

402 **Supplementary Material**

403 **Contents**

404	A More Experimental Results for Pragmatic Identification and Reasoning	12
405	B Annotation Details	12
406	B.1 Details For Automatic Selection	12
407	B.2 Details For Fine-grained Annotation	12
408	B.3 Details on Human Refinements	13
409	C Experimental Detail	13
410	C.1 Pragmatic Identification and Reasoning (PIR)	13
411	C.2 Conversational Question Answering (CQA)	14
412	D More Detail on DiPlomat	15
413	E Grice Maxims and Pragmatic Reasoning	22
414	F Computational Resources	22
415	G Limitations & Negative Societal Impacts	22
416	H Ethics Concern	22
417	I Responsibility & Dataset Liscence	23
418	J Datasheets for Our Dataset	23
419	J.1 Motivation	23
420	J.2 Composition	23
421	J.3 Collection Process	25
422	J.4 Preprocessing/cleaning/labeling	26
423	J.5 Uses	26
424	J.6 Distribution	26
425	J.7 Maintenance	27

426 A More Experimental Results for Pragmatic Identification and Reasoning

427 We conducted additional experiments on PIR employing various models, including RoBERTa_{large} [50],
428 DeBERTa_{base} [65], and ALBERT_{base} [66]. The training and testing procedures remained consistent
429 with the aforementioned models described in the main body. All experimental results have been
430 compiled and presented in Table 6. Analysing newly proposed result, it’s obvious to observe that our
conclusions mentioned in the main body still hold.

Table 6: Pragmatics Identification and Reasoning Results. The numerical results are accuracy scores in their percentage.

	C → P	CP → R	C → PR
Random	50	20	10
BERT _{base}	63.2 ± 1.1	91.3 ± 0.7	50.2 ± 6.8
RoBERTa _{base}	64.4 ± 1.3	92.0 ± 0.4	50.0 ± 11.28
RoBERTa _{large}	63.8 ± 0.0	60.8 ± 0.5	0.0 ± 0.0
GPT-2 _{base}	64.4 ± 0.7	90.9 ± 0.9	13.06 ± 1.1
DialoGPT _{medium}	65.0 ± 0.6	24.5 ± 1.9	3.8 ± 1.5
DeBERTa _{base}	64.9 ± 0.2	92.6 ± 0.6	43.9 ± 1.2
ALBERT _{base}	65.1 ± 0.4	90.6 ± 0.2	34.9 ± 1.8

431

432 B Annotation Details

433 B.1 Details For Automatic Selection

434 Different methodologies are employed to address various pragmatic phenomena. To leverage prior
435 advancements in the field, we begin by segmenting each dialogue into individual utterances. Sub-
436 sequently, we employ two distinct approaches, namely string matching and pretrained model clas-
437 sification, to identify these phenomena within our source data. In the case of scalar implicature,
438 which exhibits a noticeable pattern characterized by word pairs such as (*some, all*) appearing in
439 adjacent turns of dialogues, we employ string matching to annotate instances of scalar implicature in
440 conversations. Similarly, for **popcq implicature**, which often features a continuous question mark,
441 we utilize this characteristic as a means of detection. With regards to idioms, which exhibit more
442 evident patterns, we employ the idiom set proposed by Saxena and Paul [18] to conduct searches.
443 For other types of phenomena that lack obvious patterns, we leverage a pretrained RoBERTa base
444 model [50], and fine-tune it for our specific task. The sarcasm dataset by Misra [67] is used for
445 finetuning the sarcasm model, the MOVER dataset by Zhang and Wan [42] for hyperbole and the
446 ColBERT dataset by Annamoradnejad and Zoghi [20] for paronomasia. Several models have been
447 proposed for metaphor detection, thus we utilize an existing model [68] specifically designed for
448 metaphor identification.

449 **Topic Segmentation** The original dialogues employed in our study consist of lengthy and multi-
450 turn exchanges, which are **ill-suited** for our research objectives. Consequently, we implement a
451 segmentation process to break down these dialogues into shorter units. To achieve this, we employ
452 two techniques, namely BERTScore [52] and TextTiling [69]. The segmentation procedure starts
453 with computing the BERTScore between adjacent turns and subsequently applying the TextTiling
454 algorithm to the generated BERTScores.

455 B.2 Details For Fine-grained Annotation

456 AMT is integral to our process. To ensure clarity and consistency, we provide explicit instructions
457 to the workers. Additionally, to further elucidate the objectives of our study, we offer illustrative
458 examples. The task itself is presented below the instructions and examples, with the dialogue and
459 corresponding turn numbers provided for workers to select. Furthermore, as workers check a checkbox,

460 we prompt them to select a confidence score and provide a rationale. In order to strike a balance
461 between our budget, the quality of annotations, and the speed of annotation, we have determined the
462 compensation of \$0.1 per completed task. The whole view of the worker interface is presented in
463 Figure 7. After the annotation process, we collect responses that are assigned with a confidence score
464 of 4 or higher.

465 **Specifically, we surveyed 10 users to accomplish our task. All users can complete a single task within**
466 **45 seconds, leading to a wage pay of around 8 dollars per hour, which is about a dollar higher than**
467 **the federal minimum hourly wage of the United States.**

468 B.3 Details on Human Refinements

469 Disturbing choices are chosen based on the BERTScore metric [52]. The rationale with the highest
470 similarity, as determined by other dialogues, is selected and included in the pool of candidate options.
471 The instructions provided to the workers align with those used for Fine-grained Annotation, wherein
472 they are also instructed to assign a confidence score to their responses. The remuneration for workers
473 is set at \$0.05 per task. The worker interface is included in Figure 8.

474 **AMT Workers Requirements** In order to guarantee the quality of annotated data, the qualification
rules for workers are strict and can be found in Table 7.

Table 7: AMT workers requirements

Country _{In}	United States, Canada, Great Britain, Australia, Singapore, Ireland, New Zealand
# Tasks approved _{GreaterThanOrEqualTo}	1300
Tasks approved Rate _{GreaterThanOrEqualTo}	95%

475

476 C Experimental Detail

477 C.1 Pragmatic Identification and Reasoning (PIR)

478 **BERT_{base}** [70] BERT (Bidirectional Encoder Representations from Transformers) is a revolutionary
479 language representation model that has had a significant impact on natural language processing (NLP)
480 tasks. It has achieved remarkable performance across various NLP benchmarks, including question
481 answering, sentiment analysis, named entity recognition, and many others. Its birth brings profound
482 influence on pretrained language models.

483 **RoBERTa_{base} & RoBERTa_{large}** [50] RoBERTa improves upon BERT by incorporating enhance-
484 ments such as larger and more diverse training data, longer pretraining duration, dynamic masking,
485 and advanced training strategies. These improvements enable RoBERTa to achieve even better
486 performance on a wide range of NLP benchmarks. While BERT paved the way for contextualized rep-
487 resentations in NLP, RoBERTa further refines and pushes the boundaries of language understanding,
488 making it a powerful and preferred choice for many researchers and practitioners in the field.

489 **ALBERT_{base} & ALBERT_{large}** [66] ALBERT (A Lite BERT) is a highly efficient and compact
490 variant of the BERT model that addresses the computational limitations of the original architecture.
491 It incorporates parameter-reduction techniques to alleviate training time constraints and achieve
492 improved performance compared to BERT.

493 **DeBERTa_{base}** [62] DeBERTa (Decoding-enhanced BERT with Disentangled Attention) is a state-
494 of-the-art language representation model that builds upon the BERT architecture and introduces
495 several key innovations, including disentangled attention mechanism. The performance of DeBERTa
496 has been demonstrated to surpass that of BERT on a wide range of NLP tasks.

497 **GPT2_{base}** [71] Leveraging transformers decoder, Radford et al. [71] proposed GPT2. It represents
498 a significant breakthrough in natural language processing and generation. One of the most notable

Table 8: Hyperparameters for models on $\mathbf{CP} \rightarrow \mathbf{R}$

Model	learning rate	batch size	weight decay	epochs
BERT _{base}	5e-5	12	0.001	50
BERT _{large}	5e-5	12	0.001	50
ALBERT _{base}	5e-5	12	0.001	50
ALBERT _{large}	5e-5	12	0.001	50
DeBERTa _{base}	5e-5	12	0.001	50
RoBERTa _{base}	5e-5	12	0.001	50
RoBERTa _{large}	5e-5	12	0.001	50
GPT2 _{base}	0.001	8	0.01	50
DialogGPT _{medium}	0.001	2	0.01	50

Table 9: Batch size for models on $\mathbf{C} \rightarrow \mathbf{P}$

Model	Batch Size
BERT _{base}	80
ALBERT _{base}	24
ALBERT _{large}	24
DeBERTa _{base}	24
RoBERTa _{base}	80
RoBERTa _{large}	24
GPT2 _{base}	24
DialogGPT _{medium}	8

500 features of GPT-2 is its ability to generate coherent and contextually relevant text. Through unsu-
 501 pervised pretraining on a large corpus of internet text, GPT-2 learns to predict the next word in a
 sequence of text, enabling it to generate human-like responses.

502 **DialogGPT_{medium}** [35] DialogGPT is dialogue-oriented GPT. It builds upon the GPT architecture
 503 and extends it to support interactive conversations. DialogGPT is trained in a supervised manner
 504 using a dialogue dataset, which allows it to understand and generate responses in a conversational
 505 context.

506 The PIR task encompasses three distinct settings: $\mathbf{C} \rightarrow \mathbf{P}$, $\mathbf{CP} \rightarrow \mathbf{R}$, and $\mathbf{C} \rightarrow \mathbf{PR}$. In the $\mathbf{C} \rightarrow \mathbf{P}$
 507 setting, models are trained for 20 epochs, employing a batch size as indicated in Table 9, a learning
 508 rate of $2e - 5$, and weight decay of 0.01. As for $\mathbf{CP} \rightarrow \mathbf{R}$, the hyperparameters adopted are listed
 509 in Table 8. For the $\mathbf{C} \rightarrow \mathbf{PR}$ setting, there is no training required; instead, we simply load the best
 510 checkpoint obtained from the previous training for this task. The concrete implementation is as
 511 follows: we initially flatten the test dataset of $\mathbf{C} \rightarrow \mathbf{P}$, ensuring that each instance contains both a
 512 dialogue and a pragmatic turn extracted from the same dialogue. As for the test dataset of $\mathbf{CP} \rightarrow \mathbf{R}$,
 513 no modifications are made. It should be noted that, following the processing steps, both datasets own
 514 the same dialogues and corresponding pragmatic turns, resulting in identical instance numbers. For
 515 an instance to be deemed correct, the models must successfully accomplish both component tasks *i.e.*
 516 succeed in Identification and Reasoning.

517 C.2 Conversational Question Answering (CQA)

518 **CQA** ChatGPT was instructed to generate questions for our tasks. The prompt template that starts
 519 the questions with "Which" is depicted in Table 10. Through this methodology, we collected a total
 520 of 19,482 questions. To ensure the reliability of the answers provided to these questions, AMT is
 521 utilized. The task template is demonstrated in Figure 9. In our experiment, the hyperparameters
 522 adopted are illustrated in Table 11. To assess the performance of ChatGPT, we conducted testing
 523 using the template outlined in Table 13.

524 **Zero-Shot Natural Language Inference** Details are provided as follows. T5-XXL, and DeBERTa-
 525 v3 are tested with the pragmatic turn as premise and implied meaning as a hypothesis. The context

Table 10: ChatGPT question generation template: using "Which" to start the question.

You are sensitive and always view others' words as having some implied meanings.
For the dialogue between "A" and "B" in this task, we have offered a statement that is the implied meaning of a turn, please only offer one reading comprehension question that can be answered with only one word based on the dialogue and mostly focuses on the turn the statement mentions.
The question will be tested by only by viewing the dialogue, so please make the question hard enough that it's impossible to answer without viewing the statement.
Use "Which" to ask the question!
Following is the dialogue:
{dialogue}
Following is the statement:
{statement}
Use "Which" to ask the question! And please make the question hard enough that it's impossible to answer without viewing

Table 11: Hyperparameters for models on CQA.

Training Epoch	50
Learning Rate	$5.6e - 5$
Batch Size	24
Weight Decay	0.001

Table 12: Test ChatGPT: answer questions with only one word.

For the dialogue between "A" and "B" in this task, please answer a question according to the dialogue with only one word
Following is the dialogue:
{dialogue}
Following is the question: {question}

Table 13: ChatGPT test template of Zero-Shot CoT

This is a natural language inference task. Given the dialogue context: {context} Does {pragmatic turn} entails {implied meaning}? Reply 'entails' or 'not entails'.

Think step by step.

526 is out of reach for these models. In contrast, as shown in Table 10, ChatGPT is given the context,
527 and the red line labeled "Think step by step" represents two distinct configurations: one with
528 step-by-step and one without it.

529 **D More Detail on DiPlomat**

530 In this section, we will propose more examples of our dataset in Table 14, Table 15, Table 16, Table 17,
531 and Table 18.

Table 14: Contextual reasoning examples of **DiPlomat**

<p>A: Yeah. They say that he's the fastest pitcher there ever was. It's just he really couldn't find home plate. I mean, some of the stories you learn about this guy, it reads like fiction. When he was - I think this is around 1960. He's pitching in the minor leagues, and he pitched so fast he ripped the man's ear off.</p> <p>B: Oh.</p> <p>A: Yeah.</p>	<p>Rationale: The literal meaning is a simple expression of agreement, while the implied meaning is that the speaker is amazed by the story of Steve Dalkowski's feats.</p>
<p>B: We're talking about 2.8 million people. Has the rise of temporary workers figured into, at least, the statistical improvement of the U. S. economy for some people?</p> <p>A: It has. Overall, about one seventh of the total job growth has been in the temp sector. The temp sector is growing nine times faster than the overall private sector as a whole. And the 2.9 million workers represents a record number, both in the number of temp workers and in the percentage of the economy that they make up.</p> <p>B: You know in "Harvest Of Shame," Edward R. Murrow very famously said, the people we're showing you in this documentary have picked your Thanksgiving bounty with their bare hands, and this is how they live.</p>	<p>Rationale: The implied meaning of this turn is to reflect on our reliance on temporary workers in our day-to-day lives.</p>
<p>A: And so I got up and ran. And it wasn't too far. But I just - at that moment, I thought, I don't want to be shot in the back, and I need to find some cover. And there's really no place to hide. But there are these</p> <p>B: You found a little, like, alcove that you could duck into.</p> <p>A: There was a little alcove, yeah. And I just made myself as small as I could in that little corner.</p>	<p>Rationale: The speaker tried to protect itself from danger.</p>
<p>A: Well, there's a big argument in the United States about this. There's one group of folks who think that engagement policy failed. We engaged with China from 1979 until about 2013 when Xi Jinping came into power. And the idea of engagement was that coevolution was in the American interest as well as in China's interest. And you could bring China along to be a responsible player to some degree.</p> <p>A: Many hardliners in the United States government - and outside and including in the expert community - now claim that engagement was a sucker's game and that we have raised up a tiger which could now devour us. But there are different schools of thought about this, and many of us think that we still need to engage with China, albeit more strategically.</p> <p>B: That image of raising a tiger that will devour us is very dramatic.</p>	<p>Rationale: The situation is not necessarily an 'either/or' between China and the United States.</p>

Identify Implicature in Dialogue

Instructions

For the dialogue between "A" and "B" in this task, follow these steps :

1. Read through the dialogue
2. For **each turn** of the dialogue , identify whether its **actual meaning** is different from its **literal meaning**, such as:
 - o *Bob is a couch potato.* **implies** that "Bob sits on the sofa all day for watching TV" but not Bob is a potato.
 - o *I am so hungry that I can eat ten elephants* **implies** that "I am extremely hungry but not I will eat ten elephants".
 - o *Zombies eat brains. You're safe* **implies** that "you do not have a brain".
 - o *I'd agree with you, but then we'd both be wrong* **implies** that "you are wrong".
3. **Check the checkbox** below the dialogue corresponding to the turns that meet the condition.
4. (Confident score)When a checkbox is checked we will ask you to choose a number(1 to 5) to represent how confident you are about choosing the turn.(Higher score , more confidence)
5. Write a brief but more than 8 words implied meaning. If you can't find one , simply write **None**

Please complete the HIT carefully , and note that :

- Please read our examples!
- There may be several turns that meet the conditions , select **all of them**.
- You must **at least choose one** of the turns!
- Some actual meanings are **hard to find**, so please read patiently and carefully!
- Confident score, implied meaning **should not** leave to blank!
- Confident score doesn't represent how confident you are about your **answer**
- Confident score marks how confident you are that the turn **has implied meanings**:

Warning : Choosing too many answers randomly will cause us to mark you as unqualified worker!

Warning : Writing reasons irrelevant will also cause us to mark you as unqualified worker!

Warning : Using ChatGPT or AI methods will cause us to mark you as unqualified worker (this is strict, I will block you if there is a single suspicious hit)

Examples

(1)	A:	Did you drink the milk I kept on the table?
(2)	B:	The cat seems to be happy.

1

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

(1)
 (2)

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write a implied meaning (More than 8 words!) :

The cat seems happy implies that B thinks that the cat drinks the milk.

(1)	A:	Who made these donuts?
(2)	B:	I made some of these donuts.
(3)	A:	Ok,would you like to send some of them to Mr.Potter?
(4)	B:	I have homework to do.

2

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

(1)
 (2)

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write a implied meaning (More than 8 words!) :

"some" represents not all, B means that he has only make some of the donuts not all of the donuts.

(1)	A:	Bob, are you sure you can take care of yourself this weekend?
(2)	B:	Mom, can a duck swim?

3

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

(1)
 (2)

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write implied meaning (More than 8 words!) :

Duck can swim is for sure impling that I can take care of myself is for sure.

(1)	A:	Do you like her?
(2)	B:	She's like cream in my coffee.

4

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

(1)
 (2)

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write a implied meaning (More than 8 words!) :

Cream is wonderful impling that I like her a lot.

Task

Following is the **dialogue**:

(0)	A:	...
(1)	B:	...
(2)	B:	...
...

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

- (0)
 (1)
 ...

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write a implied meaning (More than 8 words!) :

Write reason here

Figure 7: Fine-grained annotation worker interface

Multiple Choice For Implicature In Dialogue

Instructions

For the dialogue between "A" and "B" in this task, follow these steps :

1. Read through the dialogue
2. There may be turns in dialogue that their **actual meanings** are different from their **literal meanings**, such as:
 - o *Bob is a couch potato.* **implies** that "Bob sits on the sofa all day for watching TV" but not Bob is a potato.
 - o *I am so hungry that I can eat ten elephants* **implies** that "I am extremely hungry but not I will eat ten elephants".
 - o *Zombies eat brains. You're safe* **implies** that "you do not have a brain".
 - o *I'd agree with you, but then we'd both be wrong* **implies** that "you are wrong".
3. Under each dialogue, there may be several statements. Each statement will give you a turn number telling you the corresponding turn's actual meaning is **different** from its literal meaning and offer you the implied meaning or reason.
4. **Check the checkbox** before statements which you consider it reasonable.
5. (Confident score)When a checkbox is checked we will ask you to choose a number(1 to 5) to represent how confident you are about choosing the statement.(Higher score , more confidence)

Please complete the HIT carefully , and note that :

- Please read our examples!
- There may be several statements that meet the conditions , select **all of them**.
- There are **disturbance statements**, so please read carefully.
- You must **at least choose one** of the statements!
- Some actual meanings are **hard to find**, so please read patiently and carefully!
- Confident score **should not** leave to blank!

Warning : Choosing disturbance choices too many times will cause us to mark you as unqualified worker, so please read carefully!

Warning : Choosing only one or two choices with much more choices presented to avoid selecting disturbance ones will also cause us to mark you as unqualified workers!

Examples

(1)	A:	Did you drink the milk I kept on the table?
(2)	B:	The cat seems to be happy.

1

Please choose the **statements that mark the turns with implied meaning correctly and have reasonable reason:**

turn 2 : "The cat seems to be happy" implies that the cat is delighted.

turn 2 : "The cat seems to be happy" implies that I did not drink the milk and I think the cat might drink it

Please choose a confidence score :

5 : You are totally sure that this statement is correct and believe that everyone will agree with you

(1)	A:	Who made these donuts?
(2)	B:	I made some of these donuts.
(3)	A:	Ok,would you like to send some of them to Mr.Potter?
(4)	B:	I have homework to do.

2

Please choose the **statements that mark the turns with implied meaning correctly and have reasonable reason:**

turn 1 : "Who made these donuts?" implies that A wants to eat donuts, A is hungry.

turn 2 : "I made some of these donuts" implies that I did not make all of these donuts I only make a part of them.

Please choose a confidence score :

5 : You are totally sure that this statement is correct and believe that everyone will agree with you

turn 4 : "homework" implies that the work is done at home not at school

turn 4 : B has homework to do, so he/she is not able to help A.

(1)	A:	Bob, are you sure you can take care of yourself this weekend?
(2)	B:	Mom, can a duck swim?

3

Please choose the **statements that mark the turns with implied meaning correctly and have reasonable reason:**

turn 1 : Mom is not sure whether B can take care of itself.

turn 2 : Duck can swim is for sure implying that I can take care of myself

Please choose a confidence score :

5 : You are totally sure that this statement is correct and believe that everyone will agree with you

(1)	A:	Do you like her?
(2)	B:	She's like cream in my coffee.

4

Please choose the **statements that mark the turns with implied meaning correctly and have reasonable reason:**

turn 1 : She looks like cream in my coffee

turn 2 : Cream is wonderful implying that I like her a lot.

Please choose a confidence score :

5 : You are totally sure that this statement is correct and believe that everyone will agree with you

Task

Following is the **dialogue**:

(0)	A:	...
(1)	B:	...
(2)	B:	...
...

Please choose the turns whose **actual meanings** are different from their **literal meanings**.

- turn 0 : [...]
- turn 1 : [...]
- ...

Please choose a confidence score :

5 : You are totally sure that this turn has implied meanings and believe that everyone will agree with you

Please write a implied meaning (More than 8 words!) :

Write reason here

Figure 8: Human refinements worker interface

Dialogue Question Answering

Instructions

For the dialogue between "A" and "B" in this task, follow these steps:

1. Read our example!
2. Read through the dialogue
3. Under each dialogue, there may be several questions.
4. There is a useful reference statement that may help you answer the corresponding question, please read it carefully
5. Answer each question with only one word!

Please complete the HIT carefully, and note that :

- i. *Reference Statements are useful, but don't just rely on it, read the dialogue as well!*
- ii. *Our turn number start from 0*
- iii. *There are incorrect reference statements.*

Warning : If your answer is close to incorrect reference statement, but far away from the dialogue, we will mark you as unqualified workers!

Examples

Following is the **dialogue**:

(0)	B:	How did you feel when you found out, this summer, about the abuse of children that was going on in Pennsylvania not far from places you knew?
(1)	A:	Yeah, it hurt reading the report because I was reading about these parishes that I went to growing up. I was born in '96, so growing up, you would hear kids joke about, oh, you know, priests molesting kids and whatnot. But I never knew that there was actually a - anything behind that. I just sort of thought it was people making fun of a religion. And then I didn't learn, honestly, until recently that the abuse scandal is something that was real and something that the world has known about since - what? - I think it was the early 2000s, whenever the Boston Globe or whatever that newspaper is broke the story. I didn't know that that was a thing. It wasn't something I'd ever been exposed to, so I wasn't really aware of the fact that this is a problem that was going on in my church. You know, it's unthinkable.
(2)	B:	And it's hurt your faith?
(3)	A:	I think that's a difficult thing to answer. I will say that it has hurt my faith in the Catholic Church. I don't think it has actually hurt my personal, religious faith. I'm just starting to see less of a connection between what I believe and the teachings of the Catholic Church.

Reference Statement	turn 3 : The speaker's faith in the Catholic Church has been damaged by the abuse scandal, but their personal faith remains intact.
Question	How has the abuse scandal affected the speaker's faith in the Catholic Church?
Answer	Damage.

Task

Following is the **dialogue**:

(0)	A:	...
(1)	B:	...
(2)	B:	...
...

Reference Statement	...
Question	...
Answer	[...] Write one word answer here.

Figure 9: Answer collecting worker interface

Table 15: Figurative language reasoning examples of **DiPlomat**

<p>A: Thank you. How are you? B: I'm pretty good. Thank you. You must be stuck like glue on this, but, you know, you've played in three World Cups, including one of the wins for the U. S. team in 1999. How would you describe what it's like to be out there on that field in that final game?</p>	<p>Rationale: Stuck like glue means to be attached to something, which is a particular issue or a person.</p>
<p>B: So in terms of what to do about it, we've said Twitter and Facebook have shut down these accounts, which prompts me to wonder - does shutting down a fake account do that much? Can't the Chinese government, if it's determined to go down this path, just open up two new ones in place of the one that was closed? A: It is a cat-and-mouse game, and the companies are constantly trying to get ahead of it. [...] As you said, they can always set up new accounts.</p>	<p>Rationale: Mice are constantly trying to get away from cats and cats are constantly trying to catch mice. In the same way, the Chinese government will always be trying to escape restrictions on social media accounts and media companies will always be trying to find fake accounts.</p>
<p>A: I really didn't feel safe because the Turkish government is very famous for hunting down those who oppose Erdogan. So, I mean, I just didn't want to really risk my life by going to Europe. But, you know, I talked to my team. I told them all, like, how many times I want to come because I want to be with you guys there, and I want to get a win with you guys. And then, later on, they came back with the news and said, you know what? I think the best decision is if you don't come. Let's just not risk it for one game. B: Do you feel safe in New York and elsewhere in the U. S. ? A: I have been getting last two, three days hundreds death threats, but I think I feel safe in America. But anywhere else in the world, I wouldn't really feel safe.</p>	<p>Rationale: He is implying that he is still not safe.</p>

Table 16: Commonsense reasoning examples of **DiPlomat**

<p>B: Yeah - African-American mayor from Tallahassee. A: Yes. So this is sort of a test of whether real progressive candidates can win in these sort of purplish states. [...]</p>	<p>Rationale: "Purplish" states are not really colored. They refer to US states that are neither clearly Republican (red) nor Democrat (blue) in their voting.</p>
<p>B: He wrote a lot of letters by hand, didn't he? A: He wrote tons of letters. I bet there are a hundred thousand - hundreds out there[...]</p>	<p>Rationale: tons of letters implies a very large number and not to full a ton.</p>
<p>B: Well, Pluto's official designation is a dwarf planet. And I have to tell you the people who sent this probe all the way out to Pluto are a little angry about that because when they launched it a decade ago, Pluto was still a planet. A: (Laughter) B: It got downgraded in the intervening years. A: That seems so unfair.</p>	<p>Rationale: A is expressing sympathy for the people who sent the probe, showing that they understand why they feel so disappointed.</p>

Table 17: External knowledge reasoning examples of **DiPlomat**

<p>B: Inside of his house, family pictures decorate the walls and the fridge. Les has 15 great grandchildren. He grew up in an orphanage, and he couldn't wait to leave to join the military. And so in early 1944, he boarded a ship and crossed the Atlantic Ocean to go to the frontline.</p> <p>A: I loved that sailing on, of course. It was so dramatic. You could see all these ships bobbing up and down on the ocean. And destroyers were weaving in and out of them to make sure they uncovered any mines or anything.</p>	<p>Rationale: Sailing across the ocean during wartime was a perilous experience.</p>
<p>A: . . . equivalent to a nuclear bomb?</p> <p>B: Well, it's about - its equivalent - the energy in that explosion is about 10 times the energy in the first atomic bomb. . .</p>	<p>Rationale: The energy released in the explosion is incredibly powerful.</p>
<p>B: So in your polling, in your research, do you find that it's going to come down to maybe a couple thousand votes from these unaffiliated voters and on what issues? Or will they vote?</p> <p>A: It is likely at the moment to be a very narrow victory. President Bush won in 2004 with five percent. That was 100,000 votes. In other words, if it is one percent, that would be 20,000 votes, and right now, the polls are moving around in just single percentage points. So it could be that narrow.</p> <p>B: Now, I have read that Colorado is going to be this year's Florida and Ohio, that this is going to be the state that decides the election.</p> <p>A: I think it could be, and the interesting thing is that Obama and Palin were both in Jefferson County a couple of days ago, indicating that there may be actually even a county that could be looked at to be beyond an entire state.</p>	<p>Rationale: The turn is suggesting that the county of Jefferson in Colorado could be a key factor in deciding the election, despite the fact that it is only one of many counties in the state and there are other swing states in the election.</p>

Table 18: Others examples of **DiPlomat**

<p>A: There's that feeling - I mean, so many of us have parents in the industry. I mean, that's what this region is about, especially around Detroit, and Wayne State's in Detroit, the heart of Detroit. So, it's nerve-racking. Everyone is nervous. Everyone doesn't know what's going to happen next. We're all watching the news very closely. But at the same time, it's interesting, because with my generation, we almost seem to, kind of, not be as directly impacted. I mean, our family is, it puts stress on us, but the day to day of the university and the day to day at school doesn't seem to have changed that much.</p> <p>B: I understand you have friends there who are engineering majors. Do they have any sense of what their future looks like, and will be it there in Michigan?</p> <p>A: Everybody is secure in their choices and secure in their decision. Everybody thinks that the industry will come around, especially now with the news that GM is getting money from the government. And everybody is more hopeful, and I mean, the auto industry has always been one of the largest industries and a staple in America, and to think that that industry is just going to vanish, nobody is willing to concede that.</p>	<p>Rationale: A believes that the auto industry will not vanish despite the current situation</p>
<p>B: In the meantime, what more have you learned in your reporting about the death of Carlos Hernandez Vasquez?</p> <p>A: Well, a couple of things. One thing that really stands out is that Carlos Hernandez Vasquez died in a Border Patrol station. The previous migrant children who died were taken to the hospital first; Hernandez Vasquez was not even though immigration authorities clearly knew that he was sick. He was diagnosed with the flu by a nurse practitioner.</p>	<p>Rationale: The death of Carlos Hernandez Vasquez could have been prevented if he had been taken to the hospital.</p>
<p>B: So, how do you and the retired general, James Jones, know each other?</p> <p>A: My gosh, I think - I can't even remember when I first met him. It's been so long ago. I'm sure I met him when he was head of the legislative liaison over the Senate. But I really became acquainted with him when he became a brigadier general, and, of course, I followed his career. Of course, he served very ably as a commandant in the marine corps and then as the European commander, just been with him from time to time. And I just consider him a very good friend.</p>	<p>Rationale: A has a high opinion of James Jones' character and career.</p>

533 E Grice Maxims and Pragmatic Reasoning

534 The Gricean maxims have garnered substantial attention as a foundational theory within the domain
 535 of pragmatics. This theoretical framework comprises four distinct maxims: (1) The Maxim of Quality,
 536 (2) The Maxim of Quantity, (3) The Maxim of Relevance, and (4) The Maxim of Manner [16, 15]. In
 537 contrast to rigid rules or theorems, the Gricean maxims, which capture the prevalent dynamics of
 538 conversations, are susceptible to frequent breaches in the context of human communication. These
 539 breaches, stemming from the intricacies of real-world interaction, notably manifest in the violation
 540 of one or more of these maxims. Such breaches, aligned with the cooperative principle, give rise to
 541 pragmatic phenomena that necessitate the engagement of pragmatic reasoning by recipients of the
 542 communication [15].

543 F Computational Resources

544 For our experiment, we utilized two A100s and one 3090. The majority of our experiments were
 545 conducted on the A100s, while for practical reasons, only Unified-QA-base, BART-base, and T5-
 546 small were tested on the 3090. It is important to mention that each experiment was run on a single
 547 GPU. We record the training time of models in Appendix F.

Model	C → P	CP → R	Device
BERT _{base}	0.8min/epoch	0.9min/epoch	A100
RoBERTa _{base}	0.8min/epoch	0.9min/epoch	A100
RoBERTa _{large}	2.5min/epoch	2.8min/epoch	A100
GPT-2 _{base}	5.8min/epoch	6.2min/epoch	A100
DialoGPT _{medium}	2.4min/epoch	4.2min/epoch	A100
DeBERTa _{base}	0.9min/epoch	0.9min/epoch	A100
ALBERT _{base}	0.5min/epoch	0.8min/epoch	A100

Table 19: Training Time of Models

548 G Limitations & Negative Societal Impacts

549 We acknowledge two limitations in our study: bias and subjectivity. Since our dialogues primarily
 550 stem from an interview dataset, a considerable focus is placed on political topics. This is reasonable,
 551 as pragmatic phenomena frequently emerge in the statements of politicians to advance their specific
 552 goals. However, this focus introduces a certain degree of bias into our dataset. The second limitation
 553 relates to the absence of subjectivity. In our methodology, the data undergoes two stages of human
 554 annotation, ensuring higher quality and objectivity. However, pragmatic reasoning is inherently
 555 subjective, and prioritizing objectivity compromises the preservation of subjectivity, resulting in a
 556 limitation in terms of subjectivity coverage. Our dataset exhibits minimal negative societal impacts.
 557 This is primarily due to the fact that our dialogues are transcriptions of publicly available TV shows,
 558 which inherently limits the potential for negative effects.

559 H Ethics Concern

560 **Were any ethical review processes conducted (e.g., by an institutional review board)?** No official
 561 processes were done, as our research is not on human subjects, but our data comes from published
 562 dataset.

563 **Does the dataset contain data that might be considered confidential?** No, our data comes from an
 564 existing public interview dataset.

565 **Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening,**
566 **or might otherwise cause anxiety? If so, please describe why.** Few of the dialogues may talk about
567 offensive topics. **Does the dataset identify subpopulations (e.g., by age or gender)?** Not explicitly.

568 **Is it possible to identify individuals (i.e., one or more natural persons) directly or indirectly (i.e.,**
569 **in combination with other data) from the dataset?** Yes, our data contains names of famous people.

570 **I Responsibility & Dataset Liscence**

571 We bear all responsibility in case of violation of rights and our dataset is under the license of CC
572 BY-NC-SA (Attribution-NonCommercial-ShareAlike).

573 **J Datasheets for Our Dataset**

574 **J.1 Motivation**

575 1. For what purpose was the dataset created? (Was there a specific task in mind? Was there a
576 specific gap that needed to be filled? Please provide a description.)

577 This dataset was created to study pragmatic reasoning in dialogues, a specific gap is men-
578 tioned above in Appendix G.

579 2. Who created this dataset (e.g., which team, research group) and on behalf of which entity
580 (e.g., company, institution, organization)?

581 This dataset was created by the authors of this paper.

582 3. Who funded the creation of the dataset? (If there is an associated grant, please provide the
583 name of the grantor and the grant name and number.)

584 The institute of the authors funded the creation of the dataset.

585 4. Any other comments?

586 None.

587 **J.2 Composition**

588 5. What do the instances that comprise the dataset represent (e.g., documents, photos, people,
589 countries)? (Are there multiple types of instances (e.g., movies, users, and ratings; people
590 and interactions between them; nodes and edges)? Please provide a description.)

591 An instance of our dataset represent a piece of dialogue. Description is provided in our
592 paper.

593 6. How many instances are there in total (of each type, if appropriate)?

594 We answer the question in our paper. Our datasets owns 4,177 dialogues.

595 7. Does the dataset contain all possible instances or is it a sample (not necessarily random) of
596 instances from a larger set? (If the dataset is a sample, then what is the larger set? Is the
597 sample representative of the larger set (e.g., geographic coverage)? If so, please describe
598 how this representativeness was validated/verified. If it is not representative of the larger set,
599 please describe why not (e.g., to cover a more diverse range of instances, because instances
600 were withheld or unavailable).)

601 It is a sample of all possible cases. As pragmatic phenomena aren't proved to be limited, we
602 can't guarantee a full sampling of them.

603 8. What data does each instance consist of?

604 We mention it in our paper.

605 9. Is there a label or target associated with each instance? If so, please provide a description.

606 Yes. The description is in our paper.

- 607 10. Is any information missing from individual instances? (If so, please provide a description,
608 explaining why this information is missing (e.g., because it was unavailable). This does not
609 include intentionally removed information, but might include, e.g., redacted text.)
610 No. We leverage the original dialogues.
- 611 11. Are relationships between individual instances made explicit (e.g., users' movie ratings,
612 social network links)? (If so, please describe how these relationships are made explicit.)
613 No. Instances are weakly related, but focus on the same phenomenon.
- 614 12. Are there recommended data splits (e.g., training, development/validation, testing)? (If so,
615 please provide a description of these splits, explaining the rationale behind them.)
616 Yes. We provide it.
- 617 13. Are there any errors, sources of noise, or redundancies in the dataset? (If so, please provide
618 a description.)
619 Yes. **Some workers try to finish the work as quickly as possible, therefore when we ask them**
620 **to offer a rationale for choosing a certain turn as a pragmatic turn, they simply type an "a" in**
621 **the box. However, the situation is rare, and we blocked the workers and clean the data out of**
622 **our dataset.**
- 623 14. Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g.,
624 websites, tweets, other datasets)? (If it links to or relies on external resources, a) are there
625 guarantees that they will exist, and remain constant, over time; b) are there official archival
626 versions of the complete dataset (i.e., including the external resources as they existed at the
627 time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with
628 any of the external resources that might apply to a future user? Please provide descriptions
629 of all external resources and any restrictions associated with them, as well as links or other
630 access points, as appropriate.)
631 It's self-contained.
- 632 15. Does the dataset contain data that might be considered confidential (e.g., data that is protected
633 by legal privilege or by doctor-patient confidentiality, data that includes the content of
634 individuals' non-public communications)? (If so, please provide a description.)
635 No.
- 636 16. Does the dataset contain data that, if viewed directly, might be offensive, insulting, threaten-
637 ing, or might otherwise cause anxiety? (If so, please describe why.)
638 Yes. Some of the topic are big events, they may be offensive for some people. However,
639 we consider our dataset's offensiveness to be limited, for the source dataset is a TV show
640 transcript.
- 641 17. Does the dataset relate to people? (If not, you may skip the remaining questions in this
642 section.)
643 Yes.
- 644 18. Does the dataset identify any subpopulations (e.g., by age, gender)? (If so, please describe
645 how these subpopulations are identified and provide a description of their respective distri-
646 butions within the dataset.)
647 No. This is not explicitly identified
- 648 19. Is it possible to identify individuals (i.e., one or more natural persons), either directly or
649 indirectly (i.e., in combination with other data) from the dataset? (If so, please describe
650 how.)
651 Yes; their names are given in running text.
- 652 20. Does the dataset contain data that might be considered sensitive in any way (e.g., data that
653 reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or
654 union memberships, or locations; financial or health data; biometric or genetic data; forms of
655 government identification, such as social security numbers; criminal history)? (If so, please
656 provide a description.)

657 Yes. Our dataset may have dialogues talking about religious, politics and so on.

658 21. Any other comments?

659 None.

660 J.3 Collection Process

661 1. How was the data associated with each instance acquired? (Was the data directly observable
662 (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly
663 inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or
664 language)? If data was reported by subjects or indirectly inferred/derived from other data,
665 was the data validated/verified? If so, please describe how.)

666 The data all comes from an interview dataset already published. (See our paper)

667 2. What mechanisms or procedures were used to collect the data (e.g., hardware apparatus
668 or sensor, manual human curation, software program, software API)? (How were these
669 mechanisms or procedures validated?)

670 Software program and manual human curation (2 times). See our paper for details.

671 3. If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic,
672 probabilistic with specific sampling probabilities)?

673 Randomly.

674 4. Who was involved in the data collection process (e.g., students, crowdworkers, contractors)
675 and how were they compensated (e.g., how much were crowdworkers paid)?

676 Crowdworkers. They are paid nicely. See Appendix for detail.

677 5. Over what timeframe was the data collected? (Does this timeframe match the creation
678 timeframe of the data associated with the instances (e.g., recent crawl of old news articles)?
679 If not, please describe the timeframe in which the data associated with the instances was
680 created.)

681 The dataset was collected in the early Spring of 2023, which does not necessarily reflect the
682 timeframe of the data collected.

683 6. Were any ethical review processes conducted (e.g., by an institutional review board)? (If so,
684 please provide a description of these review processes, including the outcomes, as well as a
685 link or other access point to any supporting documentation.)

686 No review processes were conducted with respect to the collection and annotation of this
687 data (though review was done for other aspects of this work; see the paper linked at the top
688 of the datasheet).

689 7. Does the dataset relate to people? (If not, you may skip the remaining questions in this
690 section.)

691 Yes.

692 8. Did you collect the data from the individuals in question directly, or obtain it via third parties
693 or other sources (e.g., websites)?

694 Other sources. By curating a published dataset.

695 9. Were the individuals in question notified about the data collection? (If so, please describe (or
696 show with screenshots or other information) how notice was provided, and provide a link or
697 other access point to, or otherwise reproduce, the exact language of the notification itself.)

698 No.

699 10. Did the individuals in question consent to the collection and use of their data? (If so, please
700 describe (or show with screenshots or other information) how consent was requested and
701 provided, and provide a link or other access point to, or otherwise reproduce, the exact
702 language to which the individuals consented.)

703 No. All data are public.

- 704 11. If consent was obtained, were the consenting individuals provided with a mechanism to
705 revoke their consent in the future or for certain uses? (If so, please provide a description, as
706 well as a link or other access point to the mechanism (if appropriate).)
707 N/A.
- 708 12. Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a
709 data protection impact analysis) been conducted? (If so, please provide a description of this
710 analysis, including the outcomes, as well as a link or other access point to any supporting
711 documentation.)
712 No. We consider our dataset having a limited negative effect, for all of our data has been
713 published for more than a year.
- 714 13. Any other comments? None.

715 **J.4 Preprocessing/cleaning/labeling**

- 716 1. Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing,
717 tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, pro-
718 cessing of missing values)? (If so, please provide a description. If not, you may skip the
719 remainder of the questions in this section.)
720 No.

721 **J.5 Uses**

- 722 1. Has the dataset been used for any tasks already? (If so, please provide a description.)
723 Yes. See our paper for details.
- 724 2. Is there a repository that links to any or all papers or systems that use the dataset? (If so,
725 please provide a link or other access point.)
726 No.
- 727 3. What (other) tasks could the dataset be used for?
728 Many more. Such as generation of implied meanings.
- 729 4. Is there anything about the composition of the dataset or the way it was collected and
730 preprocessed/cleaned/labeled that might impact future uses? (For example, is there anything
731 that a future user might need to know to avoid uses that could result in unfair treatment
732 of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable
733 harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything
734 a future user could do to mitigate these undesirable harms?)
735 No.
- 736 5. Are there tasks for which the dataset should not be used? (If so, please provide a description.)
737 No.
- 738 6. Any other comments?
739 None.

740 **J.6 Distribution**

- 741 1. Will the dataset be distributed to third parties outside of the entity (e.g., company, institution,
742 organization) on behalf of which the dataset was created? (If so, please provide a description.)
743 Yes, the dataset is freely available.
- 744 2. How will the dataset will be distributed (e.g., tarball on website, API, GitHub)? (Does the
745 dataset have a digital object identifier (DOI)?)
746 On our website.
- 747 3. When will the dataset be distributed?
748 It's already been distributed.

749 4. Will the dataset be distributed under a copyright or other intellectual property (IP) license,
750 and/or under applicable terms of use (ToU)? (If so, please describe this license and/or ToU,
751 and provide a link or other access point to, or otherwise reproduce, any relevant licensing
752 terms or ToU, as well as any fees associated with these restrictions.)

753 The dataset is licensed under a CC license.

754 5. Have any third parties imposed IP-based or other restrictions on the data associated with the
755 instances? (If so, please describe these restrictions, and provide a link or other access point
756 to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with
757 these restrictions.)

758 Not to our knowledge.

759 6. Do any export controls or other regulatory restrictions apply to the dataset or to individual
760 instances? (If so, please describe these restrictions, and provide a link or other access point
761 to, or otherwise reproduce, any supporting documentation.)

762 Not to our knowledge.

763 7. Any other comments?

764 None.

765 J.7 Maintenance

766 1. Who is supporting/hosting/maintaining the dataset?

767 The authors.

768 2. How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

769 We will post our email address.

770 3. Is there an erratum? (If so, please provide a link or other access point.)

771 Currently, no. As errors are encountered, future versions of the dataset may be released (but
772 will be versioned). They will all be provided in the same location.

773 4. Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete
774 instances)? (If so, please describe how often, by whom, and how updates will be communi-
775 cated to users (e.g., mailing list, GitHub)?)

776 Yes. However, the frequency isn't determined, and we'll publish the updated dataset on the
777 same website if a renewal occurs, and we'll announce it on the website.

778 5. If the dataset relates to people, are there applicable limits on the retention of the data
779 associated with the instances (e.g., were individuals in question told that their data would be
780 retained for a fixed period of time and then deleted)? (If so, please describe these limits and
781 explain how they will be enforced.)

782 No.

783 6. Will older versions of the dataset continue to be supported/hosted/maintained? (If so, please
784 describe how. If not, please describe how its obsolescence will be communicated to users.)

785 Yes. The older versions of the dataset will be available on the website.

786 7. If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for
787 them to do so? (If so, please provide a description. Will these contributions be validated/ver-
788 ified? If so, please describe how. If not, why not? Is there a process for communicating/dis-
789 tributing these contributions to other users? If so, please provide a description.)

790 Yes. They can email us.

791 8. Any other comments?

792 None.

793 **References**

- 794 [1] Jens Ambrasat, Christian von Scheve, Markus Conrad, Gesche Schauenburg, and Tobias Schröder. Con-
795 sensus and stratification in the affective meaning of human sociality. *Proceedings of the National Academy*
796 *of Sciences*, 111(22):8001–8006, 2014. 2
- 797 [2] Alan P Fiske. The four elementary forms of sociality: framework for a unified theory of social relations.
798 *Psychological review*, 99(4):689, 1992. 2
- 799 [3] John A Bargh and Tanya L Chartrand. The unbearable automaticity of being. *American psychologist*, 54
800 (7):462, 1999. 2
- 801 [4] Edward Finegan. *Language: Its structure and use*. Cengage Learning, 2014. 2
- 802 [5] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan,
803 Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-
804 level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378
805 (6624):1067–1074, 2022. 2
- 806 [6] Silke Anders, Roos de Jong, Christian Beck, John-Dylan Haynes, and Thomas Ethofer. A neural link
807 between affective understanding and interpersonal attraction. *Proceedings of the National Academy of*
808 *Sciences*, 113(16):E2248–E2257, 2016. 2
- 809 [7] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang,
810 Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with
811 human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022. 2
- 812 [8] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind
813 Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss,
814 Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens
815 Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack
816 Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language
817 models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- 818 [9] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of
819 thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022. 2
- 820 [10] Laura Ruis, Akbir Khan, Stella Biderman, Sara Hooker, Tim Rocktäschel, and Edward Grefenstette. Large
821 language models are not zero-shot communicators. *arXiv preprint arXiv:2210.14986*, 2022. 2
- 822 [11] Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta,
823 Noah A. Smith, and Yejin Choi. We’re afraid language models aren’t modeling ambiguity. 2023. 2
- 824 [12] Zilong Zheng, Shuwen Qiu, Lifeng Fan, Yixin Zhu, and Song-Chun Zhu. GRICE: A grammar-based
825 dataset for recovering implicature and conversational rEasoning. In *Annual Meeting of the Association*
826 *for Computational Linguistics (ACL)*, pages 2074–2085, Online, August 2021. Association for Computa-
827 tional Linguistics. doi: 10.18653/v1/2021.findings-acl.182. URL [https://aclanthology.org/2021.](https://aclanthology.org/2021.findings-acl.182)
828 [findings-acl.182](https://aclanthology.org/2021.findings-acl.182). 2, 3, 4
- 829 [13] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2, 3, 9
- 830 [14] Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh,
831 Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. Improving alignment of dialogue agents via
832 targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022. 2
- 833 [15] Alan Cruse. Meaning in language: An introduction to semantics and pragmatics. 2, 4, 22
- 834 [16] Herbert P Grice. Logic and conversation. 1975. 2, 4, 22
- 835 [17] Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. Mermaid: Metaphor generation
836 with symbolism and discriminative decoding. In *Annual Meeting of the Association for Computational*
837 *Linguistics (ACL)*, 2021. 2, 4
- 838 [18] Prateek Saxena and Soma Paul. Epie dataset: A corpus for possible idiomatic expressions. In *Text, Speech,*
839 *and Dialogue - 23rd International Conference, TSD 2020, Brno, Czech Republic, September 8-11, 2020,*
840 *Proceedings*, 2020. 2, 4, 5, 12
- 841 [19] Tosin P. Adewumi, Roshanak Vadoodi, Aparajita Tripathy, Konstantina Nikolaidou, Foteini Liwicki, and
842 Marcus Liwicki. Potential idiomatic expression (pie)-english: Corpus for classes of idioms. In *Proceedings*
843 *of the Thirteenth Language Resources and Evaluation Conference*, 2022. 2, 4

- 844 [20] Issa Annamoradnejad and Gohar Zoghi. Colbert: Using bert sentence embedding in parallel neural
845 networks for computational humor, 2022. 2, 4, 12
- 846 [21] Bodhisattwa Prasad Majumder, Shuyang Li, Jianmo Ni, and Julian McAuley. Interview: Large-scale
847 modeling of media dialog with discourse patterns and knowledge grounding. In *Annual Conference on*
848 *Empirical Methods in Natural Language Processing (EMNLP)*, pages 8129–8141, Online, November
849 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.653. URL <https://aclanthology.org/2020.emnlp-main.653>. 2, 5
- 851 [22] Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. The Ubuntu dialogue corpus: A large dataset
852 for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting*
853 *of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic,
854 September 2015. Association for Computational Linguistics. doi: 10.18653/v1/W15-4640. URL <https://aclanthology.org/W15-4640>. 3
- 856 [23] Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. Race: Large-scale reading
857 comprehension dataset from examinations. In *Annual Conference on Empirical Methods in Natural*
858 *Language Processing (EMNLP)*, 2017. 3
- 859 [24] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind
860 Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. 3
- 861 [25] Adina Williams, Nikita Nangia, and Samuel R. Bowman. A broad-coverage challenge corpus for sentence
862 understanding through inference. In *North American Chapter of the Association for Computational*
863 *Linguistics: Human Language Technologies (NAACL-HLT)*, 2018. 3, 10
- 864 [26] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing
865 dialogue agents: I have a dog, do you have pets too? In *Annual Meeting of the Association for Computational*
866 *Linguistics (ACL)*, 2018. 3
- 867 [27] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. Swag: A large-scale adversarial dataset for
868 grounded commonsense inference. In *Annual Conference on Empirical Methods in Natural Language*
869 *Processing (EMNLP)*, 2018. 3
- 870 [28] Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Cosmos qa: Machine reading
871 comprehension with contextual commonsense reasoning. In *Annual Conference on Empirical Methods in*
872 *Natural Language Processing (EMNLP)*, 2019. 3
- 873 [29] Siva Reddy, Danqi Chen, and Christopher D. Manning. Coqa: A conversational question answering
874 challenge. In *Transactions of the Association for Computational Linguistics (TACL)*, 2019. 3
- 875 [30] Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. Dream: A challenge dataset and
876 models for dialogue-based reading comprehension. In *Transactions of the Association for Computational*
877 *Linguistics (TACL)*, 2019. 3
- 878 [31] Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. Dialogue natural language inference. In
879 *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019. 3
- 880 [32] Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. Drop:
881 A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *North American*
882 *Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*,
883 2019. 3
- 884 [33] Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. Mutual: A dataset for multi-turn dialogue
885 reasoning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020. 3, 4
- 886 [34] Paloma Jeretic, Alex Warstadt, Suvrat Bhooshan, and Adina Williams. Are natural language inference
887 models impressive? learning implicature and presupposition. In *Annual Meeting of the Association for*
888 *Computational Linguistics (ACL)*, 2020. 3
- 889 [35] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing
890 Liu, and Bill Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation.
891 In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020. 3, 14
- 892 [36] Baolin Peng, Michel Galley, Pengcheng He, Chris Brockett, Lars Liden, Elnaz Nouri, Zhou Yu, Bill
893 Dolan, and Jianfeng Gao. Godel: Large-scale pre-training for goal-directed dialog. *arXiv preprint*
894 *arXiv:2206.11309*, 2022. 3

- 895 [37] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng,
896 Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin
897 Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen,
898 Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching
899 Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-
900 Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben
901 Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina,
902 Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya
903 Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui,
904 Marian Croak, Ed Chi, and Quoc Le. Lamda: Language models for dialog applications. *arXiv preprint*
905 *arXiv:2201.08239*, 2022. 3
- 906 [38] Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang,
907 Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. Towards a human-like open-domain
908 chatbot. *arXiv preprint arXiv:2001.09977*, 2020. 3
- 909 [39] Chulaka Gunasekara, Seokhwan Kim, Luis Fernando D’Haro, Abhinav Rastogi, Yun-Nung Chen, Mihail
910 Eric, Behnam Hedayatnia, Karthik Gopalakrishnan, Yang Liu, Chao-Wei Huang, Dilek Hakkani-Tür,
911 Jinchao Li, Qi Zhu, Lingxiao Luo, Lars Liden, Kaili Huang, Shahin Shayandeh, Runze Liang, Baolin
912 Peng, Zheng Zhang, Swadheen Shukla, Minlie Huang, Jianfeng Gao, Shikib Mehri, Yulan Feng, Carla
913 Gordon, Seyed Hossein Alavi, David Traum, Maxine Eskenazi, Ahmad Beirami, Eunjoon, Cho, Paul A.
914 Crook, Ankita De, Alborz Geramifard, Satwik Kottur, Seungwhan Moon, Shivani Poddar, and Rajen Subba.
915 Overview of the ninth dialog system technology challenge: Dstc9, 2020. 3
- 916 [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz
917 Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing*
918 *Systems (NeurIPS)*, 2017. 3
- 919 [41] Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak
920 Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey,
921 Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson,
922 Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan
923 Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele
924 Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément
925 Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer,
926 Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann,
927 Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey
928 Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy,
929 Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li,
930 YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni,
931 Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussaleem, Zachary Nado, John Nham, Eric
932 Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan
933 Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar
934 Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine,
935 Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang,
936 John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang,
937 Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. Palm 2 technical
938 report, 2023. 4
- 939 [42] Yunxiang Zhang and Xiaojun Wan. MOVER: Mask, over-generate and rank for hyperbole generation. In
940 *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6018–6030, Seattle, United
941 States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.440.
942 URL <https://aclanthology.org/2022.naacl-main.440>. 4, 12
- 943 [43] Orion Weller and Kevin Seppi. The rJokes dataset: a large scale humor collection. In *Proceedings of the*
944 *Twelfth Language Resources and Evaluation Conference*, pages 6136–6141, Marseille, France, May 2020.
945 European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.753>. 4
- 947 [44] Michael C. Frank and Noah D. Goodman. Predicting pragmatic reasoning in language games. *Science*, 336
948 (6084):998–998, 2012. doi: 10.1126/science.1218633. URL <https://www.science.org/doi/abs/10.1126/science.1218633>. 4
- 950 [45] Kevin Stowe, Prasetya Utama, and Iryna Gurevych. IMPLI: Investigating NLI models’ performance on
951 figurative language. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages
952 5375–5388, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/
953 2022.acl-long.369. URL <https://aclanthology.org/2022.acl-long.369>. 4

- 954 [46] Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. It’s not rocket science : Interpreting figurative language
955 in narratives. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022. 4
- 956 [47] Deepanway Ghosal, Siqi Shen, Navonil Majumder, Rada Mihalcea, and Soujanya Poria. Cicero: A
957 dataset for contextualized commonsense inference in dialogues. In *Annual Meeting of the Association for*
958 *Computational Linguistics (ACL)*, 2022. 4
- 959 [48] Deepanway Ghosal, Pengfei Hong, Siqi Shen, Navonil Majumder, Rada Mihalcea, and Soujanya Poria.
960 Cider: Commonsense inference for dialogue explanation and reasoning. In *Proceedings of the 22nd Annual*
961 *Meeting of the Special Interest Group on Discourse and Dialogue*, 2021.
- 962 [49] Pei Zhou, Pegah Jandaghi, Bill Yuchen Lin, Justin Cho, Jay Pujara, and Xiang Ren. Probing commonsense
963 explanation in dialogue response generation. In *Annual Conference on Empirical Methods in Natural*
964 *Language Processing (EMNLP)*, 2021. 4
- 965 [50] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis,
966 Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. 2019.
967 URL <http://arxiv.org/abs/1907.11692>. 5, 12, 13
- 968 [51] Yunxiang Zhang and Xiaojun Wan. Mover: Mask, over-generate and rank for hyperbole generation. In
969 *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022. 5
- 970 [52] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text
971 generation with bert. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020. 5,
972 7, 12, 13
- 973 [53] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual
974 commonsense reasoning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
975 6
- 976 [54] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In
977 *Annual Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for
978 Computational Linguistics, 11 2019. URL <http://arxiv.org/abs/1908.10084>. 7
- 979 [55] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves
980 Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language
981 generation, translation, and comprehension. In *Annual Meeting of the Association for Computational*
982 *Linguistics (ACL)*, 2019. 9
- 983 [56] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou,
984 Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. In
985 *Journal of Machine Learning Research (JMLR)*, 2020. 9, 10
- 986 [57] Daniel Khoshabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh
987 Hajishirzi. Unifiedqa: Crossing format boundaries with a single qa system. In *Annual Conference on*
988 *Empirical Methods in Natural Language Processing (EMNLP)*, 2020. 9
- 989 [58] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua,
990 and Colin Raffel. mT5: A massively multilingual pre-trained text-to-text transformer. In *North American*
991 *Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*,
992 pages 483–498, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.
993 naacl-main.41. URL <https://aclanthology.org/2021.naacl-main.41>. 9
- 994 [59] Bill MacCartney and Christopher D. Manning. Modeling semantic containment and exclusion in natural
995 language inference. In *International Conference on Computational Linguistics (COLING)*, pages 521–528,
996 Manchester, UK, August 2008. Coling 2008 Organizing Committee. URL <https://aclanthology.org/C08-1066>. 9
- 998 [60] Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen,
999 Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. TRUE: Re-evaluating fac-
1000 tual consistency evaluation. In *North American Chapter of the Association for Computational Lin-*
1001 *guistics: Human Language Technologies (NAACL-HLT)*, pages 3905–3920, Seattle, United States,
1002 July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.287. URL
1003 <https://aclanthology.org/2022.naacl-main.287>. 10
- 1004 [61] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated
1005 corpus for learning natural language inference. In *Annual Conference on Empirical Methods in Natural*
1006 *Language Processing (EMNLP)*, 2015. 10

- 1007 [62] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced bert with
1008 disentangled attention. In *International Conference on Learning Representations (ICLR)*, 2021. 10, 13
- 1009 [63] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language
1010 models are zero-shot reasoners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
1011 10
- 1012 [64] Anya Ji, Noriyuki Kojima, Noah Rush, Alane Suhr, Wai Keen Vong, Robert Hawkins, and Yoav Artzi.
1013 Abstract visual reasoning with tangram shapes. In *Annual Conference on Empirical Methods in Natural
1014 Language Processing (EMNLP)*, pages 582–601, Abu Dhabi, United Arab Emirates, December 2022.
1015 Association for Computational Linguistics. URL <https://aclanthology.org/2022.emnlp-main.38>.
1016 10
- 1017 [65] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut.
1018 ALBERT: A lite BERT for self-supervised learning of language representations. 2020. 12
- 1019 [66] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut.
1020 Albert: A lite bert for self-supervised learning of language representations. In *International Conference on
1021 Learning Representations (ICLR)*, 2020. 12, 13
- 1022 [67] Rishabh Misra. News headlines dataset for sarcasm detection, 2022. 12
- 1023 [68] Lennart Wachowiak, Dagmar Gromann, and Chao Xu. Drum up SUPPORT: Systematic analysis of
1024 image-schematic conceptual metaphors. In *Proceedings of the 3rd Workshop on Figurative Language
1025 Processing (FLP)*, pages 44–53, Abu Dhabi, United Arab Emirates (Hybrid), December 2022. Association
1026 for Computational Linguistics. URL <https://aclanthology.org/2022.flp-1.7>. 12
- 1027 [69] Marti A. Hearst. Texttiling: Segmenting text into multi-paragraph subtopic passages. *Comput. Linguist.*,
1028 23(1):33–64, mar 1997. ISSN 0891-2017. 12
- 1029 [70] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirec-
1030 tional transformers for language understanding. In *Annual Meeting of the Association for Computational
1031 Linguistics (ACL)*, 2019. 13
- 1032 [71] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models
1033 are unsupervised multitask learners. 2019. 13