

Benchmarking Machine Reading Comprehension: A Psychological Perspective

Saku Sugawara¹, Pontus Stenetorp², Akiko Aizawa¹

¹National Institute of Informatics, ²University College London
{saku, aizawa}@nii.ac.jp, p.stenetorp@cs.ucl.ac.uk

Abstract

Machine reading comprehension (MRC) has received considerable attention as a benchmark for natural language understanding. However, the conventional task design of MRC lacks explainability beyond the model interpretation, i.e., reading comprehension by a model cannot be explained in human terms. To this end, this position paper provides a theoretical basis for the design of MRC datasets based on psychology as well as psychometrics, and summarizes it in terms of the prerequisites for benchmarking MRC. We conclude that future datasets should (i) evaluate the capability of the model for constructing a coherent and grounded representation to understand context-dependent situations and (ii) ensure substantive validity by shortcut-proof questions and explanation as a part of the task design.

1 Introduction

Evaluation of natural language understanding (NLU) is a long-standing goal in the field of artificial intelligence. Machine reading comprehension (MRC) is a task that tests the ability of a machine to read and understand unstructured text and could be the most suitable task for evaluating NLU because of its generic formulation (Chen, 2018). Recently, many large-scale datasets have been proposed, and deep learning systems have achieved human-level performance for some of these datasets.

However, analytical studies have shown that MRC models do not necessarily achieve human-level understanding. For example, Jia and Liang (2017) use manually crafted adversarial examples to show that successful systems are easily distracted. Sugawara et al. (2020) show that a significant part of already solved questions is solvable even after shuffling the words in a sentence or dropping content words. These studies demonstrate that we cannot explain what type of understanding is

required by the datasets and is actually acquired by models. Although benchmarking MRC is related to the intent behind questions and is critical to test hypotheses from a top-down viewpoint (Bender and Koller, 2020), its theoretical foundation is poorly investigated in the literature.

In this position paper, we examine the prerequisites for benchmarking MRC based on the following two questions: (i) *What* does reading comprehension involve? (ii) *How* can we evaluate it? Our motivation is to provide a theoretical basis for the creation of MRC datasets. As Gilpin et al. (2018) indicate, interpreting the internals of a system is closely related to only the system’s architecture and is insufficient for explaining how the task is accomplished. This is because even if the internals of models can be interpreted, we cannot explain what is *measured* by the datasets. Therefore, our study focuses on the explainability of the task rather than the interpretability of models.

We first overview MRC and review the analytical literature that indicates that existing datasets might fail to correctly evaluate their intended behavior (Section 2). Subsequently, we present a psychological study of human reading comprehension in Section 3 for answering the *what* question. We argue that the concept of *representation levels* can serve as a conceptual hierarchy for organizing the technologies in MRC. Section 4 focuses on answering the *how* question. Here, we implement psychometrics to analyze the prerequisites for the task design of MRC. Furthermore, we introduce the concept of *construct validity*, which emphasizes validating the interpretation of the task’s outcome. Finally, in Section 5, we explain the application of the proposed concepts into practical approaches, highlighting potential future directions toward the advancement of MRC. Regarding the *what* question, we indicate that datasets should evaluate the capability of the *situation model*, which refers to the construction

Question	Foundation	Requirements	Future direction
What is reading comprehension?	Representation levels in human reading comprehension: (A) surface structure, (B) textbase, and (C) situation model.	(A) Linguistic-level sentence understanding, (B) comprehensiveness of skills for inter-sentence understanding, and (C) evaluation of coherent representation grounded to non-textual information.	(C) Dependence of context on defeasibility and novelty, and grounding to non-textual information with a long passage.
How can we evaluate reading comprehension?	Construct validity in psychometrics: (1) content, (2) substantive, (3) structural, (4) generalizability, (5) external, and (6) consequential aspects.	(1) Wide coverage of skills, (2) evaluation of the internal process, (3) structured metrics, (4) reliability of metrics, (5) comparison with external variables, and (6) robustness to adversarial attacks and social biases.	(2) Creating shortcut-proof questions by filtering and ablation, and designing a task for validating the internal process.

Table 1: Overview of theoretical foundations, requirements, and future directions of MRC discussed in this paper.

of a coherent and grounded representation of text based on human understanding. Regarding the *how* question, we argue that among the important aspects of the construct validity, *substantive validity* must be ensured, which requires the verification of the internal mechanism of comprehension.

Table 1 provides an overview of the perspectives taken in this paper. Our answers and suggestions to the *what* and *how* questions are summarized as follows: (1) Reading comprehension is the process of creating a situation model that best explains given texts and the reader’s background knowledge. The situation model should be the next focal point in future datasets for benchmarking the human-level reading comprehension. (2) To evaluate reading comprehension correctly, the task needs to provide a rubric (scoring guide) for sufficiently covering the aspects of the construct validity. In particular, the substantive validity should be ensured by creating shortcut-proof questions and by designing a task formulation that is explanatory itself.

2 Task Overview

2.1 Task Variations and Existing Datasets

MRC is a task in which a machine is given a document (*context*) and it answers the questions based on the context. Burges (2013) provides a general definition of MRC, i.e., *a machine comprehends a passage of text if, for any question regarding that text that can be answered correctly by a majority of native speakers, that machine can provide a string which those speakers would agree both answers that question.* We overview various aspects of the task along with representative datasets as follows. Existing datasets are listed in Appendix A.

Context Styles A context can be given in various forms with different lengths such as a single pas-

sage (MCTest (Richardson et al., 2013)), a set of passages (HotpotQA (Yang et al., 2018)), a longer document (CBT (Hill et al., 2016)), or open domain (Chen et al., 2017). In some datasets, a context includes non-textual information such as images (RecipeQA (Yagcioglu et al., 2018)).

Question Styles A question can be an interrogative sentence (in most datasets), a fill-in-the-blank sentence (cloze) (CLOTH (Xie et al., 2018)), knowledge base entries (QAngaroo (Welbl et al., 2018)) and search engine queries (MSMARCO (Nguyen et al., 2016)).

Answering Styles An answer can be (i) chosen from a text span of the given document (*answer extraction*) (NewsQA (Trischler et al., 2017)), (ii) chosen from a candidate set of answers (*multiple choice*) (MCTest (Richardson et al., 2013)), or (iii) generated as a free-form text (*description*) (NarrativeQA (Kočický et al., 2018)). Some datasets optionally allow answering by a *yes/no* reply (BoolQ (Clark et al., 2019)).

Sourcing Methods Initially, questions in small-scale datasets are created by experts (QA4MRE (Sutcliffe et al., 2013)). Later, fueling the development of neural models, most published datasets have more than a hundred thousand questions that are automatically created (CNN/Daily Mail (Hermann et al., 2015)), crowdsourced (SQuAD v1.1 (Rajpurkar et al., 2016)), and collected from examinations (RACE (Lai et al., 2017)).

Domains The most popular domain is Wikipedia articles (Natural Questions (Kwiatkowski et al., 2019)), but news articles are also used (Who-did-What (Onishi et al., 2016)). CliCR (Suster and Daelemans, 2018) and emrQA (Pampari et al., 2018) are datasets in the clinical domain. DuRC

(Saha et al., 2018) uses movie scripts.

Specific Skills Several recently proposed datasets require specific skills including unanswerable questions (SQuAD v2.0 (Rajpurkar et al., 2018)), dialogues (CoQA (Reddy et al., 2019)), DREAM (Sun et al., 2019)), multiple-sentence reasoning (MultiRC (Khashabi et al., 2018)), multi-hop reasoning (HotpotQA (Yang et al., 2018)), mathematical and set reasoning (DROP (Dua et al., 2019)), commonsense reasoning (CosmosQA (Huang et al., 2019)), coreference resolution (QuoRef (Dasigi et al., 2019)), and logical reasoning (ReClor (Yu et al., 2020)).

2.2 Benchmarking Issues

In some datasets, the performance of machines has already reached human-level performance. However, Jia and Liang (2017) indicate that models can easily be fooled by manual injection of distracting sentences. Their study revealed that questions simply gathered by crowdsourcing without careful guidelines or constraints are insufficient to evaluate precise language understanding.

This argument is supported by further studies across a variety of datasets. For example, Min et al. (2018) find that more than 90% of the questions in SQuAD (Rajpurkar et al., 2016) require obtaining an answer from a single sentence despite being provided with a passage. Sugawara et al. (2018) show that large parts of twelve datasets are easily solved only by looking at a few first question tokens and attending the similarity between the given questions and the context. Similarly, Feng et al. (2018) and Mudrakarta et al. (2018) demonstrate that models trained on SQuAD do not change their predictions even when the question tokens are partly dropped. Kaushik and Lipton (2018) also observe that question- and passage-only models perform well for some popular datasets. Min et al. (2019) and Chen and Durrett (2019) concurrently indicate that for multi-hop reasoning datasets, the questions are solvable only with a single paragraph and thus do not require multi-hop reasoning over multiple paragraphs. Zellers et al. (2019b) report that their dataset unintentionally contains stylistic biases in the answer options which are embedded by a language-based model.

Overall, these investigations highlight a grave issue of the task design, i.e., even if the models achieve human-level accuracies, we cannot prove that they successfully perform reading comprehen-

sion. This issue may be attributed to the low interpretability of black-box neural network models. However, a problem is that we cannot explain what is measured by the datasets even if we can interpret the internals of models. We speculate that this benchmarking issue in MRC can be attributed to the following two points: (i) we do not have a comprehensive theoretical basis of reading comprehension for specifying what we should ask (Section 3) and (ii) we do not have a well-established methodology for creating a dataset and for analyzing a model based on it (Section 4).¹ In the remainder of this paper, we argue that these issues can be addressed by using insights from the psychological study of reading comprehension and by implementing psychometric means of validation.

3 Reading Comprehension from Psychology to MRC

3.1 Computational Model in Psychology

Human text comprehension has been studied in psychology for a long time (Kintsch and Rawson, 2005; Graesser et al., 1994; Kintsch, 1988). Connectionist and computational architectures have been proposed for such comprehension including a mechanism pertinent to knowledge activation and memory storing. Among the computational models, the construction-integration (CI) model is the most influential and provides a strong foundation of the field (McNamara and Magliano, 2009).

The CI model assumes three different representation levels as follows:

- *Surface structure* is the linguistic information of particular words, phrases, and syntax obtained by decoding the raw textual input.
- *Textbase* is a set of propositions in the text, where the propositions are locally connected by inferences (*microstructure*).
- *Situation model* is a situational and coherent mental representation in which the propositions are globally connected (*macrostructure*), and it is often grounded to not only texts but also to sounds, images, and background information.

The CI model first decodes textual information (i.e., the surface structure) from the raw textual

¹These two issues loosely correspond to the plausibility and faithfulness of explanation (Jacovi and Goldberg, 2020). The plausibility is linked to what we expect as an explanation, whereas the faithfulness refers to how accurately we explain models' reasoning process.

input, then creates the propositions (i.e., textbase) and their local connections occasionally using the reader’s knowledge (*construction*), and finally constructs a coherent representation (i.e., situation model) that is organized according to five dimensions including time, space, causation, intentionality, and objects (Zwaan and Radvansky, 1998), which provides a global description of the events (*integration*). These steps are not exclusive, i.e., propositions are iteratively updated in accordance with the surrounding ones with which they are linked. Although the definition of successful text comprehension can vary, Hernández-Orallo (2017) indicates that comprehension implies the process of creating (or searching for) a situation model that best explains the given text and the reader’s background knowledge (Zwaan and Radvansky, 1998). We use this definition to highlight that the creation of a situation model plays a vital role in human reading comprehension.

Our aim in this section is to provide a basis for explaining what reading comprehension is, which requires terms for explanation. In the computational model above, the representation levels appear to be useful for organizing such terms. We ground existing NLP technologies and tasks to different representation levels in the next section.

3.2 Skill Hierarchy for MRC

Here, we associate the existing NLP tasks with the three representation levels introduced above. The biggest advantage of MRC is its general formulation, which makes it the most general task for evaluating NLU. This emphasizes the importance of the requirement of various *skills* in MRC, which can serve as the units for the explanation of reading comprehension. Therefore, our motivation is to provide an overview of the skills as a hierarchical taxonomy and to highlight the missing aspects in existing MRC datasets that are required for comprehensively covering the representation levels.

Existing Taxonomies We first provide a brief overview of the existing taxonomies of skills in NLU tasks. For recognizing textual entailment (Dagan et al., 2006), several studies present a classification of reasoning and commonsense knowledge (Bentivogli et al., 2010; Sammons et al., 2010; LoBue and Yates, 2011). For scientific question answering, Jansen et al. (2016) categorize knowledge and inference for an elementary-level dataset. Similarly, Boratko et al. (2018) propose types of

knowledge and reasoning for scientific questions in MRC (Clark et al., 2018). A limitation of both these studies is that the proposed sets of knowledge and inference are limited to the domain of elementary-level science. Although some existing datasets for MRC have their own classifications of skills, they are coarse and only cover a limited extent of typical NLP tasks (e.g., word matching and paraphrasing). In contrast, for a more generalizable definition, Sugawara et al. (2017) propose a set of 13 skills for MRC. Rogers et al. (2020) pursue this direction by proposing a set of questions with eight question types. In addition, Schlegel et al. (2020) propose an annotation schema to investigate requisite knowledge and reasoning. Dunietz et al. (2020) propose a template of understanding that consists of spatial, temporal, causal, and motivational questions to evaluate precise understanding of narratives with reference to human text comprehension.

In what follows, we describe the three representation levels that basically follow the three representations of the CI model but are modified for MRC. The three levels are shown in Figure 1. We emphasize that we do not intend to create exhaustive and rigid definitions of skills. Rather, we aim to place them in a hierarchical organization, which can serve as a foundation to highlight the missing aspects in the current MRC.

Surface Structure This level broadly covers the linguistic information and its semantic meaning, which can be based on the raw textual input. Although these features form a proposition according to psychology, it should be viewed as sentence-level semantic representation in computational linguistics. This level includes part-of-speech tagging, syntactic parsing, dependency parsing, punctuation recognition, named entity recognition (NER), and semantic role labeling (SRL). Although these basic tasks can be accomplished by some recent pretraining-based neural language models (Liu et al., 2019), they are hardly required in NLU tasks including MRC. In the natural language inference task, McCoy et al. (2019) indicate that existing datasets (e.g., Bowman et al. (2015)) may fail to elucidate the syntactic understanding of given sentences. Although it is not obvious that these basic tasks should be included in MRC and it is not easy to circumscribe linguistic knowledge from concrete and abstract knowledge (Zaenen et al., 2005; Manning, 2006), we should always care about the capabilities of basic tasks (e.g., use of checklists

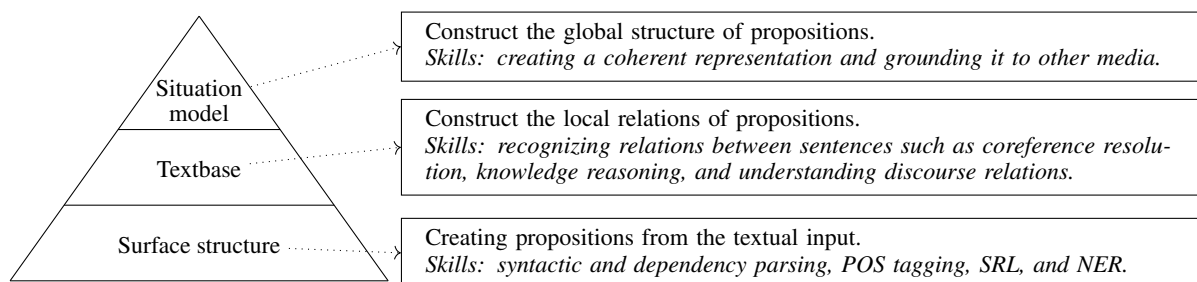


Figure 1: Representation levels and corresponding natural language understanding skills.

(Ribeiro et al., 2020)) when the performance of a model is being assessed.

Textbase This level covers local relations of propositions in the computational model of reading comprehension. In the context of NLP, it refers to various types of relations linked between sentences. These relations not only include the typical relations between sentences (discourse relations) but also the links between entities. Consequently, this level includes coreference resolution, causality, temporal relations, spatial relations, text structuring relations, logical reasoning, knowledge reasoning, commonsense reasoning, and mathematical reasoning. We also include multi-hop reasoning (Welbl et al., 2018) at this level because it does not necessarily require a coherent global representation over a given context. For studying the generalizability of MRC, Fisch et al. (2019) propose a shared task featuring training and testing on multiple domains. Talmor and Berant (2019) and Khashabi et al. (2020) also find that training on multiple datasets leads to robust generalization. However, unless we make sure that datasets require various skills with sufficient coverage, it might remain unclear whether we evaluate a model’s transferability of the reading comprehension ability.

Situation Model This level targets the global structure of propositions in human reading comprehension. It includes a coherent and situational representation of a given context and its grounding to the non-textual information. A coherent representation has well-organized sentence-to-sentence transitions (Barzilay and Lapata, 2008), which are vital for using procedural and script knowledge (Schank and Abelson, 1977). This level also includes characters’ goals and plans, meta perspective including author’s intent and attitude, thematic understanding, and grounding to other media. Most existing MRC datasets seem to struggle to target the situation model. We discuss further in Section 5.1.

Passage: *The princess climbed out the window of the high tower and climbed down the south wall when her mother was sleeping. She wandered out a good way. Finally, she went into the forest where there are no electric poles.*

Q1: *Who climbed out of the castle?* A: *Princess*
 Q2: *Where did the princess wander after escaping?*
 A: *Forest*
 Q3: *What would happen if her mother was not sleeping?*
 A: *the princess would be caught soon (multiple choice)*

Figure 2: Example questions of the different representation levels. The passage is taken from MCTest.

Example The representation levels in the example shown in Figure 2 are described as follows. Q1 is at the surface-structure level where a reader only needs to understand the subject of the first event. We expect that Q2 requires understanding of relations among described entities and events at the textbase level; the reader may need to understand who *she* means using coreference resolution. *Escaping* in Q2 also requires the reader’s commonsense to associate it with the first event. However, the reader might be able to answer this question only by looking for a place (specified by *where*) described in the passage, thereby necessitating the validity of the question to correctly evaluate the understanding of the described events. Q3 is an example that requires imagining a different situation at the situation-model level, which could be further associated with a grounding question such as *which figure best depicts the given passage?*

In summary, we indicate that the following features might be missing in existing datasets:

- Considering the capability to acquire basic understanding of the linguistic-level information.
- Ensuring that the questions comprehensively specify and evaluate textbase-level skills.
- Evaluating the capability of the situation model in which propositions are coherently organized

and are grounded to non-textual information.

Should MRC models mimic human text comprehension? In this paper, we do not argue that MRC models should mimic human text comprehension. However, when we design an NLU task and create datasets for testing human-like linguistic generalization, we can refer to the aforementioned features to frame the intended behavior to evaluate in the task. As Linzen (2020) discusses, the task design is orthogonal to how the intended behavior is realized at the implementation level (Marr, 1982).

4 MRC on Psychometrics

In this section, we provide a theoretical foundation for the evaluation of MRC models. When MRC measures the capability of reading comprehension, validation of the measurement is crucial to obtain a reliable and useful explanation. Therefore, we focus on psychometrics—a field of study concerned with the assessment of the quality of psychological measurement (Furr, 2018). We expect that the insights obtained from psychometrics can facilitate a better task design. In Section 4.1, we first review the concept of validity in psychometrics. Subsequently, in Section 4.2, we examine the aspects that correspond to construct validity in MRC and then indicate the prerequisites for verifying the intended explanation of MRC in its task design.

4.1 Construct Validity in Psychometrics

According to psychometrics, *construct validity* is necessary to validate the interpretation of outcomes of psychological experiments.² Messick (1995) report that construct validity consists of the six aspects shown in Table 2.

In the design of educational and psychological measurement, these aspects collectively provide verification questions that need to be answered for justifying the interpretation and use of test scores. In this sense, the construct validation can be viewed as an empirical evaluation of the meaning and consequence of measurement. Given that MRC is intended to capture the reading comprehension ability, the task designers need to be aware of these validity aspects. Otherwise, users of the task cannot justify the score interpretation, i.e., it cannot be confirmed that successful systems actually perform intended reading comprehension.

²In psychology, a construct is an abstract concept, which facilitates the understanding of human behavior such as vocabulary, skills, and comprehension.

4.2 Construct Validity in MRC

Table 2 also raises MRC features corresponding to the six aspects of construct validity. In what follows, we elaborate on these correspondings and discuss the missing aspects that are needed to achieve the construct validity of the current MRC.

Content Aspect As discussed in Section 3, sufficiently covering the skills across all the representation levels is an important requirement of MRC. It may be desirable that an MRC model is simultaneously evaluated on various skill-oriented examples.

Substantive Aspect This aspect appraises the evidence for the consistency of model behavior. We consider that this is the most important aspect for explaining reading comprehension, a process that subsumes various implicit and complex steps. To obtain a consistent response from an MRC system, it is necessary to ensure that the questions correctly assess the internal steps in the process of reading comprehension. However, as stated in Section 2.2, most existing datasets fail to verify that a question is solved by using an intended skill, which implies that it cannot be proved that a successful system can actually perform intended comprehension.

Structural Aspect Another issue in the current MRC is that they only provide simple accuracy as a metric. Given that the substantive aspect necessitates the evaluation of the internal process of reading comprehension, the structure of metrics needs to reflect it. However, a few studies have attempted to provide a dataset with multiple metrics. For example, Yang et al. (2018) not only ask for the answers to questions but also provide sentence-level supporting facts. This metric can also evaluate the process of multi-hop reasoning whenever the supporting sentences need to be understood for answering a question. Therefore, we need to consider both substantive and structural aspects.

Generalizability Aspect The generalizability of MRC can be understood from the reliability of metrics and the reproducibility of findings. For the reliability of metrics, we need to take care of the reliability of gold answers and model predictions. Regarding the accuracy of answers, the performance of the model becomes unreliable when the answers are unintentionally ambiguous or impractical. Because the gold answers in most datasets are only decided by the majority vote of crowd workers, the ambiguity of the answers is not considered. It

Validity aspects	Definition in psychometrics	Correspondence in MRC
1. Content	Evidence of content relevance, representativeness, and technical quality.	Questions require reading comprehension skills with sufficient coverage and representativeness over the representation levels.
2. Substantive	Theoretical rationales for the observed consistencies in the test responses including task performance of models.	Questions correctly evaluate the intended intermediate process of reading comprehension and provide rationales to the interpreters.
3. Structural	Fidelity of the scoring structure to the structure of the construct domain at issue.	Correspondence between the task structure and the score structure.
4. Generalizability	Extent to which score properties and interpretations can be generalized to and across population groups, settings, and tasks.	Reliability of test scores in correct answers and model predictions, and applicability to other datasets and models.
5. External	Extent to which the assessment scores' relationship with other measures and non-assessment behaviors reflect the expected relations.	Comparison of the performance of MRC with that of other NLU tasks and measurements.
6. Consequential	Value implications of score interpretation as a basis for the consequences of test use, especially regarding the sources of invalidity related to issues of bias, fairness, and distributive justice.	Considering the model vulnerabilities to adversarial attacks and social biases of models and datasets to ensure the fairness of model outputs.

Table 2: Aspects of the construct validity in psychometrics and corresponding features in MRC.

may be useful if such ambiguity can be reflected in the evaluation (e.g., using the item response theory (Lalor et al., 2016)). As for model predictions, an issue may be the reproducibility of results (Bouthillier et al., 2019), which implies that the reimplementation of a system generates statistically similar predictions. For the reproducibility of models, Dror et al. (2018) emphasize statistical testing methods to evaluate models. For the reproducibility of findings, Bouthillier et al. (2019) stress it as the transferability of findings in a dataset/task to another dataset/task. In open-domain question answering, Lewis et al. (2021) point out that successful models might only memorize dataset-specific knowledge. To facilitate this transferability, we need to have units of explanation that can be used in different datasets (Doshi-Velez and Kim, 2018).

External Aspect This aspect refers to the relationship between a model’s scores on different tasks. Yogatama et al. (2019) point out that current models struggle to transfer their ability from a task originally trained on (e.g., MRC) to different unseen tasks (e.g., SRL). To develop a general NLU model, one would expect that a successful MRC model should show sufficient performance on other NLU tasks as well. To this end, Wang et al. (2019) propose an evaluation framework with ten different NLU tasks in the same format.

Consequential Aspect This aspect refers to the actual and potential consequences of test use. In

MRC, this refers to the use of a successful model in practical situations other than tasks, where we need to ensure the robustness of a model to adversarial attacks and the accountability for unintended model behaviors. Wallace et al. (2019) highlight this aspect by showing that existing NLP models are vulnerable to adversarial examples, thereby generating egregious outputs.

Summary: Design of Rubric Given the validity aspects, our suggestion is to design a *rubric* (scoring guide used in education) of what reading comprehension we expect is evaluated in a dataset; this helps to inspect detailed strengths and weaknesses of models that cannot be obtained only by simple accuracy. The rubric should not only cover various linguistic phenomena (the *content* aspect) but also involve different levels of intermediate evaluation in the reading comprehension process (the *substantive* and *structural* aspects) as well as stress testing of adversarial attacks (the *consequential* aspect). The rubric is in a similar motivation with dataset statements (Bender and Friedman, 2018; Gebru et al., 2018); however, taking the validity aspects into account would improve its substance.

5 Future Directions

This section discusses future potential directions toward answering the *what* and *how* questions in Sections 3 and 4. In particular, we infer that the *situation model* and *substantive validity* are critical for benchmarking human-level MRC.

5.1 What Question: Situation Model

As mentioned in Section 3, existing datasets fail to fully assess the ability of creating the situation model. As a future direction, we suggest that the task should deal with two features of the situation model: context dependency and grounding.

5.1.1 Context-dependent Situations

A vital feature of the situation model is that it is conditioned on a given text, i.e., a representation is constructed distinctively depending on the given context. We elaborate it by discussing the two key features: defeasibility and novelty.

Defeasibility The defeasibility of a constructed representation implies that a reader can modify and revise it according to the newly acquired information (Davis and Marcus, 2015; Schubert, 2015). The defeasibility of NLU has been tackled in the task of if-then reasoning (Sap et al., 2019a), abductive reasoning (Bhagavatula et al., 2020), counterfactual reasoning (Qin et al., 2019), or contrast sets (Gardner et al., 2020). A possible approach in MRC is that we ask questions against a set of modified passages that describe slightly different situations, where the same question can lead to different conclusions.

Novelty An example showing the importance of contextual novelty is *Could a crocodile run a steeplechase?* by Levesque (2014). This question poses a novel situation where the solver needs to combine multiple commonsense knowledge to derive the correct answer. If non-fiction documents, such as newspaper and Wikipedia articles, are only used, some questions require only the reasoning of facts already known in web-based corpus. Fictional narratives may be a better source for creating questions on novel situations.

5.1.2 Grounding to Other Media

In MRC, grounding texts to non-textual information is not fully explored yet. Kembhavi et al. (2017) propose a dataset based on science textbooks, which contain questions with passages, diagrams, and images. Kahou et al. (2018) propose a figure-based question answering dataset that requires the understanding of figures including line plots and bar charts. Although another approach could be vision-based question answering tasks (Antol et al., 2015; Zellers et al., 2019a), we cannot directly use them for evaluating NLU because they focus on understanding of images rather than

texts. Similarly to the textbook questions (Kembhavi et al., 2017), a possible approach would be to create questions for understanding of texts through showing figures. We might also need to account for the *scope* of grounding (Bisk et al., 2020), i.e., ultimately understanding human language in a social context beyond simply associating texts with perceptual information.

5.2 How Question: Substantive Validity

Substantive validity requires us to ensure that the questions correctly assess the internal steps of reading comprehension. We discuss two approaches for this challenge: creating shortcut-proof questions and ensuring the explanation by design.

5.2.1 Shortcut-proof Questions

Gururangan et al. (2018) reveal that NLU datasets can contain unintended dataset biases embedded by annotators. If machine learning models exploit such biases for answering questions, we cannot evaluate the precise NLU of models. Following Geirhos et al. (2020), we define *shortcut-proof* questions as ones that prevent models from exploiting dataset biases and learning decision rules (*shortcuts*) that perform well only on i.i.d. test examples with regard to its training examples. Gardner et al. (2019) also point out the importance of mitigating shortcuts in MRC. In this section, we view two different approaches for this challenge.

Removing Unintended Biases by Filtering

Zellers et al. (2018) propose a model-based adversarial filtering method that iteratively trains an ensemble of stylistic classifiers and uses them to filter out the questions. Sakaguchi et al. (2020) also propose filtering methods based on both machines and humans to alleviate *dataset-specific* and *word-association* biases. However, a major issue is the inability to discern knowledge from bias in a closed domain. When the domain is equal to a dataset, patterns that are valid only in the domain are called *dataset-specific* biases (or annotation artifacts in the labeled data). When the domain covers larger corpora, the patterns (e.g., frequency) are called *word-association* biases. When the domain includes everyday experience, patterns are called *commonsense*. However, as mentioned in Section 5.1, commonsense knowledge can be *defeasible*, which implies that the knowledge can be false in unusual situations. In contrast, when the domain is our real world, indefeasible patterns are

called *factual knowledge*. Therefore, the distinction of bias and knowledge depends on where the pattern is recognized. This means that a dataset should be created so that it can evaluate reasoning on the intended knowledge. For example, to test defeasible reasoning, we must filter out questions that are solvable by usual commonsense only. If we want to investigate the reading comprehension ability without depending on factual knowledge, we can consider counterfactual or fictional situations.

Identifying Requisite Skills by Ablating Input Features Another approach is to verify shortcut-proof questions by analyzing the human answerability of questions regarding their key features. We speculate that if a question is still answerable by humans even after removing the intended features, the question does not require understanding of the ablated features (e.g., checking the necessity of resolving pronoun coreference after replacing pronouns with dummy nouns). Even if we cannot accurately identify such necessary features, by identifying partial features of them in a sufficient number of questions, we could expect that the questions evaluate the corresponding intended skill. In a similar vein, Geirhos et al. (2020) argue that a dataset is useful only if it is a good proxy for the underlying ability one is actually interested in.

5.2.2 Explanation by Design

Another approach for ensuring the substantive validity is to include explicit explanation in the task formulation. Although gathering human explanations is costly, the following approaches can facilitate the explicit verification of a model’s understanding using a few test examples.

Generating Introspective Explanation Inoue et al. (2020) classify two types of explanation in text comprehension: *justification explanation* and *introspective explanation*. The justification explanation only provides a collection of supporting facts for making a certain decision, whereas the introspective explanation provides the derivation of the answer for making the decision, which can cover linguistic phenomena and commonsense knowledge not explicitly mentioned in the text. They annotate multi-hop reasoning questions with introspective explanation and propose a task that requires the derivation of the correct answer of a given question to improve the explainability. Rajani et al. (2019) collect human explanations for commonsense reasoning and improve the system’s

performance by modeling the generation of the explanation. Although we must take into account the faithfulness of explanation, asking for introspective explanations could be useful in inspecting the internal reasoning process, e.g., by extending the task formulation so that it includes auxiliary questions that consider the intermediate facts in a reasoning process. For example, before answering Q2 in Figure 2, a reader should be able to answer *who escaped?* and *where did she escape from?* at the surface-structure level.

Creating Dependency Between Questions Another approach for improving the substantive validity is to create dependency between questions by which answering them correctly involves answering some other questions correctly. For example, Dalvi et al. (2018) propose a dataset that requires a procedural understanding of scientific facts. In their dataset, a set of questions corresponds to the steps of the entire process of a scientific phenomenon. Therefore, this set can be viewed as a single question that requires a complete understanding of the scientific phenomenon. In CoQA (Reddy et al., 2019), it is noted that questions often have pronouns that refer back to nouns appearing in previous questions. These mutually-dependent questions can probably facilitate the explicit validation of the models’ understanding of given texts.

6 Conclusion

In this paper, we outlined current issues and future directions for benchmarking machine reading comprehension. We visited the psychology study to analyze *what* we should ask of reading comprehension and the construct validity in psychometrics to analyze *how* we should correctly evaluate it. We deduced that future datasets should evaluate the capability of the situation model for understanding context-dependent situations and for grounding to non-textual information and ensure the substantive validity by creating shortcut-proof questions and designing an explanatory task formulation.

Acknowledgments

The authors would like to thank Xanh Ho for helping create the dataset list and the anonymous reviewers for their insightful comments. This work was supported by JSPS KAKENHI Grant Number 18H03297, JST ACT-X Grant Number JPM-JAX190G, and JST PRESTO Grant Number JPMJPR20C4.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433.
- Max Bartolo, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. 2020. [Beat the AI: Investigating adversarial human annotation for reading comprehension](#). *Transactions of the Association for Computational Linguistics*, 8:662–678.
- Regina Barzilay and Mirella Lapata. 2008. [Modeling local coherence: An entity-based approach](#). *Computational Linguistics*, 34(1):1–34.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Luisa Bentivogli, Elena Cabrio, Ido Dagan, Danilo Giampiccolo, Medea Lo Leggio, and Bernardo Magnini. 2010. Building textual entailment specialized data sets: a methodology for isolating linguistic phenomena relevant to inference. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *International Conference on Learning Representations*.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.
- Michael Boratko, Xiang Li, Tim O’Gorman, Rajarshi Das, Dan Le, and Andrew McCallum. 2020. [ProtoQA: A question answering dataset for prototypical common-sense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1122–1136, Online. Association for Computational Linguistics.
- Michael Boratko, Harshit Padigela, Divyendra Mikkineni, Pritish Yuvraj, Rajarshi Das, Andrew McCallum, Maria Chang, Achille Fokoue-Nkoutche, Pavan Kapanipathi, Nicholas Mattei, Ryan Musa, Kartik Talamadupula, and Michael Witbrock. 2018. A systematic classification of knowledge, reasoning, and context within the ARC dataset. In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 60–70. Association for Computational Linguistics.
- Xavier Bouthillier, César Laurent, and Pascal Vincent. 2019. [Unreproducible research is reproducible](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 725–734, Long Beach, California, USA. PMLR.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. Association for Computational Linguistics.
- Christopher J.C. Burges. 2013. Towards the machine comprehension of text: An essay. Technical report, Microsoft Research Technical Report MSR-TR-2013-125.
- Vittorio Castelli, Rishav Chakravarti, Saswati Dana, Anthony Ferritto, Radu Florian, Martin Franz, Dinesh Garg, Dinesh Khandelwal, Scott McCarley, Michael McCawley, Mohamed Nasr, Lin Pan, Cezar Pendus, John Pitrelli, Saurabh Pujar, Salim Roukos, Andrzej Sakrajda, Avi Sil, Rosario Uceda-Sosa, Todd Ward, and Rong Zhang. 2020. [The TechQA dataset](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1269–1278, Online. Association for Computational Linguistics.
- Danqi Chen. 2018. *Neural Reading Comprehension and Beyond*. Ph.D. thesis, Stanford University.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879. Association for Computational Linguistics.
- Jifan Chen and Greg Durrett. 2019. [Understanding dataset design choices for multi-hop reasoning](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4026–4032, Minneapolis, Minnesota. Association for Computational Linguistics.
- Michael Chen, Mike D’Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 2019. [CODAH: An adversarially-authored question answering dataset](#)

- for common sense. In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 63–69, Minneapolis, USA. Association for Computational Linguistics.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020. [HybridQA: A dataset of multi-hop question answering over tabular and textual data](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1026–1036, Online. Association for Computational Linguistics.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. [QuAC: Question answering in context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium. Association for Computational Linguistics.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. 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.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer.
- Bhavana Dalvi, Lifu Huang, Niket Tandon, Wen-tau Yih, and Peter Clark. 2018. [Tracking state changes in procedural text: a challenge dataset and models for process paragraph comprehension](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1595–1604. Association for Computational Linguistics.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasovic, Noah A. Smith, and Matt Gardner. 2019. [Quoref: A reading comprehension dataset with questions requiring coreferential reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5927–5934, Hong Kong, China. Association for Computational Linguistics.
- Ernest Davis and Gary Marcus. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103.
- Bhuwan Dhingra, Kathryn Mazaitis, and William W. Cohen. 2017. [Quasar: Datasets for question answering by search and reading](#).
- Finale Doshi-Velez and Been Kim. 2018. *Considerations for Evaluation and Generalization in Interpretable Machine Learning*, 1st edition. Springer International Publishing.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. [The hitchhiker’s guide to testing statistical significance in natural language processing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jesse Dunietz, Greg Burnham, Akash Bharadwaj, Owen Rambow, Jennifer Chu-Carroll, and Dave Ferrucci. 2020. [To test machine comprehension, start by defining comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7839–7859, Online. Association for Computational Linguistics.
- Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. [SearchQA: A new Q&A dataset augmented with context from a search engine](#).
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. [Pathologies of neural models make interpretations difficult](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3719–3728. Association for Computational Linguistics.
- James Ferguson, Matt Gardner, Hannaneh Hajishirzi, Tushar Khot, and Pradeep Dasigi. 2020. [IIRC: A dataset of incomplete information reading comprehension questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1137–1147, Online. Association for Computational Linguistics.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. [MRQA 2019 shared task: Evaluating generalization in reading comprehension](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 1–13, Hong Kong, China. Association for Computational Linguistics.

- R Michael Furr. 2018. *Psychometrics: an introduction*. Sage Publications.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models' local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, and Sewon Min. 2019. [On making reading comprehension more comprehensive](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 105–112, Hong Kong, China. Association for Computational Linguistics.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. ArXiv preprint 1803.09010.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.
- Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*, pages 80–89. IEEE.
- Arthur C. Graesser, Murray Singer, and Tom Trabasso. 1994. [Constructing inferences during narrative text comprehension](#). *Psychological review*, 101(3):371.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112. Association for Computational Linguistics.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. [The argument reasoning comprehension task: Identification and reconstruction of implicit warrants](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1930–1940, New Orleans, Louisiana. Association for Computational Linguistics.
- Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu, Qiaoqiao She, Xuan Liu, Tian Wu, and Haifeng Wang. 2018. [DuReader: a Chinese machine reading comprehension dataset from real-world applications](#). In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 37–46, Melbourne, Australia. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1693–1701. Curran Associates, Inc.
- José Hernández-Orallo. 2017. *The measure of all minds: evaluating natural and artificial intelligence*. Cambridge University Press.
- Daniel Hewlett, Alexandre Lacoste, Llion Jones, Illia Polosukhin, Andrew Fandrianto, Jay Han, Matthew Kelcey, and David Berthelot. 2016. [WikiReading: A novel large-scale language understanding task over wikipedia](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1545, Berlin, Germany. Association for Computational Linguistics.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The goldilocks principle: Reading children's books with explicit memory representations. In *International Conference on Learning Representations*.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Naoya Inoue, Pontus Stenetorp, and Kentaro Inui. 2020. [R4C: A benchmark for evaluating RC systems to get the right answer for the right reason](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6740–6750, Online. Association for Computational Linguistics.
- Alon Jacovi and Yoav Goldberg. 2020. [Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?](#) In *Proceedings of the*

- 58th Annual Meeting of the Association for Computational Linguistics, pages 4198–4205, Online. Association for Computational Linguistics.
- Peter Jansen, Niranjan Balasubramanian, Mihai Surdeanu, and Peter Clark. 2016. [What’s in an explanation? characterizing knowledge and inference requirements for elementary science exams](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2956–2965, Osaka, Japan. The COLING 2016 Organizing Committee.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2011–2021. Association for Computational Linguistics.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611. Association for Computational Linguistics.
- Samira Ebrahimi Kahou, Adam Atkinson, Vincent Michalski, Ákos Kádár, Adam Trischler, and Yoshua Bengio. 2018. [FigureQA: An annotated figure dataset for visual reasoning](#). In *International Conference on Learning Representations Workshop Track*.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How much reading does reading comprehension require? a critical investigation of popular benchmarks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015. Association for Computational Linguistics.
- Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *the IEEE Conference on Computer Vision and Pattern Recognition*.
- Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 252–262. Association for Computational Linguistics.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. [UnifiedQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. 2020. [QASC: A dataset for question answering via sentence composition](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8082–8090. AAAI Press.
- Walter Kintsch. 1988. The role of knowledge in discourse comprehension: A construction-integration model. *Psychological review*, 95(2):163.
- Walter Kintsch and Katherine A Rawson. 2005. Comprehension. *The Science of Reading: A Handbook*, pages 211–226.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Igor Labutov, Bishan Yang, Anusha Prakash, and Amos Azaria. 2018. [Multi-relational question answering from narratives: Machine reading and reasoning in simulated worlds](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 833–844. Association for Computational Linguistics.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale reading comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 796–805. Association for Computational Linguistics.

- John Lalor, Hao Wu, and Hong Yu. 2016. [Building an evaluation scale using item response theory](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 648–657, Austin, Texas. Association for Computational Linguistics.
- Hector J. Levesque. 2014. [On our best behaviour](#). *Artificial Intelligence*, 212:27 – 35.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. Question and answer test-train overlap in open-domain question answering datasets. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, Online. Association for Computational Linguistics.
- Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. 2019. [Reasoning over paragraph effects in situations](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 58–62, Hong Kong, China. Association for Computational Linguistics.
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. [Logiqa: A challenge dataset for machine reading comprehension with logical reasoning](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3622–3628. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter LoBue and Alexander Yates. 2011. [Types of common-sense knowledge needed for recognizing textual entailment](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 329–334, Portland, Oregon, USA. Association for Computational Linguistics.
- Kaixin Ma, Tomasz Jurczyk, and Jinho D. Choi. 2018. [Challenging reading comprehension on daily conversation: Passage completion on multiparty dialog](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2039–2048, New Orleans, Louisiana. Association for Computational Linguistics.
- Christopher D. Manning. 2006. Local textual inference: It’s hard to circumscribe, but you know it when you see it—and NLP needs it. Unpublished manuscript.
- David Marr. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Danielle S McNamara and Joe Magliano. 2009. Toward a comprehensive model of comprehension. *Psychology of learning and motivation*, 51:297–384.
- Samuel Messick. 1995. Validity of psychological assessment: Validation of inferences from persons’ responses and performances as scientific inquiry into score meaning. *American psychologist*, 50(9):741.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. [Compositional questions do not necessitate multi-hop reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4249–4257, Florence, Italy. Association for Computational Linguistics.
- Sewon Min, Victor Zhong, Richard Socher, and Caiming Xiong. 2018. [Efficient and robust question answering from minimal context over documents](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1725–1735. Association for Computational Linguistics.
- Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. [LSDSem 2017 shared task: The story cloze test](#). In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pages 46–51. Association for Computational Linguistics.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. [Did the model understand the question?](#) In *Proceedings of the 56th Annual Meeting of the Association*

- for *Computational Linguistics (Volume 1: Long Papers)*, pages 1896–1906. Association for Computational Linguistics.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. *CoRR*, abs/1611.09268.
- Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. TORQUE: A reading comprehension dataset of temporal ordering questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.
- Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. Who did What: A large-scale person-centered cloze dataset. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2230–2235. Association for Computational Linguistics.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018. MCScript: 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 Resources Association (ELRA).
- Simon Ostermann, Michael Roth, and Manfred Pinkal. 2019. MCScript2.0: A machine comprehension corpus focused on script events and participants. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, pages 103–117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. 2018. emrQA: A large corpus for question answering on electronic medical records. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2357–2368, Brussels, Belgium. Association for Computational Linguistics.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany. Association for Computational Linguistics.
- Lianhui Qin, Antoine Bosselut, Ari Holtzman, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi. 2019. Counterfactual story reasoning and generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5043–5053, Hong Kong, China. Association for Computational Linguistics.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 193–203. Association for Computational Linguistics.
- A Rogers, O Kovaleva, M Downey, and A Rumshisky. 2020. Getting closer to AI complete question answering: A set of prerequisite real tasks. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 8722–8731.
- Marzieh Saeidi, Max Bartolo, Patrick Lewis, Sameer Singh, Tim Rocktäschel, Mike Sheldon, Guillaume Bouchard, and Sebastian Riedel. 2018. Interpretation of natural language rules in conversational machine reading. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2087–2097, Brussels, Belgium. Association for Computational Linguistics.
- Amrita Saha, Rahul Aralikkatte, Mitesh M. Khapra, and Karthik Sankaranarayanan. 2018. DuoRC: Towards

- complex language understanding with paraphrased reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1683–1693. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. WINOGRANDE: an adversarial winograd schema challenge at scale. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 8732–8430.
- Mark Sammons, V.G.Vinod Vydiswaran, and Dan Roth. 2010. “Ask not what textual entailment can do for you...”. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1199–1208. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019a. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3027–3035.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019b. Social IQa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Roger C. Schank and Robert P. Abelson. 1977. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Lawrence Erlbaum.
- Viktor Schlegel, Marco Valentino, Andre Freitas, Goran Nenadic, and Riza Batista-Navarro. 2020. A framework for evaluation of machine reading comprehension gold standards. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5359–5369, Marseille, France. European Language Resources Association.
- Lenhart K Schubert. 2015. What kinds of knowledge are needed for genuine understanding? In *IJCAI 2015 Workshop on Cognitive Knowledge Acquisition and Applications (Cognitum 2015)*.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. What makes reading comprehension questions easier? In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219. Association for Computational Linguistics.
- Saku Sugawara, Yusuke Kido, Hikaru Yokono, and Akiko Aizawa. 2017. Evaluation metrics for machine reading comprehension: Prerequisite skills and readability. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 806–817. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. Assessing the benchmarking capacity of machine reading comprehension datasets. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 8918–8927.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. DREAM: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231.
- Simon Suster and Walter Daelemans. 2018. CliCR: a dataset of clinical case reports for machine reading comprehension. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1551–1563. Association for Computational Linguistics.
- Richard Sutcliffe, Anselmo Peñas, Eduard Hovy, Pamela Forner, Álvaro Rodrigo, Corina Forascu, Yassine Benajiba, and Petya Osenova. 2013. Overview of QA4MRE main task at CLEF 2013. *Working Notes, CLEF*.
- Alon Talmor and Jonathan Berant. 2019. MultiQA: An empirical investigation of generalization and transfer in reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4911–4921, Florence, Italy. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. NewsQA: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200. Association for Computational Linguistics.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong

- Kong, China. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems* 32, pages 3261–3275. Curran Associates, Inc.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. Towards AI-complete question answering: a set of prerequisite toy tasks. In *International Conference on Learning Representations*.
- Qizhe Xie, Guokun Lai, Zihang Dai, and Eduard Hovy. 2018. Large-scale cloze test dataset created by teachers. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2344–2356, Brussels, Belgium. Association for Computational Linguistics.
- Semih Yagcioglu, Aykut Erdem, Erkut Erdem, and Nazli Ikizler-Cinbis. 2018. RecipeQA: A challenge dataset for multimodal comprehension of cooking recipes. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1358–1368, Brussels, Belgium. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380. Association for Computational Linguistics.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. Learning and evaluating general linguistic intelligence. *CoRR*, abs/1901.11373.
- Weihao Yu, Zihang Jiang, Yanfei Dong, and Jiashi Feng. 2020. ReClor: A reading comprehension dataset requiring logical reasoning. In *International Conference on Learning Representations*.
- Annie Zaenen, Lauri Karttunen, and Richard Crouch. 2005. Local textual inference: Can it be defined or circumscribed? In *Proceedings of the ACL Workshop on Empirical Modeling of Semantic Equivalence and Entailment*, pages 31–36, Ann Arbor, Michigan. Association for Computational Linguistics.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019a. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6720–6731.
- 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*, pages 93–104. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019b. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. 2018. ReCoRD: Bridging the gap between human and machine commonsense reading comprehension. ArXiv preprint 1810.12885.
- Rolf A Zwaan and Gabriel A Radvansky. 1998. Situation models in language comprehension and memory. *Psychological bulletin*, 123(2):162.

A Machine Reading Comprehension Datasets

Tables 3, 4, 5, and 6 list machine reading comprehension and related datasets along with their answer styles, dataset size, type of corpus, sourcing methods, and focuses.

Dataset	Answer style	Size	Corpus	Question source	Focus
QA4MRE (Sutcliffe et al., 2013)	multiple-choice	240	technical document	expert	exam-level questions
MCTest (Richardson et al., 2013)	multiple-choice	2.6K	written story	crowd	children-level narrative
bAbI (Weston et al., 2015)	descript	10K * 20	generated text	automated	toy tasks for prerequisite skills
CNN/ DailyMail (Hermann et al., 2015)	extract	1.4M	news article	automated	entity cloze
Children’s Book Test (Hill et al., 2016)	extract	688K	narrative	automated	large-scale automated
SQuAD 1.1 (Rajpurkar et al., 2016)	extract	100K	Wikipedia	crowd	large-scale crowdsourced
LAMBADA (Paperno et al., 2016)	descript	10K	narrative	crowd	hard language modeling
WikiReading (Hewlett et al., 2016)	descript	18m	Wikipedia	automated	Wikidata articles
Who did What (Onishi et al., 2016)	multiple-choice	200K	news article	automated	cloze of person names
MS MARCO (Nguyen et al., 2016)	descript	100K	web snippet	query	description on web snippets
NewsQA (Trischler et al., 2017)	extract	120K	news article	crowd	blindly created questions
SearchQA (Dunn et al., 2017)	extract	140K	web snippet	trivia	49.6 snippets on average
RACE (Lai et al., 2017)	multiple-choice	100K	language exam	expert	middle and high school English exam in China
Story Cloze Test (Mostafazadeh et al., 2017)	multiple-choice	3.7K	written story	crowd	98,159 stories for training
TriviaQA (Joshi et al., 2017)	extract	650K	web snippet	trivia	trivia questions
Quasar (Dhingra et al., 2017)	extract	80K	web snippet	query	search queries
TextbookQA (Kembhavi et al., 2017)	multiple-choice	26K	textbook	expert	figures included
AddSent SQuAD (Jia and Liang, 2017)	extract	3.6K	Wikipedia	crowd	distracting sentences injected

Table 3: Machine reading comprehension datasets published until 2017. In the *answer style* column, *descript* represents description (free-form answering) and *extract* denotes answer extraction by selecting a span in given texts. *Size* indicates the size of the whole dataset including training, development, and test sets. In the *question source* column, *crowd* indicates questions written by crowdworkers and *query* indicates questions collected from search-engine queries.

Dataset	Answer style	Size	Corpus	Question source	Focus
ARCT (Habernal et al., 2018)	multiple-choice	2.0K	debate article	crowd expert	reasoning on argument
QAngaroo (Welbl et al., 2018)	multiple-choice	50K	Wikipedia, MEDLINE	automated	multi-hop reasoning
CLOTH (Xie et al., 2018)	multiple-choice	99K	various	expert	cloze in exam texts
NarrativeQA (Kočíský et al., 2018)	descript	45K	movie script	crowd	summary and full story tasks
MCScript (Ostermann et al., 2018)	multiple-choice	30K	written story	crowd	commonsense reasoning, script knowledge
ClICR (Suster and Daelemans, 2018)	extract	100K	clinical case text	automated	cloze style queries
ARC (Clark et al., 2018)	multiple-choice	8K	science exam	expert	retrieved documents from textbooks
DuoRC (Saha et al., 2018)	extract	186K	movie script	crowd	commonsense reasoning, multi-sentence reasoning
ProPara (Dalvi et al., 2018)	extract	2K	science exam	automated	procedural understanding
DuReader (He et al., 2018)	descript	200K	web snippet	query crowd	Chinese, Baidu Search/Knows
MultiRC (Khashabi et al., 2018)	multiple-choice	6K	various documents	crowd	multi-sentence reasoning
Multi-party Dialog (Ma et al., 2018)	extract	13K	TV show transcript	automated	1.7k crowd dialogues, cloze query
SQuAD 2.0 (Rajpurkar et al., 2018)	extract no answer	100K	Wikipedia	crowd	unanswerable questions
ShARC (Saeidi et al., 2018)	yes/no/irrelevant	32K	web snippet	crowd	reasoning on rules from government documents
QuAC (Choi et al., 2018)	extract yes/no	100K	Wikipedia	crowd	dialogue-based, 14k dialogues
Textworlds QA (Labutov et al., 2018)	extract	1.2M	generated text	automated	simulated worlds, logical reasoning
SWAG (Zellers et al., 2018)	multiple-choice	113K	video captions	language-model	commonsense reasoning
emrQA (Pampari et al., 2018)	extract	400K	clinical documents	automated	using annotated logical forms on i2b2 dataset
HotpotQA (Yang et al., 2018)	extract yes/no	113K	Wikipedia	crowd	multi-hop reasoning
OpenbookQA (Mihaylov et al., 2018)	multiple-choice	6.0K	textbook	crowd	commonsense reasoning
RecipeQA (Yagcioglu et al., 2018)	multiple-choice	36K	recipe script	automated	multimodal questions
ReCoRD (Zhang et al., 2018)	extract	120K	news article	crowd	commonsense reasoning, cloze query

Table 4: Machine reading comprehension datasets published in 2018. In the *answer style* column, *descript* represents description (free-form answering) and *extract* denotes answer extraction by selecting a span in given texts. *Size* indicates the size of the whole dataset including training, development, and test sets. In the *question source* column, *crowd* indicates questions written by crowdworkers and *query* indicates questions collected from search-engine queries.

Dataset	Answer style	Size	Corpus	Question source	Focus
CoQA (Reddy et al., 2019)	extract yes/no	127K	Wikipedia	crowd	dialogue-based, 8k dialogues
Commonsense QA (Talmor et al., 2019)	multiple-choice	12K	ConceptNet	crowd	commonsense reasoning
Natural Questions (Kwiatkowski et al., 2019)	extract yes/no	323K	Wikipedia	query crowd	short or long answer styles
DREAM (Sun et al., 2019)	multiple-choice	10K	language exam	expert	dialogue-based, 6.4k multi-party dialogues
DROP (Dua et al., 2019)	descript	96K	Wikipedia	crowd	discrete reasoning
SocialIQA (Sap et al., 2019b)	multiple-choice	38K	crowd	crowd	commonsense reasoning about social situation
BoolQ (Clark et al., 2019)	yes/no	16K	Wikipedia	query crowd	boolean questions, subset of Natural Questions
MSCript 2.0 (Ostermann et al., 2019)	multiple-choice	20K	narrative	crowd	commonsense reasoning, script knowledge
HellaSWAG (Zellers et al., 2019b)	multiple-choice	70K	web snippet	language-model	commonsense reasoning, WikiHow and ActivityNet
CODAH (Chen et al., 2019)	multiple-choice	2.8K	written prompt	expert	adversarial collection
Quoref (Dasigi et al., 2019)	extract	24K	Wikipedia	crowd	coreference resolution
CosmosQA (Huang et al., 2019)	multiple-choice	36K	narrative	crowd	commonsense reasoning
PubMedQA (Jin et al., 2019)	yes/no	273.5K	PubMed	expert automated	biomedical domain, 1k expert questions
ROPES (Lin et al., 2019)	extract	14K	textbook Wikipedia	crowd	paragraph effects in situations

Table 5: Machine reading comprehension datasets published in 2019. In the *answer style* column, *descript* represents description (free-form answering) and *extract* denotes answer extraction by selecting a span in given texts. *Size* indicates the size of the whole dataset including training, development, and test sets. In the *question source* column, *crowd* indicates questions written by crowdworkers and *query* indicates questions collected from search-engine queries.

Dataset	Answer style	Size	Corpus	Question source	Focus
QuAIL (Rogers et al., 2020)	multiple-choice	15K	various	crowd	prerequisite real tasks
QASC (Khot et al., 2020)	multiple-choice	10K	textbook	crowd	knowledge composition
AdversarialQA (Bartolo et al., 2020)	extract	36K	Wikipedia	crowd	adversarial collection
ReClor (Yu et al., 2020)	multiple-choice	6.1K	exam	expert	logical reasoning
R ⁴ C (Inoue et al., 2020)	extract descript	5K	Wikipedia	crowd	multi-hop reasoning
TechQA (Castelli et al., 2020)	descript	1.4K	tech documents	crowd	tech forum questions
LogiQA (Liu et al., 2020)	multiple-choice	8.7K	exam	expert	logical reasoning
ProtoQA (Boratko et al., 2020)	descript	9.8K	web snippet	crowd	commonsense reasoning over prototypical situations
IIRC (Ferguson et al., 2020)	descript	13K	Wikipedia	crowd	incomplete information
HybridQA (Chen et al., 2020)	extract	70K	Wikipedia	crowd	understanding tabular data
TORQUE (Ning et al., 2020)	extract	21K	TempEval-3	crowd	temporal ordering
2WikiMultiHopQA (Ho et al., 2020)	extract yes/no descript	200K	Wikipedia	automated	multi-hop reasoning

Table 6: Machine reading comprehension datasets published in 2020. In the *answer style* column, *descript* represents description (free-form answering) and *extract* denotes answer extraction by selecting a span in given texts. *Size* indicates the size of the whole dataset including training, development, and test sets. In the *question source* column, *crowd* indicates questions written by crowdworkers and *query* indicates questions collected from search-engine queries.