

Machine Reading Comprehension using Case-based Reasoning

Dung Thai¹, Dhruv Agarwal¹, Mudit Chaudhary¹, Wenlong Zhao¹, Rajarshi Das^{2*},
Manzil Zaheer², Jay-Yoon Lee⁴, Hannaneh Hajishirzi⁵, Andrew McCallum¹

¹University of Massachusetts Amherst

²AWS AI Labs, ³Google Research, ⁴Seoul National University, ⁵University of Washington
{dthai, dagarwal, mchaudhary, wenlongzhao, mccallum}@cs.umass.edu

Abstract

We present an accurate and interpretable method for answer extraction in machine reading comprehension that is reminiscent of case-based reasoning (CBR) from classical AI. Our method (CBR-MRC) builds upon the hypothesis that contextualized answers to similar questions share semantic similarities with each other. Given a test question, CBR-MRC first retrieves a set of similar cases from a non-parametric memory and then predicts an answer by selecting the span in the test context that is most similar to the contextualized representations of answers in the retrieved cases. The semi-parametric nature of our approach allows it to attribute a prediction to the specific set of evidence cases, making it a desirable choice for building reliable and debuggable QA systems. We show that CBR-MRC provides high accuracy comparable with large reader models and outperforms baselines by 11.5 and 8.4 EM on NaturalQuestions and NewsQA, respectively. Further, we demonstrate the ability of CBR-MRC in identifying not just the correct answer tokens but also the span with the most relevant supporting evidence. Lastly, we observe that contexts for certain question types show higher lexical diversity than others and find that CBR-MRC is robust to these variations while performance using fully-parametric methods drops. ¹

1 Introduction

Machine reading comprehension (MRC) aims to measure the ability of models to understand and reason over a specified sequence of text. In the extractive setting, the task requires a model to answer a question by reading a set of one or more passages, referred to as the context, and identifying a span of text from that context as the answer.

*Work done while at University of Washington.

¹Our code is available at <https://github.com/dungtn/cbr-txt>




Figure 1: **Overview of CBR-MRC.** *Top:* given a test question, CBR-MRC first retrieves similar cases from a memory of known (question, context, answer) triples. *Bottom:* CBR-MRC scores each candidate answer span by comparing its contextualized representation with the answer spans of the retrieved cases to output a score. The candidate with the highest aggregate score with the answer spans from the cases is output as the prediction (“Charles Babbage”).

Current state-of-the-art machine readers (Devlin et al., 2019; Choi et al., 2018; Liu et al., 2019; Lewis et al., 2020; Izacard and Grave, 2021) use fine-tuned transformer-based models (Vaswani et al., 2017). However, the prediction mechanisms in these models are largely opaque to humans, limiting their interpretability and maintainability in real-world scenarios (Ribeiro et al., 2016; Yang et al., 2018; Camburu et al., 2018). This becomes particularly relevant when practitioners need to determine

the cause of erroneous predictions, for instance, to patch models deployed in production systems. Moreover, to build reliable QA systems, it is essential for models to not only predict the correct answer but also provide evidence to support the predictions (Thayaparan et al., 2020; Zeng et al., 2020).

To address these limitations, we introduce **CBR-MRC** — a semi-parametric approach for MRC that makes predictions by explicitly reasoning over a set of retrieved answers for similar questions and their contexts from a non-parametric memory. Our key idea is that contextualized answers to similar questions share representational similarities. We posit that a (question, passage, answer) triple captures a set of *latent relations*, which can be understood to represent the *types* or *intents* present within a question given the accompanying context. We use the intersection of latent relation sets between questions to learn a similarity function that can then be used to predict answers to questions with similar intents. More concretely, our approach first compares the contextualized representations of candidate answer spans in the target context with gold answer spans of similar questions from memory. Then, the candidate with the highest aggregate similarity with the gold case answer spans is selected as the predicted answer. For example, in Figure 1, the question “*Who invented the computer?*” shares the same latent relations with the question “*Who is the inventor of the telephone?*”, and the contextualized span similarity with the case answer “*Graham Bell*” is maximized when compared with the correct answer “*Charles Babbage*” and low when compared with other candidate spans such as “*mechanical computer*”.

CBR-MRC adapts the case-based reasoning (Schank, 1983) paradigm from classical AI (Kolodner, 1983; Risland, 1983; Leake, 1996), which typically consists of four steps — retrieve, reuse, revise, and retain. CBR-MRC uses the first two — **retrieval** from a memory of cases, or *casebase*, using the embedding of the test question followed by **reusing** reasoning patterns encoded in the contextualized answer embeddings of the retrieved cases to select the most similar span from the test context as the predicted answer. To the best of our knowledge, we are the first to propose a novel adaptation of the CBR framework for the task of question answering over unstructured text. We empirically verify CBR-MRC on NaturalQuestions and NewsQA, two

common MRC datasets, and find our method outperforms strong baselines by up to 11.53 EM. The inference procedure of CBR-MRC is not just accurate but also interpretable — we know which cases the model uses for its prediction and how much each case contributes. Further, the semi-parametric nature of CBR-MRC allows for efficient and accurate domain adaptation by simply adding new cases to the casebase without the need to modify any parameters.² Lastly, we also find that our method shows higher gains compared to fully-parametric models when the training set contains lexically diverse contexts for questions with similar latent relations, demonstrating the robustness of CBR-MRC when questions may be answered using several different passages.

In summary, our contributions are as follows:

1. We propose CBR-MRC — a novel semi-parametric approach for MRC that outperforms strong baselines on multiple datasets and settings by up to 11.53 EM on NaturalQuestions and 8.4 EM on NewsQA.
2. We show that CBR-MRC makes predictions by identifying the correct supporting evidence. When presented with gold and noisy evidence sentences, CBR-MRC predicts the correct answer span with 83.2 Span-F1 and 74.6 Span-EM compared to BLANC (Seonwoo et al., 2020a) with 79.0 Span-F1 and 66.8 Span-EM on NaturalQuestions.
3. We show the ability of CBR-MRC for efficient and accurate few-shot domain adaptation. Compared to a strong baseline (Friedman et al., 2021), we achieve a 5.5 EM gain on RelEx (Levy et al., 2017) and at-par performance on BioASQ (Tsatsaronis et al., 2015).
4. We show that questions with different sets of latent relations can be expressed with varying degrees of lexical diversity in their accompanying contexts and analyze model behavior on such questions. As lexical diversity increases, fully-parametric methods drop by up to 15.4 F1, while CBR-MRC shows robust performance with drops of only up to 4.1 F1.

²Case augmentation can be extended to make point-fixes for specific inference errors without expensive re-training.

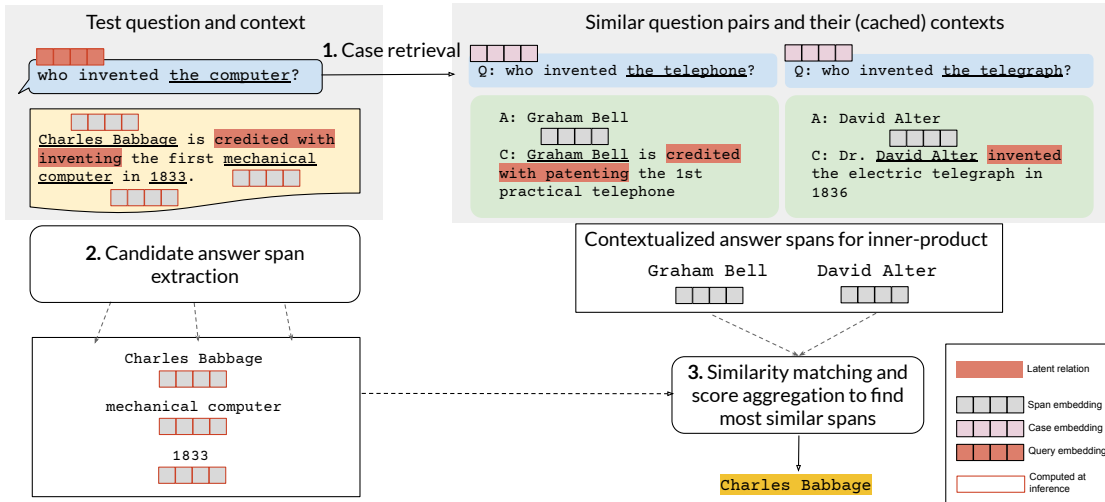


Figure 2: **Inference with CBR-MRC.** For a given target query, (1) other similar queries (and their contexts) are retrieved; (2) candidate answer spans are extracted from the target context; (3) the candidate spans in the target context are ranked w.r.t. answer spans of the retrieved queries by comparing their contextualized embeddings. Finally, the span with the highest inner-product similarity is selected as the prediction. Note that the casebase questions and their corresponding answer span embeddings are pre-computed and cached.

2 Related Work

Machine Reading Comprehension. MRC evaluates machine ability to reason over natural language by training it to answer questions on a given passage. Several variations of this task exist, such as cloze-style (Hermann et al., 2015; Yagcioglu et al., 2018), multiple-choice (Richardson et al., 2013; Lai et al., 2017), extractive (Yang et al., 2015; Trischler et al., 2016), and generative question-answering (Bajaj et al., 2016; Kočiský et al., 2017). Our work focuses on the extractive MRC setting, which aims to identify the correct answer span within a passage given a test question.

Interpretable methods for MRC have also seen much recent interest. Thayaparan et al. (2020) highlights the challenges facing MRC in terms of explainability and its impact on MRC performance. Several methods (Yu et al., 2014; Min et al., 2018; Gravina et al., 2018) have attempted to improve interpretability in MRC by identifying the supporting sentence for the answer. However, these methods rely on the availability of gold-supporting sentences, which may not always be present. Other methods utilize graph networks (Ye et al., 2019; Tu et al., 2019) and attention mechanisms to capture the attended parts of text or relations (Seo et al., 2016; Shao et al., 2020). In this work, we improve interpretability by providing a novel inference procedure that allows us to clearly attribute predictions to previously seen instances that were used in the decision-making process.

Case-based Reasoning. Previous work (Das et al., 2020a,b, 2022; Thai et al., 2022; Yang et al., 2023) has demonstrated strong performance using CBR on knowledge graph reasoning tasks. CBR has also been successfully applied to the task of semantic parsing (Das et al., 2021; Pasupat et al., 2021; Awasthi et al., 2023), where methods typically retrieve the logical forms of similar cases and pass them through a seq2seq model as input along with a target query. Both settings, however, operate over structured spaces, where similarities between instances can easily be determined by collecting and comparing symbolic patterns in the data, e.g., paths between nodes in a KG, or syntactic similarities in semantic parsing. In this work, however, we adapt CBR to operate over unstructured text, and it is not obvious how previous approaches can be re-used to work in this setting out-of-the-box.

A concurrent work (Iyer et al., 2023) performs answer sentence selection by using Graph Neural Networks to model interactions among questions and answers, assuming that for each question-answer, there exist similar question-answers in the training data. We instead focus on the more fine-grained task of answer span prediction from contexts and use transformer models to score question-question and answer-answer similarities between each retrieved case and a test example.

In-Context Learning (ICL). ICL is an emergent property observed in large language models (LLMs) (Brown et al., 2020; Min et al.,

2022) showing remarkable success in several NLP tasks (Cheng et al., 2022; Lewkowycz et al., 2022; Dunn et al., 2022). To make predictions, a model is simply prompted with demonstrations for a target task and a test instance. Recent work (Rubin et al., 2021; Liu et al., 2021; Li et al., 2023; Khattab et al., 2022) has shown strong performance by retrieving relevant demonstrations. This setting can be understood as a specific instantiation of the case-based reasoning framework, where the “reuse” step uses ICL instead of a fine-tuned model. ICL, however, is sensitive to the prompt (Chen et al., 2022), leading to black-boxed inference and reducing reliability of predictions. Problem decomposition and iterative prompting methods, such as chain-of-thought (Wei et al., 2022) and least-to-most (Zhou et al., 2022; Drozdov et al., 2022), attempt to address this issue by converting complex problems into intermediate steps. In contrast to methods reliant on emergent pre-training behavior, our work *explicitly* encodes the relational similarity hypothesis in a trained model to reuse exemplars for making predictions.

3 Method

We now formally describe CBR-MRC (Figure 2). A *case* in CBR consists of a problem and its solution. Formally, we define a case c as a tuple $\{q, \mathcal{A}, p\}$, representing a question, its corresponding answer set³, and the passage context that contains each element in the answer set. The context p is a span of text that not only contains all $a \in \mathcal{A}$ but also captures one or more latent relations between the question and the answer. We refer to a collection of all cases as the *casebase* \mathcal{B} , which is typically set to the full training data and may continually be augmented to incorporate new data patterns. In the MRC setting, annotations for p for all $c \in \mathcal{B}$ are assumed to be available.

3.1 Case Retrieval

Given an input question q , CBR-MRC retrieves *similar* cases from the casebase, where similarity is determined by the overlap of latent relations encoded by the query and the casebase questions. Formally, let \mathcal{R}_q be the set of latent relations expressed in q . For example, for the question “*who invented the telephone?*”, the latent relation expressed is the concept “*invented by*”.⁴ A case $c_k := \{q_k, \mathcal{A}_k, p_k\}$

³ \mathcal{A} is a singleton if every question has a unique answer.

⁴Note that CBR-MRC does not require explicit annotations of question relation types, e.g., mapping to KB schema

is similar to q , if the latent relations expressed in q and q_k overlap, i.e., $\mathcal{R}_q \cap \mathcal{R}_{q_k} \neq \emptyset$. For the above example, a similar question may be “*Who was the inventor of the television?*”. We hypothesize that these latent relations are captured in dense representations obtained from language models; thus, relational overlap can be assessed via vector similarity scores. In our experiments, we use a pre-trained BERT (Devlin et al., 2019) model to encode questions and use the representation of the [CLS] token for similarity calculations.

To ensure that a question representation faithfully captures the expressed relations and is not influenced by the entities in the question, we replace mentions of all entities with a [MASK] token, e.g., “*who invented the telephone?*” becomes “*who invented the [MASK]?*”. Previous work (Soares et al., 2019) has shown this technique to be useful in learning entity-independent relational representations. In CBR-MRC, masking removes spurious similarities, where the presence of a common entity results in two relationally divergent questions being adjudged as similar, for e.g., “*who invented the telephone?*” and “*what is the telephone used for?*” should be considered dissimilar in our setup. In practice, we use T-NER (Ushio and Camacho-Collados, 2021), an automatic entity recognizer, to identify mentions of all entities in the question.

The similarity score between two masked queries is computed as the inner product between their normalized vector representations (cosine similarity). The representations for all questions in the casebase are pre-computed and cached. During inference, we retrieve the top- k similar cases for a query by performing a fast nearest-neighbor search over the cached representations.

Case retrieval during training. To minimize the number of retrieved cases with a low overlap in latent relations during training, we additionally employ two filtering heuristics — (1) we use a minimum threshold of 0.95 for the cosine similarity score to filter out dissimilar cases, and (2) we use a *wh-filter* to remove retrieved cases where the question keyword (e.g., “*who*”, “*where*”, “*when*”) present in the case question does not match the keyword present in the input query.

3.2 Case Reuse

We hypothesize that contexts associated with similar questions also share relational similarities. For relations. It instead leverages such semantics implicitly.

instance, Figure 2 shows the context surrounding the answer span of the retrieved case “*Graham Bell is credited with patenting...*” and the target context surrounding the gold answer “*Charles Babbage is credited with inventing...*” In CBR-MRC, we use the contextualized embeddings of the answer spans from retrieved cases to find this semantic overlap within the target context. For each candidate answer span in the context, a similarity score is explicitly computed with answer representations from the retrieved cases.

For efficient processing, we employ a dense retriever that uses contextualized answer spans from past cases as queries and searches for the target answer over the input passage embeddings. Given an input question q and an input passage p , we search for a text span $s \in \mathcal{S}_p$, the set of all possible spans in p , whose representation is the most similar to representations of answer spans in similar question contexts. We employ a standard BERT-base model as the passage encoder of the dense retriever. Due to the combinatorial complexity of considering all $s \in \mathcal{S}_p$, we follow Seo et al. (2019) and consider only a subset of candidate spans in p likely to be the answer. In particular, we consider a set of candidates $\mathcal{S}_{\text{cand}} \subseteq \mathcal{S}_p$ such that $\mathcal{S}_{\text{cand}}$ is composed of only entity mentions, date-time strings, numbers, and quoted strings that can be extracted from p . To increase span coverage, we include all n-grams up to length three. See A.1 for more details regarding the candidate span extraction process.

We now formulate the task to predict the answer \hat{a} to the question q from a context passage p given its relevant past cases. Let $c_k := \{q_k, \mathcal{A}_k, p_k\}$ be the k -th retrieved case for the given question q . For an answer $a_k \in \mathcal{A}_k$, let $\text{Enc}(a_k)$ be the contextualized answer span embedding pooled from the embeddings of the answer tokens in p_k . Similarly, for all candidate answers $s \in \mathcal{S}_{\text{cand}}$, let $\text{Enc}(s)$ be their contextualized embeddings pooled from the encoded tokens in context p . The predicted answer span based on the k -th retrieved case is given by

$$\hat{a}^{(k)} = \arg \max_{s \in \mathcal{S}_{\text{cand}}} \max_{a_k \in \mathcal{A}_k} \text{Enc}(s)^T \text{Enc}(a_k),$$

i.e., the candidate answer span most similar to the answer set \mathcal{A}_k based on similarity scores computed using inner products. At inference, the similarity scores of each candidate span are aggregated across the top- k retrieved cases by taking the maximum scores. The candidate span with the highest aggregated score is predicted as the answer.

3.2.1 Training

The passage encoder must be fine-tuned to maximize similarity scores between answer span embeddings of similar questions. We fine-tune a pre-trained BERT model with a contrastive loss (Khosla et al., 2020) function using answer and non-answer spans.

Specifically, consider a train question q_t , its set of answers \mathcal{A}_t , and its set of candidate spans $\mathcal{S}_{\text{cand},t} \subseteq \mathcal{S}_{p_t}$ derived from the passage p_t . We use the gold answer set \mathcal{A}_t as the set of positive candidate spans, and define the set of negative candidate spans as $\mathcal{S}_{\text{cand},t} \setminus \mathcal{A}_t$. Let $\text{sim}(\cdot, \cdot) := \text{Enc}(\cdot)^T \text{Enc}(\cdot)$ be the similarity function (dot product or cosine similarity) between vector representations. For the top- k similar questions as described in the previous section, we fine-tune the passage encoder with the following objective derived from the soft nearest-neighbor loss (Salakhutdinov and Hinton, 2007; Frosst et al., 2019):

$$\begin{aligned} \mathcal{L}_t = & -\log \sum_{s \in \mathcal{A}_t} \max_k \max_{a_k \in \mathcal{A}_k} \exp(\text{sim}(a_k, s)/\tau) \\ & + \log \sum_{s \in \mathcal{S}_{\text{cand},t} \setminus \mathcal{A}_t} \max_k \max_{a_k \in \mathcal{A}_k} \exp(\text{sim}(a_k, s)/\tau), \end{aligned}$$

where τ is a temperature hyperparameter. The loss for the entire training set is simply the mean of \mathcal{L}_t over all $q_t \in \mathcal{Q}$, the training set. Among the top- k retrieved cases, we pick the spans with the highest similarity scores. This loss function encourages the similarity scores of the gold answer spans with respect to the answers in the top- k cases to be higher than the similarity of other spans of texts in the input context that are not the answer.

4 Experiments

We evaluate CBR-MRC on two machine reading comprehension tasks — supervised answer span extraction and few-shot domain adaptation. For three MRC datasets, we use the training data to both construct the casebase as well as train our model and report performance on the in-domain test sets. We report the few-shot domain adaptation performance on two additional datasets.

4.1 Datasets

NaturalQuestions (NaturalQ) consists of over 323K questions and answers, representative of real queries people make on the Google search engine (Kwiatkowski et al., 2019). The answers are sourced from Wikipedia and annotated by

crowd workers, including both a long and short version. Our experiments utilize short answers as the gold answers and use the long answers as context. **NewsQA** has 120K crowd-sourced questions based on CNN news articles (Trischler et al., 2016). Annotation is obtained by two sets of crowd workers: one creating questions based on article highlights and the other answering them using spans from the full article. **BioASQ** is a large-scale biomedical dataset that features question-answer pairs crafted by domain experts (Tsatsaronis et al., 2015). Questions are linked to relevant scientific literature found on PubMed, with the abstracts of these articles serving as the context for MRC. **RelationExtraction (RelEx)** comprises question-answer pairs that express Wikidata relations between entities mentioned in Wikipedia articles (Levy et al., 2017). The dataset is derived from the WikiReading (Hewlett et al., 2016) slot-filling dataset through predefined templates.

We utilize the curated version of all the previously mentioned datasets from the MRQA shared task (Fisch et al., 2019), including annotations for the gold answer spans within their context. The gold context for constructing a case, as described in Section 3, is a short paragraph surrounding the gold answer span taken from the originally provided text. Additionally, note that we use BioASQ and RelEx to *only evaluate* in the few-shot setting and not train our models.

4.2 Experiment Setup

We refer the reader to Appendix A.1 for more details on the models used in this work.

Baselines. We compare CBR-MRC to BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2019), and SpanBERT (Joshi et al., 2020) as well as the SOTA baselines for our settings – BLANC (Seonwoo et al., 2020b) and MADE (Friedman et al., 2021).

Evaluation Metrics.⁵ We report exact match (EM) and F1 scores to measure model performance. EM calculates the percentage of predictions that match all tokens of any one of the correct answers for a given question. F1 score measures the token overlap between the prediction and the correct answer. As these metrics do not take into account the

⁵(1) All reported numbers are from single runs on the test set unless noted otherwise. (2) Hyperparameters used by our models were tuned on a development set and the details are included in the appendix.

Model	NewsQA		NaturalQ	
	EM	F1	EM	F1
BERT	50.11	65.07	64.48	76.39
ALBERT	51.18	66.02	63.81	75.89
RoBERTa	52.36	67.28	66.33	78.54
MADE	<u>56.55</u>	72.12	68.86	80.09
SpanBERT	52.85	67.93	66.60	78.31
BLANC	55.52	<u>70.31</u>	68.33	80.04
SpanBERT _{LARGE}	53.84	69.06	69.14	80.66
BLANC _{LARGE}	57.40	72.36	<u>70.59</u>	<u>81.99</u>
CBR-MRC	64.95	69.17	82.12	83.10

Table 1: **MRC performance on NewsQA and NaturalQ.** CBR-MRC outperforms all baselines in finding the exact match of the answers. On NewsQA, our model shows +8.4 EM over MADE but has lower F1 which indicates CBR-MRC focuses on getting the right answer rather than overfitting to the answer tokens.

surrounding context of the answer span (i.e., the supporting evidence), we also report Span-EM and Span-F1, following Seonwoo et al. (2020b), which consider overlapping indices between the predicted and gold answer spans and exclude answers with irrelevant supporting text.

4.3 Results

We report our main MRC results on NewsQA and NaturalQ in Table 1. We find that CBR-MRC consistently outperforms all baselines. On NaturalQ, CBR-MRC achieves the state-of-the-art, outperforming the next best model (BLANC_{LARGE}) by 11.53 EM. Similarly, on NewsQA, CBR-MRC shows an improvement of 12.56 EM over BLANC_{LARGE}, and 8.4 EM over the next best model (MADE). Note that the best baseline model (BLANC_{LARGE}) has three times the number of parameters as CBR-MRC. Furthermore, our model shows performance gains in EM but a lower F1 which indicates CBR-MRC focuses on getting the correct answer rather than over-fitting to the answer tokens.

4.3.1 Supporting Evidence Identification

It is becoming increasingly important to train models to predict the *right* answer for the *right* reasons. Since CBR-MRC identifies the answer span by comparing it with multiple answer-containing contexts of similar cases, we hypothesize that it can cut through the noise and identify the correct occurrence of the span that *exactly* answers the question. Our main baseline is BLANC (Seonwoo et al., 2020b), which trains an auxiliary model specifically for context prediction.

Dataset	Model	Span-EM	Span-F1
NaturalQ	BERT	60.63	72.92
	ALBERT	60.31	72.66
	RoBERTa	62.59	75.07
	SpanBERT	62.71	75.16
	BLANC	64.57	76.99
	SpanBERT _{LARGE}	65.28	77.62
	BLANC _{LARGE}	66.75	79.04
	CBR-MRC	74.64	83.16
NewsQA	BERT	45.53	59.18
	ALBERT	46.54	60.12
	RoBERTa	47.43	61.36
	SpanBERT	48.04	62.26
	BLANC	50.60	64.39
	SpanBERT _{LARGE}	49.03	63.43
	BLANC _{LARGE}	52.39	66.48
	CBR-MRC	54.07	68.21
NaturalQ*	RoBERTa	60.12	65.99
	SpanBERT	57.63	63.47
	BLANC	61.43	67.07
	CBR-MRC	71.30	81.04

Table 2: **Identifying the correct supporting evidence.** We report Span-F1 and Span-EM scores for CBR-MRC and baseline models, indicating how often the span with the correct supporting evidence is identified versus selecting either the incorrect answer or a spurious occurrence of the correct answer. We report baseline numbers for BLANC as reported in their paper. NaturalQ* is a subset of NaturalQuestions containing questions with answers having at least two mentions within the context.

Table 2 highlights the performance gains of CBR-MRC compared to BLANC and several strong baselines. Our method shows improvements on NaturalQ by 7.89 Span-EM and 4.12 Span-F1 points in predicting the span indices. On NewsQA, it improves performance by 1.68 Span-EM and 1.73 Span-F1 points compared to BLANC. To determine the gains due to predicting the answer span with the correct supporting evidence, we follow the approach in BLANC and evaluate performance on the NaturalQuestions (NaturalQ*) subset that contains questions with answers that have at least two mentions within the context. Our method outperforms BLANC on this subset by 6.73 and 4.05 points for Span-EM and Span-F1, respectively.

4.3.2 Robustness to Lexical Diversity

CBR-MRC performs answer extraction by explicitly comparing span similarities of the target context with the context spans in the filtered casebase. Purely parametric methods can also be seen as following a similar mechanism, albeit performing the




Figure 3: **Robustness to lexical diversity in passages.** Lexical diversity of latent relation clusters is the number of unique tokens in passages seen at training for that latent relation cluster (or question type). We plot F1 performance (averaged over 6 clusterings varying cut thresholds) on 8 buckets of increasing lexical diversity for the latent relations seen at training. CBR-MRC shows a drop in performance only up to 4.10 points, while BLANC and MADE show drops of up to 15.38 and 11.80 points, respectively.

two steps implicitly. These models, therefore, must rely only on the parameter values learned at training in order to handle several relational semantics, which may each be expressed in contexts with varying levels of diversity in language form. In practice, we find that some relations are indeed expressed with more lexical diversity than others (Appendix A.3). In Figure 3, we investigate how CBR-MRC performs at inference under varying amounts of lexical diversity observed in passage contexts at *training*. We first cluster questions according to their embeddings and then compute n-gram statistics (Li et al., 2015; Dušek et al., 2020; Tevet and Berant, 2020) across passage contexts found in each cluster to measure diversity. We use NaturalQuestions as a case study for our analysis and compare performance with BLANC and MADE, two fully-parametric baselines.

Clustering questions by latent relations. We cluster questions in the training set based on the latent relations they encode using hierarchical agglomerative clustering (HAC) (Krishnamurthy et al., 2012) over similarity scores. For a fair evaluation, we compute 6 different flat clusterings of varying cluster tightness, i.e., how similar each question is within a cluster. For each clustering, we compute the per-cluster lexical diversity as the count of the number of unique tokens in the set of passage contexts across questions in the cluster, normalized by the cluster size. We then cluster the

Model	BioASQ		RelEx	
	EM	F1	EM	F1
MADE _{ZERO}	43.3	62.9	72.0	84.9
MADE _{FT}	62.5	74.5	82.3	90.0
CBR-MRC	62.4	64.6	87.8	88.2

Table 3: **Few-shot domain adaptation.** We compare CBR-MRC with two variants of the MADE adapter networks, which are also trained on NaturalQuestions. MADE_{ZERO} demonstrates zero-shot performance, while MADE_{FT} is further fine-tuned on 256 samples for each dataset. CBR-MRC is evaluated by simply adding the samples into the casebase without any fine-tuning.

test set by assigning each question to the cluster label of its top-1 similar train question for each of the 6 clusterings, resulting in 6 *test* clusterings (see Appendix A.4 for more details).

Results. Figure 3 shows the mean F1 scores between CBR-MRC and the baselines on test queries for varying values of lexical diversity observed at training for the same latent relations as the test queries. We group clusters by lexical diversity scores into buckets and report the aggregated F1 scores for each bucket.⁶ On average, the performance gap between CBR-MRC and baselines tends to increase as lexical diversity increases. The min-max F1 difference, which captures the drop in performance as lexical diversity increases, is 4.10 points for CBR-MRC while BLANC and MADE show significantly larger differences of 15.38 and 11.80 points, respectively. We posit that as diversity in contexts for the same latent relation increases, it is harder for fully parametric models to learn parameters that are suitable for all lexical variations. On the other hand, the semi-parametric nature of CBR-MRC allows it to partially shift this burden to test time and ground itself using only a few retrieved cases that are explicitly compared.

4.3.3 Few-shot Domain Adaptation

It is important for models to generalize to unseen datasets. We test CBR-MRC trained on NaturalQ on two diverse out-of-domain datasets - BioASQ for MRC (Tsatsaronis et al., 2015) and RelEx for relation extraction (Levy et al., 2017). We compare our model with the MADE adapter networks (Friedman et al., 2021) that are trained on NaturalQuestions, and further fine-tuned on 256 samples from

⁶(1) We use k -means over lexical diversity scores of all clusters to construct $B = 8$ buckets. (2) We report the F1 performance on all 6 clusterings in Appendix A.5.

Cases (k)	1	5	10	20
EM	82.03	82.12	77.95	77.39
F1	82.13	83.10	78.00	78.45

Table 4: **Performance with different number of case retrievals** $k \in \{1, 5, 10, 20\}$. Both EM and F1 display the same trend. Both performance metrics increase as k increases from 1 to 5. However, larger sets of retrievals ($k > 10$) reduce performance to levels worse than even $k = 1$, which indicates the introduction of noise.

the target dataset. We follow the same practice in (Friedman et al., 2021) and evaluate CBR-MRC on a held-out subset of 400 samples. For both datasets, we use the trained model obtained from the experiment on the NaturalQuestions dataset, and the casebase for each dataset is built from the same training set that MADE also uses.

Table 3 presents the performance of CBR-MRC and MADE adapters. Our model outperforms the zero-shot adapter network (MADE_{ZERO}) by simply putting some samples from the target dataset into the casebase. When compared to the fine-tuned adapter (MADE_{FT}), our model gains 5.5 EM points on RelEx while being comparable to MADE_{FT} on BioASQ. These results confirm that CBR-MRC can transfer to a new domain by simply collecting a handful of sample cases, which is significantly less expensive than model fine-tuning.

We notice a drop in F1 scores on both datasets. F1 scores evaluate the partial matching of answers. The inference procedure used by previous work, including our baselines, independently predicts the start and end spans of the answer, which encourages the inclusion of partial matches. In CBR-MRC, on the other hand, the entire candidate span is considered when making a prediction by comparing against the representations of gold spans from the casebase. This results in higher similarity scores for candidate spans that map to the same answer type, rather than matching the correct answer partially, which explains lower F1 scores for CBR-MRC.

4.3.4 Effect of Retrieval Quantity

We explore how the number of retrieved cases k affects the reading comprehension performance, via varying k on the NaturalQuestions dataset. We find that while increasing k does enrich the set of contexts for referencing in the reuse step, too many retrievals may introduce irrelevant cases that degrade performance.

4.3.5 Effect of Modeling

We analyze CBR-MRC performance with different base models and show that CBR-MRC can improve upon them. We run the experiments on NQ with the same set of hyper-parameters.

Model	EM	F1	Span-EM	Span-F1
BERT	64.48	76.39	60.63	72.92
CBR-MRC_{BERT}	82.12	83.1	74.64	83.16
DeBERTa	43.11	74.07	36.03	63.7
CBR-MRC_{DeBERTa}	80.30	82.22	76.09	80.39

Table 5: **Performances of CBR-MRC variants on NQ.** We compare two variants of CBR-MRC with BERT and DeBERTa as its base model.

5 Conclusion

We present CBR-MRC, a semi-parametric model for machine reading comprehension that is simple, accurate, and interpretable. Our model stores a collection of cases, retrieves the most relevant cases for a given test question, and then explicitly reuses the reasoning patterns encoded in the embeddings of these cases to predict an answer. We show that our model performs well for both extracting answers and identifying supporting evidence on several MRC tasks compared to fully-parametric baselines. We also demonstrate the ability of our model to transfer to new domains with limited labeled data. Finally, we analyze our model under varying conditions of lexical diversity and find that it is robust to high lexical diversity, whereas fully-parametric models show a drop in performance.

6 Limitations

CBR-MRC, and the CBR framework more generally, relies on the existence of past cases to make predictions for a new problem. This can pose a challenge for composite questions, which require multi-hop reasoning, since the likelihood of enumerating each algebraic combination of reasoning patterns in the casebase is impractical. Models, thus, may only be able to match a portion of the question, resulting in partially correct reasoning. To address these limitations, future work could explore methods to explicitly encourage such compositional generalization, such as question decomposition with a recursive application of CBR.

Ethics Statement

The objective of our research and the proposed methodology is to enhance performance on the Ma-

chine Reading Comprehension task. The datasets utilized in this study have been extensively employed in previous research and, as far as our knowledge extends, are not associated with any privacy or ethical concerns.

Acknowledgements

We thank members of UMass IESL for helpful discussions and feedback. This work was supported in part by Amazon Digital Services, in part by IBM Research AI through the AI Horizons Network, in part by the Chan Zuckerberg Initiative under the project "Scientific Knowledge Base Construction", and in part using high-performance computing equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Technology Collaborative. Any opinions, findings, and conclusions / recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor(s).

References

- Abhijeet Awasthi, Soumen Chakrabarti, and Sunita Sarawagi. 2023. Structured case-based reasoning for inference-time adaptation of text-to-sql parsers. *arXiv preprint arXiv:2301.04110*.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2016. *Ms marco: A human generated machine reading comprehension dataset*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.
- Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. 2022. On the relation between sensitivity and accuracy in in-context learning. *arXiv preprint arXiv:2209.07661*.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, et al. 2022. Binding language models in symbolic languages. *arXiv preprint arXiv:2210.02875*.

- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. *arXiv preprint arXiv:1808.07036*.
- Rajarshi Das, Ameya Godbole, Shehzaad Dhuliawala, Manzil Zaheer, and Andrew McCallum. 2020a. A simple approach to case-based reasoning in knowledge bases. In *AKBC*.
- Rajarshi Das, Ameya Godbole, Nicholas Monath, Manzil Zaheer, and Andrew McCallum. 2020b. Probabilistic case-based reasoning for open-world knowledge graph completion. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4752–4765, Online. Association for Computational Linguistics.
- Rajarshi Das, Ameya Godbole, Ankita Naik, Elliot Tower, Robin Jia, Manzil Zaheer, Hannaneh Hajishirzi, and Andrew McCallum. 2022. Knowledge base question answering by case-based reasoning over subgraphs. In *ICML*.
- Rajarshi Das, Manzil Zaheer, Dung Ngoc Thai, Ameya Godbole, Ethan Perez, Jay Yoon Lee, Lizhen Tan, Lazaros Polymenakos, and Andrew McCallum. 2021. Case-based reasoning for natural language queries over knowledge bases. In *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Andrew Drozdov, Nathanael Schärli, Ekin Akyürek, Nathan Scales, Xinying Song, Xinyun Chen, Olivier Bousquet, and Denny Zhou. 2022. Compositional semantic parsing with large language models. *arXiv preprint arXiv:2209.15003*.
- Alexander Dunn, John Daggdelen, Nicholas Walker, Sanghoon Lee, Andrew S Rosen, Gerbrand Ceder, Kristin Persson, and Anubhav Jain. 2022. Structured information extraction from complex scientific text with fine-tuned large language models. *arXiv preprint arXiv:2212.05238*.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2020. Evaluating the state-of-the-art of end-to-end natural language generation: The e2e nlg challenge. *Computer Speech & Language*, 59:123–156.
- 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 2nd Machine Reading for Reading Comprehension (MRQA) Workshop at EMNLP*.
- Dan Friedman, Ben Dodge, and Danqi Chen. 2021. Single-dataset experts for multi-dataset qa. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Nicholas Frosst, Nicolas Papernot, and Geoffrey E. Hinton. 2019. Analyzing and improving representations with the soft nearest neighbor loss. In *International Conference on Machine Learning*.
- Alessio Gravina, Federico Rossetto, Silvia Severini, and Giuseppe Attardi. 2018. Cross attention for selection-based question answering. In *NL4AI@AI*IA*.
- John A Hartigan and Manchek A Wong. 1979. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend.
- 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.
- Roshni Iyer, Thuy Vu, Alessandro Moschitti, and Yizhou Sun. 2023. Question-answer sentence graph for joint modeling answer selection. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 968–979, Dubrovnik, Croatia. Association for Computational Linguistics.
- Gautier Izacard and Edouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *EACL*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673.
- Janet L. Kolodner. 1983. Maintaining organization in a dynamic long-term memory. *Cognitive Science*, 7(4):243–280.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2017. The narrativeqa reading comprehension challenge.

- Akshay Krishnamurthy, Sivaraman Balakrishnan, Min Xu, and Aarti Singh. 2012. Efficient active algorithms for hierarchical clustering. In *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML'12*, page 267–274, Madison, WI, USA. Omnipress.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*.
- 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. *CoRR*, abs/1909.11942.
- David B Leake. 1996. Case-based reasoning: experiences, lessons, and future directions.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Neurips*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. *arXiv preprint arXiv:2206.14858*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Tianle Li, Xueguang Ma, Alex Zhuang, Yu Gu, Yu Su, and Wenhui Chen. 2023. Few-shot in-context learning for knowledge base question answering. *arXiv preprint arXiv:2305.01750*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*.
- Sewon Min, Victor Zhong, Richard Socher, and Caiming Xiong. 2018. Efficient and robust question answering from minimal context over documents.
- Panupong Pasupat, Yuan Zhang, and Kelvin Guu. 2021. Controllable semantic parsing via retrieval augmentation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7683–7698, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- 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, Seattle, Washington, USA. Association for Computational Linguistics.
- Edwina L Rissland. 1983. Examples in legal reasoning: Legal hypotheticals. In *IJCAI*, pages 90–93.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2021. Learning to retrieve prompts for in-context learning. *arXiv preprint arXiv:2112.08633*.
- Ruslan Salakhutdinov and Geoff Hinton. 2007. Learning a nonlinear embedding by preserving class neighbourhood structure. In *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, pages 412–419, San Juan, Puerto Rico. PMLR.
- Roger C. Schank. 1983. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge University Press, USA.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension.
- Minjoon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur P Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. Real-time open-domain question answering with dense-sparse phrase index. In *ACL*.

- Yeon Seonwoo, Ji-Hoon Kim, Jung-Woo Ha, and Alice Oh. 2020a. Context-aware answer extraction in question answering. In *EMNLP*.
- Yeon Seonwoo, Ji-Hoon Kim, Jung-Woo Ha, and Alice Oh. 2020b. Context-aware answer extraction in question answering. In *EMNLP*.
- Nan Shao, Yiming Cui, Ting Liu, Shijin Wang, and Guoping Hu. 2020. [Is graph structure necessary for multi-hop question answering?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *ACL*.
- Guy Tevet and Jonathan Berant. 2020. Evaluating the evaluation of diversity in natural language generation. *arXiv preprint arXiv:2004.02990*.
- Dung Thai, Srinivas Ravishankar, Ibrahim Abdelaziz, Mudit Chaudhary, Nandana Mihindukulasooriya, Tahira Naseem, Rajarshi Das, Pavan Kapanipathi, Achille Fokoue, and Andrew McCallum. 2022. Cbr-ikb: A case-based reasoning approach for question answering over incomplete knowledge bases. *arXiv preprint arXiv:2204.08554*.
- Mokanarangan Thayaparan, Marco Valentino, and André Freitas. 2020. [A survey on explainability in machine reading comprehension](#).
- Adam Trischler, Tong Wang, Xingdi (Eric) Yuan, Justin D. Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016. [Newsqa: A machine comprehension dataset](#).
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R. Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, Yannis Almirantis, John Pavlopoulos, Nicolas Baskiotis, Patrick Gallinari, Thierry Artières, Axel-Cyrille Ngonga Ngomo, Norman Heino, Éric Gaussier, Liliana Barrio-Alvers, Michael Schroeder, Ion Androutsopoulos, and Georgios Paliouras. 2015. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, 16.
- Ming Tu, Kevin Huang, Guangtao Wang, Jing Huang, Xiaodong He, and Bowen Zhou. 2019. [Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents](#).
- Asahi Ushio and Jose Camacho-Collados. 2021. T-NER: An all-round python library for transformer-based named entity recognition. In *EACL (demo)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Neurips*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Semih Yagcioglu, Aykut Erdem, Erkut Erdem, and Nazli Ikizler-Cinbis. 2018. [Recipeqa: A challenge dataset for multimodal comprehension of cooking recipes](#).
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. [WikiQA: A challenge dataset for open-domain question answering](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018, Lisbon, Portugal. 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, Brussels, Belgium. Association for Computational Linguistics.
- Zonglin Yang, Xinya Du, Erik Cambria, and Claire Cardie. 2023. End-to-end case-based reasoning for commonsense knowledge base completion. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3491–3504.
- Deming Ye, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, and Maosong Sun. 2019. [Multi-paragraph reasoning with knowledge-enhanced graph neural network](#).
- Lei Yu, Karl Moritz Hermann, Phil Blunsom, and Stephen Pulman. 2014. [Deep learning for answer sentence selection](#).
- Changchang Zeng, Shaobo Li, Qin Li, Jie Hu, and Jianjun Hu. 2020. [A survey on machine reading comprehension: Tasks, evaluation metrics and benchmark datasets](#).
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A Appendix

A.1 Experiments: Implementation Details

All CBR-MRC models employ the BERT-base architecture and the cased vocabulary. We use a fixed set of hyperparameters to train all our models. More specifically, we use an Adam optimizer with a learning rate of $2e - 5$, epsilon $1e - 8$, and max gradient norm value of 5.0. We set the number of warm-up steps to be 8% of the training data size. All models are trained for 10 epochs, which takes 2-4 days on a single RTX-8000 GPU each, depending on the size of the dataset. The standard models we report have parameter counts of 108M (BERT, MADE, SpanBERT, BLANC), 17M (ALBERT), 124M (RoBERTa), and the large variants have a parameter count of 333M.

Our approach follows prior QA work (Seo et al., 2019) and selects a broad set of spans that may be answers to factoid-type questions, such as entity mentions, date-time strings, numbers, and quoted strings. We also include all n-grams up to three words. For date-time string extraction, we use the *datefinder* Python package, and for entity mentions, we use the T-NER library (Ushio and Camacho-Collados, 2021).

A.2 Experiments: Dataset Statistics

For our experiments, we use the curated versions of NaturalQuestions, NewsQA, BioASQ, and RelEx datasets from the MRQA shared task (Fisch et al., 2019). Table 6 presents the number of examples present in each of the datasets used in the paper.

Dataset	Train	Dev
NaturalQ	104,071	12,836
NewsQA	74,160	4,212
BioASQ	-	1,504
RelEx	-	2,948

Table 6: Dataset statistics

A.3 Lexical Diversity Analysis: Distribution in NaturalQuestions

We evaluate the distribution of lexical diversity in the NaturalQuestions dataset. Out of the set of test questions under consideration in our case study on lexical diversity, we find that 7.46%, or 504, unique (question, answer) pairs contain more than one context, which is used to extract the answer, highlighting the relevance of lexical diversity as a

consideration for models to address. A question in this set may contain up to 4 different contexts that can answer the same exact question. In Table 8, we show examples from the test set of questions that can be answered using multiple, diverse contexts.

A.4 Lexical Diversity Analysis: Clustering questions by latent relations

We cluster questions in the training set based on the latent relations they contain using hierarchical agglomerative clustering (HAC) (Krishnamurthy et al., 2012) over masked similarity scores⁷. Importantly, since we do not have ground-truth latent relations for questions, we run our analysis on *multiple* flat clusterings, where each cluster in a clustering should represent the same semantic relations. To obtain flat clusters in a fair manner, we cut the HAC tree using C thresholds, computed using k -means (Hartigan and Wong, 1979) over similarity scores across all questions in the training set. This yields C clusterings with increasing levels of cluster tightness, i.e., how similar each question is in a cluster. We set C to 6 and lower-bound the similarity scores to 0.9 to reduce noise in the clusters. For each clustering, we then compute the per-cluster lexical diversity as the count of the number of unique tokens in the set of passage contexts across questions in the cluster, normalized by the cluster size. Finally, we cluster the test set by assigning each question to the cluster label of its top-1 similar train question for each of the C clusterings, resulting in 6 test clusterings. In our analysis, we retain only unique (q, \mathcal{A}, p) triples resulting in 7,213 instances from an original set of 12,836. We additionally drop 98 questions that cannot be answered by any of the three models under consideration, resulting in a final set of 7,115 test questions.

A.5 Lexical Diversity Analysis: F1 Performance

Figure 4 shows the absolute F1 performance of CBR-MRC and the two baselines (BLANC and MADE) on each clustering separately. Note that Figure 3 is a reduction from this, in that it takes the average difference between the F1 score of CBR-MRC and the baselines across the 6 clusters.

⁷We restrict the number of nearest-neighbors to 20 for each question.




Figure 4: **F1 performance v/s lexical diversity in passage contexts on clusterings with decreasing levels of cluster-tightness.** We cluster questions by latent relations in the training set using HAC using 6 cut thresholds and assign each test question to one of these train clusters. We then bucket the clusters by lexical diversity scores and compute F1 performance of each bucket for each of the 6 clusterings shown.

A.6 Does CBR-MRC have the correct inductive bias?

Text representations obtained from pre-trained language models have shown the ability to capture the latent semantics of text spans. In this experiment, we evaluate CBR-MRC when the parameters are set to the pre-trained BERT-base model *without* fine-tuning. As shown in Table 7, we find that our model with no fine-tuning performs surprisingly well, indicating that CBR-MRC does have the correct inductive bias. However, fine-tuning with our proposed contrastive loss leads to a significant increase in performance.

Model	NewsQA	NaturalQ
BERT	50.11	64.48
SpanBERT	52.85	66.60
CBR-MRC_{PRE-TRAINED}	18.63	46.07
CBR-MRC_{FINE-TUNED}	64.95	82.12

Table 7: **Pre-trained CBR-MRC performance.** We report the performance (EM) of CBR-MRC using both a pre-trained passage encoder as well as an encoder fine-tuned with the contrastive loss as described in §3. We additionally include scores from reasonable, *fine-tuned* baselines for comparison.

Question	Answer	Context
when does season 5 of ruby come out	October 14, 2017	<p>[...] Four seasons , referred to as “ Volumes ” , have been released , with a fifth currently ongoing since its premiere on October 14 , 2017 . As of [...]</p> <p>[...] The fifth Volume premiered on October 14 , 2017 , a date which was first announced at the RTX Austin 2017 event . The episodes [...]</p> <p>[...] <Table><Tr><Th colspan="2"> Season </Th><Th colspan="2"> Episodes </Th><Th colspan="2"> Originally aired </Th> [...] <Td colspan="1"> 5 </Td><Td colspan="2"> 14 </Td><Td colspan="1"> October 14 , 2017 (2017 - 10 - 14) </Td><Td> [...]</p>
who wrote most of the declaration of independence	Thomas Jefferson	<p>[...] Committee of Five ” to draft a declaration , consisting of John Adams of Massachusetts , Benjamin Franklin of Pennsylvania , Thomas Jefferson of Virginia , Robert R. Livingston of New York , and Roger Sherman of Connecticut . The committee [...]</p> <p>[...] John Adams , a leader in pushing for independence , had persuaded the committee to select Thomas Jefferson to compose the original draft of the document , which Congress edited to produce the final version [...]</p> <p>[...] The source copy used for this printing has been lost and may have been a copy in Thomas Jefferson ’s hand . Jefferson ’s original draft is preserved at the Library of Congress , complete with changes made by John Adams [...]</p>
who plays the mom on the tv show mom	Allison Janney	<p>[...] working as a waitress and attending Alcoholics Anonymous meetings . Her mother Bonnie Plunkett (Allison Janney) is also a recovering drug and alcohol addict . Christy ’s daughter [...]</p> <p>[...] It stars Anna Faris and Allison Janney in lead roles as dysfunctional daughter / mother duo Christy and Bonnie Plunkett [...]</p> <p>[...] her relationship with her mother Bonnie [...] Allison Janney as Bonnie Plunkett : Christy ’s mother , a joyful if cynical recovering addict who is now grateful with life . She tries [...]</p>
when was the us department of homeland security created	November 25 , 2002	<p><Tr><Th> Formed </Th><Td> November 25 , 2002 ; 15 years ago (2002 - 11 - 25) </Td></Tr></p> <p>[...] <Th colspan="2"> Agency overview </Th> </Tr><Tr><Th> Formed </Th><Td> November 25 , 2002 ; 15 years ago (2002 - 11 - 25) </Td> [...]</p> <p>[...] The Department of Homeland Security was established on November 25 , 2002 , by the Homeland Security Act of 2002 . It was intended to consolidate U.S. [...]</p>
who has won the eurovision song contest the most times	Ireland	<p>[...] Ireland has finished first seven times , more than any other country [...]</p> <p>[...] <Table><Tr><Th> Wins </Th><Th> Country </Th><Th> Years </Th></Tr><Tr><Td> 7 </Td><Td> Ireland </Td><Td> 1970 , 1980 , 1987 , 1992 , 1993 , 1994 , 1996 </Td></Tr><Tr><Td> 6 </Td><Td> Sweden </Td><Td> 1974 , 1984 , 1991 , 1999 , 2012 , 2015 </Td></Tr><Tr><Td> [...] </Td></Tr></p> <p>[...] The country with the highest number of wins is Ireland , with seven . The only person to have won more than once [...]</p>

Table 8: **Lexical Diversity in Natural Questions.** We show here unique question-answer pairs that can be answered using diverse contexts.