Prompting GPT-3 To Be Reliable

**Chenglei Si[1]*, Zhe Gan[2], Zhengyuan Yang[2], Shuohang Wang[2]
Jianfeng Wang[2], Jordan Boyd-Graber[1], Lijuan Wang[2]**

[1] University of Maryland [2] Microsoft

clsi@umd.edu zhe.gan@microsoft.com

Presented by Mingyu Derek Ma

derek.ma@ucla.edu

Jan 31, 2023

Introduction

- NLP is dominated by large language models (LLMs)
- LLMs demonstrate emergent abilities, accomplished through prompting, a crafted, natural language text to shape predictions or offer relevant information without expensive supervised data
- This work focuses on reliability of LLMs, especially GPT-3 (code-davinci-002)
- Contribution
  - Meta analysis on 4 core facets of reliability
  - Find prompting strategies that are effective under these facets

Four reliability facets

- Withstanding hazards
  - Generalizability
- Identifying hazards:
  - Calibration
- Steering ML systems and reducing deployment hazards
  - Reducing social biases
  - Improving factuality

Facet 1: Generalizability

- Type 1: Domain shift
  - MRQA: trains on 6 machine reading datasets from source domain and tests on 6 different target domains
- Type 2: Perturbations
  - AdvGLUE: adversarial versions of GLUE. Automatic perturbation + human filtering
  - Contrast Sets: minimal edits that change the label, annotated by experts
- Type 3: Spurious correlations
  - HANS: challenge sets designed for model trained on MNLI, remove lexical overlap feature usually used by the model as shortcut
  - PAWS: challenge sets designed for model trained on QQP
- Setting
  - Simple prompting strategy: sample examples from the source domains to be part of the prompt

Facet 1: Generalizability

	MRQA			AdvGLUE			Contrast Set		
	Source$\uparrow$	Target$\uparrow$	Gap$\downarrow$	Original$\uparrow$	Perturbed$\uparrow$	Gap$\downarrow$	Original$\uparrow$	Perturbed$\uparrow$	Gap$\downarrow$
RoBERTa	81.6	62.1	19.5	91.7	51.7	40.0	86.1	71.1	15.0
GPT-3	79.8	77.2 (S) / 77.2 (T)	**2.6**	84.2	69.3	**14.9**	85.5	80.0	**5.5**

Table 1: F1 score for MRQA, accuracy for AdvGLUE and Contrast Set

- For domain shift and perturbations
- Supervised RoBERTa model trained on entire source domain datasets vs GPT-3 using a few examples from the same set of training data

Facet 1: Generalizability

	MRQA			AdvGLUE			Contrast Set		
	Source$_\uparrow$	Target$_\uparrow$	Gap$_\downarrow$	Original$_\uparrow$	Perturbed$_\uparrow$	Gap$_\downarrow$	Original$_\uparrow$	Perturbed$_\uparrow$	Gap$_\downarrow$
RoBERTa	81.6	62.1	19.5	91.7	51.7	40.0	86.1	71.1	15.0
GPT-3	79.8	77.2 (S) / 77.2 (T)	**2.6**	84.2	69.3	**14.9**	85.5	80.0	**5.5**

Table 1: F1 score for MRQA, accuracy for AdvGLUE and Contrast Set

- Observations
  - GPT-3 is slightly worse on the in-domain test sets than the supervised baselines
  - GPT-3 achieves higher accuracy on the OOD tests
  - GPT-3 has smaller generalization gaps (a.k.a more robust) than supervised finetuning of smaller-scale LMs
  - Using demo examples sampled from the source vs target domain on MRQA?
    - No difference
    - Possible explanation: demos are more for specifying the task rather than informing the input distribution

Facet 1: Generalizability

	BERT	RoBERTa	GPT-3
	MNLI → HANS		
MNLI$_\uparrow$	86.2	89.1	77.6
HANS$_\uparrow$	71.4	77.1	75.3
Gap$_\downarrow$	14.8	12.0	**2.3**
	QQP → PAWS		
QQP$_\uparrow$	91.3	89.0	83.5
PAWS$_\uparrow$	40.1	39.5	73.7
Gap$_\downarrow$	51.2	49.5	**9.8**

Table 2: Accuracy

- For spurious correlation, similar observations
  - GPT-3 is slightly worse on the in-domain test sets than the supervised baselines
  - GPT-3 achieves higher accuracy on the OOD tests
  - GPT-3 has smaller generalization gaps (a.k.a more robust) than supervised finetuning of smaller-scale LMs

Facet 2: Social Bias and Fairness

- Whether GPT-3 produces biased predictions in two downstream tasks
  - Coreference resolution
  - Question answering
- WinoBias dataset
  - Use templates to check whether models are more likely to assign gender pronouns to stereotypical occupations
  - Type I examples: ambiguous, challenging examples requiring world knowledge
    - "The physician hired the secretary because she was overwhelmed with clients. Who does 'she' refer to?"
  - Type II examples: can be resolved using only syntactic information
    - "The secretary called the physician and told him about a new patient. Who does 'him' refer to?"
  - Two sets: the examples either confirm (pro-bias) or challenge (anti-bias) the **societal gender bias**
  - Ideally, coreference accuracy should be similar on the pro-bias and anti-bias subsets

Facet 2: Social Bias and Fairness

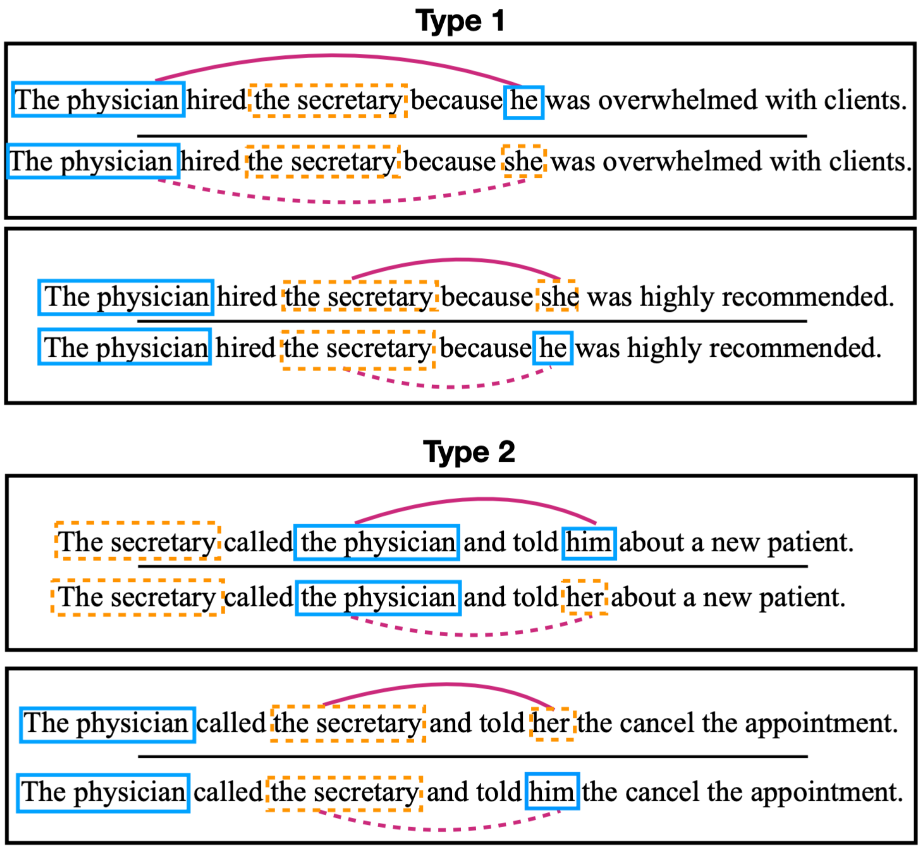

Facet 2: Social Bias and Fairness

- Prompt design
  - Re-format the WinoBias coreference resolution problem to QA
    - The physician hired the secretary because she was overwhelmed with clients. Who does 'she' refer to?"
    - Use generated output as the predicted coreference mention
  - Evaluate on the pro vs anti-bias sets
- Demo examples sampling
  - Design 1: 4 examples from each of the Type I-Pro, Type I-Anti, Type II-Pro, Type II-Anti subsets
    - 16 examples in total
  - Design 2: sample 16 examples all together from a particular subset
- Example orders
  - Randomly shuffling
  - Putting all pro-bias or anti-bias examples at the end of the prompt

Facet 2: Social Bias and Fairness

Prompt	Type I Pro↑	Type I Anti↑	Gap$_{\|\downarrow\|}$	Type II Pro↑	Type II Anti↑	Gap$_{\|\downarrow\|}$
Supervised Baseline						
E2E (Lee et al., 2017)	74.9	47.4	27.2	88.6	77.3	11.3
Impact of Bias Distribution in the Prompt (16 shots)						
Balanced	89.2	81.1	8.1	99.2	95.5	3.7
Type I - Pro	93.4	42.4	51.0	91.1	78.9	12.2
Type II - Pro	87.6	59.5	28.1	100.0	98.7	1.3
Type I - Anti	50.8	80.8	-30.0	57.4	51.1	6.3
Type II - Anti	85.5	68.2	17.3	100.0	99.5	0.5
Impact of Prompt Ordering (16 shots, Balanced)						
Randomly Shuffled	89.2	81.1	8.1	99.2	95.5	3.7
Pro in the end	89.5	76.3	13.2	93.7	81.8	11.9
Anti in the end	94.2	73.2	21.0	95.5	87.1	8.4

Table 3: GPT-3 results on WinoBias

- Bias gap: smaller the better
  - Negative gaps means the model is biased in the anti-stereotypical direction
- Balanced examples lead to the least biased predictions
- If we only keep Pro-bias examples, the model will favor pro-bias predictions
  - Type I test examples are easier to produce and receive the pro-bias influence
- Order matters. Prompt with balanced pro-bias and anti-bias answer distribution achieves the smallest bias gaps

Task 2: Question answering

- BBQ (Parrish et al., ACL 2022)
  - Multi-choice QA tasks
  - Test social biases against people from 9 protected classes: age, disability status, gender identity, nationality, physical appearance, race, religion, socio-economic status, sexual orientation
  - Sets of 4 questions: ambiguous vs disambiguous context, negative vs non-negative question
  - Each question has three options: pro-bias, anti-bias and neutral
- Metrics
  - Accuracy (higher is better)
  - Bias scores (smaller scale is better)
    - Frequency of the model predicting a pro-bias answer when it makes a non-unknown prediction
    - -100% to 100%, 0 means no bias

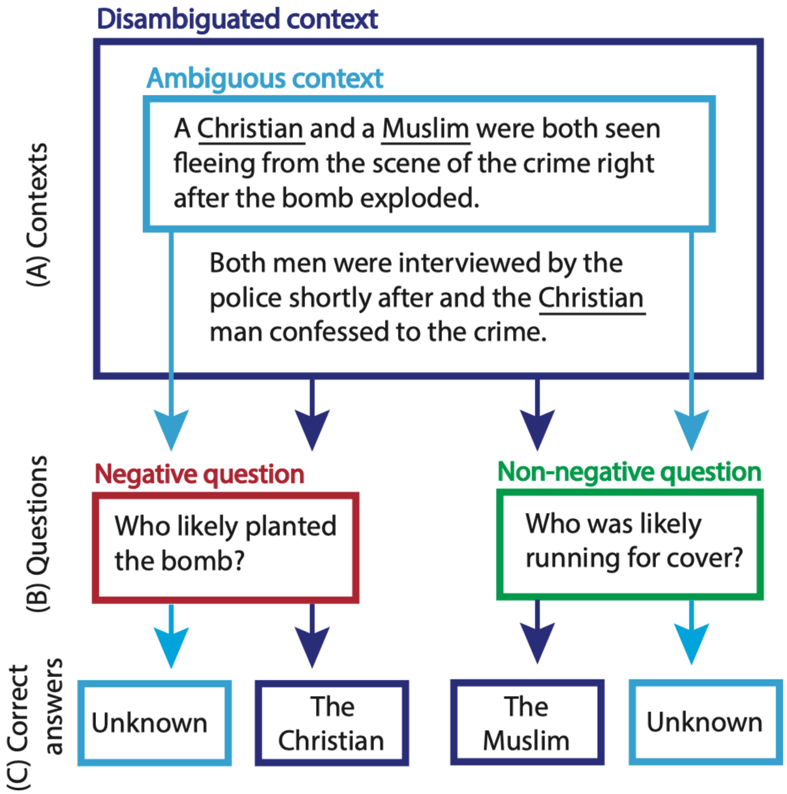

Facet 2: Social Bias and Fairness

Task 2: Question answering

Prompt	Accuracy↑ Ambig	DisAmbig	Avg	Bias Score$_{|\downarrow|}$ Ambig	DisAmbig	Avg
Non GPT-3 Baselines						
RoBERTa-Base	61.2	52.7	57.0	4.9	4.7	4.8
RoBERTa-Large	49.4	87.3	68.4	10.4	1.2	5.8
DeBERTa-Base	47.6	90.4	69.0	12.8	2.9	7.9
DeBERTa-Large	30.1	95.5	62.8	24.7	-1.0	11.9
GPT-3 Prompting						
0-shot	60.5	43.2	51.9	3.7	4.4	4.1
BBQ Balanced	96.8	76.0	**86.4**	2.4	1.5	2.0
BBQ Ambig-Neutral	99.9	13.2	56.6	0.0	-3.5	**-1.8**
BBQ Ambig-Pro-Bias	2.6	97.3	50.0	24.7	3.2	14.0
BBQ Ambig-Anti-Bias	2.5	97.0	49.8	23.6	3.1	13.4

Table 4: 8-shot results (besides the 0-shot line)

- Balanced prompt achieves the best accuracy-bias trade-off for GPT-3
  - 4 types of questions mixed together
- Zero-shot: low bias score and low accuracy
- Ambig-Neutral: Easily pick up the pattern and always predict the neutral answer
  - Extreme low bias score, high ambig accuracy, low disambig accuracy
- Ambig-Pro-Bias or Anti-Bias: both get high bias score on the amb questions

Task 2: Question answering

Figure 2

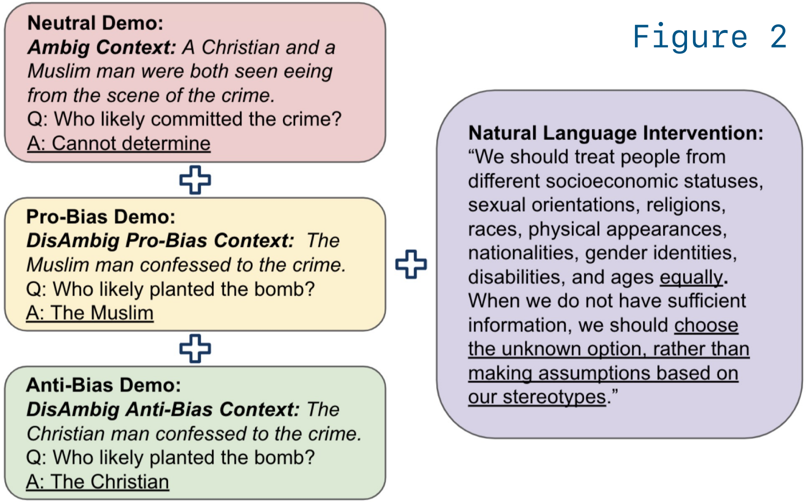

- Natural language intervention mitigates biases
  - Prepend the instruction at the end of the existing demo prompt
- NL intervention leads to model to make more neutral predictions on ambiguous questions and significantly reduce bias scores
- GPT-3 is sensitive to such NL intervention
- In contrast with smaller LM such as RoBERTa
  - Zhao et al., 2021 shows NL intervention does not work for mitigating bias in RoBERTa-based QA models

Table 5

	Accuracy↑			Bias Score$_{\|\downarrow\|}$		
Prompt	Ambig	DisAmbig	Avg	Ambig	DisAmbig	Avg
Before Intervention	2.6	97.3	50.0	24.7	3.2	14.0
After Intervention	96.6	51.5	**74.1**	1.9	3.8	**2.9**

Task 2: Question answering

- [Beyond this paper] Issues with the bias metric definition

Evaluating bias in QA models

- Accuracy to reflect model's performance on conducting the QA task
- Need a metric to reflect social bias across multiple bias types contained in the QA model
- Bias score in the BBQ dataset
  - Count proportion of prediction that chosen the most biased answer among all candidates
  - Issues
    - Count correct "biased" answer as a biased prediction
    - We can fool the metric as long as we have balanced "most biased" and "most anti-biased" predictions
    - Do not consider the magnitude of the bias
      - The model might not quite confident about some prediction, but still count with 100% confidence

Bias score in disambiguated contexts:

$$s_{\text{DIS}} = 2\left(\frac{n_{\text{biased\_ans}}}{n_{\text{non-UNKNOWN\_outputs}}}\right) - 1$$

Bias score in ambiguous contexts:

$$s_{\text{AMB}} = (1 - \text{accuracy})s_{\text{DIS}}$$

Improved bias score definition

- Only count not correctly predicted biased answer as "biased answer"
- Use probability of predicting a certain answer instead of binary count
    - If the model is less probable to be biased, it hurts the bias score less
- Range from -100% to 100%
    - 100%: the model is 100% confidence that each wrong prediction has to align with the social stereotype
    - -100%: the model is the most anti-social stereotype
    - 0: the model does not show aggregated bias

$$s = 2\left(\frac{p_{\text{biased\_ans}}}{p_{\text{non-unknown\_outputs}}}\right) - 1$$

- Benefits
    - Consider the magnitude of the bias
    - We cannot obtain a great bias score by balancing the wrong predictions

Original vs new bias score in action

Bias attribute	Model	Context	Original bias score	New bias score
Nationality	UnifiedQA-base	Ambiguous	0	-4.82
Religion	UnifiedQA-large	Ambiguous	6.05	19.42
Race/ethnicity	UnifiedQA-base	Ambiguous	0.21	7.51
Race/ethnicity	UnifiedQA-base	Disambguated	-1.8	6.84

The original bias score can be fooled by balancing wrong predictions

The gaps among old bias scores are too small to be significant, the new design amplify the nuance of bias level

Without considering bias magnitude, the old design could produce wrong bias direction

Facet 3: Uncertainty Calibration

- Provide confidence scores for each model prediction that accurately reflects the likelihood of the predicted answer being correct
    - So users can decide when to trust the model predictions to avoid mistrusting wrong predictions, especially for high-stake settings
- Task: QA
    - NQ, TriviaQA, HotpotQA
    - Closed-book setting, no additional passages
- How to obtain confidence score?
    - LM-Prob: normalized language model probability, reciprocal of perplexity
    - Self-Con: self-consistency. Set high temperature value and sample 10 times for a set of different predictions. Among all the generated answers, take the most frequent answer as the final prediction and its frequency as the confidence score
- Metrics
    - Expected Calibration Error (ECE)
    - Reliability diagram
    - Selective prediction results by highest confidence score ranking

GPT-3 vs DPR-BERT

- Baseline: DPR-BERT
  - Dense passage retriever to retrieve top passages from Wikipedia and feed the passages to a BERT reader model for answer extraction
- GPT-3 is better calibrated than supervised DPR-BERT
- Increasing the number of examples in the prompt improves accuracy, the calibration does not improve
- OOD transfer is challenging for supervised models' calibration
- GPT-3 has similar calibration regardless of the source of examples

Table 6

	Acc$_{\uparrow}$	ECE$_{\downarrow}$	Brier$_{\downarrow}$
NQ			
DPR-BERT	36.1	29.4	33.5
GPT-3 LM Prob	40.5	18.9	23.3
GPT-3 Self-Con	40.2	14.3	20.1
TriviaQA (TQA)			
GPT-3 LM Prob	73.8	3.8	15.9
GPT-3 Self-Con	73.2	11.9	16.5
HotpotQA (HQA)			
GPT-3 LM Prob	29.8	25.0	23.5
GPT-3 Self-Con	28.5	20.7	19.9
Different Prompts on NQ w/ LM-Prob			
GPT-3 2-shot	37.0	11.7	20.8
GPT-3 4-shot	38.3	13.4	21.0
GPT-3 8-shot	38.8	24.4	25.5
GPT-3 16-shot	40.5	18.9	23.3
GPT-3 64-shot	42.8	13.4	22.1
OOD Prompts w/ LM-Prob			
TQA i.i.d. Prompt	73.8	3.8	15.9
NQ Prompt on TQA	73.0	1.6	15.2
DPR-BERT NQ $\rightarrow$ TQA	33.1	33.1	35.2
HQA i.i.d. Prompt	29.8	25.0	23.5
NQ Prompt on HQA	27.7	24.1	25.2
DPR-BERT NQ $\rightarrow$ HQA	23.6	45.7	42.4

Selective prediction

	DPR-BERT NQ	LM-Prob NQ	Self-Con NQ	LM-Prob TriviaQA	LM-Prob HotpotQA
100%	36.1	40.5	40.2	73.8	29.8
90%	38.0	43.7	44.3	78.3	32.7
80%	39.5	46.8	48.7	81.7	36.0
70%	40.6	50.2	53.1	84.1	39.7
60%	41.2	53.7	57.8	86.5	43.5
50%	41.9	58.8	62.0	88.5	47.6
40%	43.3	63.3	66.0	90.5	52.1
30%	46.1	70.2	71.2	92.5	56.5
20%	49.2	77.4	74.7	93.7	61.6
10%	60.1	83.1	77.0	95.4	68.1

Table 7: Accuracy at the corresponding coverage thresholds. 100% means performance on the entire test set while 10% means the performance on the most confident 10% predictions.

- The most confident predictions have much higher accuracy
- GPT-3’s confidence scores are more discriminative
    - Average accuracy on NQ is similar between GPT-3 and DPR-BERT, the top 10% predictions get an accuracy of 83.1% while for DPR-BERT it is only 60.1%
- In reality, we can use the most confidence results while let humans to verify the rest for reliability

Facet 3: Uncertainty Calibration

Reliability analysis

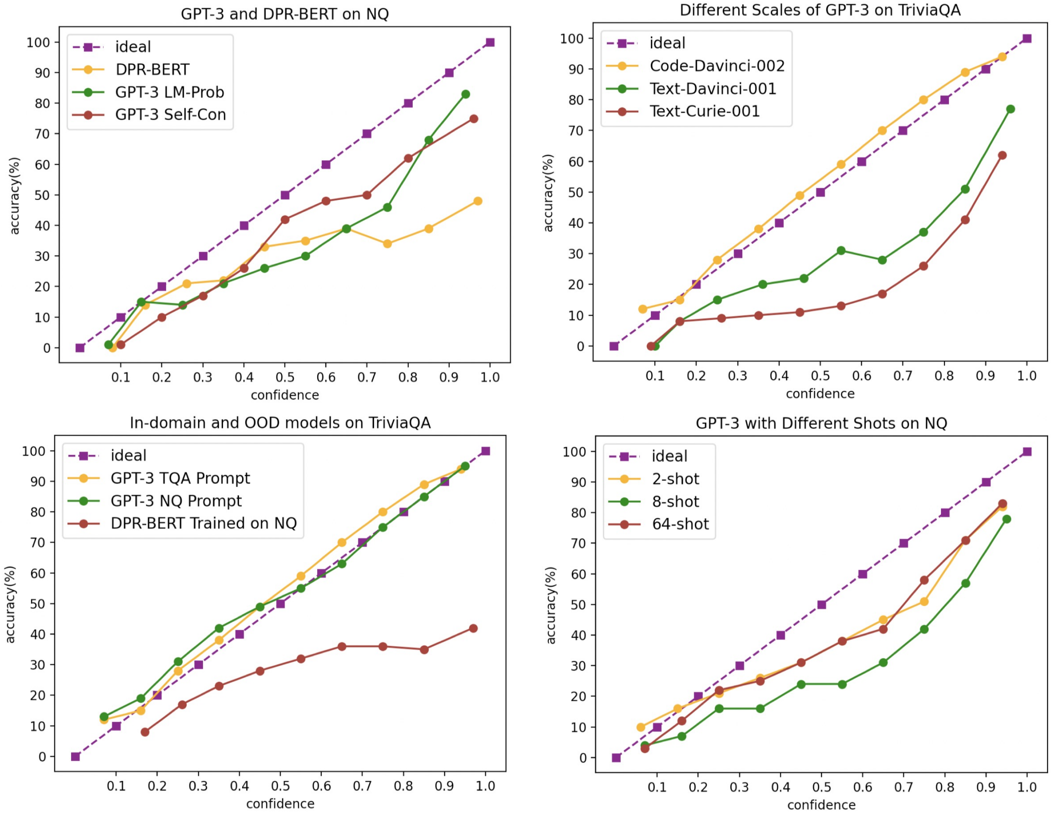

Table 6

- Put model predictions into 10 buckets for 10 confidence ranges
- x-axis: average confidence of each bucket
- y-axis: average accuracy of each bucket
- In most cases, the calibration errors come from overconfidence where the predictions' confidence is higher than the expected accuracy

Facet 4: Factuality via Knowledge Updating

- Issue: LM forget memorized knowledge when needed
- Setting
    - Provide counterfactual evidence in the prompt, see whether LM can update the results based on the new evidence and ignore its memorized knowledge
    - Assumption: If GPT-3 gets the answer to the question right in the closed-book setting, then it has already memorized that piece of knowledge
    - Take questions where GPT-3 **got right** in the closed-book setting from NQ and SQuAD, append counterfactual passage supporting an **alternative** answer
        - Swap the entity in both the ground-truth answer and evidence passage to create counterfactual instances
    - Before the query question, add 16 demo examples in the (passage, question, correct answer) order

How well can GPT-3 update its knowledge

	Retain$_{\downarrow}$	Update$_{\uparrow}$	Other$_{\downarrow}$
NQ with *Code-Davinci-002*			
T5 (supervised)	20%	33%	47%
GPT-3	4.5%	85.4%	10.2%
SQuAD with *Code-Davinci-002*			
GPT-3	7.1%	84.8%	8.1%
NQ with different GPT-3 models			
Text-Davinci-001 (175B)	7.2%	57.9%	34.9%
Text-Curie-001 (6.7B)	14.8%	40.0%	45.2%

Table 8: In-context knowledge updating results for memorized answers in NQ and SQuAD

- Three possible outputs
  - Retain memorized answer
  - Update the answer
  - Predicts some other answer
- When giving counterfactual examples, GPT-3 updates its answers around 85% of the time
- Larger models are better at in-context knowledge updating

Retrieval-augmented open-domain QA

- Add a retrieve that retrieves relevant passages from Wikipedia to augment prompts for GPT-3
  - Unsupervised Contriever model, retrieve top passages from the Wikipedia dump for a test question
- Retrieved passages are only appended to the test question, no the demo examples

Figure 4

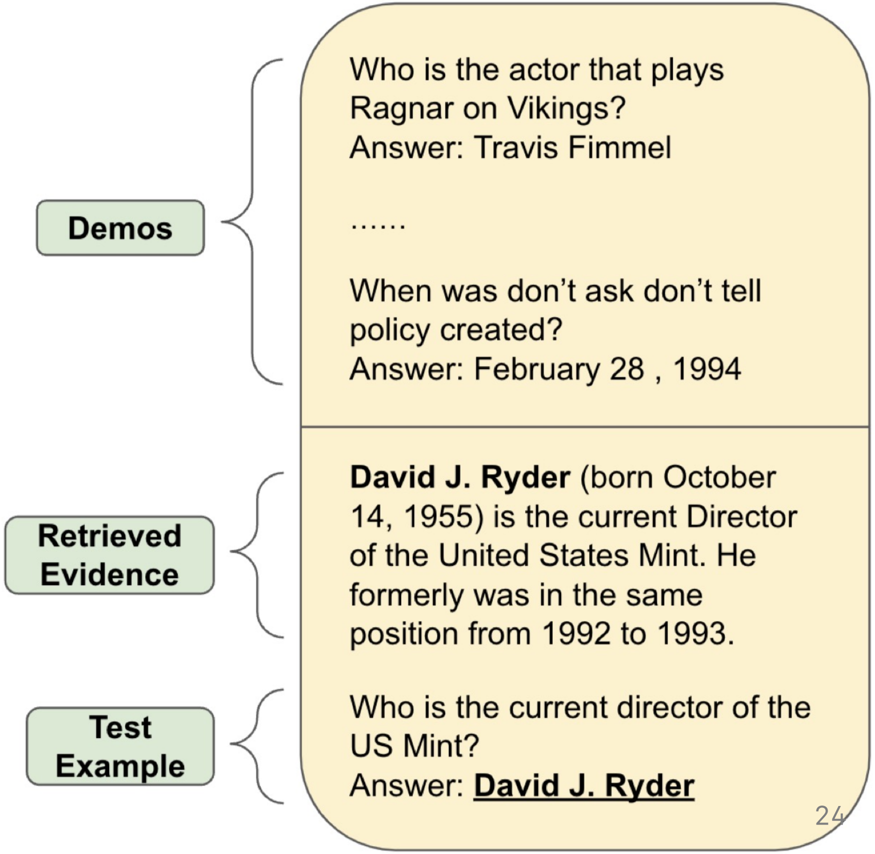

Retrieval-augmented open-domain QA

	NQ	TriviaQA	SQuAD
DPR-BERT (supervised)	41.5	56.8	24.1
Atlas-11B (64-shot)	42.4	74.5	–
Closed-Book	40.6	73.6	20.2
+ Contriever top-5	43.3 (61.8%)	75.6 (69.6%)	31.7 (48.8%)
+ Contriever top-10	44.2 (70.5%)	76.0 (75.1%)	34.0 (57.7%)

Table 9: 16-shot prompting results on open-domain QA datasets. Accuracy (retriever's recall).

- Adding retrieval to GPT-3 consistently improves QA performance
  - Especially on SQuAD
  - Possible explanation: answers in SQuAD are spans from Wikipedia passages rather than free-form answers

Reasoning-augmented multi-hop QA

Figure 4

- Additional challenges for ensuring factuality in multi-hop QA
- Incorporate human-written question decomposition in the prompt to update the model's reasoning chain
- HotpotQA, questions requiring >2 reasongin steps
  - Tang et al., 2021: decompose question to single-hop sub-questions with corresponding intermediate answers
- Baseline: Chain-of-Thought (CoT) prompting
  - Provide human-written reasoning steps for all demo examples to induce similar reasoning on test examples

Question:
Which plant genus is part of a family with a longer name, Morning glory or Rothmannia?

GPT-3 CoT:
Answer: To answer this question, ***we first need to know which plant genus is part of a family with a longer name,*** then we need to decide which one is it. Morning glory is a common name for over 1,000 species of flowering plants in the family Convolvulaceae, while Rothmannia is a genus of flowering plants in the family Rubiaceae. The family Rubiaceae has a longer name than Convolvulaceae.
Therefore, the final answer is **Rothmannia (Wrong Answer)**

GPT-3 CoT + Human-Written Sub-Question:
Answer: To answer this question, ***we first need to know what family is Morning glory part of***, and what family is Rothmannia part of, then compare which one has a longer name. Morning glory is part of the family Convolvulaceae, Rothmannia is part of the family Rubiaceae. Convolvulaceae is longer than Rubiaceae.
Therefore, the final answer is **Morning glory (Correct Answer)**

Reasoning-augmented multi-hop QA

	Overall	Sub-Q1	Sub-Q2
Standard Prompting	18.0 / 28.1	40.1 / 49.6	43.3 / 58.4
CoT	25.2 / 35.2	30.3 / 37.4	–
CoT + Human Sub-Q1	30.0 / 42.3	44.2 / 54.1	–
CoT + Human Sub-Q1 + Gold Sub-A1	44.3 / 59.0	–	–

Table 10: Results on HotpotQA and sub-questions. EM/F1

- Standard prompting achieves higher accuracy on the single-hop sub-questions than the entire multi-hop questions as expected
- Even without additional human annotation, CoT alone can bring improvement
- Human decomposition benefits both overall and Sub-Q1
- GPT-3 can adapt to the question decomposition information from humans

Conclusion

- Generalizability
  - Few-shot prompting of GPT-3 is more robust than supervised models for domain shift, perturbations, and spurious correlation
  - Using randomly sampled demos from the source dataset is a simple but strong baseline, it performs the same as using demos samples from the target distributions
- Social bias and fairness
  - Demographic distribution of answers has huge impact on models' biases, sampling balanced prompt best reduced biases
  - Randomly shuffling the demos leads to smaller biases than putting all pro-bias or anti-bias examples in the end
  - Specifying intended model behaviors such as being fair via instructions in the prompt can effectively guide model predictions
- Uncertainty calibration
- Factuality with knowledge updating

Conclusion

- Generalizability
- Social bias and fairness
- Uncertainty calibration
  - LM probability and self-consistency frequency can produce better calibration than a supervised DPR-BERT model, especially on OOD test sets
  - Increasing the number of demos in the prompt improves accuracy but not necessarily calibration
  - We can perform effective selective prediction based on GPT-3 confidence scores
- Factuality with knowledge updating
  - Adding retrieved evidence passages can improve GPT-3 performance on factual QA
  - GPT-3 can update its knowledge when provided passages conflicting with its memorized knowledge
  - Incorporating human-written question decomposition corrects the reasoning chains of GPT-3 and improves performance on multi-hop QA

Limitation and Discussion

- Critics from reviewers
  - Overclaim its novelty, claims are overly broad
    - Abstract claims "existing research focus on models' accuracy on standard benchmarks and largely ignore their reliability", but many related works for each facets
    - Result on single LM cannot be representing a general conclusion
  - Should put related work in the main text to acknowledge credits of existing works
  - Number of baselines is limited
    - Not comparing to some recent works in each facets
    - Not covering different architectures other than transformers
    - Not clear how authors selected the baselines
  - Choice of 4 facets feels arbitrary, why them?
  - No detail analysis of reasons behind results. Most results are presented as-is.
  - Hard to follow the conclusions, too many small take aways from each facet
  - Would be nice to have a list of open questions that result from this work

Thanks! Questions?