

SemEval-2020 Task 4: Commonsense Validation and Explanation

Cunxiang Wang^{1,2*}, Shuailong Liang^{3*}, Yili Jin⁴, Yilong Wang¹,
Xiaodan Zhu⁵ and Yue Zhang²

¹Zhejiang University, China; ²School of Engineering, Westlake University, China

³Singapore University of Technology and Design, Singapore;

⁴Sun Yat-Sen University, China; ⁵Queen’s University, Canada

wangcunxiang, zhangyue@westlake.edu.cn,
shuailong_liang@mymail.sutd.edu.sg, zhu2048@gmail.com

Abstract

In this paper, we present SemEval-2020 Task 4, **Commonsense Validation and Explanation (ComVE)**, which includes three subtasks, aiming to evaluate whether a system can distinguish a natural language statement that *makes sense* to humans from one that does not, and provide the reasons. Specifically, in our first subtask, the participating systems are required to choose from two natural language statements of similar wording the one that *makes sense* and the one does not. The second subtask additionally asks a system to select the key reason from three options why a given statement does not make sense. In the third subtask, a participating system needs to generate the reason. We finally attracted 39 teams participating at least one of the three subtasks. For Subtask A and Subtask B, the performances of top-ranked systems are close to that of humans. However, for Subtask C, there is still a relatively large gap between systems and human performance. The dataset used in our task can be found at <https://github.com/wangcunxiang/SemEval2020-Task4-Commonsense-Validation-and-Explanation>; The leaderboard can be found at <https://competitions.codalab.org/competitions/21080#results>.

1 Introduction

In the past decades, computer’ ability in processing natural language has significantly improved. However, its intelligence for understanding common sense expressed in language is still limited. For example, it is straightforward for humans to judge that the following sentence is plausible, or makes sense: “*John put a turkey into a fridge*” while “*John put an elephant into the fridge*” does not, but it is non-trivial for a computer to tell the difference. Arguably, commonsense reasoning plays a central role in a natural language understanding system (Davis, 2017). It is essential to gauge how well computers can understand whether a given statement *makes sense*. In our task, we take an operational definition of *making sense* by asking human subjects to generate natural language statements that obey or violate their commonsense knowledge about the world.¹

Many existing tasks embed the evaluation of commonsense understanding in other tasks such as coreference resolution (Levesque et al., 2012; Morgenstern and Ortiz, 2015), subsequent event prediction (Roemmele et al., 2011), ordinal common-sense inference (Zhang et al., 2017), situations with adversarial generations (Zellers et al., 2018), event validation (Wang et al., 2018), reading comprehension (Mostafazadeh et al., 2016; Ostermann et al., 2018b; Ostermann et al., 2018a), dialogue (Cui et al., 2020) and QA (Davis, 2016; Talmor et al., 2018; Mihaylov et al., 2018). They verify whether a system is equipped with common sense by testing whether the system can give a correct answer when the input does not contain such knowledge. The above tasks do not directly evaluate commonsense validation and they do not explicitly identify the key factor required in a commonsense validation process.

The SemEval-2020 Task 4 includes three subtasks on testing whether a system can distinguish natural language statements that make sense from those that do not, and probe the reasons. In the first subtask, a

* Equal Contribution

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

¹Note that the term “makes sense” may also be used to describe if a statement is meaningful, regardless of whether it is plausible or conforms to common sense. In this task the term refers to statements that obey common sense.

system needs to choose the against-common-sense statement from two natural language statements of similar wordings, e.g., “*John put an elephant into the fridge*” and “*John put a turkey into the fridge*”, respectively. The second task aims to find the key reason from three provided options why a given nonsensical statement does not make sense. For example, for the nonsensical statement, “*John put an elephant into the fridge*”, the three options are “*An elephant is much bigger than a fridge*”, “*Elephants are usually white while fridges are usually white*”, and “*An elephant cannot eat a fridge.*” A system needs to identify the correct reason. In addition, the third task requires the participating systems to generate the reason automatically. We hope that the task and datasets can facilitate studies on commonsense validation, its interpretability, and the related natural language understanding and generation problems.

There are 39 teams submitting valid systems to at least one subtask. In Subtask A and Subtask B, top-performing systems achieve performances closed to that of human subjects. However, for Subtask C, there is still a relatively large between system and human performances.

2 Task Setup

2.1 Task Definition

Formally, each instance in our dataset is composed of eight sentences: $\{s_1, s_2, o_1, o_2, o_3, r_1, r_2, r_3\}$. s_1 and s_2 are two similar statements that differ by only a few words; one of them makes sense (i.e., conforms to common sense) while the other does not. They are used in our Subtask A: the **Validation** subtask, which requires a model to identify which one makes sense. For the statement that does not make sense, we have three candidate reasons, i.e., three options o_1 , o_2 , and o_3 ; one of them explains why the statement does not make sense. So, in our Subtask B, the **Explanation (Multi-Choice)** subtask, a model is required to find the correct reason from the three options. For the same nonsensical statement, in Subtask C, the **Explanation (Generation)** subtask, a participating system needs to generate the reason why it does not make sense. Three references, r_1 , r_2 , and r_3 , are used for evaluating Subtask C. Below we give an example for each subtask, in which we introduce some notations we will use in the paper.

- **Subtask A: Validation**

Task: *Select the statement of the two that does not make sense.*

s_1 : *John put a turkey into a fridge.*

s_2 : *John put an elephant into the fridge.*

In this example, s_1 is a sensical statement, also denoted as s_c , while s_2 is the nonsensical statement, which is also denoted as s_n .

- **Subtask B: Explanation (Multi-Choice)**

Task: *Select the best reason that explains why the given statement does not make sense.*

Nonsensical statement (s_n): *John put an elephant into the fridge.*

o_1 : *An elephant is much bigger than a fridge.*

o_2 : *Elephants are usually white while fridges are usually white.*

o_3 : *An elephant cannot eat a fridge.*

In this example, the option o_1 is the correct reason, which is also denoted also as o_c , while o_2 and o_3 are not the reason, which are also denoted as o_{n1} and o_{n2} .

- **Subtask C: Explanation (Generation)**

Task: *Generate the reason why this statement does not make sense.*

Nonsensical statement (s_n): *John put an elephant into the fridge.*

Reference reasons (used for calculating the BLEU score):

r_1 : *An elephant is much bigger than a fridge.*

r_2 : *A fridge is much smaller than an elephant.*

r_3 : *Most of the fridges aren't large enough to contain an elephant.*

Score	Description
0	The reason is not grammatically correct, or not comprehensible at all, or not related to the statement at all.
1	The reason is just the negation of the statement or a simple paraphrase. Obviously, a better explanation can be made.
2	The reason is relevant and appropriate, though it may contain a few grammatical errors or unnecessary parts. Or like case 1, but it’s hard to write a proper reason.
3	The reason is appropriate and is a solid explanation of why the statement does not make sense.

Table 1: Rubrics used in human evaluation in Subtask C.

2.2 Evaluation Metrics

The Subtasks A and B are evaluated using accuracy. Subtask C is evaluated with the BLEU score (Papineni et al., 2002). In addition, for Subtask C, we further perform human evaluation. We randomly select 100 instances from the test set and evaluate system outputs on Amazon Mechanical Turk. We ask three different crowd-sourcing workers to score each generated reason with a scale ranging from 0 to 3, inclusively, according the rubrics listed in Table 1.

Then we calculate the average score of the three scores as our final human evaluation score. Formally, the human evaluation score of system k is

$$score_k = \frac{\sum_{i=1}^{100} \sum_{j=1}^3 score_{ijk}}{100 * 3}, \quad (1)$$

where $score_{ijk}$ means the score from the j_{th} annotator for system k on the i_{th} instance.

3 Data Construction

Our data construction is mainly performed on Amazon Mechanical Turk, which consists of two steps:

- Step 1: In this step, we construct datasets for Subtask A and Subtask B. Specifically, we ask a crowd-sourcing worker to write a sensical statement s_c and a nonsensical statement s_n . For the nonsensical statement s_n , the worker further writes three sentences, o_1, o_2, o_3 ; one of them, denoted as o_c , explains why the nonsensical statement does not make sense; two of them, denoted as o_{n1} and o_{n2} , serve as the confusing choices. (Refer to Section 3.1 for details.)
- Step 2: We then make three reference reasons, r_1, r_2, r_3 for Subtask C. We use o_c as one of three references, and collect two more references in this step. We ask two different crowd-sourcing workers to write each of them. Note that instead of letting the same worker in step 1 to write these two references, we asked two more workers. The reason is to encourage diversity of the reference. (Refer to Section 3.2 for details.)

Finally, each instance of the dataset have 8 sentences: $\{s_1, s_2, o_1, o_2, o_3, r_1, r_2, r_3\}$. Note that one sentence in o_1, o_2, o_3 is repeated in r_1, r_2, r_3 , but for convenience of description, we denote it differently.

3.1 Step 1: Collecting Data for Subtask A and B

Annotation Guidelines. When writing instances, workers were asked to follow several principles: (1) Try to avoid complex knowledge and focus on daily common sense. Make the questions as understandable as possible, so that a literate person is able to give the right answers. (2) The confusing reason options, o_{n1} and o_{n2} , should better contain more content words or information such as entities and activities in the nonsensical statements s_n . For example, the confusing reasons of “*John put an elephant into the fridge*” should better contain both “*elephant*” and “*fridge*”. (3) The confusing reasons, o_{n1} and o_{n2} , should be related to the statements s_n and the correct reason o_c and not deviate from the context; otherwise it may be easily captured by pretrained models like BERT (Talmor et al., 2018). (4) The three option reasons, $o_1, o_2,$

Type of Sentences	Training Set	Dev Set	Test Set
Sensical Statements	7.67	7.12	7.25
Nonsensical Statements	7.69	7.16	7.36
Correct Reasons	8.13	7.96	8.09
Confusing Reasons	7.80	7.14	7.29
Referential Reasons	8.08	7.92	8.06

Table 2: Average length of different types of sentences of Training/Dev/Test set

and o_3 should only be related to the incorrect statements s_n rather than the correct statements s_c , because we want further studies to be able to estimate nonsensical statements s_n without the correct statement s_c . (5) The confusing reasons, o_{n1} and o_{n2} , should make sense themselves. Otherwise, the models may simply ignore the incorrect options o_{n1} , o_{n2} without considering the casual semantics. This concern is raised from and motivated by the fact that models can achieve high performance in the ROC Story Cloze Task, when only looking at the alternative endings and ignoring the story content (Schwartz et al., 2017). (6) We ask the annotators to make the nonsensical statement s_n contain about the same number of words as the sensical statement s_c , and the correct reason o_c have similar length with other two options. We drop the instances which do not meet such requirements.

Use of Inspirational Materials. It is not easy for all crowd-sourcing workers to write instances from scratch. To address this issue, we also provide them with external reading materials to stimulate inspiration, such as the sentences of the Open Mind Common Sense (OMCS) project (Havasi et al., 2010). For example, “*he was sent to a (restaurant)/(hospital) for treatment after a car crash*” can be inspired by the two sentences “*restaurants provide food*” and “*hospitals provide medical care*”.

Quality Control. To ensure the quality of the data, we manually check the instances and drop or request a rewriting of the low-quality ones. If one worker writes too many low-quality instances, we will remove her or him from our annotator pool. With such process, we finally accept around 30% submitted instances.

3.2 Step 2: Collecting Data for Subtask C

Annotation Guidelines. To collect data for Subtask C, each worker is given a nonsensical statement s_n and a sensical statement s_c and asked to write a reason to explain why the nonsensical statement s_n does not make sense. They shall follow the following rules: (1) Do not explain why the sensical statement s_c makes sense. (2) Avoid mentioning the sensical statement s_c . (3) Write the reason, rather than simply add the word “*not*” or “*can’t*” to the nonsensical statement s_n to form an explanation. (4) Write the reason, don’t use patterns like “*XXX is not for YYY*” to create an explanation. (5) Do not try to justify why the nonsensical statement s_n makes sense. (6) Write only one sentence, do not be overly formal. (7) Refrain from using “*because*” at the beginning of a sentence. (8) Do not try to correct the statement s_n , but just give the reason.

Quality Control. As the same as in Step 1, after the annotators write the reasons in Step 2, the first two authors of the paper perform the check process again. We reject low-quality reasons (that violate the rules significantly) and low-quality annotators (who write many low-quality reasons with the number above a threshold).

3.3 Data Summary and Analysis

For SemEval-2020, we created 11,997 instances (i.e., 11,997 8-sentence tuples). We further split the instances into three subsets with 10,000 (the training set), 997 (the development set), and 1,000 (the test set) instances, respectively. We randomly assign the label of the correct options in subtask A and B to avoid unbalanced correct labels. We conduct three more data analysis experiments to evaluate data quality, including sentence length, common words and repetition.

Average Length. In Table 2, we present the average length of each type of sentence in the training/dev/test set. The sentences in the development and test set have shorter lengths than those in the training set. This is because we check the development and test more carefully and more strictly, thus removing longer and more incomprehensible instances, which lowers the average lengths of the dev/test set. The

Types of Sentences	Word:Word Frequency(‰)				
Sensical Statements	can:1.54	his:1.45	with:1.285	my:1.247	people:0.912
Nonsensical Statements	can:1.604	his:1.411	with:1.299	my:1.254	people:0.873
Correct Reasons	not:4.545	cannot:1.731	people:1.579	can:1.389	no:1.337
Confusing Reasons	can:2.095	people:1.671	not:1.49	have:1.152	than:0.81
Referential Reasons	not:5.711	cannot:1.65	can:1.09	people:1.088	have:0.978
(a) Training set					
Types of Sentences	Word:Word Frequency(‰)				
Sensical Statements	can:1.343	my:1.322	his:1.291	put:1.259	with:0.978
Nonsensical Statements	can:1.458	my:1.335	put:1.304	his:1.15	with:0.975
Correct Reasons	not:2.325	can:1.709	no:1.614	people:1.362	than:1.31
Confusing Reasons	can:2.332	people:1.619	not:1.609	have:1.094	than:1.079
Referential Reasons	not:4.264	cannot:1.212	can:1.203	it:1.179	have:1.093
(b) Dev+Test set					

Table 3: Top-5 common words and their frequencies in different types of sentences in the training and dev+test set. 1.000‰ means this word appear once in every 1000 words.

sensical statements and nonsensical statements almost have the same average lengths in the three sets (the differences are equal or smaller than 1%), which is balanced. However, there is an obvious gap between the correct reasons and confusing reasons in terms of the average lengths (roughly 4% in the training set and 10% in the dev/test set).

Common Word Analysis. The most common words are important for showing the differences between sentences. We only present those words which have obvious different frequencies between sensical statements and nonsensical statements or between correct/referential reasons and confusing reasons. So, we skip most uninformative words, including ‘a’, ‘an’, ‘the’, ‘to’, ‘in’, ‘on’, ‘of’, ‘for’, ‘and’, ‘is’, ‘are’ and ‘be’. After removing those words, we can list the top-5 common words in each type of sentence in the training/dev+test sets. For sensical statements s_c and nonsensical statements s_n , there are no significant differences between the training, dev, and test set. However, there is an obvious gap in the correct reasons o_c and confusing reasons o_n in negative words such as “not”, “no”, and “cannot”. In the training data, negative words are about 3 times more common in the correct option o_c than in the confusing options o_n . In the dev+test data, the gap is about 40%, which indicates that the dev+test data has a higher quality than the training data. However, as discussed in (Niven and Kao, 2019), *spurious statistical cues* can affect BERT’s results. We conjecture that the negative words are also spurious effective clues, which make the Subtask B potentially easier.

Repetition. The dev+test set have 12 instances (0.6%) that repeat the same nonsensical statements in the training data and 36 instances (1.8%) that repeat the same correct reasons with the training data.

3.4 Cautions of using the data

The following advice is given to all task participants and future users: (1) Feel free to use whatever additional data they deem appropriate for the tasks to train their model. (2) Do not use the input of Subtask B/C to help Subtask A and do not use the option o of Subtask B to help Subtask C. Otherwise the task will be artificially easy. This is because of two reasons: a) The nonsensical statements s_n of Subtask B and Subtask C is exactly the nonsensical statements s_c of Subtask A and, participants can use the input of the Subtask B/C to directly obtain the answer of Subtask A and the option answers o of Subtask B will also reduce the difficulty of Subtask A; b) the correct reason o_c of Subtask B is also one of the reference reason o_c in Subtask C.

4 Systems and Results

In this section, we show the evaluation results of all the submitted systems for the three subtasks. Since most systems share similar model architecture for subtasks A and B, we discuss the two subtasks together.

4.1 Subtask A and Subtask B

The formal evaluation results of Subtask A and B are shown in Table 4 and 5. There are in total 39 valid submissions for Subtask A and 27 valid submissions for Subtask B. Most top-performing submissions

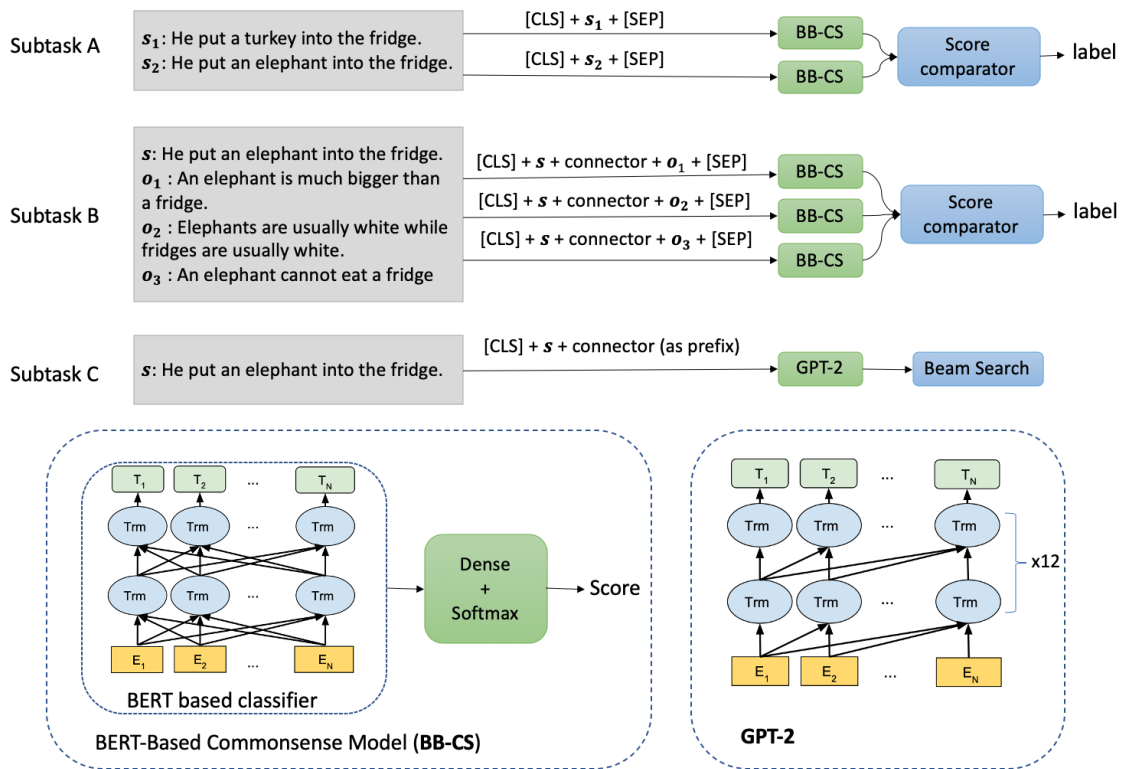


Figure 1: The most commonly used model architectures used in the three subtasks. This figure is mostly based on Team Solomon’s system. For Subtask B and C, the connector can be simply “No, ”, to help in constraining the model to learn a choice that explains the unreasonability of the statement. For Subtask A and B, the pretrained models are finetuned on the task-specific data with MLM-objective, and then trained as a binary classification task to score each input. For Subtask C, the cross-entropy loss of next-token-prediction is used to train the model, and beam search is used at inference.

adopted the pretrained language models such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019c), XLNET (Yang et al., 2019) and ALBERT (Lan et al., 2019) as the encoder of the model, and then finetune on the training set of the task. See Figure 1 for the most commonly-used model architectures for Subtask A and B. Also, the top-performing systems take advantage of external knowledge graphs such as ConceptNet (Speer et al., 2017), or unstructured text containing commonsense knowledge. Below we introduce in detail several top-performing systems and their main features.

- **CN-HIT-IT.NLP** (Zhang et al., 2020) ranks top in Subtask A. They use a variant of K-BERT (Liu et al., 2019a) as the encoder to enhance language representations through knowledge graphs. K-BERT is a Transformer-based model, which enhances the language representations of the text by injecting relevant triples from a knowledge graph to form a knowledge-rich sentence tree, and then uses a mask-Transformer to make the triples visible only to the corresponding entity. They use ConceptNet as the commonsense repository to extract the triples for the statements.
- **ECNU-SenseMaker** (Zhao et al., 2020) ranks top in Subtask B. They use Knowledge-enhanced Graph Attention Network to leverage heterogeneous knowledge from both the structured knowledge base (i.e. ConceptNet) and the unstructured text to better improve the commonsense understanding. Like **CN-HIT-IT.NLP**, their model is also based on K-BERT. In addition, they use unstructured text from ConceptNet and Subtask C to pretrain the language model.
- **IIE-NLP-NUT** (Xing et al., 2020) uses RoBERTa as the encoder, and conduct a second pretraining on the original RoBERTa model with the textual corpus from Open Mind Common Sense (Singh et

Team	Acc.	Rank	Team	Acc.	Rank	Team	Acc.	Rank
Human	99.1	-						
CN-HIT-IT.NLP	97.0	1	panaali*	92.5	14	Lijunyi	83.0	27
ECNU-SenseMaker	96.7	2	ZhengxianFan*	92.4	15	ehsantaher*	82.5	28
IIE-NLP-NUT	96.4	3	LMVE	90.4	16	TakeLab*	81.2	29
nlpX*	96.4	3	Warren*	90.4	16	Vicki*	79.8	30
Solomon	96.0	5	TMLab*	89.2	18	TR	79.7	31
Qiaoning	95.9	6	UAICS	89.1	19	KDE SenseForce	79.6	32
BUT-FIT	95.8	7	JUST	89.1	19	Hitachi*	78.4	33
olenet*	95.5	8	eggy*	89.0	21	CUHK	72.4	34
KaLM	95.3	9	UI	88.2	22	paramitamirza*	69.2	35
CS-NET	94.8	10	Armins*	87.1	23	UoR	67.6	36
fkerem*	94.4	11	DEEPYANG	85.1	24	chengguang*	62.3	37
JUSTers	92.9	12	WUY*	84.2	25	praveenjoshi007*	55.9	38
CS-NLP	92.7	13	YNU-oxz	83.6	26	dania*	21.6	39

Table 4: Subtask A results of all the submitted systems. Those marked with * did not submit system description paper. Human performance are based on the trial data instead of the test data used by the participating systems.

Team	Acc.	Rank	Team	Acc.	Rank	Team	Acc.	Rank
Human	97.8	-						
ECNU-SenseMaker	95.0	1	JBNU	91.4	10	Masked Reasoner	73.5	19
CN-HIT-IT.NLP	94.8	2	Qiaoning	90.8	11	KDE SenseForce	72.8	20
IIE-NLP-NUT	94.3	3	CS-NET	89.0	12	SSN-NLP	68.3	21
Solomon	94.0	4	WUY*	85.3	13	TakeLab*	66.8	22
LMVE	93.8	5	SWAGex	84.6	14	UoR	65.9	23
CS-NLP	93.7	6	TMLab*	82.0	15	dania*	55.5	24
KaLM	93.2	7	UI	80.5	16	CUHK	51.2	25
BUT-FIT	93.1	8	ehsantaher*	79.3	17	bhu*	36.4	26
JUSTers	92.3	9	uzh*	75.8	18	praveenjoshi007*	32.6	27

Table 5: Subtask B results of all the submitted systems. Those marked with * did not submit system description paper. Human performance are based on the trial data instead of the test data used by the participating systems.

al., 2002). They also explore several prompt templates to constructs as the inputs to the model.

- **Solomon** (Srivastava et al., 2020), **KaLM** (Wan and Huang, 2020), **CS-NET** (Dash et al., 2020), **JUSTers** (Fadel et al., 2020), **CS-NLP** (Saeedi et al., 2020), **UI** (Doxolodeo and Mahendra, 2020), **TR** (Teo, 2020) **UoR** (Markchom et al., 2020), **Masked Reasoner** (Lu, 2020) have similar model architecture, with RoBERTa as the encoder. In addition, **UoR** finetunes the pretrained language model on NLI and STS dataset, and **UI** finetunes on MNLI data. **TR** combines RoBERTa features with additional features from text-to-image generation using Gradient Boosted Decision Tree, and give better results in post-evaluation.
- **Qiaoning** (Liu, 2020) and **JUST** (Mohammed and Abdullah, 2020) use several ensembles of BERT, ALBERT, XLNet and RoBERTa.
- **BUT-FIT** (Jon et al., 2020), **LMVE** (Liu et al., 2020), **Lijunyi** (Li et al., 2020) use ALBERT as the encoder. **BUT-FIT** uses back-translation from Czech for data augmentation, and **LMVE** uses hint sentences, back-translation from French and intra-subtask transfer learning between Subtasks A and B to enhance their system.
- **UAICS** (Cusmulic et al., 2020), **DEEPYANG** (Bai and Zhou, 2020), **YNU-oxz** (Ou et al., 2020), **KDE-SenseForce** (Mendbayar and Aono, 2020), **CUHK** (Wang et al., 2020) **JBNU** (Na and Lee, 2020), **SWAGex** (Rim and Okazaki, 2020) are BERT-based. **JBNU** puts an BiLSTM on top of BERT, and **SWAGex** finetunes BERT with SWAG data. **CUHK** uses a Multitask Learning framework MTDNN (Liu et al., 2019b), adopting the “Explain, Reason and Predict” system.

Team	BLEU	Rank	Human	Rank	Team	BLEU	Rank	Human	Rank
Human	-	-	2.58	-					
BUT-FIT	22.4	1	1.84	4	CN-HIT-IT.NLP+	9.7	10	1.74	10
Solomon	19.3	2	1.84	4	SWAGex	7.1	11	1.75	8
KaLM	18.5	3	2.08	2	UI	5.5	12	0.73	14
panaali*	17.2	4	1.22	11	TMLab*	5.4	13	1.05	12
JUSTers	16.1	5	1.94	3	CUHK	4.3	14	0.58	16
cdjhz*	16.0	6	1.75	8	SSN-NLP	2.2	15	0.59	15
JBNU	15.9	7	1.80	6	UoR+	0.9	16	0.53	17
ANA	15.7	8	2.10	1	Masked Reasoner+	0.6	17	0.81	13
LMVE+	12.9	9	1.78	7					

Table 6: Subtask C results of all the submitted systems. Those marked with * did not submit a system description paper, and those marked with + means they do not include Subtask C in their system description paper.

It can be seen from the results that pretrained language models such as RoBERTa can achieve rather high performance, e.g., the team Solomon achieves 96.0% and 94.0% on Subtask A and Subtask B, respectively, without using further resources. This shows that large-scale pretrained language models do contain commonsense knowledge to deal with the Subtask A and the Subtask B in this challenge. Additionally finetuning the pretrained language models on commonsense-related text such as OMCS, which we use as inspirational materials, can push the results even higher, close to human performance. The best-performing teams on Subtask A and Subtask B both adopt K-BERT, which incorporates the external knowledge base (i.e. ConceptNet) to complement the pretrained language models with knowledge triples. This shows that knowledge-graph-enhanced approaches, such as K-BERT can effectively incorporate external knowledge. However, the high number may also indicate data leaking to some extent, since in the data creation stage, both ConceptNet and OMCS are used as references for the annotator to write the data instances.

4.2 Subtask C

The results for Subtask C are shown in Table 6. There are in total 17 valid submissions for Subtask C. There are generally two approaches: (1) sequence-to-sequence approach, where the source side is the non-sensical statement, and the reason is the target sequence. (2) language model generation approach, which uses large-scale pretrained auto-regressive language models such as GPT-2 (Radford et al., 2019) for reason generation, where the non-sensical sentence acts as prompt. An example of the language model generation approach is shown in Figure 1, which is most commonly used and achieves relatively good results. Below we describe in detail the systems and their main features.

- **BUT-FIT** (Jon et al., 2020) experiments with both the sequence-to-sequence approach and the language generation approach. For the sequence-to-sequence approach, they use BART (Lewis et al., 2019) with beam-search decoding to achieves the highest BLEU among all the teams. For the language generation approach, the nonsensical statement is used as a prompt. At the training stage, the statement and the explanation are concatenated together, and a GPT-2 is trained on these sequences with a next token prediction objective. At the test time, based on the statement, the model generates the reason tokens until the end-of-sentence token is generated.
- **KaLM** (Wan and Huang, 2020) uses the sequence-to-sequence architecture BART. To enhance the source side statement, they extract keywords from the statement and search for evidence from Wiktionary.² After that, they concatenate the evidence along with the original statement as the source sentence for the generation. This approach proves effective and makes their system second-best for human evaluations.
- **ANA** (Konar et al., 2020) has the highest human evaluation score with a multitask learning framework. Specifically, they use a decoder-only transformer based on GPT-2 as the backbone model, and train

²Wiktionary version: enwiktionary-20200220

the model with two self-attention heads: one for language models and another for classification. They then use data from both task B and task C to calculate language model loss and classification loss. Furthermore, they use OMCS at the pretraining stage and use CoS-E (Rajani et al., 2019) and OpenBook (Mihaylov et al., 2018) at the task-specific training stage.

- **Solomon** (Srivastava et al., 2020), **JUSTers** (Fadel et al., 2020), **SWAGex** (Rim and Okazaki, 2020), **UI** (Doxolodeo and Mahendra, 2020) and **CUHK** (Wang et al., 2020) use GPT or GPT-2 finetuned on the task training data. **JBNU** (Na and Lee, 2020) uses UniLM, which incorporates three LM tasks: unidirectional LM, bidirectional LM and sequence-to-sequence prediction LM, and only use one of the reference correct reasons. **UI** does not use the training data and treats the generation as a Cloze task. **SSN-NLP** (S, 2020) uses the seq2seq NMT framework without a pretrained LM.

Large-scale pretrained language models such as BART and GPT-2 dominates the submissions. The two systems with the highest human evaluations, namely **ANA** and **KaLM**, use additional resources such as Wiktionary, OMCS, and other commonsense datasets. This again shows that additional knowledge from structured databases can help with the generation of the reasons. From Table 6 we can see that BLEU does not correlate well with Human Evaluation, especially for the top-performing systems. According to a further experiment of BUT-FIT, the naive baseline of “copying source sentence as the reason” can give a BLEU of 17.23, which can rank No. 4 among all the submissions. This indicates that BLEU, which focuses on the surface token overlap, has difficulty in evaluating the generated text reliably. The top-performed system achieves the human evaluation score of 2.10, showing the power of pretrained language models, but considering the human performance of 2.58, we still have a long way to go to generate human acceptable reasons.

5 Related Work

Commonsense reasoning in natural language has been studied in different forms of tasks and has recently attracted extensive attention. In the Winograd Schema Challenge (WSC) (Levesque et al., 2012; Morgenstern and Ortiz, 2015), a model needs to solve hard co-reference resolution problems based on commonsense knowledge. For example, “*The trophy would not fit in the brown suitcase because it was too big. What was too big (trophy or suitcase)?*” The Choice of Plausible Alternatives (COPA) (Roemmele et al., 2011) emphasizes on events and consequences. Each question in COPA aims to find the suitable cause or result of the premise from two given alternatives. All premises and alternatives are simple sentences. For example, the premise can be “*The man broke his toe. What was the CAUSE of this?*” and the two candidate answers are “(1) *He got a hole in his sock.*” and “(2) *He dropped a hammer on his foot.*” Several subsequent datasets are inspired by COPA. The JHU Ordinal Common-sense Inference (JOCI) (Zhang et al., 2017) aims to label the plausibility from 5 (very likely) to 1 (impossible) of human response after a particular situation. Situations with Adversarial Generations (SWAG) (Zellers et al., 2018) request a system to choose the most likely-to-happen alternative after a specific situation. Those datasets emphasize the pre-situations and/or the after-situations of certain situations, but not on the reasons why they occur or are caused. Besides, our dataset is not limited to events or situations. It concerns a broader commonsense setting, which includes events, descriptions, assertion etc.

Some datasets are inspired by reading comprehension. The Story Cloze Test and ROCStories Corpora (Mostafazadeh et al., 2016; Sharma et al., 2018) aim to figure out the right ending from two candidate sentences after a four-sentence story. For a narrative text, MCScript (Ostermann et al., 2018a) gives various types of questions and pairs of answer candidates for each question. Most questions require knowledge beyond the facts mentioned in the text. Compared to those reading comprehension tasks, our benchmark encourages people to use any external resources they want.

Some other datasets evolve from QA problems and care more about factual commonsense knowledge. SQUABU (Davis, 2016) provides a small hand-constructed test of commonsense and scientific questions. CommonsenseQA (Talmor et al., 2018) asks crowd workers to create questions from ConceptNet (Speer et al., 2017), which is a large graph of commonsense knowledge, where each question discriminates its answer candidates between three target concepts that all share the same relationship to a single source

drawn from ConceptNet. OpenBookQA (Mihaylov et al., 2018) provides questions and answer candidates, as well as thousands of diverse facts about elementary level science that are related to the questions. The AI2 Reasoning Challenge (ARC) (Clark et al., 2018) gives thousands of questions with different knowledge types, as well as a relevant 14M-sentence corpus, mixed with science facts and other narrative sentences. MuTual provides a dataset for Multi-Turn dialogue reasoning in the commonsense area (Cui et al., 2020). Those questions are not easy to answer without specializing certain domain knowledge, while our questions are based on daily common sense.

Some datasets focus on non-sentential eventual plausibility (Wang et al., 2018; Porada et al., 2019), such as “gorilla-ride-camel”. In contrast, our dataset is based on *statements* which includes events, descriptions, assertion etc, not merely events, such as “China’s territory is larger than Japan’s”. And some datasets concentrate on limited attributes or actions of world knowledge, such as physics (Forbes and Choi, 2017). Our dataset concerns general commonsense knowledge beyond just physical common sense, the sentence in our task “Tom’s mom become (happy)/(upset) when Tom gets high grades in the exam” is about social and emotional common sense. For our first task, those statements that conforms to commonsense can also be phrased as being plausible. Thus our first task is similar to plausibility tests, despite that plausibility has a broader scope while our focus is on commonsense only.

More importantly, compared with our work, the **above** tasks do not directly estimate general common sense or ask the logical reasons behind the correct answers and questions. In recent years, some large-scale commonsense inference knowledge resources have been developed, which may be helpful in commonsense reasoning tasks. Atomic (Sap et al., 2018) presents a large-scale everyday commonsense knowledge graph, which has nine *if-then* relations with variables, including causes, effects, and so on. Event2Mind (Rashkin et al., 2018) proposes a new corpus and task, aiming to find out the mentioned/unmentioned people’s intents and reactions under various daily circumstances. These datasets are not directly useful for our benchmark since they focus only on a small domain. ConceptNet is a seminal knowledge graph that has been upgraded over time (Liu and Singh, 2004; Havasi et al., 2007; Speer and Havasi, 2013; Speer et al., 2017). ConceptNet constructs triples using labeled edges as relations and various words and/or phrases as entities. It also has the sentences describing the corresponding triples. In contrast to these datasets, we investigate the evaluation of common sense, rather than building a resource.

Before organizing this shared-task, a pilot study (Wang et al., 2019) has been performed, showing that there is still a significant gap between human and machine performance when no training data is provided, despite that the models have already been pretrained with over 100 million natural language sentences. In our task here, we also provide training data with human annotations.

6 Summary

This paper summarizes SemEval-2020 Task 4: Commonsense Validation and Explanation. In this task, we construct a dataset that consists of 11,997 instances and 83,986 sentences. The task attracted around 40 participating teams, out of which 31 teams submit their system papers. The pretrained models are shown to be very effective in Subtask A and Subtask B, but there is still a large room to improve system performances in Subtask C. Contextualized embedding such as RoBERTa and BART play a central role in the success of the top-performing models, demonstrating that such methods contain commonsense information to a good extent.

We attribute the high performance on Subtask A and B to several main reasons: 1) Subtask A is a relatively easy question by definition: a model needs only to detect a relatively less plausible content among the two candidate sentences. 2) Pretrained models are obtained on billion-words large corpora such as Wikipedia data, which help obtain commonsense knowledge (Zhou et al., 2019), which helps achieve considerably better performance. 3) As described in the annotation process, we use the sentences from OMCS to inspire crowd-sourcing workers. The top-3 systems also use OMCS, which potentially help them to attain better performances. 4) For Subtask B, as discussed in our data analysis section, the data has some flaws in the average length and common words, which reduces the difficulty. 5) Some instances have obvious patterns. For example, there are tens of instances that contain “put XXX into YYY”, and “XXX is bigger than YYY”, making the problems simpler. 6) Hundreds of crowd-sourcing

workers write instances. It is likely for workers to think about the shared commonsense knowledge, such as “XXX is bigger/shorter/quicker/slower than YYY”.

We consider future works in four directions: 1) We observe that there is still a gap between machine performance and human performance in Subtask C, and the reason generation task still needs further investigation. 2) The artifacts or spurious correlations in the datasets can be further removed, e.g., by making different candidate sentences in subtask B be the same, removing instances with shared commonsense knowledge, removing artifacts in common words, and filtering out common patterns. 3) Subtask A can be turned into a more difficult form. Instead of comparing which statement makes more sense, we can form it into a classification task, validating if one statement makes sense or not. 4) We notice that the BLEU score does not closely align with human evaluation for systems with high performances, and it is desirable to develop an auto-metric for comparing the semantic correlation between two reasons.

Acknowledgements

This work is supported by the National Science Foundation of China (Grant No. 61976180), the Westlake University, and the Bright Dream Joint Institute for Intelligent Robotics. The research of Xiaodan Zhu is supported by NSERC. Yue Zhang is the corresponding author.

References

- Yang Bai and Xiaobing Zhou. 2020. Deepyang at semeval-2020 task 4: Using the hidden layer of bert model for differentiating common sense. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the AI2 reasoning challenge. *CoRR*, abs/1803.05457.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. Mutual: A dataset for multi-turn dialogue reasoning. In *Proceedings of the 58th Conference of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Ciprian-Gabriel Cusmuluc, Lucia-Georgiana Coca, and Adrian Iftene. 2020. Uaics at semeval-2020 task 4 - using a bidirectional transformer for task a. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Ranjan Soumya Dash, Sandeep Routray, Prateek Varshney, and Ashutosh Modi. 2020. Cs-net at semeval-2020 task 4: Siamese bert for comve. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Ernest Davis. 2016. How to write science questions that are easy for people and hard for computers. *AI Magazine*, 37:13–22, 04.
- Ernest Davis. 2017. Logical formalizations of commonsense reasoning: a survey. *Journal of Artificial Intelligence Research*, 59:651–723.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Kerenza Doxolodeo and Rahmad Mahendra. 2020. Ui at sem-eval 2020 task 4 : Commonsense validation and explanation by exploiting contradiction. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Ali Fadel, Mahmoud Al-Ayyoub, and Erik Cambria. 2020. Justers at semeval-2020 task 4: Evaluating transformer models against commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Maxwell Forbes and Yejin Choi. 2017. Verb physics: Relative physical knowledge of actions and objects. In *ACL*.
- Catherine Havasi, Robert Speer, and Jason Alonso. 2007. Conceptnet 3: a flexible, multilingual semantic network for common sense knowledge. In *Recent advances in natural language processing*, pages 27–29. Citeseer.

- Catherine Havasi, Robert Speer, Kenneth Arnold, Henry Lieberman, Jason Alonso, and Jesse Moeller. 2010. Open mind common sense: Crowd-sourcing for common sense. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- Josef Jon, Martin Fajcik, Martin Docekal, and Pavel Smrz. 2020. But-fit at semeval-2020 task 4: Multilingual commonsense. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Anandh Konar, Chenyang Huang, Amine Trabelsi, and Osmar Zaiane. 2020. Ana at semeval-2020 task 4: multi-task learning for commonsense reasoning (union). In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12*, pages 552–561. AAAI Press.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Junyi Li, bin wang, and Haiyan Ding. 2020. Lijunyi at semeval-2020 task 4: An albert model based on different training sizes and depths for commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Hugo Liu and Push Singh. 2004. Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2019a. K-bert: Enabling language representation with knowledge graph. *arXiv preprint arXiv:1909.07606*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Shilei Liu, Yu Guo, BoChao Li, and Feiliang Ren. 2020. Lmve at semeval-2020 task 4: Commonsense validation and explanation using pretraining language model. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Pai Liu. 2020. Qiaoning at semeval-2020 task 4: Commonsense validation and explanation system based on ensemble of language model. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Daming Lu. 2020. Masked reasoner at semeval-2020 task 4: Fine-tuning roberta for commonsense reasoning. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Thanet Markchom, Bhuvana Dhruva, Chandresh Pravin, and Huizhi Liang. 2020. Uor at semeval-2020 task 4: Pre-trained sentence transformer models for commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Khanddorj Mendbayar and Masaki Aono. 2020. Kde senseforce. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. 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*, pages 2381–2391. Association for Computational Linguistics.
- Roweida Mohammed and Malak Abdullah. 2020. Teamjust at semeval-2020 task 4: Commonsense validation and explanation using ensembling techniques. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.

- Leora Morgenstern and Charles L. Ortiz. 2015. The winograd schema challenge: Evaluating progress in commonsense reasoning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15*, pages 4024–4025. AAAI Press.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James F. Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *HLT-NAACL*.
- Seung-Hoon Na and Jong-Hyeon Lee. 2020. Jbnu at semeval-2020 task 4: Bert and unilm for commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. *CoRR*, abs/1907.07355.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018a. Mscript: A novel dataset for assessing machine comprehension using script knowledge. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*. European Language Resource Association.
- Simon Ostermann, Michael Roth, Ashutosh Modi, Stefan Thater, and Manfred Pinkal. 2018b. Semeval-2018 task 11: Machine comprehension using commonsense knowledge. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 747–757. Association for Computational Linguistics.
- Xiaozhi Ou, Xiaobing Zhou, and hongling Li. 2020. Ynu-oxz at semeval-2020 task 4: Commonsense validation using bert with bidirectional gru. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Ian Porada, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. Can a gorilla ride a camel? learning semantic plausibility from text. *arXiv preprint arXiv:1911.05689*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *Technical report*.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. *arXiv preprint arXiv:1906.02361*.
- Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A. Smith, and Yejin Choi. 2018. Event2mind: Commonsense inference on events, intents, and reactions. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 463–473. Association for Computational Linguistics.
- Wiem Ben Rim and Naoaki Okazaki. 2020. Swagex at semeval-2020 task 4: Commonsense explanation as next event prediction. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S. Gordon. 2011. Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning. In *AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*, Stanford University, March.
- Kayalvizhi S. 2020. Ssn-nlp at semeval-2020 task 4: Text classification and generation on common sense context using neural networks. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Sirwe Saeedi, Aliakbar Panahi, Seyran Saeedi, and Alvis C Fong. 2020. Cs-nlp team at semeval-2020 task 4: Evaluation of state-of-the-art nlp deep learning architectures on commonsense reasoning task. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Maarten Sap, Ronan LeBras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. 2018. ATOMIC: an atlas of machine commonsense for if-then reasoning. *CoRR*, abs/1811.00146.

- Roy Schwartz, Maarten Sap, Ioannis Konstas, Leila Zilles, Yejin Choi, and Noah A. Smith. 2017. The effect of different writing tasks on linguistic style: A case study of the roc story cloze task. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 15–25. Association for Computational Linguistics.
- Rishi Sharma, James Allen, Omid Bakhshandeh, and Nasrin Mostafazadeh. 2018. Tackling the story ending biases in the story cloze test. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 752–757. Association for Computational Linguistics.
- Push Singh, Thomas Lin, Erik T Mueller, Grace Lim, Travell Perkins, and Wan Li Zhu. 2002. Open mind common sense: Knowledge acquisition from the general public. In *OTM Confederated International Conferences” On the Move to Meaningful Internet Systems”*, pages 1223–1237. Springer.
- Robert Speer and Catherine Havasi. 2013. Conceptnet 5: A large semantic network for relational knowledge. In *The People’s Web Meets NLP*, pages 161–176. Springer.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Vertika Srivastava, Sudeep Sahoo, Kumar, YEON KIM, HYANG, Rohit R.R, Mayank Raj, and AJAY JAISWAL. 2020. Team solomon at semeval-2020 task 4: Be reasonable: Exploiting large-scale language models for commonsense reasoning. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *CoRR*, abs/1811.00937.
- Don Teo. 2020. Tr at semeval-2020 task 4: Exploring the limits of language-model-based common sense validation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Jiajing Wan and Xinting Huang. 2020. Kalm at semeval-2020 task 4: Knowledge-aware language models for comprehension and generation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Su Wang, Greg Durrett, and Katrin Erk. 2018. Modeling semantic plausibility by injecting world knowledge. *arXiv preprint arXiv:1804.00619*.
- Cunxiang Wang, Shuailong Liang, Yue Zhang, Xiaonan Li, and Tian Gao. 2019. Does it make sense? and why? a pilot study for sense making and explanation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4020–4026, Florence, Italy, July. Association for Computational Linguistics.
- Hongru Wang, Xiangru Tang, Sunny Lai, and Sak Kwong Leung. 2020. Cuhk at semeval-2020 task 4: Commonsense explanation, reasoning and prediction with multi-task learning. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Luxi Xing, Yuqiang Xie, Yue Hu, and Wei Peng. 2020. Iie-nlp-nut at semeval-2020 task 4: Guiding plm with prompt template reconstruction strategy for comve. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Sheng Zhang, Rachel Rudinger, Kevin Duh, and Benjamin Van Durme. 2017. Ordinal common-sense inference. *Transactions of the Association for Computational Linguistics*, 5:379–395.
- Yice Zhang, Jiaxuan Lin, Yang Fan, Peng Jin, Yuanchao Liu, and Bingquan Liu. 2020. Cn-hit-it.nlp at semeval-2020 task 4: Enhanced language representation with multiple knowledge triples. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Qian Zhao, Siyu Tao, Jie Zhou, Linlin Wang, and Xin Lin. 2020. Ecnu-sensemaker at semeval-2020 task 4: Leveraging heterogeneous knowledge resources for commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.

Xuhui Zhou, Yue Zhang, Leyang Cui, and Dandan Huang. 2019. Evaluating commonsense in pre-trained language models. *arXiv preprint arXiv:1911.11931*.