

Feeding What You Need by Understanding What You Learned

Xiaoqiang Wang^{1*}, Bang Liu^{2*†}, Fangli Xu³, Bo Long⁴, Siliang Tang^{1‡} and Lingfei Wu^{5‡}

¹Zhejiang University, ²Université de Montréal & Mila, ³Squirrel AI Learning

⁴JD.COM, ⁵JD.COM Silicon Valley Research Center

{xq.wang, siliang}@zju.edu.cn

bang.liu@umontreal.ca, lili@yixue.us

{bo.long, lingfei.wu}@jd.com

Abstract

Machine Reading Comprehension (MRC) reveals the ability to understand a given text passage and answer questions based on it. Existing research works in MRC rely heavily on large-size models and corpus to improve the performance evaluated by metrics such as Exact Match (EM) and F_1 . However, such a paradigm lacks sufficient interpretation to model capability and can not efficiently train a model with a large corpus. In this paper, we argue that a deep understanding of model capabilities and data properties can help us feed a model with appropriate training data based on its learning status. Specifically, we design an MRC capability assessment framework that assesses model capabilities in an explainable and multi-dimensional manner. Based on it, we further uncover and disentangle the connections between various data properties and model performance. Finally, to verify the effectiveness of the proposed MRC capability assessment framework, we incorporate it into a curriculum learning pipeline and devise a Capability Boundary Breakthrough Curriculum (CBBC) strategy, which performs a model capability-based training to maximize the data value and improve training efficiency. Extensive experiments demonstrate that our approach significantly improves performance, achieving up to an 11.22% / 8.71% improvement of EM / F_1 on MRC tasks.

1 Introduction

A competency assessment is used to measure someone's capabilities against the requirements of their job (Cheryl Lasse, 2020). In other words, it measures how (behaviors) someone does the what (task or skill). By showing what it looks like to be good

in a job, a competency assessment can effectively empower and engage people who want to understand and improve their unique skill profile and tell them what action to take to close any gaps so they can own their development. A natural question that arises here is: can we develop competency assessments for machine learning models to help better understand their capabilities and improve their performance on a given task?

In this paper, we focus on competency assessments for machine reading comprehension (MRC). MRC is a core task in natural language processing (NLP) that aims to teach machines to understand human languages and answer questions (Zeng et al., 2020; Chen et al., 2019). Recently, pre-trained language models (LMs) (Mikolov et al., 2013; Peters et al., 2018; Pennington et al., 2014; Devlin et al., 2018) have demonstrated superior performance on MRC tasks by pre-training on large amounts of unlabeled corpus and fine-tuning on MRC datasets. The performance is usually evaluated by metrics such as Exact Match (EM) and F_1 score, lacking interpretability to the capabilities of a model. That is to say, such metrics only tell how good a model performs overall on a specific dataset, but uncovers little about what specific skills a model has gained and the level of each skill.

We argue that the value of each data sample varies during the training process of a model, depending on the model's current capabilities. A deep understanding of the model's intrinsic capabilities can help us estimate each data sample's learning value and better manage the training process to improve the training efficiency. Take student learning as an example. There is no doubt that a college student can do well in solving primary school level exercises, but such exercises do not help improve his/her ability. On the contrary, a primary school student can not acquire knowledge efficiently from college-level exercises due to the big gap between his/her current knowledge or skills and the require-

*Equal contribution.

†Canada CIFAR AI Chair.

‡Corresponding authors.

Capability c_i	Subclasses	Metrics m_i^j	
Reading words	Recognize vocabulary	Intra-n (Gu et al., 2018) and Ent-n (Serban et al., 2017).	
	Recognize function words	Frequency of function words.	
Reading sentences	Recognize grammaticality	Height and width of a question’s constituency parsing tree.	
	Readability	Readability metrics.	
Understanding words	Arithmetic operation	Frequency of numerical expressions (CD tag).	
	Logical operation	Frequency of logically qualified words such as <i>any</i> , <i>all</i> and <i>every</i> .	
Understanding sentences	Linguistic reasoning	Syntactic and semantic overlap	BLEU-n (Papineni et al., 2002), BERTScore (Zhang* et al., 2020) and MoverScore (Zhao et al., 2019) between the context and question.
		Coreference resolution	Frequency of personal and possessive pronouns, such as PRP and PRP\$ tags.
		Con/Dis-junction, negation	Frequency of coordinating junctions, such as <i>and</i> and <i>or</i> .
	Factual reasoning	Causality	Frequency of causal clauses, such as <i>because</i> and <i>the reason for</i> .
		Spatial/Temporal	Frequency of spatial/temporal expressions, such as <i>before</i> , <i>after</i> , and <i>in front of</i> .
	Multi-hop reasoning	Number of supporting evidences.	

Table 1: Example set of MRC model capabilities $\{c_i\}$ and corresponding metrics $\{m_i^j\}$. See Section 2.1 for details.

ment to solve the exercises. We need to measure the ability of a student and then choose the appropriate exercises accordingly.

Existing works on interpreting MRC model capabilities concentrate on analyzing a model’s behavior with adversarial data (Jia and Liang, 2017), or defining the prerequisite skills to solve a specific dataset (Sugawara et al., 2017). However, these works require costly human annotation efforts or ignore the fact that model capabilities change during the training progresses.

In this paper, we design a competency assessment framework for MRC model capabilities. Specifically, we define four major capability dimensions for understanding text and solving MRC tasks: *reading words*, *reading sentences*, *understanding words* and *understanding sentences*, which are inspired by the computational models of human text comprehension in psychology (Kintsch, 1988). Based on the proposed framework, we can obtain a more appropriate assessment of model capabilities than the regular EM or F_1 metrics.

Furthermore, we analyze a variety of data properties to estimate how good a model has to be to solve a specific data sample and identify the relationships between data properties and model performance. This greatly helps us estimate the learning value of each training sample. Based on this analysis, we discover a very common situation: if a sample is scored as a high value in one capability dimension, the other dimensions have the same tendency as well, and vice versa. To alleviate these inevitable correlations, we utilize data whitening to quantify each sample as four capability-specific scores in a decorrelated fashion.

Finally, to reveal the potential usefulness of our proposed competency assessment framework and

evaluate its efficiency, we employ it in a curriculum learning pipeline and design a Capability Boundary Breakthrough Curriculum (CBBC) strategy. This strategy gradually enlarges the model capability boundary by picking samples around the boundary and breaking through it. Based on the analysis of model capabilities and data properties, we feed the model with training samples that are neither too simple nor too hard for it to solve. Extensive experiments on four benchmark datasets demonstrate that our approach significantly improves the performance of existing MRC models, achieving up to an 11.22% / 8.71% improvement of EM / F_1 on MRC tasks. These results show the reasonability and effectiveness of our proposed assessment framework and provide a widely applicable measurement for dealing with the relationship between the model capability and data quality.

2 Competency Assessment of MRC Capabilities

In this section, we first formulate our competency assessment framework of 4-dimensional MRC capabilities. Based on this framework, the data properties related to each capability dimension are described as corresponding heuristic metrics. We then uncover the relationship between various data properties and model performance in a decorrelated manner, quantifying each sample as 4-dimensional capability-specific scores with little correlation.

2.1 Assessment Framework Formulation

Human text comprehension has been studied in psychology for a long time. Constructionist, landscape model, and computational architectures have been proposed for such comprehension (McNamara and Magliano, 2009). Among them, the construction-

Context: James is a trouble making Turtle . One day, James went to the grocery store and pulled all the pudding off the shelves and ate two jars. Then he walked to the fast food restaurant and ordered 15 bags of fries.

Q1: Who is the trouble making turtle?

A1: James

Q2: Where did James go after he went to the grocery store?

A2: A fast food restaurant

Required capabilities: syntactic matching, temporal relation, semantic overlap

Figure 1: Two example questions Q1 and Q2 with different difficulties require different capabilities.

integration (CI) model (Kintsch, 1988) is one of the most basic and influential theories. The CI model assumes three different representation levels (surface structure, textbase, and situation model) and a two-step process (construction and integration) to understand text comprehensively. It first constructs the propositions (*i.e.* textbase) from the raw textual input (*i.e.* surface structure), then integrates the local connections into a globally coherent representation (*i.e.* situation model). Based on this situation model, a given text is understood comprehensively and can even be grounded to other modalities. Inspired by the two-step process of the CI model, we formulate our assessment framework by 4-dimensional capabilities as summarized in Table 1. We sketch out the meaning of each MRC capability $\{c_i\}_{i=1}^4$ and highlight some heuristic metrics $\{m_i^j\}_{j=1}^{n(i)}$ (where $n(i)$ means the number of metrics to measure a sample’s learning value to capability c_i) as follows.

Reading words. To formulate the surface structure of the CI model in our framework, we first highlight the text representation at the verbal or linguistic level. Theoretically, the units at the linguistic level are the words that make up the text and the hierarchical sentence constituents to which these words belong. Empirically, Sugawara et al. (2018) has shown that some questions are answered correctly by just reading the first k tokens. Similarly, the perturbation-based experiments of Nema and Khapra (2018) have demonstrated the significant influence of four types of words (*i.e.* content words, named entities, question types, and function words) on an MRC question. Therefore, the dimension of *reading words* is defined as recognizing the observed vocabulary and the special words’ appearance (*i.e.* function words). In this study, The former is implemented as *Intra-n* (Gu et al., 2018)

and *Ent-n* (Serban et al., 2017) to measure vocabulary distribution, while the latter is computed as the frequency of corresponding words.

Reading sentences. The rules that are used to form a sentence using the aforementioned linguistic units are conventional phrase-structure grammars. Consequently, before understanding the information contained in a text, an MRC system inevitably requires capturing the sentence structure and handling the possible obscure words. We define the dimension of *reading sentences* as recognizing grammaticality and readability, and they are implemented by constituency parsing tree statistics and readability metrics¹, respectively.

Understanding words. The semantic representation of text is usually established by local and global links according to the linguistic units at word-level and sentence-level, respectively. To reflect the local semantic structure, we design the dimension of *understanding words* to assess how well an MRC model understands the relationships between words. In this work, we exemplify two relations (*i.e.* the arithmetic operations and logical items) that usually have salient patterns in the text. The former directly focuses on statistical and operational reasoning from the text, while the latter deals with the reasoning of predicate logic, *e.g.* conditionals and qualifiers. Inspired by the human annotation process (Boratko et al., 2018; Schlegel et al., 2020), where the annotators are asked to label as many reasoning skills as possible by paying more attention to corresponding indicative words, the sub-capabilities of this dimension is quantified as the frequency of those words.

Understanding sentences. Integrating the local structures into a global representation requires not only the text itself but also specific knowledge. To simplify the forms of knowledge, we divide the dimension of *understanding sentences* into two subclasses, linguistic and factual reasoning. They respectively mean understanding the relationship between sentences based on linguistics and the events (*i.e.* five dimensions including time, space, causation, intentionality, and objects). Among metrics of this dimension, *BERTScore* (Zhang* et al., 2020), *MoverScore* (Zhao et al., 2019) and *LS_score* (Wu et al., 2020) are used to measure semantic overlap between the context and question and multi-hop reasoning is an extra particular subclass on the HotpotQA (Yang et al., 2018; Cheng et al., 2021)

¹<https://py-readability-metrics.readthedocs.io/>

	Value	SQuADv1		SQuADv2		HotpotQA		RACE		
		r	p	r	p	r	p	r	p	
F_1	v_1	-0.131	0.000	-0.135	0.000	-0.146	0.007	-0.129	0.010	
		-0.120	0.018	-0.124	0.026	-0.135	0.009	-0.118	0.017	
	v_2	-0.162	0.002	-0.152	0.027	-0.174	0.022	-0.173	0.003	
		-0.154	0.000	-0.144	0.025	-0.166	0.000	-0.165	0.005	
	v_3	-0.134	0.000	-0.130	0.000	-0.141	0.029	-0.135	0.011	
		-0.163	0.016	-0.159	0.026	-0.170	0.006	-0.164	0.017	
	v_4	-0.166	0.001	-0.155	0.018	-0.182	0.000	-0.181	0.019	
		-0.198	0.020	-0.187	0.026	-0.214	0.018	-0.213	0.006	
		-0.208	0.015	-0.197	0.020	-0.224	0.013	-0.223	0.002	
		-0.206	0.000	-0.195	0.002	-0.222	0.010	-0.221	0.001	
		-0.168	0.023	-0.157	0.022	-0.184	0.023	-0.183	0.006	
		-0.168	0.010	-0.157	0.004	-0.184	0.002	-0.183	0.000	
	F_{logits}	v_1	-0.144	0.012	-0.151	0.000	-0.165	0.007	-0.142	0.010
			-0.133	0.018	-0.140	0.026	-0.154	0.009	-0.131	0.017
		v_2	-0.188	0.002	-0.175	0.000	-0.205	0.022	-0.203	0.003
			-0.180	0.000	-0.167	0.025	-0.197	0.028	-0.195	0.005
v_3		-0.163	0.000	-0.157	0.000	-0.172	0.000	-0.163	0.011	
		-0.192	0.016	-0.186	0.026	-0.201	0.006	-0.192	0.017	
v_4		-0.206	0.001	-0.192	0.018	-0.228	0.023	-0.226	0.019	
		-0.238	0.020	-0.224	0.026	-0.260	0.018	-0.258	0.006	
		-0.248	0.015	-0.234	0.020	-0.270	0.013	-0.268	0.002	
		-0.246	0.023	-0.232	0.002	-0.268	0.010	-0.266	0.001	
		-0.208	0.023	-0.194	0.022	-0.230	0.023	-0.228	0.006	
		-0.208	0.010	-0.194	0.004	-0.230	0.002	-0.228	0.000	

Table 2: The Pearson’s correlation (r) between capability-specific value v_i and model performance. Stronger correlations are marked in **bold**. All the correlations are with the p -value < 0.05 .

dataset. For the other sub-capabilities of this dimension, we consider lessons of the ablation operations performed by Sugawara et al. (2020) to observe the performance change of the MRC model and quantify them using the corresponding indicative structures.

Consider the two examples questions shown in Figure 1. To solve Q_1 , an MRC system just needs to match the words between the question and context. However, Q_2 requires understanding temporal relations among the events (went to the grocery store \rightarrow walked to the fast-food restaurant) and the verb semantics (walk to means go to). Therefore, Q_2 is more challenging to the MRC system than Q_1 . Please refer to Appendix B for more detailed examples and descriptions of our employed metrics.

2.2 Relationship Between Data Properties and Model Performance

Based on our assessment framework, the learning value of each sample is also decomposed into four dimensions, namely capability-specific values. In this section, we first uncover the connection between the capability-specific values and model performance from four dimensions and then recalibrate the connection by removing the inter-dimension correlations.

Capability-specific value. Given a sample x , we represent it by four capability-specific value (denoted as $\{v_i(x)\}_{i=1}^4$) to reflect its learning value for

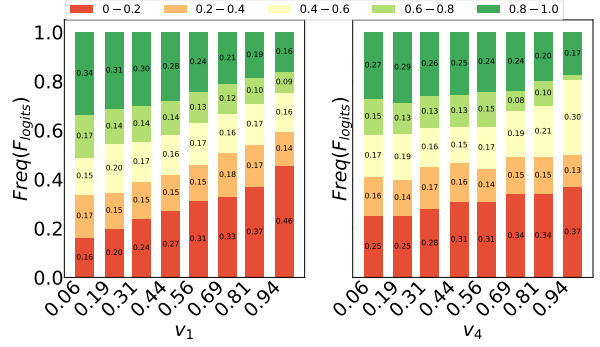


Figure 2: Bar diagram illustrating the relationship between the distribution of model performance and different ranges of v_i . Horizontal axes represent the different score ranges of v_i of samples, and the vertical axis means the performance distribution by the frequency of F_{logits} on five levels (plotted in five colors).

each capability dimension. According to our assessment framework, $v_i(x)$ can be computed by merging the corresponding metrics $\{m_i^j(x)\}_{j=1}^{n(i)}$. Specifically, considering the sensitivity of capability-specific value to different ranges of the metric score, we normalize each raw metric $m_i^j(x)$ from its original scale to range $[0, 1]$ by the cumulative density function (CDF) as Platanios et al. (2019), which is denoted as $\widetilde{m}_i^j(x)$. In this work, the normalization computes the cumulative density from a higher model performance to ensure that the normalized metric and model performance are negatively correlated. The capability-specific score $v_i(x)$ is formulated as: $v_i(x) = \frac{1}{n(i)} \sum_{j=1}^{n(i)} \widetilde{m}_i^j(x)$.

Analysis between capability-specific values and model performance. For each sample x , we obtain a 4-dimensional score $\{v_i(x)\}_{i=1}^4$. It is necessary to explore the relationship between samples’ $v_i(x)$ and model performance for knowing about what specific capabilities a model has gained and the level of each capability. In this work, we employ BERT-base (Devlin et al., 2018) as the MRC model and train it respectively on training split of datasets SQuADv1 (Rajpurkar et al., 2016), SQuADv2 (Rajpurkar et al., 2018), HotpotQA (Yang et al., 2018) and RACE (Lai et al., 2017). We then analyze the correlations between four capability-specific scores and the model’s overall performance on the corresponding dev split. In addition to F_1 , we also report the results of scaled F_1 (denoted as F_{logits}) by taking the model’s confidence to an answer span or candidate into account.

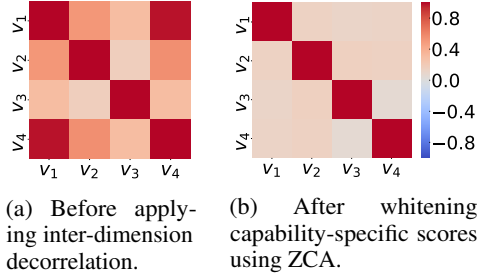


Figure 3: Pairwise correlations of capability-specific scores before and after inter-dimension decorrelation.

F_{logits} is computed as:

$$F_{logits} = \begin{cases} F_1 * \ln(slog) * \ln(elog) & \text{or} \\ F_1 * \ln(candlog) \end{cases} \quad (1)$$

where *slog* and *elog* mean the model output logits for start and end token in answer extraction style questions, and *candlog* represents the largest logits among all candidate answers.

Table 2 quantitatively shows the Pearson’s correlation coefficients (r) between capability-specific values and model performance. From the results, we have the following observations: First, each capability-specific score has a relatively strong correlation with the model performance under a statistically significant guarantee, showing the reasonability of our capability-based assessment framework. Second, F_{logits} shows better relevancy than F_1 , which indicates that F_{logits} is a more appropriate performance measurement in our framework.

We further explore the distribution of model performance over different ranges of v_i . The distribution diagrams of v_1 and v_4 are shown in Figure 2. There are two inspiring characteristics in this diagram: First, among all the bins of v_i , the frequency of prediction results within the intermediate range (0.4 ~ 0.6) are similar ($\approx 50\%$). Second, as the v_i increases, the frequency of prediction results within a low range (0.0 ~ 0.2) also increases, while the one of a high range (0.8 ~ 1.0) decreases. These observations reveal that the samples with high v_i can be used in indicative measurements to the corresponding model capability c_i . Please refer to Appendix C for more diagrams illustrating this relationship.

Inter-dimension decorrelation. Let $\mathcal{V} = \{v_i | i = 1, \dots, 4\}$. Pairwise correlations of \mathcal{V} are illustrated in Figure 3a in a heatmap fashion. The results show a common situation where if a sample is difficult (scored as high capability-specific value) in a dimension, the other dimensions have the same tendency and vice versa. To alleviate the inevitable

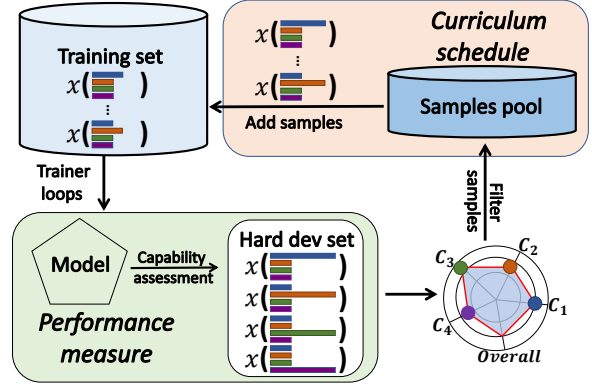


Figure 4: Illustration of our capability boundary breakthrough curriculum learning (CBBC). Different capabilities and their levels are represented in different colors and the length of the color bar, respectively. A longer color bar indicates a stronger capability.

correlations and construct a clear value representation for our following specific application scenario (*i.e.* CBBC), we eliminate the 4-dimensional capabilities by decorrelation. Specifically, we employ zero-phase component analysis (ZCA) whitening (Bell and Sejnowski, 1996) to diagonalize the covariance matrix while keeping the local information of the samples as much as possible.

As shown in Figure 3b, the 4-dimensional capabilities are not highly correlated after inter-dimension decorrelation, which can be in favor of constructing clear indicators for our following data sampling in CBBC.

3 Improve Learning Efficiency with Competency Assessment

In this section, our universal assessment framework of the model capability is adapted into a specific MRC training scenario to evaluate its usefulness and efficiency. Specifically, we embed our proposed assessment framework into a curriculum learning pipeline and make a capability boundary breakthrough curriculum learning (CBBC) strategy. Based on the assessment framework, our CBBC can guide a model to learn according to its capability boundary by understanding what the model has learned from data (*i.e.* capability-specific value v_i) and choosing appropriate samples with comparable learning values from four dimensions. It is worth noting that our competency assessment framework is also applicable to other training pipelines that balance the relationship between the model capabilities and data properties, such as active learning (Settles, 2009) and self-training (Mihalcea, 2004) (Appendix E).

Figure 4 shows an illustration of the pipeline of our CBBC. Following the original formulation of curriculum learning (Bengio et al., 2009), our CBBC organizes all samples by a sequence of ordered training stages $\{s\}_{s=1}^S$ and corresponding training sets $\{D^s\}_{s=1}^S$ with an easy-to-difficult fashion. The classic curriculum learning works (Soviany et al., 2021) usually consist of two essential components: the performance measurer and the curriculum scheduler. In general, the measurer is used to determine the learning status of a model by evaluating performance, while the scheduler is responsible for deciding when and how to update the curriculum by selecting the input samples.

In our work, the measurer and scheduler are implemented by analyzing the multi-dimensional capability levels of the model interpretably and measuring the capability-specific values of the data in a decorrelated way, respectively. That is to say, the only difference between our CBBC and the original curriculum learning design is incorporating MRC capability assessment into the curriculum learning. Without significantly increasing the complexity of the pipeline, our proposed assessment framework can generally empower the MRC training pipeline in a plug-and-play manner.

Performance measurer. Recall what we have discussed in Section 2.2 that the samples with high v_i can be used in indicative measurements to the corresponding model capability c_i . In this work, we use samples scored in the top- k of each capability-specific value to assess the corresponding model capability. More precisely, we first evaluate the model on the dev set and obtain an average F_{logits} for each capability on the corresponding top- k subset. Then partial correlation (Baba et al., 2004) (denoted as ρ_i) between dimension v_i and F_{logits} is computed to mask the contributions of the other dimensions $\mathcal{V} \setminus \{v_i\}$. After that, each model capability on stage s is quantified as: $c_i^s = \frac{\rho_i}{\sum_{j=1}^4 \rho_j} F_{logits}$. Empirically, we set k in top- k as 32.

Curriculum scheduler. Following the most works (Xu et al., 2020; Platanios et al., 2019), we schedule the curriculum at a linear pace (every 1,000 training iterations). During each curriculum schedule, we enlarge the training set two times until it includes all the samples. The capability upper bound $\overline{c_i^{s+1}}$ for $s + 1$ stage by exponential growth: $\overline{c_i^{s+1}} = \max\{\gamma c_i^s, 1.0\}$. After that, we use criterion $v_i(x) < \overline{c_i^{s+1}}$ to construct candidate set D_i^{s+1} for the i -th capability on the state $s + 1$, and use

absolute contribution of v_i to F_{logits} as sampling ratio (*i.e.* $\rho_1 : \rho_2 : \rho_3 : \rho_4$) to construct D^{s+1} .

4 Experiments

Datasets. We employ two question styles to evaluate our CBBC: answer span extraction and multiple choice. The former consists of SQuADv1 (Rajpurkar et al., 2016), SQuADv2 (Rajpurkar et al., 2018) and HotpotQA (Yang et al., 2018), while the latter adopts RACE (Lai et al., 2017). For each dataset, we train and evaluate the model on official training and dev split, respectively.

Implementation details. The source code and hyperparameters are included in the supplementary material. We use BERT-base (Devlin et al., 2018) as our backbone model, which is initialized by pre-trained parameters from cased BERT. AdamW (Loshchilov and Hutter, 2017) optimizer with weight decay $5e - 4$ and epsilon 8 is used to finetune the model with max sequence length 384, document stride 128. The learning rate warms up over the first 10% steps and then decays linearly to 0 for all experiments with training batch size 16 and maximum iteration 40,000.

Baseline models. In addition to the BERT-base model, we also consider the following ten baselines. The first two baselines are trained through a pre-defined curriculum learning strategy, which sorts the samples, then feeds them to the model stage-by-stage. “B+CL+ \mathcal{V} (M_2)” sorts the samples by four capability-specific scores in an easy-to-difficult order. “B+antiCL+ \mathcal{V} (M_3)” does like “ M_2 ”, but in a reverse difficult-to-easy order. The following five baselines are trained using our CBBC strategy to maximize the data value in each dimension, respectively. “B+C+ v_1 (M_4)”, “B+C+ v_2 (M_5)”, “B+C+ v_3 (M_6)” and “B+C+ v_4 (M_7)” use the corresponding v_1 , v_2 , v_3 and v_4 respectively to perform the competency test and filter samples. “B+C+ \mathcal{V}_{corr} (M_8)” is trained using four correlated scores through CBBC. The following three baselines are devised by embedding other instance scoring methods into our CBBC pipeline. “B+C+DatasetMap (M_9)”, “B+C+Forgetting (M_{10})” and “B+C+Predictability (M_{11})” substitute the capability-specific scores with the confidence score of true answer span (Swayamdipta et al., 2020), number of “forgotten” events (Toneva et al., 2018) and predictability score (Le Bras et al., 2020), respectively. The last two baselines (denoted as M_{12} and M_{13}) are

Name	Method	SQuADv1		SQuADv2		HotpotQA		RACE
		EM	F ₁	EM	F ₁	EM	F ₁	Acc.
M_1	B	81.25	88.41	77.32	80.31	63.53	76.42	63.02
M_2	B+CL+ \mathcal{V}	84.21	90.12	81.16	84.28	66.44	78.71	66.68
M_3	B+antiCL+ \mathcal{V}	79.80	87.31	75.43	79.19	63.01	75.03	61.80
M_4	B+C+ v_1	82.47	89.29	78.56	81.62	64.18	76.88	63.92
M_5	B+C+ v_2	82.67	89.30	78.59	81.98	64.12	77.01	64.11
M_6	B+C+ v_3	83.66	89.68	80.47	83.23	65.36	78.75	65.82
M_7	B+C+ v_4	85.03	89.90	81.60	84.21	66.30	79.13	65.73
M_8	B+C+ \mathcal{V}_{corr}	87.25	91.65	84.98	87.64	68.39	80.42	68.66
M_9	B+C+DatasetMap	82.04	89.05	78.04	81.11	63.86	76.57	63.51
M_{10}	B+C+Forgetting	83.51	89.77	79.62	82.99	64.75	77.62	64.91
M_{11}	B+C+Predictability	84.65	90.67	81.47	84.22	66.36	79.75	66.82
M_{12}	B+DRCA	85.05	90.59	82.19	85.30	67.07	79.32	67.48
M_{13}	B+CBCL	86.15	90.89	83.17	86.85	67.72	79.65	67.79
	Ours	89.71	93.18	87.64	90.51	69.82	82.57	71.01

Table 3: Quantitative results on four benchmark datasets. B and C represent the BERT backbone and our CBBC strategy, respectively. The best results are highlighted in **bold**.

start-of-the-art curriculum learning pipelines consisting of DRCA (Xu et al., 2020) and CBCL (Platanios et al., 2019). Finally, our full model is trained using four decorrelated scores through CBBC instead. The critical difference between the full model and M_8 is the decorrelation operation.

4.1 Experimental Results

Quantitative Results. We present a summary of our quantitative results in Table 3. As shown in the table, we have the following key observations.

On the one hand, our proposed competency framework does benefit the MRC learning efficiency in either a single or multiple dimensions. For example, when using a pre-defined curriculum strategy, M_2 achieves EM and F_1 far beyond M_1 , highlighting that our quantification to data properties properly estimates the learning value contained in the data. M_3 degrades performance w.r.t. M_1 , demonstrating that the learning strategy from easy to difficult samples is more reasonable than the reverse. When equipped with our CBBC, all models of M_4 , M_5 , M_6 and M_7 achieve improvements w.r.t. M_1 on four datasets, which indicates the significant contribution of each capability dimension on gradually increasing the model capability. In particular, among the four different dimensions, M_7 has the best result, indicating that understanding sentences is a relatively more important capability for MRC. M_8 outperforms all the models except for ours. This demonstrates that our CBBC can maximize the learning value of the data sample to increase an MRC model’s capability.

On the other hand, our framework wins other scoring methods and curriculum learning pipelines by a considerable margin. Although M_9 , M_{10} , M_{11} , M_{12} and M_{13} achieve substantial improvements on four datasets w.r.t. M_1 , their perfor-





Figure 5: Illustration of performance (smoothed by averaging F_{logits} every 32 steps) of various baseline models on HotpotQA dev split as training progresses. Ours (denoted as red plot) outperforms the other models from the start of training.



(a) Plots illustrating changes of each MRC capability during training. (b) A 5-dimensional map of MRC model capabilities on step 50, 1k, 2k, 20k.

Figure 6: Illustration of MRC model capabilities on different training stages.

mances are still worse than our full model. These results verify that our proposed framework can assess the model capability more correctly and make better use of the learning value within data.

Finally, our full model achieves significantly higher EM , F_1 and $Acc.$ compared to all other baselines, demonstrating the necessity of the decorrelation between capability-specific scores. Its superior performance roots from constructing a decorrelated value representation of each dimension for our CBBC learning strategy. Overall, compared to M_1 , our full model achieves tremendous improvement of EM / F_1 up to 11.22% / 8.71% on the average of three answer extraction style datasets.

Qualitative Results. Figure 5 shows the performance of baselines on the HotpotQA dev set. There are two observations worth noting here. First, the performance of our full model lies consistently on top of the other baseline models during the whole training stage. This result shows that CBBC can make the model more prepared for complex samples by enlarging its capability boundary step by step. Second, the performance plot of the baseline model with v_4 sits on top of other baselines with v_1 , v_2 , and v_3 from the beginning of training to the end. This result highlights the main contribution of v_4 (understanding sentences) to the final performance.

Furthermore, the capability map after max-min normalization of the model is shown in Fig-

Value	SQuADv1	SQuADv2	HotpotQA	RACE
v_1	0.602	0.536	0.630	0.589
	0.625	0.559	0.653	0.612
v_2	0.696	0.683	0.755	0.620
	0.714	0.701	0.773	0.638
v_3	0.730	0.573	0.550	0.797
	0.674	0.517	0.494	0.741
v_4	0.553	0.674	0.672	0.819
	0.488	0.609	0.607	0.754
	0.468	0.589	0.587	0.734
	0.472	0.593	0.591	0.738
	0.547	0.668	0.666	0.813
	0.503	0.612	0.609	0.761

Table 4: Pearson’s correlation coefficients between human judgments and data properties for four capabilities and their subclasses.

ure 6. First, among 4-dimensional capability, the c_3 (*i.e.* understanding words) has the largest initial value. A possible explanation is that pre-trained BERT has a fair amount of prior knowledge obtained from unlabeled corpus, which concentrates more on semantic understanding of words. Second, the capability c_1 increases at the fastest speed as the training progresses. Interestingly, the model M_4 based on v_1 does not seem improving accordingly as the capability c_1 increases. The possible reason could be that the superficial structure is easy to learn from samples but makes a limited contribution to the final performance. *Please refer to Appendix D for the results of other MRC models.*

4.2 Human Annotation

Annotation specification. we ask three annotators to answer ($100 \times 4 = 400$) questions randomly sampled from four datasets, consisting of SQuADv1, SQuADv2, HotpotQA, and RACE. Using only our proposed four capabilities, they first read the context, question, and gold standard answer (the correct candidate answer under multiple-choice situation), and then choose the evidence sentences in context. After that, they respectively label the subclasses of four major capabilities as 1 (required) or 0 (not required). *Please refer to Appendix A for more details about human annotation.*

Annotation results. In the annotation of required capabilities, the inter-annotator agreement is 75.33% for all 400 samples. We use the average of three corresponding annotator labels as the final human judgments for a specific sub-capability required by the question. Finally, a sample will be annotated ($2 + 2 + 2 + 6 = 12$) human ratings. Table 4 summarizes the correlations between human

judgments and capability-specific scores of samples. The relatively strong correlations on all four dimensions indicate that our employed heuristic metrics can reasonably approximate the learning value contained in the samples.

5 Related Work

Analytic approaches to MRC capability. Some works performed skill-based analyses for the MRC model. In the scientific question domain, Clark et al. (2018) constituted the ARC benchmark, which requires far more powerful knowledge and reasoning than previous benchmarks. In a generalizable definition, Sugawara et al. (2017) proposed a set of 10 skills for MCTest (Richardson et al., 2013). The others focused more on the analysis of the MRC dataset itself. For example, Sugawara et al. (2020) proposed a semi-automated, ablation-based methodology to assess the capacities of datasets. Rajpurkar et al. (2016) analyzed their proposed datasets using several types of reasoning, *e.g.* lexical and syntactic variation, and multiple sentence reasoning. Nevertheless, they require too costly human efforts and ignore that the model capability changes as training progresses.

Data selection for debiased representations. Some works proposed different criteria to score instances according to the model response to input. Swayamdipta et al. (2020) built data maps using training dynamics measures for scoring data samples. Toneva et al. (2018) also employed the number of “forgotten” events to measure a sample, which was misclassified during a later epoch of training, despite being classified correctly earlier. The others (Le Bras et al., 2020) used adversarial filtering algorithms to rank instances based on their “predictability”. However, these approaches require training a model once in advance on the dataset to obtain the corresponding training dynamics, which is computationally expensive, especially when using a large model.

6 Conclusion

We design a competency assessment framework for MRC capabilities, which describes model skills in an explainable and multi-dimensional manner. By leveraging the framework, we further uncover and disentangle the connections between various data properties and model performance on a specific task, as well as propose a capability boundary breakthrough curriculum (CBBC) strategy to maxi-

mize the data value and improve training efficiency. The experiments performed on four benchmark datasets verified that our approach can significantly improve the performance of existing MRC models. Our work shows a deep understanding of model capabilities and data properties helps monitor the model skills during training and improves learning efficiency. Our framework and learning strategy are also generally applicable to other NLP tasks.

Acknowledgements

This work has been supported in part by the National Key Research and Development Program of China (2018AAA0101900), Zhejiang NSF (LR21F020004), Key Research and Development Program of Zhejiang Province, China (No. 2021C01013), Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies, Chinese Knowledge Center of Engineering Science and Technology (CKCEST).

References

- Kunihiro Baba, Ritei Shibata, and Masaaki Sibuya. 2004. Partial correlation and conditional correlation as measures of conditional independence. *Australian & New Zealand Journal of Statistics*, 46(4):657–664.
- Anthony Bell and Terrence J Sejnowski. 1996. Edges are the ‘independent components’ of natural scenes. *Advances in neural information processing systems*, 9:831–837.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Michael Boratko, Harshit Padigela, Divyendra Mikilineni, Pritish Yuvraj, Rajarshi Das, Andrew McCallum, Maria Chang, Achille Fokoue-Nkoutche, Pavan Kapanipathi, Nicholas Mattei, et al. 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.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2019. Graphflow: Exploiting conversation flow with graph neural networks for conversational machine comprehension. *arXiv preprint arXiv:1908.00059*.
- Yi Cheng, Siyao Li, Bang Liu, Ruihui Zhao, Sujian Li, Chenghua Lin, and Yefeng Zheng. 2021. Guiding the growth: Difficulty-controllable question generation through step-by-step rewriting. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5968–5978.
- Cheryl Lasse. 2020. [What is a competency assessment?](#)
- 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. *ArXiv*, abs/1803.05457.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#).
- Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. 2019. Uncertainty-guided continual learning with bayesian neural networks. In *International Conference on Learning Representations*.
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR.
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR.
- Mingfei Gao, Zizhao Zhang, Guo Yu, Serkan Ö Arık, Larry S Davis, and Tomas Pfister. 2020. Consistency-based semi-supervised active learning: Towards minimizing labeling cost. In *European Conference on Computer Vision*, pages 510–526. Springer.
- Xiaodong Gu, Kyunghyun Cho, Jung-Woo Ha, and Sunghun Kim. 2018. DialogWAE: Multimodal response generation with conditional wasserstein auto-encoder.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- 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 2021–2031.
- Walter Kintsch. 1988. The role of knowledge in discourse comprehension: A construction-integration model. *Psychological review*, 95(2):163.
- Andreas Kirsch, Joost Van Amersfoort, and Yarin Gal. 2019. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Advances in neural information processing systems*, 32:7026–7037.

- 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 785–794.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *International Conference on Machine Learning*, pages 1078–1088. PMLR.
- David D Lewis and William A Gale. 1994. A sequential algorithm for training text classifiers. In *SIGIR'94*, pages 3–12. Springer.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Danielle S McNamara and Joe Magliano. 2009. Toward a comprehensive model of comprehension. *Psychology of learning and motivation*, 51:297–384.
- Rada Mihalcea. 2004. Co-training and self-training for word sense disambiguation. In *Proceedings of the Eighth Conference on Computational Natural Language Learning (CoNLL-2004) at HLT-NAACL 2004*, pages 33–40.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Preksha Nema and Mitesh M. Khapra. 2018. Towards a better metric for evaluating question generation systems. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3950–3959, Brussels, Belgium. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237.
- Robert Pinsler, Jonathan Gordon, Eric Nalisnick, and José Miguel Hernández-Lobato. 2019. Bayesian batch active learning as sparse subset approximation. *Advances in neural information processing systems*, 32:6359–6370.
- Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabas Poczos, and Tom Mitchell. 2019. Competence-based curriculum learning for neural machine translation. 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 1162–1172, Minneapolis, Minnesota. 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.
- 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.
- Matthew Richardson, Christopher JC 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.
- Viktor Schlegel, Marco Valentino, André Freitas, Goran Nenadic, and Riza Batista-Navarro. 2020. A framework for evaluation of machine reading comprehension gold standards. *arXiv preprint arXiv:2003.04642*.
- Iulian Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Burr Settles. 2009. Active learning literature survey.
- Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2021. Curriculum learning: A survey. *arXiv preprint arXiv:2101.10382*.
- 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.

- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. Assessing the benchmarking capacity of machine reading comprehension datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8918–8927.
- Saku Sugawara, Hikaru Yokono, and Akiko Aizawa. 2017. Prerequisite skills for reading comprehension: Multi-perspective analysis of mctest datasets and systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. 2018. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 189–198.
- Kai Wei, Rishabh Iyer, and Jeff Bilmes. 2015. Submodularity in data subset selection and active learning. In *International Conference on Machine Learning*, pages 1954–1963. PMLR.
- Hanlu Wu, Tengfei Ma, Lingfei Wu, Tariro Manyumwa, and Shouling Ji. 2020. Unsupervised reference-free summary quality evaluation via contrastive learning. *arXiv preprint arXiv:2010.01781*.
- Benfeng Xu, Licheng Zhang, Zhendong Mao, Quan Wang, Hongtao Xie, and Yongdong Zhang. 2020. [Curriculum learning for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6095–6104, Online. 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.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. In *International Conference on Learning Representations*.
- Changchang Zeng, Shaobo Li, Qin Li, Jie Hu, and Jianjun Hu. 2020. A survey on machine reading comprehension—tasks, evaluation metrics and benchmark datasets. *Applied Sciences*, 10(21):7640.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M Meyer, and Steffen Eger. 2019. Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. 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 563–578.

A Annotation Details

We ask three annotators to answer ($100 \times 4 = 400$) questions randomly sampled from four datasets, consisting of SQuADv1, SQuADv2, HotpotQA, and RACE. They are graduate students majoring in Computer Science or Electronic Engineering and competent in English. They voluntarily offer to help without being compensated in any form. Before annotation, they are informed of the detailed annotation instruction with the following three steps.

- **Step 1.** Make a hypothesis using a question statement and gold standard answer or the correct candidate answer under the multiple-choice situation.

Example 1:

Q: Why did Tom look angry? A: His sister ate his cake.

→ *Hypothesis*: Tom looked angry because his sister ate his cake.

Example 2:

Q: When did French Revolution occur? A: In 1789

→ *Hypothesis*: French Revolution occurred in 1789.

- **Step 2.** Select sentences (from the context) required to provide the hypothesis.

Example 1:

Context: (C1) Tom is a student. (C2) Tom looks annoyed because his sister ate his cake. (C3) His sister's name is Sylvia.

Hypothesis: Tom looks angry because his sister ate his cake.

→ *Select*: C2

- **Step 3.** Select capabilities required for understanding an entailment from selected context sentences to hypothesis and label the corresponding capability as 1 (required).

Example 1:

C2: Tom looks annoyed because his sister ate his cake.

Hypothesis: Tom looks angry because his sister ate his cake.

→ *Capability*: causal relation (“because”), semantic overlap (lexical knowledge of “annoyed = angry”)

Then, we describe our annotation schema in greater detail. We present the respective phenomenon, give a short description, and present an example illustrating the corresponding feature.

- **Reading words**

Recognize vocabulary. We annotate this as “1” if repetition of some word rarely occurs in a sentence (less than two times in every ten words).

Question with label 0: The creek of which Paradise Creek is a tributary of what river?

Context: Paradise Creek is a 9.6 mi tributary of Brodhead Creek in the Poconos of eastern Pennsylvania in the United States. Brodhead Creek is a 21.9 mi tributary of the Delaware River in the Poconos of eastern Pennsylvania in the United States.

Question with label 1: Of these two publications—Báiki and Sick—what type of publication is the one that was published most frequently?

Context: Báiki: The International Sámi Journal (“Báiki” means a place in Sami) is a biannual English-language publication that covers Sami culture, history, and current affairs. The coverage also includes the community affairs of the Sami in North America, estimated at some 30,000 people. Sick was a satirical-humor magazine published from 1960 to 1980, lasting 134 issues.

Recognize function words. We annotate this as “1” if a sentence consists of lots of the structural relationships between words signaled by function words (more than five times in every ten words).

Question: What drug is among the list of illegal drugs in the Philippines and can be taken by mouth or by injection?

Context: [...] Ephedrine and methylenedioxy methamphetamine are also among the list of illegal drugs that are of great concern to the authorities. Ephedrine is a medication and stimulant. [...]

- **Reading sentences**

Recognize grammaticality. We annotate this as “1” the sentence pattern and grammar involved are relatively complex, such as multiple nested subordinate clause structures.

Question: Sudha Kheterpal, who is a musician best known as the percussionist in Faithless, has played with what singer who is recognized as the highest-selling Australian artist of all time by the Australian Recording Industry Association?

Readability. We annotate this as “1” if there are a lot of obscure words in the question or context (more than five times in every ten words).

Question: The creature HNoMS Draug is named after comes from what kind of mythology?

Context: Two ships of the Royal Norwegian Navy have borne the name HNoMS “Draug”, after the sea revenant Draugr: The draugr or draug (Old Norse: “draugr”, plural draugar ; modern Icelandic: “draugur”, Faroese: “dreygur” and Danish, Swedish, and Norwegian: “draug”), also called aptrganga or aptrgandr , literally “again-walker” (Icelandic: “afturganga”) is an undead creature from Norse mythology.

- **Understanding words**

Arithmetic operation. We annotate this as “1” if an arithmetic operation needs to be performed to answer the question, such as addition, subtraction, ordering, and counting.

Question: How many points were the Giants behind the Dolphins at the start of the 4th quarter?

Context: New York was down 17-10 behind two rushing touchdowns.

Logical operation. We annotate this as “1” if it is required to understand the concept of quantification (existential and universal) in order to determine the correct answer.

Question: How many presents did Susan receive?

Context: On the day of the party, all five friends showed up. Each friend Quantification had a present for Susan.

- **Understanding sentences**

Syntactic and semantic overlap. We annotate this as “1” if some part of the context and the question overlap semantically.

Question: Is it freezing today?

Context: The weather is cold today.

Coreference resolution. We annotate this as “1” if inter-sentence coreference and anaphora need to be resolved in order to retrieve the expected answer.

Question: What is the name of the psychologist who is known as the originator of social learning theory?

Context: Albert Bandura OC (born December 4, 1925) is a psychologist who is the David Starr Jordan Professor Emeritus of Social Science in Psychology at Stanford University. [...] He is known as the originator of social learning theory and the theoretical construct of self-efficacy and is also responsible for the influential 1961 Bobo doll experiment.

Con/Dis-junction, negation. We annotate this as “1” if logical conjunction, disjunction, or negation needs to be resolved in order to conclude the answer.

Question: Is dad in the living room?

Context: Dad is either in the kitchen or the living room.

Question: How many percent are not Marriage couples living together?

Context: [...] 46.28% were Marriage living together. [...]

Causality. We annotate this as “1” if causal (*i.e.* cause-effect) reasoning between events, entities, or concepts is required to answer a question correctly.

Question: Why did Sam stop Mom from making four sandwiches?

Context: [...] There are three of us, so we need three sandwiches. [...]

Spatial/Temporal relationship. We annotate this as “1” understanding about directions, environment, spatiality, and succession is required in order to derive an answer.

Question: What is the 2010 population of the city 2.1 miles southwest of Marietta Air Force Station?

Context: Marietta Air Force Station is located 2.1 mi northeast of Smyrna, Georgia.

Multi-hop reasoning. We annotate this as “1” if information to answer the question needs to be gathered from multiple supporting facts,

“Multi-hop” by commonly mentioned entities, concepts, or events. This phenomenon is also known as “Bridging” in literature.

Question: What show does the host of The 2011 Teen Choice Awards ceremony currently star on?

Context: [...] The 2011 Teen Choice Awards ceremony, hosted by Kaley Cuoco, aired live on August 7, 2011, at 8/7c on Fox. [...] Kaley Christine Cuoco is an American actress. Since 2007, she has starred as Penny on the CBS sitcom “The Big Bang Theory”, for which she has received Satellite, Critics’ Choice, and People’s Choice Awards. [...]

B Examples of Our Employed Metrics

We first present a brief overview of employed metrics, especially those that are adapted from other studies. In the following descriptions, the question Q and corresponding context C are denoted as the sequence of n-grams $Q^n = \{q_i^n\}$ and $C^n = \{c_j^n\}$, respectively.

Intra-n and Ent-n. *Intra-n* (Gu et al., 2018) and *Ent-n* (Serban et al., 2017) are originally designed to evaluate the diversity of neural dialogue responses. The former calculates the ratio of distinct unigrams (Intra-1) and bigrams (Intra-2) in generated responses, while the latter measures the information entropy of n-grams. Specifically, in our work, they are formulated as:

$$\text{Intra-n} = \frac{\text{Unique}(Q^n)}{|Q^n|} \quad (2)$$

$$\text{Ent-n} = \sum_{i=1} -\frac{\text{Count}(q_i^n)}{|Q^n|} \log \frac{\text{Count}(q_i^n)}{|Q^n|} \quad (3)$$

Tree statistics. Empirically, most of the complicated sentences have a relatively high and wide constituency parsing tree. In this paper, we calculate the height and width of the constituency parsing tree by Stanford CoreNLP API (Manning et al., 2014).

Readability metrics. Readability is the ease with which a reader can understand a written text. Readability metrics produce an approximate representation of the US grade level needed to comprehend the text and are widely used in the field of education to assess the English proficiency of non-native English speakers. We employ the py-readability package (<https://pypi.org/>

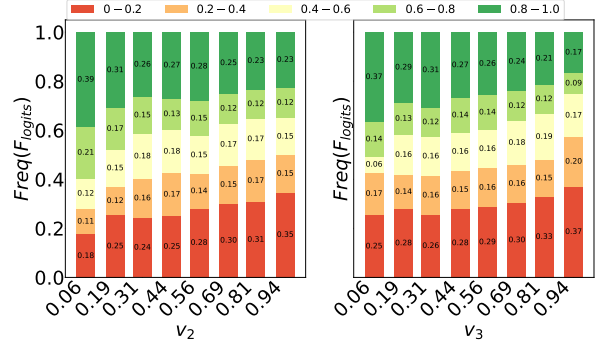


Figure 7: Bar diagram illustrating the relationship between the distribution of model performance and different ranges of v_i . Horizontal axes represent the different score ranges of v_i of samples, and the vertical axis means the performance distribution by the frequency of F_{logits} on five levels (plotted in five colors).

[project/py-readability-metrics/](#)) to calculate the readability of a question based on two metrics, including *Flesch Kincaid Grade Level* and *Automated Readability Index (ARI)*.

BERTScore and MoverScore. To measure the semantic overlap between the question and corresponding context, we slightly modify the contextualized embedding-based similarity metrics of text generation task, comprising *BERTScore* (Zhang* et al., 2020) and *MoverScore* (Zhao et al., 2019). Unlike the original implementation, the question and the context rather than the gold standard reference are fed into the metrics computation.

We exemplify two samples from HotpotQA dev set to show the difference of specific metrics in Table 5 and Table 6.

C Additional Diagrams of Samples’ Capability-specific Value

The distribution diagrams of v_2 and v_3 are shown in Figure 7. They present a conclusion consistent with v_1 and v_4 discussed in Section 2.2. That is, among all the bins of v_i , the frequency of prediction results within the intermediate range (0.4 ~ 0.6) are similar ($\approx 50\%$). Furthermore, as the v_i increases, the frequency of prediction results within a low range (0.0 ~ 0.2) also increases, while the one of a high range (0.8 ~ 1.0) decreases.

In addition to the distribution of model performance over different ranges of v_i , the mean value and standard deviation of model performance (in F_{logits}) over them and their subclasses are illustrated in Figure 8. As shown in the figure, it qualitatively shows the relatively strong correlation be-

ID	Q1(5a7322a25542991f9a20c634)	Q2(5a72bd0b5542992359bc318f)
Context	<p>The Metropolitan Life Insurance Company Tower, colloquially known as the Met Life Tower, is a landmark skyscraper located on Madison Avenue near the intersection with East 23rd Street, across from Madison Square Park in Manhattan, New York City. Designed by the architectural firm of Napoleon LeBrun & Sons and built by the Hedden Construction Company, the tower is modeled after the Campanile in Venice, Italy. The hotel located in the clock tower portion of the building has the address 5 Madison Avenue, while the office building covering the rest of the block, occupied primarily by Credit Suisse, is referred to as 1 Madison Avenue. 15 Hudson Yards is a residential building currently under construction on Manhattan’s West Side. Located in Chelsea near Hell’s Kitchen Penn Station area, the building is a part of the Hudson Yards project, a plan to redevelop the Metropolitan Transportation Authority’s West Side Yards. The tower started construction on December 4, 2014.</p>	<p>Andrea Louise Martin (born January 15, 1947) is an American actress, singer, author, and comedian, best known for her work in the television series "SCTV". She has appeared in films such as "Black Christmas" (1974), "Wag the Dog" (1997), "Hedwig and the Angry Inch" (2001), "My Big Fat Greek Wedding" (2002), and "My Big Fat Greek Wedding 2" (2016), and lent her voice to the animated films "Anastasia" (1997), "The Rugrats Movie" (1998) and " (2001). Mark S. Hoplamazian is an American businessman who is the President and CEO of Hyatt Hotels Corporation. He received his A.B. in economics from Harvard College and his M.B.A. from the University of Chicago Booth School of Business.</p>
Question	Was the Metropolitan Life Insurance Company Tower [Met Life Tower] or the 15 Hudson Yards building designed by the firm of Napoleon LeBrun & Sons?	Who achieved more academically, Andrea Martin or Mark Hoplamazian?

Table 5: Two samples from HotpotQA dev set. The answer span in the context is marked in green.

Value	Metrics	Q1(5a7322a25542991f9a20c634)		Q2(5a72bd0b5542992359bc318f)	
		raw	normalized	raw	normalized
v_1	intra1	0.521	0.252	0.608	0.681
	entropy1	6.243	0.530	6.363	0.636
	ntopwrods	0.268	0.509	0.274	0.575
v_2	height	9.000	0.904	7.000	0.809
	flesch kincaid	16.748	0.913	10.574	0.196
	ari	17.659	0.895	9.835	0.119
v_3	nnums	0.039	0.722	0.051	0.861
	nlogicals	0.002	0.506	0.009	0.735
v_4	BERTScore	0.862	0.043	0.698	0.959
	MoverScore	0.154	0.022	-0.210	0.900
	ncoreferences	0.005	0.079	0.071	0.991
	njunctions	0.012	0.069	0.071	0.957
	ncausals	0.001	0.275	0.009	0.619
	nspatialtemporals	0.078	0.779	0.035	0.512
	nfacts	2.000	0.674	3.000	0.913

Table 6: Two samples with our raw and normalized metrics. The higher normalized scores are marked in bold.

Method	SQuADv1		SQuADv2		HotpotQA		RACE
	EM	F ₁	EM	F ₁	EM	F ₁	Acc.
R-Net	76.59	85.83	71.87	75.10	60.78	73.17	59.49
R-Net+CBBC+v ₁	78.11	86.36	73.37	77.15	61.75	73.80	60.20
R-Net+CBBC+v ₂	78.23	86.91	73.40	76.52	61.03	73.80	59.92
R-Net+CBBC+v ₃	79.71	87.63	74.98	78.40	61.92	74.86	61.32
R-Net+CBBC+v ₄	80.27	87.77	76.35	79.15	62.85	76.29	62.63
R-Net+CBBC+V _{corr}	82.49	88.47	78.50	81.15	64.41	77.28	63.33
R-Net+CBBC+V	83.87	89.75	80.85	83.56	65.87	77.96	65.47

Table 7: Quantitative results of R-Net backbone on four benchmark datasets. The best results are highlighted in **bold**.

Method	SQuADv1		SQuADv2		HotpotQA		RACE
	EM	F ₁	EM	F ₁	EM	F ₁	Acc.
QANet	77.85	86.54	73.42	76.62	61.72	74.09	60.69
QANet+CBBC+v ₁	79.37	87.07	74.91	78.68	62.69	74.72	61.40
QANet+CBBC+v ₂	79.50	87.62	74.94	78.05	61.97	74.72	61.12
QANet+CBBC+v ₃	80.97	88.34	76.52	79.93	62.86	75.77	62.52
QANet+CBBC+v ₄	81.54	88.48	77.89	80.67	63.79	77.21	63.83
QANet+CBBC+V _{corr}	83.75	89.18	80.04	82.68	65.35	78.20	64.53
QANet+CBBC+V	85.13	90.46	82.40	85.08	66.81	78.88	66.66

Table 8: Quantitative results of QANet backbone on four benchmark datasets. The best results are highlighted in **bold**.

tween the model performance and each capability-specific score.

D Additional Experiments Using Other Models

In addition to the Transformer-based MRC model, we also perform ablation analysis using the following systems to further verify the effectiveness of our proposed assessment framework, whose training setting is consistent with that of our BERT-based MRC model.

- R-Net (Wang et al., 2017) matches the question and passage with gated recurrent neural networks (RNNs) to obtain the question-aware passage representation and employs the pointer networks to locate the positions of answer span from the passages.
- QANet (Yu et al., 2018) encodes the local and global interactions with the convolution and self-attention, respectively. It achieves higher training efficiency while obtaining the equivalent accuracy to the recurrent models.

The quantitative results of R-Net and QANet are summarized in Table 7 and Table 8, respectively. Our assessment framework also provides a significant performance improvement to these weaker

backbones than BERT, such as RNN-based R-Net and convolution-based QANet.

E Additional Experiments Using Other Pipelines

To further verify the effectiveness of our proposed MRC competency assessment framework and reveal more available application scenarios for it, we embed it into the active learning pipeline besides curriculum learning.

CBBC in active learning. Given the training state of a model, active learning aims to select the most valuable samples from the unlabeled dataset and hand it over to the oracle (*e.g.* human annotator) for labeling so as to reduce the cost of labeling as much as possible while still maintaining performance. Take the most common pool-based active learning (Lewis and Gale, 1994; Gal and Ghahramani, 2016) as an example, which queries the best sample based on the confidence evaluation and ranking of the entire dataset. This query strategy is usually implemented by the uncertainty-based sampling (Ebrahimi et al., 2019; Gal et al., 2017; Houlsby et al., 2011; Kirsch et al., 2019) and distribution-based sampling (Pinsler et al., 2019; Wei et al., 2015). Compared to the original pool-based active learning (shown in Figure 9a), CBBC-guided active learning (shown in Figure 9b) pro-

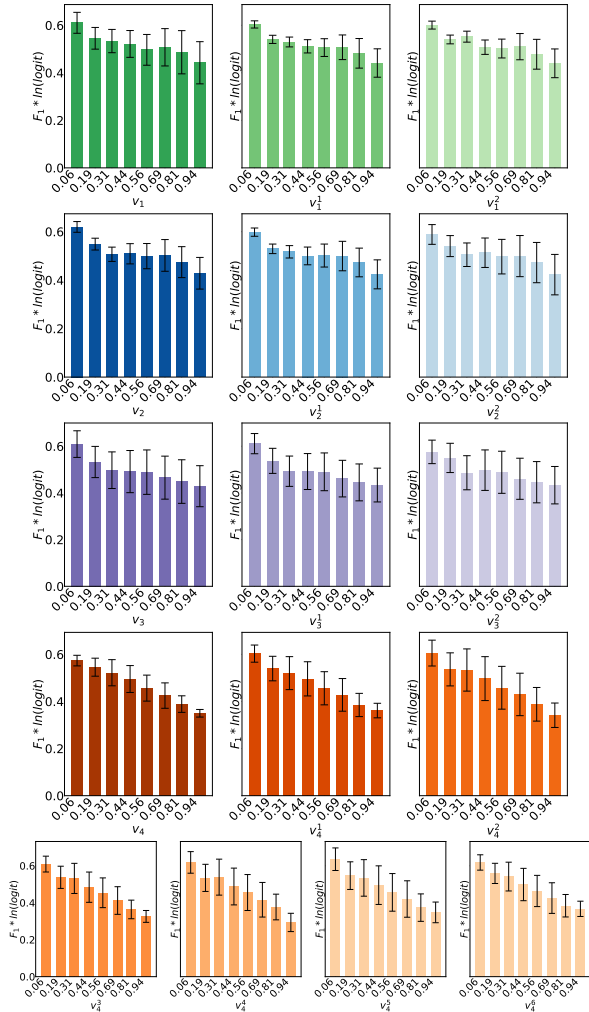



Figure 8: Bar diagram illustrating the relationship between the mean value and standard deviation of model performance and different ranges of v_i and its subclasses. The height of each bar and its error line represent the mean value and standard deviation of model performance, respectively.


vides a novel and interpretable query strategy by assessing the model in a 4-dimensional capability. In addition to the pool-based baseline, we employ a more recent active learning pipeline ICAL (Gao et al., 2020) as a comparison, which selects samples with the high inconsistency of predictions over a set of data augmentations.

We employ BERT as our MRC backbone. In each active learning cycle, we continue to train the MRC model by adding 5% labeled data points by simulating the oracle annotating process. The initial training set is randomly sampled from the HotpotQA train split and follows Gao et al. (2020) for the setting of initial training set size and active learning budget.

Figure 10 illustrates the results of different meth-



(a) Pipeline of the pool-based active learning.



(b) Pipeline of CBBC-guided active learning.

Figure 9: Architecture comparison between the typical pool-based active learning pipeline and our CBBC-guided pipeline.

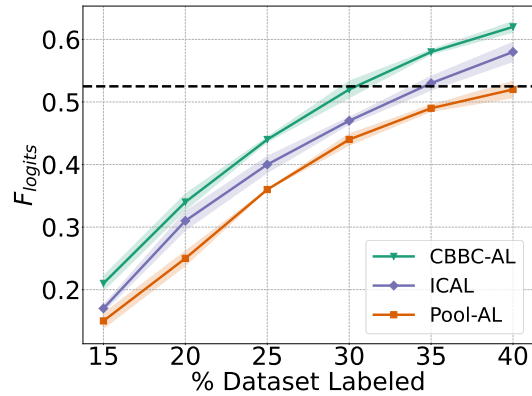


Figure 10: Performance comparison under active learning pipeline on the HotpotQA dev split. The solid lines indicate the results averaged over five trials, and shadows represent the standard deviation. The dotted line at the top represents the performance of the BERT-based MRC model with the whole training set labeled.

ods at each active learning cycle qualitatively. Our CBBC-guided active learning (denoted as “CBBC-AL”) achieve a higher MRC performance than the pool-based active learning (denoted as “Pool-AL”) and ICAL from the start of training, demonstrating that the assessment of MRC model capability derived by our CBBC is reasonable and can also make a substantial difference to active learning beside curriculum learning. When using only 35% labeled samples, our CBBC-guided active learning outperforms the baseline model normally trained on the entire dataset by a considerable margin.