

How Ambiguous Are the Rationales for Natural Language Reasoning? A Simple Approach to Handling Rationale Uncertainty

Hazel Kim

University of Oxford
Classting AI Research
hazel.kim@cs.ox.ac.uk

Abstract

The quality of rationales is essential in the reasoning capabilities of language models. Rationales not only enhance reasoning performance in complex natural language tasks but also justify model decisions. However, obtaining impeccable rationales is often impossible. Our study aims to investigate how ambiguous rationales play in model performances of natural language reasoning. We first assess the ambiguity of rationales through the lens of entropy and uncertainty in model prior beliefs, exploring its impact on task performance. We then propose a simple way to guide models to choose between two different reasoning paths depending on the ambiguity of the rationales. Our empirical results demonstrate that this approach leads to robust performance, particularly in adversarial scenarios where rationale quality is inconsistent.¹

1 Introduction

Language models (LMs) have achieved significant progress on sophisticated reasoning tasks requiring commonsense knowledge or selecting the best answer among tricky multiple choices (Wei et al., 2022; Liu et al., 2022). Recent advances have enabled LMs to generate explicit free-text rationales and use them to guide task predictions with better search space for candidate answers without injecting additional knowledge (Zelikman et al., 2022; Kojima et al., 2022).

However, it is impossible to consistently obtain perfect rationales from models or even from humans. (Jung et al., 2022; Wang et al., 2023; Chen et al., 2023). Despite the advantage of rationales that can explain model decisions, the human utility of those generated by large-scale LMs is far from satisfactory. Joshi et al. (2023) has recently reported that only 20% of them are useful for humans

¹The work was mostly done at Classting AI Research. The author is currently affiliated with the University of Oxford.

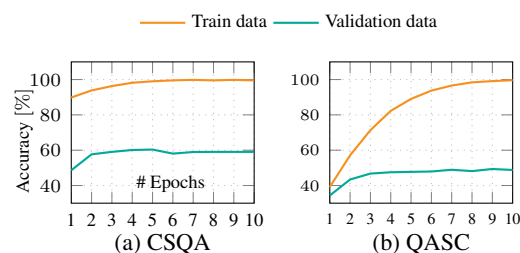


Figure 1: **The huge performance mismatch between training and validation sets by epochs.**

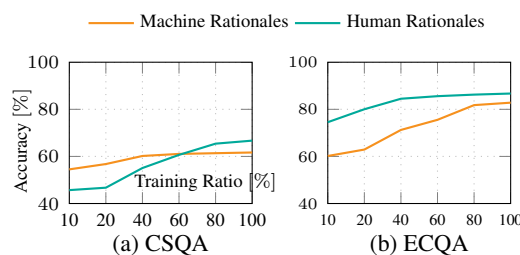


Figure 2: **Performance of human or machine rationales by training ratios.** Though we add up the data points, the performance of the reasoning model remains consistent once it meets certain performance thresholds.

to gain additional information to answer questions. Human annotation may enhance the quality of rationales but is often costly yet does not guarantee perfect conditions, either.

Above all, learning the patterns of rationales that are generated based on a tremendous number of different statements is almost improbable. By nature, rationales carry varying normative concepts, corresponding to various kinds of reasoning to justify thinking or action in a certain way (Wedgwood, 2017). The diversity leads models to face the aleatoric uncertainty (i.e., data uncertainty) that identifies ambiguity or randomness inherent in the observation, causing the difficulty of learning the rationales, as shown in Figure 1. The Figure 2 shows that the performance does not significantly improve with additional data due to the noisiness in the dataset even with human rationales.

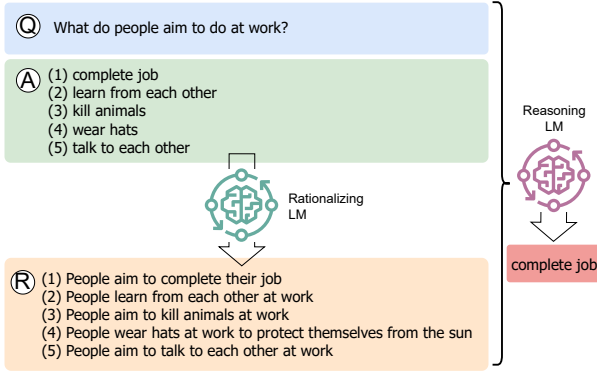


Figure 3: **An example of reasoning task on generated rationales.** The rationalizing LM generates rationales of the answer choices and the reasoning LM predicts answers on the given question, answer choices, and the generated rationales.

Our study aims to investigate how the ambiguous rationales play in model performances of natural language reasoning. We explore how differently human-annotated and machine-generated rationales impact model reasoning performances. We then introduce a method of simply dealing with unclear machine-generated rationales, proposing a suitable two-system reasoning mechanism depending on the level of ambiguity. Because the quantity of the data does not improve the quality of the data, we target to work with the suboptimal quality of the rationales for faithful enough reasoning rather than improving the quality of the rationales itself.

2 Method

We study how the model should deal with ambiguous rationales for reasoning tasks.

2.1 Natural Language Reasoning with Rationales

Following Shwartz et al. (2020); Wiegrefe et al. (2021); Wang et al. (2023), we explore multiple-choice question-answering tasks in the form of question q and a set of answer choices $A = \{a_i\}$. For such a task, the model predicts an answer $\hat{a} = \arg \max_{a_i \in A} \rho(q, a_i)$ and aims to match the predicted answer with a correct answer choice $a^* \in A$. Note that $\rho(q, a_i)$ is a plausibility score for each (q, a_i) pair.

Given rationales, the reasoning model learns to output plausibility score $\rho(q, A, a_i, r_i)$ such that a question q , the answer choices A , answer candidate $a_i \in A$, and rationale r_i . The reasoning model computes ρ_i by aggregating the probabilities P of to-

kens t_i^j where a_i 's tokens $t_i = [t_i^1, t_i^2, t_i^3, \dots, t_i^{|a_i|}]$:

$$\rho_i = \sum_{j=1}^{|a_i|} \log P(t_i^j | t_i^{j-1}, \dots, t_i^1, t_i^1, q, A, r_i). \quad (1)$$

The aggregated probabilities $P(a_i | q, A, R)$ normalize ρ_i by using softmax function as $P(a_i | q, A, R) = e^{\rho_i} / \sum_{j=1}^{|A|} e^{\rho_j}$. The standard training objective is to maximize the likelihood of the correct answer choice using cross-entropy loss, computed as:

$$L = -\sum_{a_i \in A} Q(a_i | q, A) \log P(a_i | q, A, R) \quad (2)$$

where $Q(a_i | q, A)$ is 1 if $a_i = a^*$ and 0 otherwise.

We mainly explore the reasoning module that produces the ideal outcome given rationales with various qualities.

2.2 Ambiguous Rationales

We calculate the entropy values of the rationales, using model prior beliefs as informativeness on multiple choices A , and define those with high entropy values as ambiguous.

Entropy. In information theory (Shannon, 1948), the entropy measures information or uncertainty inherent to the possible outcome:

$$H = -\sum_x p(x) \log p(x) \quad (3)$$

where p_i is the probability of the event occurring and $\log p$ shows the informativeness of the event.

In the rationale entropy estimation, we use the prior beliefs as informativeness, $\log P(a_i | x_i, \hat{\theta})$, and the posterior beliefs as the probability of the event occurring, $P(a_i | x_i, \hat{\theta}_{MLE})$ as illustrated:

$$H(x) = -P(a_i | x_i, \hat{\theta}_{MLE}) \log P(a_i | x_i, \hat{\theta}) \quad (4)$$

We aim to measure how the posterior beliefs are uncertain about what prior beliefs have high informativeness.

Model Prior and Posterior Beliefs. LMs have shown that their parameters might contain a significant body of factual and commonsense knowledge as they get trained on a massive volume of text (Wei et al., 2022; Cohen et al., 2023). We call the internal knowledge that the models have had before being fine-tuned with the training set a prior belief of the model $\hat{\theta}$. After the same model gets trained with a training set, we call their knowledge a posterior belief of the maximum likelihood model $\hat{\theta}_{MLE}$. The prior and posterior belief probabilities

come from the model plausibility score as in Equation 1. We normalize the probability scores over the answer choices for each question. We do not particularly measure model prior and posterior beliefs in the context of Bayesian settings but rather a deterministic approach with a single model.

Rationale Ambiguity. We use the expected value of entropies with the entire data points as a threshold to determine ambiguous and unambiguous rationales within the distribution of how data points interact with model training.

$$R(x) = \begin{cases} r_{\text{unambiguous}}, & \text{if } h(x_i) < \tau \\ r_{\text{ambiguous}}, & \text{if } h(x_i) \geq \tau \end{cases} \quad (5)$$

where τ is $\mathcal{E}[H(x)]$. The randomness and ambiguity mainly come from their indecisive behaviors between prior and posterior beliefs, comparing before and after the model learning.

2.3 AURA: Reasoning with Ambiguous Rationales.

We utilize a two-stage reasoning system depending on the ambiguity scores of rationales from Equation 5. We first use a pre-trained model to train a reasoning module on the entire dataset (Reasoning 1). We use the pre-trained model $\hat{\theta}$ as prior beliefs and the trained model $\hat{\theta}_{\text{MLE}}$ as posterior beliefs to calculate rationale entropy. We then train the same pre-trained model again on the selected ambiguous rationales to learn unlearned or underlearned data points that have high entropy scores (reasoning 2). It works as ensemble learning as illustrated in Equation 6.

$$\{P(y|x^*, \theta^{(t)})\}_{t=1}^{T=2} \rightarrow \{P(y|W^{(t)})\}_{t=1}^{T=2}, \quad (6)$$

$$W^{(t)} \sim P(W|x^t, D)$$

where $t = 1$ and $t = 2$ indicate Reasoning 1 and 2 procedures, respectively.

3 Experiments

In this section, we observe the superiority of AURA which deals with aleatoric uncertainty in rationales for reasoning tasks.

3.1 Experimental Setup

We use five different datasets and seven types of baseline approaches to explore AURA. We conducted the experiments with 5 other random seeds and reported the average scores for the results.

Datasets We experiment with four commonsense question-answering datasets that require machine rationales in addition to one benchmark that includes human rationales: (a) CSQA (Talmor et al., 2019) a 5-choice dataset that requires knowledge extracted from ConceptNet. (a-1) ECQA (Aggarwal et al., 2021) explains each choice to questions from CSQA but with human rationales. (b) StrategyQA (Geva et al., 2021) is a dataset that requires an advanced level of reasoning strategies to produce binary (yes/no) answers. (c) OpenbookQA (Mihaylov et al., 2018) is a 4-choice dataset to test model ability as a form of open book testing. (d) QASC (Khot et al., 2020) is a complicated 8-choice question set.

Baselines We consider four rationalize-then-reason pipeline approaches similar to AURA as baselines. We also compare our approach with those with self-rationales or without any rationales. (1) Without rationales: fine-tune T5-base model without providing rationales. (2) Answer upon self-rationales using a medium-sized LM: make GPT-neox (20B) (Black et al., 2022) learn from a few examples to generate machine rationales and the following answers. (3) Answer upon self-rationales using a large-scale LM: same mechanism as (2) but use GPT3 (175B) Wei et al. (2022). (4) NILE (Kumar and Talukdar, 2020): a method of language inference on label-specific explanations. Following the standard setup, we fine-tune NILE with a T5 (3B) model, unlike other baselines that use a T5-base model. (5) Train on prompted rationales with a standard training method: use GPT-neox but obtain answers not from the same model but from a separate reasoning model, T5-base in our experiment. (6) Dropout context is the same as (5) but randomly drops out from the input question while training the reasoning model (Wang et al., 2023). (7) same as (5) but trained with counterfactual regularization objective. We report the results presented by Wang et al. (2023).

Implementation Details For machine-generated rationales, we follow the same setup with Wang et al. (2023) for a fair comparison, 7 manually-annotated prompt examples to GPT-neox (20B) for each dataset. We adopt the T5-base model for the reasoning module.

3.2 Evaluation

We test AURA on In-Distribution (ID) and Out-of-Distribution (OoD) settings and different training

Method	Dataset				
	CSQA	StrategyQA	OBQA	QASC	Average
w/o Rationales ... (1)	58.68	58.12	55.85	35.58	52.06
Self-Rationalization					
GPT-neox, 20B ... (2)	38.41	55.31	33.80	32.61	45.89
GPT3, 175B (Chain-of-Thought) ... (3)	63.50	64.40	-	-	-
Pipeline Rationalization					
NILE ... (4)	57.60	57.31	-	-	-
Standard Training ... (5)	59.48	57.11	56.65	37.50	52.69
Dropout Context ... (6)	59.64	51.45	57.55	35.37	51.00
PINTO Counterfactual Regularization [†] ... (7)	61.67	60.87	58.85	37.82	54.80
AURA [‡]	67.13 (+5.46)	76.35 (+15.48)	65.89 (+7.04)	40.82 (+3.00)	62.55 (+7.04)

Table 1: **Comparison of task performances reported with accuracy.** We bold the best results and report the performance gain by AURA from the best results among existing pipeline rationalization methods with (+ sign). [†]: a current state-of-the-art method in the pipeline rationalization, [‡]: our proposed method.

Method	Dataset (<i>source</i> \rightarrow <i>target</i>)				
	CSQA \rightarrow OBQA	CSQA \rightarrow QASC	OBQA \rightarrow CSQA	QASC \rightarrow CSQA	QASC \rightarrow OBQA
w/o Rationales ... (1)	32.05	39.17	24.87	45.74	34.90
Pipeline Rationalization					
NILE ... (4)	32.40	40.93	-	-	-
Standard Training.... (5)	31.05	40.04	25.37	47.71	34.50
Dropout Context ... (6)	32.30	38.85	23.01	44.27	32.90
PINTO Counterfactual Regularization [†] ... (7)	34.90	42.25	27.66	48.03	35.75
AURA [‡]	49.80 (+14.90)	43.20 (+0.95)	42.91 (+15.25)	52.48 (+4.45)	51.40 (+15.65)

Table 2: **OoD Test.** Performance on a target dataset after training on a source dataset, denoted as *source* \rightarrow *target*.

ratios to investigate the superiority of AURA.

Comparison against Other Baseline Methods.

We first evaluate AURA in ID settings where we use the target dataset for rationalizing and reasoning modules and test the performance on unseen examples of the same dataset. We have observed results as shown in Table 1. AURA achieves strong task performance compared to other baseline methods on almost all four benchmark datasets. For CSQA and StrategyQA, AURA outperforms the self-rationalizing method with chain-of-thought (GPT3, 175B) by 3.63% and 11.95% accuracy points, respectively. AURA outperforms the state-of-the-art approach with the counterfactual regularization method on four out of four benchmarks and achieved the new state-of-the-art by increasing the accuracy score by 7.75 % points on average.

Out-of-distribution (OoD) Performance We now evaluate AURA in OoD settings where we train on source datasets and test on different target datasets to investigate how AURA helps models generalize well to reason on rationales rather than memorize the concepts or answers. Table 2 shows that AURA outperforms all existing methods on the entire benchmark datasets by an average score

of 10.24% and at most 15.65% of accuracy. Following previous work, we adopt the intuition that the models produce better OoD results than others if they utilize rationales faithfully.

Training Ratios. We test AURA on different training ratios to see its performance in low-resource settings. We observe that AURA is strong on limited resource settings on all four benchmarks as shown in Figure 4. We observe AURA outperforms PINTO, the state-of-the-art rationalize-then-reason method, in both ID and OoD settings. Due to the nature of aleatoric uncertainty in rationales, the performance gains by AURA are pretty consistent over different training ratios.

3.3 Analysis

RQ1. Does the quality rationales contribute more than the good reasoning mechanisms to the overall performance? Robust reasoning method matters more than obtaining quality rationales for predicting right answers. As shown in Figure 5, the performance gain by AURA (+7.65%) is greater than that of better rationales, using human rationales instead of machine rationales (+4.53%) when tested on machine rationales. Similarly, when

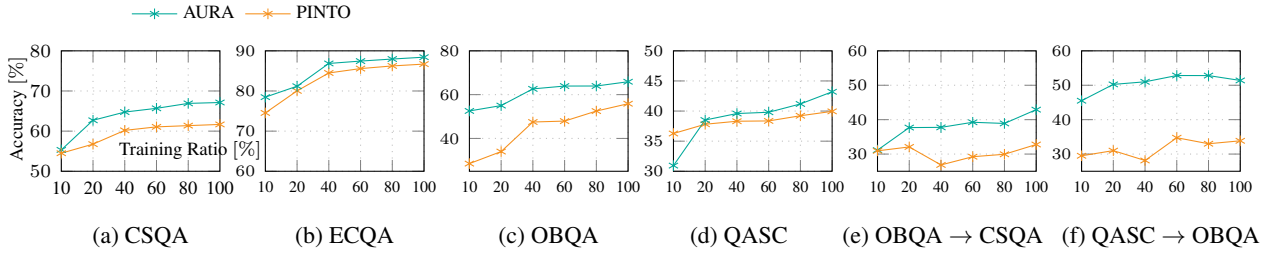


Figure 4: **Performance changes depending on different training ratios.** The performance of in-distribution settings with (a) – (d) and of out-of-distribution settings with (e) – (f).

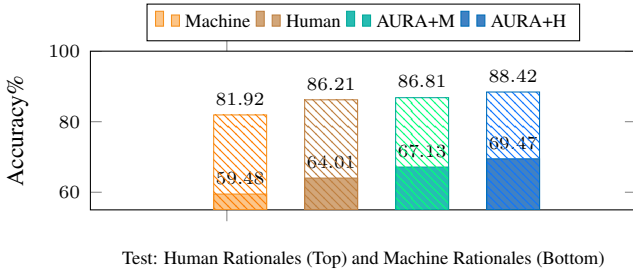


Figure 5: **Performance gains by AURA depending on training and testing rationales.** M: machine-generated rationales, H: human-written rationales. The legend shows types of training rationales with or without AURA. Without AURA is standard training.




Figure 6: **Performance comparison between human and machine rationales in OoD settings.** For the competitive comparison, we used PINTO, the state-of-the-art method, to compare with AURA as the reasoning method in this experiment.

tested on human rationales, AURA (+4.89%) works slightly better than better rationales (+4.29%).

In OoD settings, the statement above is more solid than in ID settings as shown in Figure 6. Superior rationales in ID settings do not guarantee satisfaction in OoD settings (34.91% → 32.2%, 42.25% → 34.82%). On the other hand, the performance gains by better rationales (-2.71%, -7.42%) underperform those by AURA (+14.89%, +0.95%). It is worth noting that AURA always guarantees slight performance gains without any performance degradation. This shows that better reasoning is capable of covering unsatisfactory rationalization.

RQ2. Do quality rationales provide a better influence on training or inference? We have observed in Figure 5 that the quality of rationales does not improve the model learnability much compared to what it does to model inference ability. Standard training on human rationales has increased model performance by 4.53% or 4.29% from machine rationales. Both models trained on machine and human rationales have improved their performance accuracy by 22.44% and 22.2% each when tested on human rationales instead of machine rationales.

If we have a limited amount of quality rationales, we better use them while inferring the answers to obtain the most decent performance level among other methods using them. However, in realistic scenarios, we cannot tell if the seemingly ideal rationales are still satisfactory to the question-answer pairs that we want the answers to. As discussed in RQ1, if the inferring data are OoD, the results would be disappointing.

Hence our research statement is that it is very important to study how we deal with ambiguous, confusing rationales for natural language reasoning. AURA indeed is strong on both OoD settings and unsatisfactory rationales.

4 Conclusion

Generating impeccable rationales is often impossible, and learning the patterns of rationales is almost improbable. We first explore how ambiguous it is for language models to learn rationales for reasoning tasks, both with human-annotated and machine-generated rationales. We then introduce a method of simply dealing with unclear rationales, proposing a suitable two-system reasoning mechanism depending on the level of ambiguity. We empirically argue that AURA produces robust performance superiority against the adversarial quality of rationales and low-resource settings.

5 Limitations

Our approach has been tested mainly on medium-sized LMs due to computational cost. We have also focused on commonsense reasoning tasks rather than those requiring domain-specific knowledge. This is due to our approach of exploring the benefits of model prior knowledge from pre-trained LMs.

Acknowledgements

This research was supported and funded by Classting Inc. The author thanks the members of the Classting AI Research, Jaesang Yoo, Yohaun Yoon, Cheoleui Hong, and Heeseok Jung, for their support for this work in various ways.

References

- Shourya Aggarwal, Divyanshu Mandowara, Vishwajeet Agrawal, Dinesh Khandelwal, Parag Singla, and Dinesh Garg. 2021. Explanations for commonsenseqa: New dataset and models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3050–3065.
- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usven Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. **GPT-NeoX-20B: An open-source autoregressive language model**. In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin. Association for Computational Linguistics.
- Hanjie Chen, Faeze Brahman, Xiang Ren, Yangfeng Ji, Yejin Choi, and Swabha Swayamdipta. 2023. **REV: information-theoretic evaluation of free-text rationales**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 2007–2030. Association for Computational Linguistics.
- Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. **Crawling the internal knowledge-base of language models**. In *Findings of the Association for Computational Linguistics: EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 1811–1824. Association for Computational Linguistics.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. **Did aristotle use a laptop? A question answering benchmark with implicit reasoning strategies**. *Trans. Assoc. Comput. Linguistics*, 9:346–361.
- Brihi Joshi, Ziyi Liu, Sahana Ramnath, Aaron Chan, Zhewei Tong, Shaoliang Nie, Qifan Wang, Yejin Choi, and Xiang Ren. 2023. **Are machine rationales (not) useful to humans? measuring and improving human utility of free-text rationales**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 7103–7128. Association for Computational Linguistics.
- Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le Bras, and Yejin Choi. 2022. **Maieutic prompting: Logically consistent reasoning with recursive explanations**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 1266–1279. Association for Computational Linguistics.
- Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. 2020. **QASC: A dataset for question answering via sentence composition**. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8082–8090. AAAI Press.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. **Large language models are zero-shot reasoners**. In *NeurIPS*.
- Sawan Kumar and Partha P. Talukdar. 2020. **NILE: Natural language inference with faithful natural language explanations**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8730–8742. Association for Computational Linguistics.
- Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh Hajishirzi. 2022. **Generated knowledge prompting for commonsense reasoning**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3154–3169. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. **Can a suit of armor conduct electricity? A new dataset for open book question answering**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2381–2391. Association for Computational Linguistics.
- Claude E. Shannon. 1948. **A mathematical theory of communication**. *Bell Syst. Tech. J.*, 27(3):379–423.

- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Unsupervised commonsense question answering with self-talk](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4615–4629. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [Commonsenseqa: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4149–4158. Association for Computational Linguistics.
- Peifeng Wang, Aaron Chan, Filip Ilievski, Muhao Chen, and Xiang Ren. 2023. [PINTO: faithful language reasoning using prompt-generated rationales](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Ralph Wedgwood. 2017. *The value of rationality*. Oxford University Press.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *NeurIPS*.
- Sarah Wiegrefe, Ana Marasovic, and Noah A. Smith. 2021. [Measuring association between labels and free-text rationales](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 10266–10284. Association for Computational Linguistics.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). In *NeurIPS*.