

MEASURING AND IMPROVING BERT’S MATHEMATICAL ABILITIES BY PREDICTING THE ORDER OF REASONING

Piotr Piekos*
University of Warsaw

Henryk Michalewski
Google, University of Warsaw

Mateusz Malinowski
DeepMind

ABSTRACT

Imagine you are in a supermarket. You have two bananas in your basket and want to buy four apples. How many fruits do you have in total? This seemingly straightforward question can be challenging for data-driven language models, even if trained at scale. However, we would expect such generic language models to possess some mathematical abilities in addition to typical linguistic competence. Towards this goal, we investigate if a commonly used language model, BERT, possesses such mathematical abilities and, if so, to what degree. For that, we fine-tune BERT on a popular dataset for word math problems, AQUA-RAT, and conduct several tests to understand learned representations better. Since we teach models trained on natural language to do formal mathematics, we hypothesize that such models would benefit from training on semi-formal steps that explain how math results are derived. To better accommodate such training, we also propose new pretext tasks for learning mathematical rules. We call them (Neighbor) Reasoning Order Prediction (ROP or NROP). With this new model, we achieve significantly better outcomes than data-driven baselines and even on-par with more tailored models.

1 INTRODUCTION

Automatically solving math word problems has a long history dating back to the middle sixties Bobrow (1964). Early approaches were rule-based matching systems that solve the problem symbolically. Even though there are some impressive symbolic systems that operate in a relatively narrow domain, the inability to successfully scale them up is sometimes presented as a critique of the good-old-fashioned AI, or GOF AI Dreyfus et al. (1992). One issue is to create a formalism that covers all the aspects needed to solve these problems. On the other hand, deep learning LeCun et al. (2015) aims to develop artificial general intelligence that scales better to various problems.

However, despite many successes in computer vision and natural language processing Devlin et al. (2018); He et al. (2016); Krizhevsky et al. (2012); Lan et al. (2019); Mikolov et al. (2013), data-driven methods evade our dream of building a system with basic, every-day, mathematical skills. As large-scale natural language models become more common Devlin et al. (2018); Brown et al. (2020), we would expect them to also reason mathematically.

Since natural language understanding also involves symbolic manipulation Liang (2016), we treat mathematical reasoning as a language understanding and revisit the data-driven paradigm. For that, we rely on a recent language model, BERT Devlin et al. (2018), and challenge it with math word problems Ling et al. (2017). Even though such language models have initially shown promising results, more recent investigation shows they may rely on various biases in their predictions Hendricks et al. (2018); Brown et al. (2020); Bhardwaj et al. (2020); Kurita et al. (2019). Here, we also follow that line of investigation and show these models can answer correctly without an understanding of the rationale behind it.

Furthermore, as directly predicting answers to challenging math problems often requires multiple steps of reasoning, we show that we can improve BERT’s generalization by exposing it to rationales Ling

*piotrpiekos@gmail.com

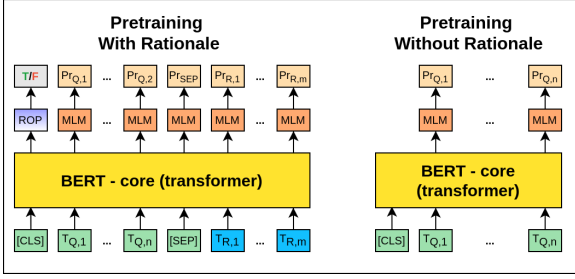


Figure 1: Architectures for different methods. We use shared architecture but we separate question tokens (green blocks) from rationales (blue blocks) using different segment encodings. For question-answering, we fine-tune the whole model with a classification layer. We use the embedding that corresponds to the [CLS] token as the input representation.

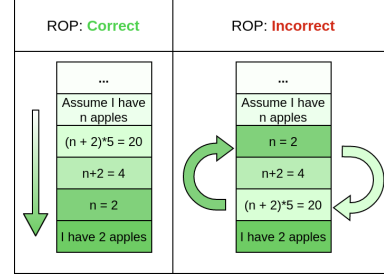


Figure 2: ROP or NROP with positive (left) and negative (right) labels. We randomly swap two rationales and classify if that change has happened.

et al. (2017); Hendricks et al. (2016); Lei et al. (2016). These are, however, only used during training akin to a teacher that shows a student a justification for each answer. But then, the student is evaluated only on the ability to answer these questions during the college exam correctly with no access to rationales. Finally, to learn a better representation from rationales and to improve the generalization even further, we introduce novel pretext tasks and corresponding losses, which we name (Neighbor) Reasoning Order Prediction (ROP or NROP). With that, we achieved significantly better results than data-driven baselines, even on-par with methods that are more tailored to math-word problems and the AQuA-RAT dataset.

2 METHODS

We use the following methods, each initialized with BERT-base pre-trained on Wikipedia and Books Corpus Devlin et al. (2018); Zhu et al. (2015). Note that all use the same number of parameters.

- 1) **BERT-base** - we directly fine-tune BERT to predict the correct answer and show its transfer to math word problems.
- 2) **BERT-AQuA** - we use the MLM loss on the AQuA-RAT questions before training to predict correct answer.
- 3) **BERT-AQuA-RAT** - we use the MLM loss on the AQuA-RAT questions and rationales and show if we can inject knowledge from rationales into BERT.
- 4) **BERT-(N)ROP** - we use the MLM loss and the novel (N)ROP loss for coherence prediction (defined later) and show if we can improve the results by focusing the model on rationales.

We base our architecture on BERT Devlin et al. (2018) that has 12 transformer blocks Vaswani et al. (2017). As the core, we use the standard configuration described in Devlin et al. (2018). We use three self-supervised losses. One is the standard Masked Language Modelling (MLM) but extended to work on rationales. Other two are our new losses, (Neighbour) Reasoning Order Prediction (ROP or NROP). Figure 1 shows two variants of our models. Note that, during fine-tuning, rationales and all the self-supervised losses are discarded.

MLM is the Masked Language Modelling Devlin et al. (2018). We randomly mask 15% of the input tokens by a special token [MASK]. The objective is to predict the masked token using its context casted as a classification problem over the tokenizer vocabulary.

Interestingly, there are parallels between masking numbers and solving mathematical equations, where it can be seen as solving the equation with unknown. For example, $2 + [\text{MASK}] = 4$ becomes $2 + x = 4$. As a consequence, models during training organically deal with mathematical calculations without defining a specific loss for mathematics allowing soft transitions between natural and more formal languages.

ROP is our novel coherence loss. Since rationales are sequences of consecutive reasoning steps, the order of the execution is critical as shown in Figure 2.

Following this intuition, we introduce Reasoning Order Prediction (ROP) that predicts whether the order of the rationale steps is preserved. Hence it encourages the network to pay more attention

Model	Accuracy
Random chance	20.0%
LSTM Ling et al. (2017)	20.8%
BERT-base (ours)	28.3(± 2.0)%
BERT-AQUA (ours)	29.1(± 1.7)%
BERT-AQuA-RAT (ours)	32.3(± 1.8)%
BERT-ROP (ours)	35.4(± 1.0)%
BERT-NROP (ours)	37.0(± 1.1)%
AQuA-RAT Ling et al. (2017)	36.4%
MathQA Amini et al. (2019)	37.9%

Table 1: Comparison of data-driven (first six rows) with two hybrid approaches that use stronger and hence more specific inductive biases (last two rows). Standard deviation estimates (over random initializations) is given in parentheses, where we see our losses can reduce the variability slightly.

Number of tokens	650K	1.3M	1.95M	2.6M	3.25M
% AQuA used	20%	40%	60%	80%	100%
BERT-AQuA Accuracy score	28.23%	27.84%	28.63%	28.23%	29.01%
% of AQuA-RAT used	6.66%	13.33%	20%	26.66%	33.33%
BERT-NROP Accuracy score	28.23%	34.11%	29.80%	36.07%	32.55%

Table 2: Accuracy scores conditioned on the number of tokens available for training. We align models trained on questions with models trained on question with rationales, so that we have comparable number of tokens in both cases. This shows that training with rationales leads to a better representation. Even better than more questions.

to rationales. The loss is similar to Sentence Order Prediction (SOP) Lan et al. (2019), but ours is focused on learning of reasoning steps.

NROP is an extension of ROP where two consecutive rationales are swapped making the prediction (swap or no swap) task more challenging and, hence, it can arguably lead to a better representation as understanding the correct ordering between neighbor steps is more nuanced.

3 RESULTS

Dataset. We use AQuA-RAT Ling et al. (2017). It has about 100k crowd-sourced math questions with five candidate answers (one is correct). Each question has a rationale – a step-by-step explanation of how the answer is computed – that is only available during training. At test time answer predictions are based on questions. The train set has roughly 100k question-answer-rationale triples, while dev and test about 250 question-answer pairs each. We observe a highly sensitive relationship between scores on dev and test sets. we propose an extended dev set consisting of 5000 randomly chosen samples from the training set extended by the whole dev set. The details are in the appendix.

Main results. Table 1 shows our main results. We see that our method is the state-of-the-art among the models with minimal inductive biases and is very competitive to the other two models that are more specific to handle word math problems (e.g., requires programs). Moreover, even though BERT is already a stronger model than LSTM, it is better to use its MLM pretext task and loss on the AQuA-RAT questions (BERT-AQuA) or even better on questions and rationales (BERT-AQuA-RAT). However, models with our novel coherence prediction losses can better learn from rationales (BERT-ROP and BERT-NROP).

Rationales. We hypothesize that rationales contain more information and are more useful for training than questions. To measure that, we want to compare training on the same amount of training data comprised only of questions to the training on questions with rationales. We have estimated that rationales have 1.71 times more tokens than questions. This means that a question combined with rationale has around 3 times more tokens than just a question. We want to prove that rationales contain more information than question. This means that if our hypothesis is valid, training on 20% questions and rationales should give better results than training on 60% questions (counting the

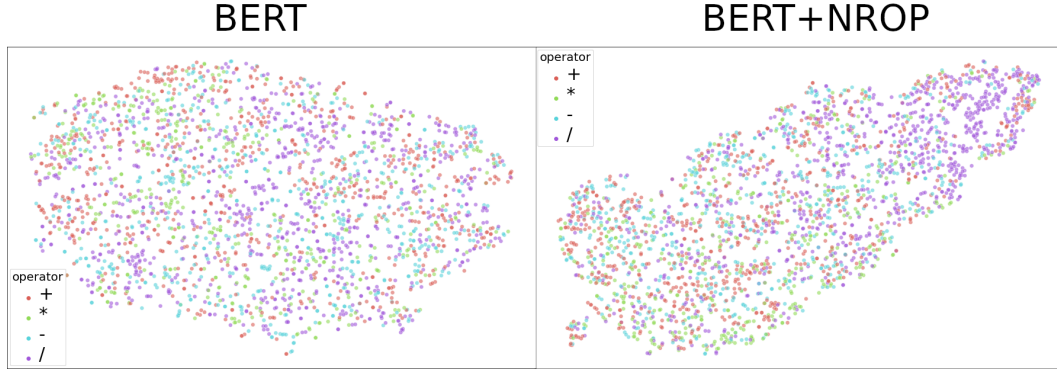


Figure 3: BERT and BERT+NROP embeddings. Colours represent different operators in rationales (T-SNE). BERT+NROP embeddings better separate operators.

number of tokens). We therefore created samples of respective sizes of just questions and questions combined with rationales. We show our results in Table 2. They confirm our hypothesis that more questions do not necessarily help in learning a better representation. On contrary, incorporating rationales can be more helpful.

Embeddings. To better understand the difference between BERT and BERT+NROP, we analyze their embeddings. For our analysis, we sample 2500 questions with a single operator in rationales, and next we visualise them with T-SNE Van der Maaten & Hinton (2008). We show both in Figure 3. We observe that BERT+NROP embeddings preserve more information about different operators.

Permutation consistency. Random guessing on AQUA-RAT yields 20%. For that and to separate from other questions that were solved by chance, we have constructed a new evaluation task – permutation consistency test – where each question gets 5 answers at different positions. Table 3a shows our procedure. Here, models only score a single point if they solve all 5 questions correctly. Hence, random chance is 0.032% in such experiments. Table 3b shows our results. BERT+NROP solves almost three times as many questions as BERT. Additionally, further inspection shows that BERT relies on choosing the answers that most stand out, e.g., numbers ending with zeros or floats while every other option is an integer. We didn’t observe that simple patterns with BERT+NROP. Questions solved by BERT+NROP usually contain one or two operations and show that BERT+NROP better understands the problem. Below, we exemplify two math problems solved by both models.

Example of a problem solved by BERT+NROP: 8 man work for 6 days to complete a work. How many men are required to complete same work in 1/2 day?

Answers: A)93, B)94, C)95, D)96, E)97

Correct Option: D

Example of a problem solved by BERT A ship went on a voyage. After it had traveled 180 miles a plane started with 10 times the speed of the ship. Find the distance when they meet from starting point.?

Answers: A)238, B)289, C)200, D)287, E)187

Correct Option: C

Drop from 37.0% to 11.02% (Table 3b) suggests that models rely strongly on the order of answers. To reduce such a bias, we test several permutation invariant losses.

1) **AUG.** We sample randomly 25 permutations of all the possible answers and use them during training. Original ordering is not used, so there is no order bias. This method can be seen as a data augmentation technique.

2) **SEP-NC.** The original models are trained on a 5-class classification task, where we build the representation by using questions and all the candidate answers, i.e., $\mathbf{BERT}(Q||P)$. Here, $||$ denotes concatenation, Q is the question and P represents the sequence of all answers. In SEP-NC, we block the path between all the candidate answers and the BERT-base. Next, we use a late-fusion to predict if the given candidate answer matches with the question. That is, we use the following formulation $f(\mathbf{BERT}(Q)||\mathbf{BERT}(C))$, where $C \in P$ is a single candidate answer and f is a multi-

Original question	
How much is 27 / 3	A)13 B) 9 C)3 D)12 E)17
Generated questions	
How much is 27 / 3	A) 9 B)13 C)3 D)12 E)17
How much is 27 / 3	A)13 B) 9 C)3 D)12 E)17
How much is 27 / 3	A)13 B)3 C) 9 D)12 E)17
How much is 27 / 3	A)13 B)12 C)3 D) 9 E)17
How much is 27 / 3	A)13 B)17 C)3 D)12 E) 9

(a) Our generation method for the permutation consistency test. Models get a point only if they solve all them.

Model	Score
Random chance	0.032%
BERT	4.33%
BERT+NROP	11.02%
BERT AUG	13.4%
BERT+NROP AUG	19.7%
BERT SEP-NC	15.0%
BERT+NROP SEP-NC	22.7%
BERT SEP-C	16.1%
BERT+NROP SEP-C	23.9%

(b) Our results for the permutation consistency test.

layer perception (with two layers). At test time, the model is prompted to score all five candidate answers and select the one with the highest score. Appendix material has more information about this method.

3) **SEP-C**. As models trained with SEP-NC do not have access to all the possible answers, their biases to answer positions are significantly reduced. However, these models cannot compare each answer to all other candidate answers. Here, we use the following formulation $f(\text{BERT}(Q||P)||\text{BERT}(C))$ to measure the compatibility of the input (question Q and all the candidate answers P) with the given candidate answer $C \in P$. Moreover, we also reset the positional encoding between every possible answer in P . In such a way, we hypothesise the network can, on the one hand, learn a less biased representation, and on the other hand, exploit relationship between the candidate answers. Table 3b shows SEP-NC and SEP-C significantly outperform the original model on the permutation consistency test. We give more details in the supplementary material.

SEP-NC and SEP-C improve results in permutation consistency tests. However, they give similar results to original methods in accuracy measuring task. They achieve respectively 33.5% (SEP-NC) and 35.4% (SEP-C).

Questions difficulty. We categorize questions by difficulty for BERT-NROP and BERT. To estimate a question’s difficulty, we have ranked the candidate answers according to the model’s uncertainties. For instance, if the correct answer has the 2nd largest probability, we assign to that question difficulty two. With that, we group questions into 5 difficulty categories, from the easiest: D_1, \dots, D_5 .

Manual inspection shows that for BERT+NROP: D_5 requires additional knowledge or implicitly defined numbers (e.g., adding first 100 numbers), D_4 requires geometry or non-linear equations and systems, D_3 requires solving linear systems with a few basic operations, D_2 requires solving simple equations, and D_1 has one or two basic operations with clearly written numbers. We show an example from each group in the supplementary material. We didn’t observe a similar pattern for BERT with the exception of the easiest group D_1 where the model chooses the answer that is somewhat different from other candidates. We provide an example of each group in the supplementary material.

Finally, we also compare the difficulty of questions with the difficulty perceived by humans. For that, we have conducted a small-group human study, where we have asked participants to solve some AQuA-RAT questions and rate their difficulty. We find a positive correlation between the difficulty measured by our models (as described above) to the difficulty judged by humans. Details are in the supplementary material.

Conclusions. We have investigated if BERT Devlin et al. (2018) – a pre-trained, large language model – can deal with mathematical reasoning. We find that its representation is biased Brown et al. (2020); Bhardwaj et al. (2020); Kurita et al. (2019) also in mathematics. Our novel pretext tasks and losses reduce that bias, but the network still finds shortcuts that do not generalize well. We hope our work will spark interest of the community in developing language models that are more capable of mathematical reasoning.

Acknowledgements. We thank Wang Ling (DeepMind) for his comments and suggestions on our draft. All experiments were performed using the Entropy cluster funded by NVIDIA, Intel, the Polish National Science Center grant UMO-2017/26/E/ST6/00622 and ERC Starting Grant TOTAL. The work of Henryk Michalewski was supported by the Polish National Science Center grant UMO-2018/29/B/ST6/02959.

REFERENCES

- Stephan Alaniz and Zeynep Akata. Explainable observer-classifier for explainable binary decisions. *arXiv preprint arXiv:1902.01780*, 2019.
- Aida Amini, Saadia Gabriel, Peter Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. *arXiv preprint arXiv:1905.13319*, 2019.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Rishabh Bhardwaj, Navonil Majumder, and Soujanya Poria. Investigating gender bias in bert. *arXiv preprint arXiv:2009.05021*, 2020.
- Daniel G Bobrow. Natural language input for a computer problem solving system. 1964.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in neural information processing systems*, 2020.
- Eugene Charniak. Computer solution of calculus word problems. In *Proceedings of the 1st international joint conference on Artificial intelligence*, pp. 303–316, 1969.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*, 2019.
- Rodney R Cocking, Rodney T Cocking, and Jose P Mestre. *Linguistic and cultural influences on learning mathematics*. Psychology Press, 1988.
- Mark G Core, H Chad Lane, Michael Van Lent, Dave Gomboc, Steve Solomon, and Milton Rosenberg. Building explainable artificial intelligence systems. 2006.
- Scott Crossley, Tiffany Barnes, Collin Lynch, and Danielle S McNamara. Linking language to math success in an on-line course. *International Educational Data Mining Society*, 2017.
- Wietse de Vries, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim. Bertje: A dutch bert model. *arXiv preprint arXiv:1912.09582*, 2019.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Hubert L Dreyfus, L Hubert, et al. *What computers still can’t do: A critique of artificial reason*. MIT press, 1992.
- John Rupert Firth. *Papers in Linguistics 1934-1951: Repr.* Oxford University Press, 1961.
- Lynn S Fuchs, Douglas Fuchs, Donald L Compton, Sarah R Powell, Pamela M Seethaler, Andrea M Capizzi, Christopher Schatschneider, and Jack M Fletcher. The cognitive correlates of third-grade skill in arithmetic, algorithmic computation, and arithmetic word problems. *Journal of Educational Psychology*, 98(1):29, 2006.
- Lynn S Fuchs, Douglas Fuchs, Karla Stuebing, Jack M Fletcher, Carol L Hamlett, and Warren Lambert. Problem solving and computational skill: Are they shared or distinct aspects of mathematical cognition? *Journal of educational psychology*, 100(1):30, 2008.
- Mor Geva, Ankit Gupta, and Jonathan Berant. Injecting numerical reasoning skills into language models. *arXiv preprint arXiv:2004.04487*, 2020.
- Caglar Gulcehre, Sarath Chandar, Kyunghyun Cho, and Yoshua Bengio. Dynamic neural turing machine with soft and hard addressing schemes. *arXiv preprint arXiv:1607.00036*, 2016.
- Hossein Hajipour, Mateusz Malinowski, and Mario Fritz. Ireen: Iterative reverse-engineering of black-box functions via neural program synthesis. *arXiv preprint arXiv:2006.10720*, 2020.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision (ECCV)*, pp. 630–645. Springer, 2016.
- Lisa Anne Hendricks, Zeynep Akata, Marcus Rohrbach, Jeff Donahue, Bernt Schiele, and Trevor Darrell. Generating visual explanations. In *European Conference on Computer Vision*, pp. 3–19. Springer, 2016.
- Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 771–787, 2018.
- Danqing Huang, Shuming Shi, Chin-Yew Lin, Jian Yin, and Wei-Ying Ma. How well do computers solve math word problems? large-scale dataset construction and evaluation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 887–896, 2016.
- W Lewis Johnson. Agents that learn to explain themselves.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. Measuring bias in contextualized word representations. *arXiv preprint arXiv:1906.07337*, 2019.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- H Chad Lane, Mark G Core, Michael Van Lent, Steve Solomon, and Dave Gomboc. Explainable artificial intelligence for training and tutoring. Technical report, UNIVERSITY OF SOUTHERN CALIFORNIA MARINA DEL REY CA INST FOR CREATIVE . . . , 2005.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. *arXiv preprint arXiv:1606.04155*, 2016.
- Percy Liang. Learning executable semantic parsers for natural language understanding. *Communications of the ACM*, 59(9):68–76, 2016.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*, 2017.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Meghann Lomas, Robert Chevalier, Ernest Vincent Cross, Robert Christopher Garrett, John Hoare, and Michael Kopack. Explaining robot actions. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, pp. 187–188, 2012.
- Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in neural information processing systems*, pp. 1682–1690, 2014.
- Mateusz Malinowski, Carl Doersch, Adam Santoro, and Peter Battaglia. Learning visual question answering by bootstrapping hard attention. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–20, 2018.

- Jose P Mestre. The role of language comprehension in mathematics and problem solving. In *Linguistic and cultural influences on learning mathematics*, pp. 201–220. Taylor and Francis, 2013.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.
- Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pp. 2204–2212, 2014.
- Richard J Murnane, John B Willett, M Jay Braatz, and Yves Duhaldeborde. Do different dimensions of male high school students’ skills predict labor market success a decade later? evidence from the nlsy. *Economics of Education Review*, 20(4):311–320, 2001.
- Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. Neuro-symbolic program synthesis. *arXiv preprint arXiv:1611.01855*, 2016.
- Maria Chiara Pasolunghi, Cesare Cornoldi, and Stephanie De Liberto. Working memory and intrusions of irrelevant information in a group of specific poor problem solvers. *Memory & Cognition*, 27(5):779–790, 1999.
- Markus N Rabe, Dennis Lee, Kshitij Bansal, and Christian Szegedy. Mathematical reasoning via self-supervised skip-tree training. *arXiv preprint arXiv:2006.04757*, 2020.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. Analysing mathematical reasoning abilities of neural models. *arXiv preprint arXiv:1904.01557*, 2019.
- Shuming Shi, Yuehui Wang, Chin-Yew Lin, Xiaojiang Liu, and Yong Rui. Automatically solving number word problems by semantic parsing and reasoning. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1132–1142, 2015.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Amber Y Wang, Lynn S Fuchs, and Douglas Fuchs. Cognitive and linguistic predictors of mathematical word problems with and without irrelevant information. *Learning and individual differences*, 52:79–87, 2016.
- Lei Wang, Dongxiang Zhang, Lianli Gao, Jingkuan Song, Long Guo, and Heng Tao Shen. Mathdqn: Solving arithmetic word problems via deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771, 2019.

Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *European Conference on Computer Vision*, pp. 451–466. Springer, 2016.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pp. 5753–5763, 2019.

Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 21–29, 2016.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pp. 19–27, 2015.

Yanyan Zou and Wei Lu. Text2math: End-to-end parsing text into math expressions. *arXiv preprint arXiv:1910.06571*, 2019.

A AQUA-RAT EXAMPLE

Example problem: *A starts a business with Rs.40,000. After 2 months, B joined him with Rs.60,000. C joined them after some more time with Rs.120,000. At the end of the year, out of a total profit of Rs.375,000, C gets Rs.150,000 as his share. How many months after B joined the business, did C join?*

Options: A) 30', B) 32', C) 35', D) 36', E) 40'

Rationale: Assume that C was there in the business for x months

$$A : B : C = 40000 * 12 : 60000 * 10 : 120000 * x$$

$$= 40 * 12 : 60 * 10 : 120x = 40 : 5 * 10 : 10x$$

$$= 8 : 10 : 2x$$

$$= 4 : 5 : x$$

$$C's\ share = 375000 * x / (9 + x) = 150000$$

$$\Rightarrow 375x / (9 + x) = 150$$

$$\Rightarrow 15x = 6(9 + x)$$

$$\Rightarrow 5x = 18 + 2x$$

$$\Rightarrow 3x = 18$$

$$\Rightarrow x = 18/3 = 6$$

It means C was there in the business for 6 months. Given that B joined the business after 2 months. Hence C joined after 4 months after B joined

Answer is B

Correct answer: B

AQuA-RAT question is a closed test math word problem with five possible answers given, and one of them is correct. Single data point consists of **Question**, **Options**, **Rationale** and **Correct answer**.

Question and **Options** are used for predicting **Correct answer** during test time. **Rationale** is a step by step derivation of the solution. Steps are separated by a newline token.

In the example above two subproblems are being solved: Finding a ratio between A, B, C shares (first four steps) and deriving C's share (steps from 6 to 10). ROP loss by swapping two rows with 50% probability forces model to be aware of the reasoning coherence. Moreover, this example shows why NROP is a more difficult task than ROP and requires more subtlety. If two rows from the same subproblem are swapped, the question of guessing whether they were swapped is more difficult than if it swapped rows from different subproblems.

Extended validation dataset We observe a highly sensitive relationship between dev and test sets (Figure 4, left), where small changes in the accuracies in the former set can lead to more dramatic changes at test time. Indeed, the correlation of results between both sets is only 0.082. As the validation set is quite small, we propose an extended dev consisting of 5000 randomly chosen samples from the training set extended by the whole dev set. Although not ideal, and the sensitive relationship

is still present, we have increased the correlation to 0.401. With such a new validation set, we report 37% test accuracy but we can also see that 40% is within the reach. Details are in the appendix in the (Figure 4, right).

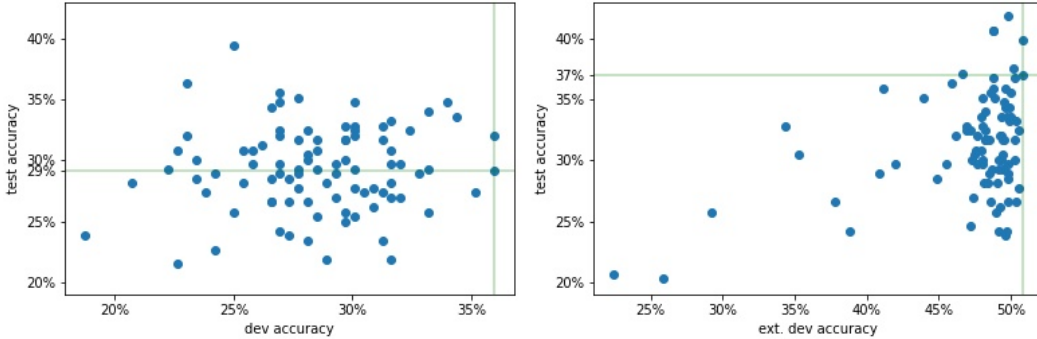


Figure 4: Accuracies for dev and test sets. Green lines show the iteration that maximizes validation accuracy. The image also shows the sensitivity of relationship between test and the original (left) or our extended (right) validation set.

B INPUT REPRESENTATION

All BERT variants use the representation that corresponds to a special token [CLS] that we put at the beginning of the whole input sequence consisting of question tokens followed by rationale tokens, and in the downstream, question-answering task, rationale tokens are replaced by the answer options. With that, the classification uses the contextual embedding of [CLS] that captures the entire input. MLM classifies over the entire vocabulary of possible words while the other two losses use a binary cross-entropy loss for the predictions.

C TRAINING PROTOCOL

We train all our architectures on AQuA-RAT using the following training phases. In all cases, we choose our best model based on the performance on the validation set (dev set), and report the final performance on the test set.

Pre-training. Each model is pre-trained on a large corpus of texts written in natural language sampled from English Wikipedia and BooksCorpus Devlin et al. (2018); Zhu et al. (2015). We use this as the base (BERT-base) model that is also used in all other variants of BERT. In practice, we initialize all the models with the weights using the HuggingFace library Wolf et al. (2019). Our model therefore has the same number of weights as BERT-base.

Self-supervision. Here, we use our newly introduced losses, where our models use questions and possibly rationales from the AQuA-RAT dataset. Both questions and rationales use the same word embeddings but to distinguish between both modalities we use different positional encodings and segment embeddings. That is, we share the latter among all the question tokens, and separately among all the rationale tokens. We use dynamic masking Liu et al. (2019). Here, tokens are randomly masked for each batch. We naturally extend this approach to other losses that we use in this phase. That is, ROP and NROP negative examples are randomly recreated every k epochs, where $k = 2$ in our case.

Fine-tuning is the last training phase. Here, once our models have learnt the representation during the self-supervised phase, we tune such a representation to the question-answering downstream task. In this task, our input consists of question tokens and possible answer options. There are five such options that comes with the dataset. Like other methods, we treat this as a five-class classification task where the classification head is added on top of the final embedding of the input. We consider the embedding corresponding to the first (from the left) [CLS] token as such the final representation.

D IMPLEMENTATION DETAILS

Thoroughly in our experiments, we use four TITAN V GPUs. We use a multi-gpu setup. In the pre-training phase, we use batch size equals to four for each GPU device. Therefore the effective batch size equals to sixteen. We use the learning rate $5 \cdot 10^{-5}$ and trained the models for 24 epochs. In the fine-tuning phase, we use early stopping criteria, based on the accuracy score on the validation set. We use the following criteria. If the model does not improve the performance in 15 consecutive epochs, we stop training, and evaluate a model that yields the highest validation performance. We use ADAM optimizer with learning rate 10^{-5} and gradient clipping that sets the maximal gradient’s norm to one. All our settings use the same hyper-parameters but they differ due to the random initialization of our self-supervised networks (during the self-supervised training phase) and the classification networks (during the fine-tuning phase). Self-supervision phase takes around 4 days on 4 gpus, whereas fine-tuning takes 8 hours on single gpu

E PERMUTATION INVARIANT METHODS

In the main paper, we have shown that typical models can use positional biases in achieving answers. This results in a low permutation consistency score (Table 3 in the main paper). To handle that issue, we have defined extra variants that do not use positional encodings for the answer options and instead they rely on the retrieval mechanics where input representations are matched against the candidate answers. Here, we describe two such variants.

E.1 ORIGINAL METHODS

Original models create an embedding of a sentence extended by possible questions. This embedding is then transformed by a linear layer to predict the correct answer. That is,

$$o_1 = f_1(\mathbf{BERT}(Q||P))$$

where o_1 is a 5-dimensional vector with probabilities for each possible answer, Q is a question, P are all possible answers, $||$ represents concatenation, f_1 is a single fully connected layer from 768-dimensional space to 5-dimensional space with the softmax activation. **BERT** is a BERT-base sentence embedding. The same approach is used for BERT+(N)ROP.

E.2 SEP-NC

In SEP-NC and SEP-C, we use separate embeddings for a question and **SE**parate embedding for a candidate answer. They differ, however, in the fact that SEP-C has access to all five possible answers, while SEP-NC has access only to one prompted candidate answer. Therefore NC stands for "no candidates", while C stands for "candidates".

We train the SEP-NC model on a binary classification task to predict whether each candidate answer C is correct. The method produces two embeddings, one for question and another one for a candidate answer $C \in P$, and next concatenates them. That is,

$$o_2 = f_2(\mathbf{BERT}(Q)||\mathbf{BERT}(C))$$

where o_2 is an estimated probability that C is a correct answer, P is the sequence of all possible answers, f_2 is a single fully connected layer from 1536 ($768 * 2$) dimensional space to 1-dimensional space with the sigmoid activation. Note that, all candidate answers are independent of the question. That is, BERT cannot use positional biases in deriving an answer. At test time, the model is prompted to score all five candidate answers and select the one with the highest score. We naturally extended that approach to BERT+ROP and BERT+NROP. Table 3 (the main paper) shows a significant improvement over the baseline method.

E.3 SEP-C

SEP-NC method could be too restrictive as it does not allow the model to compare against different answers. Therefore, we propose another approach that 1) alleviate the issue with positional biases, but 2) can compare between different answer options. We call that approach SEP-C.

Originally for each token, a positional encoding is assigned based on its position. In SEP-C, before assigning positional encoding, we artificially reset the position at the beginning of each possible answer. For example, if possible answers are: a)10, b)20, c)30, d)40, e)50 they are changed into 10; 20; 30; 40; 50 and after the tokenization, we get the following list of tokens: ['1', '0', ';', '2', '0', ';', '3', '0', ';', '4', '0', ';', '5', '0']. Modified positional encoding will assign value based only on the relative position to the beginning of the current possible answer. Therefore, in the example above, each '0' will receive the same positional encoding, and '1' will get the same positional encoding as '2', '3', and so on.

Formally, we have

$$o_3 = f_3(\text{BERT}(Q||P_m)||\text{BERT}(C))$$

where P_m is the sequence of all the possible answers but modified as explained above. Note that, in this formulation, the model can use the information for all the possible answer options, but their order is not taken into account. Table 3 (the main paper) shows a significant improvement over the baseline method.

F QUESTION DIFFICULTY

In this section we present an example from each difficulty group for BERT+NROP and BERT. We have described the grouping procedure in the main paper.

F.1 BERT+NROP

D₅: *How many ways A boy can reach the top of stairs which contain 10 steps, when he can take either one or two steps every time?*

Answers: A)88, B)89, C)90, D)91, E)92

Correct Answer: B

Model Answer: D

D₄: *A square piece of cloth is trimmed by 4 feet on one edge to form a rectangular piece, which is then cut diagonally in half to create two triangles. If the area of each of triangle is 70 square feet, what was the perimeter (in feet) of the original piece of square cloth?*

Options: A)56, B)58, C)60, D)62, E)64

Correct Answer: A

Model Answer: B

D₃: *Train A leaves a station every 16 minutes and Train B leaves every 17 minutes. If both trains just left the station simultaneously, how long until they do so again?*

Options: A)272 minutes, B)304 minutes, C)190 minutes, D)70 minutes, E)35 minutes

Correct Answer: A

Model Answer: B

D₂: *10kg of a mixture contains 30% sand and 70% clay. In order to make the mixture contain equal quantities of clay and sand how much of the mixture is to be removed and replaced with pure sand?*

Options: A)10/7, B)20/7, C)30/7, D)40/7, E)50/7

Correct Answer: B

Model Answer: C

D₁: *If one third of 3/4 of a number is 21. Then, find the number?*

Options: A)84, B)66, C)28, D)19, E)11

Correct Answer: D

Model Answer: D

F.2 BERT

D₅: *The length of the ribbon was originally 30 cm. It was reduced in the ratio 5 : 3. What is its length now?*

Answers: A)18, B)30, C)6, D)15, E)12

Correct Answer: A

Model Answer: B

D₄: *An electric pole, 14 metres high, casts a shadow of 10 metres. Find the height of a tree that casts a shadow*

of 15 metres under similar conditions.

Options: A)21, B)22, C)20, D)23, E)24

Correct Answer: A

Model Answer: C

D₃: A rope 20 meters long is cut into two pieces. If the length of one piece of rope is 3 meters shorter than the length of the other, what is the length, in meters, of the longer piece of rope?

Options: A)7.5, B)8.9, C)9.9, D)11.5, E)11.7

Correct Answer: D

Model Answer: B

D₂: Jerry purchased a 1-year \$5,000 bond that paid an annual interest rate of 12% compounded every six months. How much interest had this bond accrued at maturity?

Options: A)\$5102, B)\$618, C)\$216, D)\$202, E)\$200

Correct Answer: B

Model Answer: A

D₁: I have a money pouch containing Rs. 700. There are equal number of 25 paise coins, 50 paise coins and one rupee coins. How many of each are there?

Options: A)453, B)651, C)400, D)487, E)286

Correct Answer: C

Model Answer: C

F.3 HUMAN STUDY

We carried an initial human study on the group of 16 volunteers from University of Warsaw. Volunteers were Mathematics and Informatics students from the Faculty of Mathematics, Informatics and Mechanics. We asked the participants to solve questions sampled from the AQuA-RAT dataset. We are interested in the relation between BERTs difficulty, BERT+NRDP difficulty and human difficulty. Therefore to have a full image we would like to have 2 questions for each question difficulty pair, for example (**D₁** BERT, **D₂**: BERT+NRDP) . However, that would give 25 combinations and 50 questions if we wanted to have 2 questions per combination. That would be too much to ask from a volunteer participant. In order to reduce the number of questions, we group our 5 difficulty groups into 3 categories as follows.

- Easy: **D₁**
- Medium: **D₂** and **D₃** combined
- Hard: **D₄** and **D₅** combined

Because of that we have only 9 possible combinations and by sampling 2 questions from each combination we still have a feasible number of questions (18).

Apart from solving the question, we asked to rate question difficulty on a scale from 1 (the simplest) to 10 (the most challenging). In general, our participants were knowledgeable in math and solved all the questions correctly. With that grouping we now

The average human-rated difficulty for each of 9 combinations is presented in Figure 5. The results show that the progression of human difficulty is correlated with the difficulty judged by the models. Additionally, the human difficulty seems to be more sensitive to BERT+NRDP difficulty than to BERTs. In other words, increasing the difficulty of BERT+NRDP will increase the human difficulty more than the increasing difficulty of BERT. This observation fits our previous observations that BERT+NRDP solves the most straightforward questions while BERT is looking for some leaks, like looking for the roundest answer.

G DISTRIBUTION OF ANSWERS

Table 4 shows the distribution of the answers in the AQuA-RAT Ling et al. (2017) dataset in all the folds. Imbalance in distributions could potentially be used by models to find easy, shortcut solutions. For instance, a constant classifier that always choose the first answer (A) gets about 24% test accuracy.

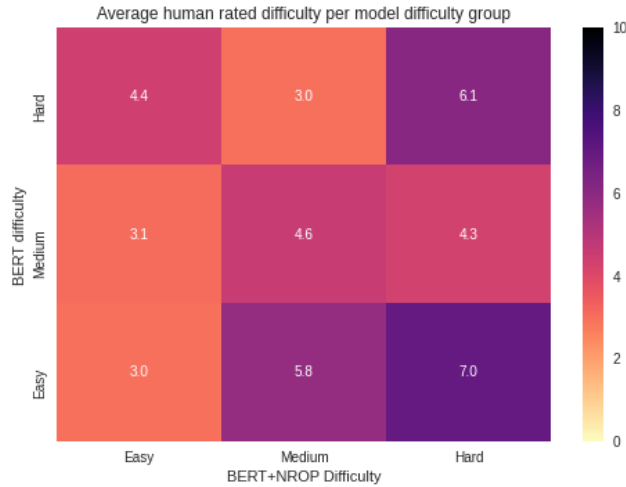


Figure 5: The average human-judged difficulty for questions from each model difficulty group.

dataset	A	B	C	D	E
train	21.03%	22%	22.87%	19.95%	14.15%
dev	27.17%	25.98%	16.93%	19.69%	10.24%
test	24.80%	22.83%	20.87%	18.11%	13.38%

Table 4: Answer distribution in each dataset.

H RELATED WORK

We are inspired by the following research.

BERTology. We use BERT Devlin et al. (2018) as our core. It uses Transformers Vaswani et al. (2017); powerful neural architectures that applies a trainable function to all the pairs of input embeddings. It also uses masking that covers a fraction of the input words and requires the network to predict the hidden words based on the context. With both ingredients, the meaning (representation) of a word emerges from the “company it keeps” Firth (1961). In practice, often, such representations are pre-trained on large textual corpora with no need for annotations, and next fine-tuned on the downstream tasks. BERT’s strong performance has resulted in the Cambrian explosion of studies of the inner working mechanisms and various modifications Clark et al. (2019); de Vries et al. (2019); Lan et al. (2019); Liu et al. (2019); Sanh et al. (2019); Radford et al.; Raffel et al. (2019); Yang et al. (2019). Finally, our Reasoning Order Prediction (ROP) is inspired by Sentence Order Prediction (SOP) Lan et al. (2019). However, ROP works with multiple rationale sentences, where by changing the order we force the network to understand the consecutive “reasoning” steps. We have also further extended ROP to a more difficult Neighbor Reasoning Order Prediction (NROP).

Language and math. Development psychologists Cocking et al. (1988); Mestre (2013) often argue for the necessity of learning languages and point out that those with limited language skills are in danger of under-performing at school. Moreover, it is also believed that language studies involve discipline in learning and manipulating formal structures, and thus may promote the development of the organization of thoughts also required in mathematical reasoning. The similarity between linguistic competence and mathematics is especially pronounced when solving math word problems Fuchs et al. (2006; 2008); Wang et al. (2016). Interestingly, attention appears to be crucial in problem solving Fuchs et al. (2006); Pasolunghi et al. (1999). Crossley et al. (2017) show that language skills are correlated with the performance in mathematical tests also among the university students. In particular, they pointed out that ability to use complex syntactic structures and cohesion devices are linked to better scores in a blended discrete mathematics course. We take inspiration from all such studies and decide to build our mathematical model based on language models.

Math word problems. Solving math word problems is a significant component of the mathematics curriculum and is taught very early, thoroughly, and universally. Such the emphasize is often motivated by that solving them is among the best predictors of employability, and is considered as a distinct area of mathematical competence Murnane et al. (2001); Wang et al. (2016). Since

solving such problems is unique to human intelligence, math word problems are also interesting for the AI community. This results in various approaches, more traditional symbolic methods, neural networks, and neuro-symbolic methods. Bobrow (1964); Charniak (1969); Shi et al. (2015); Ling et al. (2017); Amini et al. (2019); Parisotto et al. (2016); Wang et al. (2018); Zou & Lu (2019) as well as datasets Ling et al. (2017); Amini et al. (2019); Huang et al. (2016); Saxton et al. (2019). An interesting approach is proposed in Rabe et al. (2020), in which authors use self-supervised tasks on parsing trees of formal expressions. This approach requires syntax trees, and hence we would have to use an external parser. As our goal was to make an end-to-end model, we did not experiment with it, but there are no obstacles against using it in symbiosis with our methods. Geva et al. (2020) also proposes self-supervised training for improving mathematical abilities in language models. We, however, focused on a data-driven approach to exclude choice biases and therefore restricted ourselves from using generated data.

Rationales. In human communication, we always expect there is some rationale behind each decision. Hence, we set the same expectations to our artificial agents. Symbolic or semi-symbolic architectures naturally produce justifications as a sequence of formulas in some formal language Lane et al. (2005); Core et al. (2006); Lomas et al. (2012); Johnson; Liang (2016); Malinowski & Fritz (2014). Ideally, such rationales would also be shared and communicated to us through some language. The latter approach is especially appealing when applied to black-box neural networks. For instance, Hendricks et al. (2016) propose a system that classifies the input image as well as it produces a textual explanation on “why this class is suitable for the given image”.

Systems that produce explanations either in the form of the language Ling et al. (2017); Hendricks et al. (2016), attention Bahdanau et al. (2014); Mnih et al. (2014); Gulcehre et al. (2016); Malinowski et al. (2018); Xu & Saenko (2016); Yang et al. (2016), phrase selection Lei et al. (2016), distillation into programs Hajipour et al. (2020), or decision trees Alaniz & Akata (2019) can potentially increase the transparency of the black-box neural networks. However, most of these approaches create rationales posthoc where the justification is conditioned on answers or by querying the network. In our work, we use rationales to learn a finer representation that can potentially lead to better decisions. In this sense, our technique is conceptually closer to methods that derive answers based on the program and use rationales paired with questions to guide the program induction process Ling et al. (2017).