for SubjectiveResponses

Data set name: SubjectiveResponses

Citation (if available): TBD

Data set developer(s): Elisa Ferracane

Data statement author(s): Elisa Ferracane

Others who contributed to this document: N/A

A. CURATION RATIONALE

The purpose of this dataset is to capture subjective judgments of responses to questions. We choose witness testimonials in U.S. congressional hearings because they contain question-answer sessions, are often controversial and elicit subjectivity from untrained crowdsourced workers. The data is sourced from publicly available transcripts provided by the U.S. government (<https://www.govinfo.gov/app/collection/chrhg>) and downloaded using their provided APIs (<https://api.govinfo.gov/docs/>). We download all transcripts from 113th-116th congresses (available as of September 18, 2019), then use regexes to identify speakers, turns, and turns containing questions. We retain hearings with only one witness and with more than 100 question-response pairs as a signal of argumentativeness. To ensure a variety of topics and political leanings, we sample hearings from each congress and eliminate those whose topic is too unfamiliar to an average American citizen (e.g. discussing a task force in the Nuclear Regulatory Commission). This process yields a total of 20 hearings: 4 hearings from the 113th congress (CHRG-113hhrhg86195, CHRG-113hhrhg88494, CHRG-113hhrhg89598, CHRG-113hhrhg93834), 5 hearings from the 114th (CHRG-114hhrhg20722, CHRG-114hhrhg22125, CHRG-114hhrhg26003, CHRG-114hhrhg95063, CHRG-114hhrhg97630), 7 hearings from the 115th (CHRG-115hhrhg25545, CHRG-115hhrhg30242, CHRG-115hhrhg30956, CHRG-115hhrhg31349, CHRG-115hhrhg31417, CHRG-115hhrhg31504, CHRG-115hhrhg32380), and 4 hearings from the 116th (CHRG-116hhrhg35230, CHRG-116hhrhg35589, CHRG-116hhrhg36001, CHRG-116hhrhg37282). For annotation, we then select the first 50 question-response pairs from each hearing.

Code used to create the dataset is available at https://github.com/elisaF/subjective_discourse.

B. LANGUAGE VARIETY/VARIETIES

- BCP-47 language tag: en-US
- Language variety description: American English as spoken in U.S. governmental setting

C. SPEAKER DEMOGRAPHIC

- Description: The speakers are from two groups: the questioners are politicians (members of Congress) and the witnesses can be politicians, businesspeople or other members of the general public.
- Age: No specific information was collected about the ages, but all are presumed to be adults (30+ years old).
- Gender: No specific information was collected about gender, but members of Congress include both men and women. The witnesses included both men and women.
- Race/ethnicity (according to locally appropriate categories): No information was collected.
- First language(s): No information was collected.
- Socioeconomic status: No information was collected.
- Number of different speakers represented: 91 members of Congress and 20 witnesses.
- Presence of disordered speech: No information was collected but none is expected.

D. ANNOTATOR DEMOGRAPHIC

Annotators:

- Description: Workers on the Amazon Mechanical Turk platform who reported to live in the U.S. and had a >95% approval rating with >500 approved HITs were recruited during the time period of November 2019 - March 2020.
- Age: No information was collected.
- Gender: No information was collected.

- Race/ethnicity (according to locally appropriate categories): No information was collected.
- First language(s): No information was collected.
- Training in linguistics/other relevant discipline: None.

Annotation guideline developer:

- Description: Elisa Ferracane
- Age: 40.
- Gender: Female.
- Race/ethnicity (according to locally appropriate categories): Hispanic.
- First language(s): American English.
- Training in linguistics/other relevant discipline: PhD candidate in computational linguistics.

E. SPEECH SITUATION

- Description: Witness testimonials in U.S. congressional hearings spanning the 114th-116th Congresses.
- Time: 2013-2019
- Place: U.S. Congress
- Modality (spoken/signed, written): transcribed from spoken.
- Scripted/edited vs. spontaneous: mostly spontaneous, though members of Congress sometimes read questions they have written down
- Synchronous vs. asynchronous interaction: synchronous
- Intended audience: the U.S. government and the general public, as all hearings are both transcribed and televised

F. TEXT CHARACTERISTICS

The genre is political discourse in a highly structured setting where a chairperson runs the meeting, and each member of Congress is afforded 5 minutes to question the witness but can yield their time to others. Topics vary based on the congressional committee that is holding the hearing, and include oversight of other governmental bodies (e.g.,

IRS, Department of Justice) and inquiries into businesses suspected of misconduct (e.g., FaceBook, Wells Fargo).

G. RECORDING QUALITY

N/A

H. OTHER

N/A

I. PROVENANCE APPENDIX

N/A

About this document

A data statement is a characterization of a dataset that provides context to allow developers and users to better understand how experimental results might generalize, how software might be appropriately deployed, and what biases might be reflected in systems built on the software.

Data Statements are from the University of Washington. Contact: datastatements@uw.edu. This document template is licensed as [CC0](https://creativecommons.org/share-your-work/public-domain/cc0/).

This version of the markdown Data Statement is from June 4th 2020. The Data Statement template is based on worksheets distributed at the [2020 LREC workshop on Data Statements](https://sites.google.com/uw.edu/data-statements-for-nlp/), by Emily M. Bender, Batya Friedman, and Angelina McMillan-Major. Adapted to community Markdown template by Leon Dercyznski.