

# Learning with Instance Bundles for Reading Comprehension

Dheeru Dua<sup>♡</sup> Pradeep Dasigi<sup>♣</sup> Sameer Singh<sup>♡</sup> Matt Gardner<sup>♣</sup>

<sup>♡</sup>University of California, Irvine, USA

<sup>♣</sup>Allen Institute for Artificial Intelligence

ddua@uci.edu

## Abstract

When training most modern reading comprehension models, all the questions associated with a context are treated as being independent from each other. However, closely related questions and their corresponding answers are not independent, and leveraging these relationships could provide a strong supervision signal to a model. Drawing on ideas from contrastive estimation, we introduce several new supervision losses that compare question-answer scores across multiple related instances. Specifically, we normalize these scores across various neighborhoods of closely contrasting questions and/or answers, adding a cross entropy loss term in addition to traditional maximum likelihood estimation. Our techniques require bundles of related question-answer pairs, which we either mine from within existing data or create using automated heuristics. We empirically demonstrate the effectiveness of training with instance bundles on two datasets—HotpotQA and ROPES—showing up to 9% absolute gains in accuracy.

## 1 Introduction

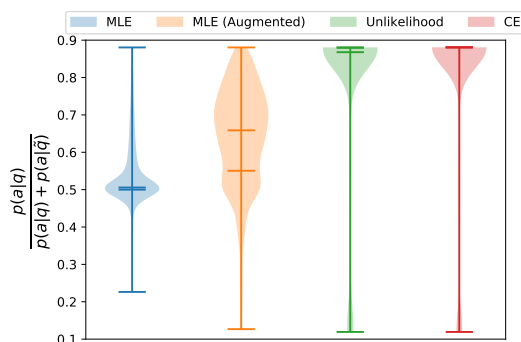
Machine learning models are typically trained with the assumption that the training instances sampled from some data distribution are independent and identically distributed. However, this assumption can cause the learner to ignore distinguishing cues (Dietterich et al., 1997) between related or minimally different questions associated with a given context, resulting in inconsistent model predictions (Jia and Liang, 2017; Ribeiro et al., 2019; Asai and Hajishirzi, 2020). In cases like ROPES, where the dataset contains only minimally different questions, we see that the performance of a competitive model (RoBERTA) is close to random (Lin et al., 2019). One potential reason for this poor performance is that the model considers each question independently, instead of looking at differences between related questions.

**Context:** Marsilea is a genus of approximately 65 species of aquatic ferns of the family Marsileaceae. The name honours Italian naturalist Luigi Ferdinando Marsili (1656-1730) .... Brabejum is a genus of a single species of large evergreen tree, Brabejum stellatifolium in the Proteaceae, commonly called wild almond, bitter almond or ghoeboontjie.

**Question,  $q$ :** Is the Marsilea or the Brabejum the genus of more individual species of plants?  
**Answer,  $a$ :** Marsilea

**Question,  $\tilde{q}$ :** Is the Marsilea or the Brabejum the genus of less individual species of plants?  
**Answer,  $\tilde{a}$ :** Brabejum

(a) Instance bundle created from HotpotQA. The QA pair,  $(\tilde{q}, \tilde{a})$  is heuristically generated from  $(q, a)$  in the dataset.



(b) Distribution over probability scores of gold QA pair,  $(a, q)$  normalized over the two questions  $\{q, \tilde{q}\}$  in each instance bundle, in HotpotQA’s comparison type dev set. The higher value indicates that correct QA pair has a higher likelihood than the wrong pair (with 0.5 indicating model has the same likelihood for the answer for both the contrastive questions).

Figure 1: Example of contrastive QA pairs and effect of different loss functions while training with such data.

This problem can be addressed by training models with sets of related question-answer (QA) pairs simultaneously, instead of having a loss function that decomposes over independent examples. We use the term *instance bundle* to refer to these sets of closely contrasting examples. Consider an instance bundle from HotpotQA in Figure 1a, containing two contrastive QA pairs, which differ in their input by only one word (changing *more* to *less*), resulting in different answers. With both examples in

the training set, traditional maximum likelihood estimation will incentivize the model to figure out the difference between the inputs that leads to the difference in the answers, but the instances are likely to be seen far apart from each other during training, giving only a weak and indirect signal about the relationship between the pair.

To learn from these instance bundles more effectively, we draw on contrastive estimation (Smith and Eisner, 2005), a method for re-normalizing an unsupervised probabilistic model using a neighborhood of related examples (originally a set of perturbations of some observed text). We extend this technique to apply to supervised reading comprehension problems by carefully selecting appropriate “neighborhoods” from instance bundles. The simplest choice of neighborhood is the set of contrasting answers from the instance bundle, resulting in a method similar to unlikelihood training (Welleck et al., 2020) or noise-contrastive estimation (Gutmann and Hyvärinen, 2010). However, there are other choices, including the set of contrasting *questions*, or combinations of questions and answers. These re-normalized loss functions are not effective on their own, which is likely why they have not been used before for training reading comprehension models, but provide substantial gains when combined with maximum likelihood training.

An intuitive explanation for the improved performance is illustrated in Figure 1b. When trained on non-contrasting data with maximum likelihood estimation, the model gives roughly equal values for both  $p(a|q)$  and  $p(a|\tilde{q})$ , even though  $q$  and  $\tilde{q}$  are opposites. Augmenting the contrasting data helps the model differentiate these two probabilities, but not as much as unlikelihood training, which itself is not as effective as contrastive estimation.

We empirically demonstrate the utility of this approach on two reading comprehension datasets: HotpotQA (Yang et al., 2018) and ROPES (Lin et al., 2019). To define the instance bundles on these datasets, we introduce various heuristics for obtaining closely related instances. We show that using contrastive estimation on the instance bundles gives up to a 9% absolute performance improvement over prior training techniques. These results strongly suggest that data should be collected in instance bundles wherever possible, to allow for stronger supervision signals during training.

## 2 Contrastive Estimation for Reading Comprehension

Reading comprehension is the task of producing an answer  $a$  given a question  $q$  about a context  $c$ . The question is tied to a particular passage, so in the discussion that follows we will typically use  $q$  as a shorthand to refer to both  $q$  and  $c$  together. Reading comprehension models are typically trained to maximize the likelihood of the answer to each training question. Given a model’s exponentiated scoring function  $\psi(q, a)$  for a QA pair,<sup>1</sup> this MLE objective normalizes the scores over all possible answer candidates  $\mathcal{A}$  for a given question:

$$\begin{aligned}\mathcal{L}_{\text{MLE}}(q_i, a_i) &= \log p(a_i|q_i) \\ &= \log \frac{\psi(q_i, a_i)}{\sum_{c \in \mathcal{A}} \psi(q_i, c)}\end{aligned}$$

In this work we use a generative model for  $\psi$ , but many other alternatives are available, and our contribution is applicable to any scoring function. Specifically, we use as  $\psi$  the (locally normalized) probability assigned by the generative model to an answer candidate for a given question.

Instead of normalizing scores over all possible answer candidates, *contrastive estimation* (CE, Smith and Eisner, 2005) normalizes scores over some *neighborhood* of closely related instances. This method was originally introduced for unsupervised linguistic structure prediction, with a neighborhood obtained by permuting observed text to get inputs that had similar content but were ungrammatical. Our contribution is to apply this general idea to supervised reading comprehension problems. In our setting, given a neighborhood  $\mathcal{N}(q, a)$  of related QA pairs, CE can be described as

$$\mathcal{L}_{\text{CE}}(q_i, a_i) = \log \frac{\psi(q_i, a_i)}{\sum_{q_j, a_k \in \mathcal{N}(q_i, a_i)} \psi(q_j, a_k)}$$

Smith and Eisner (2005) replace the MLE objective with CE, which worked well in their unsupervised learning problem. In supervised learning, MLE is a much stronger training signal, and CE on its own severely underperforms MLE. This is because CE provides no learning signal for the very large space of alternative answers to a question that are not in the neighborhood. However, CE can provide a much stronger signal than MLE for a small set of potentially confusing alternatives, as

<sup>1</sup> $\psi$  is parameterized by model parameters  $\theta$ , but we omit this in the equations for simplicity of exposition.

there are fewer ways for the model to increase the probability of the correct answer. To adapt CE to supervised settings we interpolate between the two losses, instead of replacing MLE with CE:

$$\mathcal{L} = \alpha_1 \mathcal{L}_{MLE} + \alpha_2 \mathcal{L}_{CE} \quad (1)$$

Interestingly, this can be viewed as forcing the scoring function  $\psi$  to capture different probabilistic interpretations, as both losses perform softmaxes over different sets of alternatives. Additionally, if  $\psi$  has some locally-normalized component, as is true for the generative models we work with and for many other common models (such as BIO tagging, or independent span start and span end positions), this interpolation trades off between the locally-focused MLE and the more global view of the problem that the normalization in CE provides.

The key question in applying CE to reading comprehension is how to choose a neighborhood  $\mathcal{N}$  for a given training example. We do so by making bundles of related instances, then extracting various combinations of questions and answers from a bundle to use as neighborhood. Formally, an *instance bundle*  $\mathcal{B}$  is a collection of unique questions  $\mathcal{B}_Q$  and unique answers  $\mathcal{B}_A$ , such that there is at least one QA pair where  $a$  is the correct answer to  $q$ :  $\text{ans}(q) = a$ . We refer to such pairs as  $(q_g, a_g)$  in the discussion that follows. Our assumption is that the questions in  $\mathcal{B}_Q$  and the answers in  $\mathcal{B}_A$  are related to each other in some way—often they differ in only a single word—though we do not characterize this formally. However, a good bundle creation procedure is crucial for effective model learning. We discuss several ways for creating bundles in Section 3, and discuss the limitations of CE when effective bundles cannot be created in Section 5.2. The following section discusses choices of neighborhood functions given an instance bundle.

## 2.1 Choosing a neighborhood

Given an instance bundle  $\mathcal{B}$  with questions  $\mathcal{B}_Q$  and answers  $\mathcal{B}_A$ , there are many ways to construct a neighborhood. Figure 2 shows some of these options graphically, with the bold line showing the gold QA pair, and gray lines showing the other QA pairs that make up the neighborhood. We distinguish between two kinds of neighborhood methods. A *single neighborhood* CE model is one that perturbs and normalizes over a single variable, either the question (input) or the answer (output). Similarly, *multiple neighborhood* CE models perturb

both variables jointly and normalize over the combinatorial space of both variables.

### 2.1.1 Single Neighborhood Models

These models construct a neighborhood using either the answers or the questions from the  $\mathcal{B}$ .

**Answer Conditional** This probabilistic model maximizes the probability of the correct answer  $a_i$  at the expense of the other answers candidates in the instance bundle  $\mathcal{B}_A$  (Figure 2a).

$$\mathcal{L}_{\text{CE-AC}}(q_g, a_g, \mathcal{B}) = \log \frac{\psi(q_g, a_g)}{\sum_{a_j \in \mathcal{B}_A} \psi(q_g, a_j)}$$

**Question Conditional** This model computes the normalization constant over the question neighborhood for a fixed answer. This effectively computes a probability distribution over *questions* in the bundle given the correct answer, and maximizes the probability of the correct question (Figure 2b).

$$\mathcal{L}_{\text{CE-QC}}(q_g, a_g, \mathcal{B}) = \log \frac{\psi(q_g, a_g)}{\sum_{q_j \in \mathcal{B}_Q} \psi(q_j, a_g)}$$

### 2.1.2 Multiple Neighbourhood Models

These models consider all possible combinations of questions,  $\mathcal{B}_Q$ , and answers,  $\mathcal{B}_A$ , in a bundle for normalization.

**Two Way** This method performs a weighted combination (Jacobs et al., 1991) of the answer conditional and question conditional losses.

$$\mathcal{L}_{\text{CE-TW}}(q_g, a_g, \mathcal{B}) = \lambda_1 \log \frac{\psi(q_g, a_g)}{\sum_{a_j \in \mathcal{B}_A} \psi(q_g, a_j)} + \lambda_2 \log \frac{\psi(q_g, a_g)}{\sum_{q_j \in \mathcal{B}_Q} \psi(q_j, a_g)}$$

**Full Partition** Instead of separate normalizations over questions and answers, this method uses a single normalization over the same sets as *Two Way*. This is equivalent to normalizing over  $\mathcal{B}_Q \times \mathcal{B}_A$  without correct pairings (Figure 2c).

$$\mathcal{L}_{\text{CE-FP}}(q_g, a_g, \mathcal{B}) = \log \frac{\psi(q_g, a_g)}{\psi(q_g, a_g) + \sum_{\substack{q_j \in \mathcal{B}_Q, \\ a_k \in \mathcal{B}_A, \\ \text{ans}(q_j) \neq a_k}} \psi(q_j, a_k)}$$

**Joint** This method switches from optimizing the probability of single QA pairs to optimizing the *set* of correct QA pairs in the bundle, also known as power-set label classification (Zhang and Zhou,

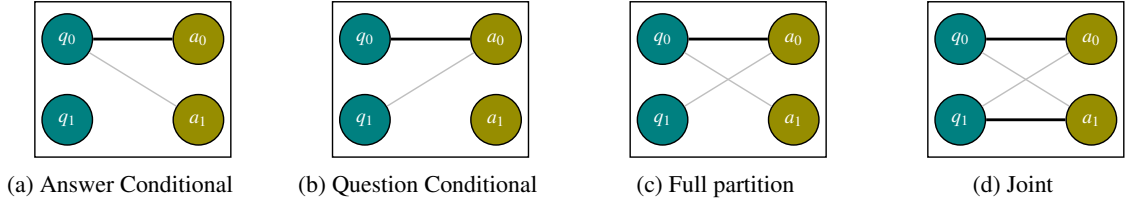


Figure 2: **Contrastive Estimation Losses** for an instance bundle of size 2, with bold lines indicating combinations whose probability is maximized at the expense of the combinations represented by gray lines, for the QA pair  $(q_i, a_i)$  in the bundle, where  $\text{ans}(q_i) = a_i$ . The CE loss is the sum of loss for each positive QA pair in the bundle.

2007) (Figure 2d). We perform this for only bundles consisting of two correct QA pairs, because the power set becomes prohibitively large for larger bundles. Let  $C(\mathcal{B})$  be a function that returns all unique subsets of size 2 from  $\mathcal{B}_Q \times \mathcal{B}_A$ , and let  $(q_{g_1}, a_{g_1})$  and  $(q_{g_2}, a_{g_2})$  be the two positive QA pairs in the bundle. The joint CE objective is

$$\mathcal{L}_{\text{CE-JT}}(\mathcal{B}) = \frac{\psi(q_{g_1}, a_{g_1})\psi(q_{g_2}, a_{g_2})}{\sum_{q_i, a_k, q_j, a_l \in C(\mathcal{B})} \psi(q_i, a_k)\psi(q_j, a_l)}.$$

## 2.2 Alternative Ways to Use Bundles

Here we briefly consider other potential baselines that make use of instance bundles in other ways.

**Data Augmentation** If the bundle  $\mathcal{B}$  contains instances that were not present in the training data (e.g., the bundle could be generated using simple heuristics; see §3), the simplest use of the bundle is to add all instances to the training data and use MLE under the standard i.i.d. assumption. This is the standard approach to using this kind of data, and it has been done numerous times previously (Andreas, 2020; Zmigrod et al., 2019). This is not applicable if the bundle was obtained by mining the existing training instances, however.

**Unlikelihood** Unlikelihood training (Welleck et al., 2020) minimizes the likelihood of carefully chosen negative examples to improve a text generation model that would otherwise assign those examples too high of a probability. Essentially, because the generative model only gets a single positive sequence in an exponentially large set, it does not get strong enough evidence to push down the probability of particularly bad generations. Unlikelihood training seeks to solve the same problem that contrastive estimation does, and it provides a natural alternative use of instance bundles. In our setting, unlikelihood training would decrease the

likelihood of negative answers in the bundle:

$$\mathcal{L}_{\text{UL}}(q_g, a_g, \mathcal{B}) = \sum_{c \in \mathcal{B}_A \setminus a_g} \log(1 - p(c|q_g))$$

Similar to (Welleck et al., 2020) and Eq. 1 we linearly interpolate  $\mathcal{L}_{\text{UL}}$  with  $\mathcal{L}_{\text{MLE}}$  to also increase likelihood of positive answer at the same time. Unlikelihood training, though easy to perform, has two drawbacks. First, it independently minimizes the likelihood of the neighborhood, which means that the probability mass is moved from negative QA pairs but may not necessarily move to the positive pair, unlike CE. Second, because it assumes a conditional probabilistic model of  $p(a|q)$ , it is not clear how to use alternative *questions* in the bundle.

## 3 Bundling Heuristics

In this section we discuss how we obtain instance bundles for use with contrastive estimation and other related baselines.

A naive way to create a bundle would be to exploit the fact that all the questions associated with a context are likely to be related, and simply make bundles consisting of all QA pairs associated with the context. However, this approach poses two problems. First, there could be many questions associated with any particular context, and smaller, more closely-related bundles are more informative. Second, and relatedly, it is likely that bundles obtained this way will have many questions whose answers can be obtained from the bundle by superficial type matching. For instance, a wh-question starting with “where” would most likely align with a location type answer. If this were bundled with a question starting with “how many”, with an answer that is a number, the bundle would be largely uninformative. We instead attempt to create bundles with minimally different questions and answers, in several different ways.

**Diverse Top-k sampling** We first discuss a method for getting alternative *answers* to a single

question. This will result in a bundle that can only be used with answer conditional CE, as there are no alternative questions in the bundle. An easy way to get answer candidates is to employ a pre-trained answering model and sample answers from the posterior distribution. However, since the model has seen all the QA pairs while training, it can easily memorize answers, resulting in a low variance, high confidence distribution. In order to achieve diverse answer samples we need to either over-generate and prune out the gold answer from the samples or induce a diversity promoting sampling. We adopt a hybrid sampling strategy where we use nucleus sampling for the first few (depending on dataset) timesteps of decoding (without replacement) and then top-k for the remaining timesteps. This forces the answer generator to consider different starting positions in the passage and then generate the best answer span (of an appropriate length) from the token produced at the first step.

**Question Mining** Some datasets, such as ROPES, are constructed with very close question pairs already in the data. When these exist, we can create instance bundles by finding natural pairings from the training set. To find these pairings, we cluster the questions with a high lexical bag-of-words overlap based on Jaccard index ( $\geq 0.8$ ), ensuring that each question in the cluster has a unique answer. In ROPES, these bundles typically result in bundles of two QA pairs that differ in one or a few words. In HotpotQA, the other dataset we focus on in this work, there are very few such pairings naturally occurring in the dataset, so we resort to heuristics to create them.

**Question Generation** HotpotQA has many questions that are phrased as multiple-choice, with answer options given in the question itself. These multiple choice questions can most often be rephrased to provide QA pairs that can be bundled with the original question. For instance, given the question, “Which animal is faster, turtle or hare?”, it is straightforward to obtain a minimally different question with the opposite answer: “Which animal is slower, turtle or hare?”. We adopt three main heuristics to generate such questions whenever possible, applicable to any dataset that has questions of this kind. All of these heuristics require identifying the two plausible answer choices from the question, which can be done with reasonably high precision using simple regular expressions.

1. We replace superlatives with their contrasting counterparts, e.g., taller/smaller, more/less, etc.
2. We negate the main verb, e.g., played  $\rightarrow$  didn’t play, by inflecting the verbs.<sup>2</sup>
3. We swap the noun phrases being compared in the question, e.g., “Are rock A’s wavelengths shorter or longer than rock B’s?” can be used to generate, “Are rock B’s wavelengths shorter or longer than rock A’s?”

## 4 Experiments

We use an encoder-decoder style T5-large model for all our experiments. The baseline models in our experiments are fine-tuned T5 models for the corresponding tasks using the MLE objective. We compare them against models that are further fine-tuned with a combination of MLE and contrastive estimation objectives, as described in Section 2.1. That is, when using various instance bundle techniques, we initialize the model with the weights from the fine-tuned MLE model, then continue training with the new loss function.<sup>3</sup> The model takes a concatenated context and question as an input to produce an answer output. We use a learning rate of  $2e-5$  for ROPES and  $5e-5$  for COMPARISON with lowercased inputs and outputs for further fine-tuning. We truncate the concatenated context and question up to a length of 650 for ROPES and 850 for COMPARISON. All the interpolation hyper-parameters ( $\alpha_l, \lambda_l$ ) are set to 1.

In addition to standard metrics on these datasets, we additionally evaluate using *consistency*. This metric evaluates to true only if all the questions in a bundle are answered correctly, and is thus a stricter version of EM.

### 4.1 Main results

We experiment with three datasets: a subset of HotpotQA containing only comparison type of questions (COMPARISON), full HotpotQA and ROPES. In general, we find that all variants of CE perform substantially better than MLE alone, with question conditional giving small improvements over other CE variants. All CE models also outperform all UL and data-augmented MLE models.

**COMPARISON** HotpotQA has several different kinds of questions, with the question category la-

<sup>2</sup><https://spacy.io/universe/project/lemminflect/>

<sup>3</sup>To control for the number of optimization steps, we also tried a baseline where we continued fine-tuning an MLE model using the same setup, but this never improved over the original MLE, so we do not include it in the tables.

Model	Dev EM/F1	Test Consis.	Test EM/F1
MLE (full HotpotQA)	68.4/72.1	36.3	57.4/65.1
<b>MLE</b>	81.8/85.4	51.2	70.8/77.9
<b>+ Aug</b>	84.9/88.0	76.7	73.9/80.7
<b>+ UL</b>	87.1/90.6	85.4	74.3/82.3
+ Answer Cond.	87.4/90.3	85.7	75.7/83.2
<b>+ Question Cond.</b>	88.3/91.5	87.3	77.1/84.3
+ Two way	87.9/91.1	83.9	74.1/81.8
+ Joint	88.3/91.1	87.0	76.8/83.8
+ Full Partition	87.2/90.7	85.9	75.8/83.5

Table 1: Results on COMPARISON subset of HotpotQA. The comparison questions from the training data are divided into train and dev set (90%-10% split) and the ones from the dev set are treated as test split.

beled in the original data. We begin by experimenting with the subset labeled as comparison questions, as they lend themselves most naturally to instance bundles. For these questions, we adopt the question generation strategy to create instance bundles. Table 1 shows a comparison of the baseline MLE model (trained on the comparison subset only) with those further fine-tuned with UL and CE over the instance bundles. Also shown is a comparison with further fine-tuning using MLE on the generated QA pairs (+Aug).

Due to unavailability of the code for best model on the HotpotQA dataset, we use a T5-large model trained on the entire HotpotQA as our baseline. Even though this model has a performance of 81.1 F1 on the whole dev set (close to the current SOTA 83.5 on the leaderboard<sup>4</sup>), on the comparison subset it performs poorly (65.1). Training an MLE model on just this subset reaches 77.9 F1, which is outperformed by unlikelihood training (82.3 F1). The best CE performance is from the question conditional model, which gets 84.3 F1.

The consistency metric evaluates if the model answers all the questions in an instance bundle correctly. All the models trained with CE are more consistent than the MLE model. We also show in Figure 1 that a model trained with answer conditional CE is still effective on the question neighborhood, as it consistently assigns higher probability to  $p(a|q)$  than to  $p(a|\tilde{q})$ , even though while training it never compared  $q$  and  $\tilde{q}$  for a given  $a$ .

**HotpotQA** We additionally experiment with the entire HotpotQA dataset. Here we use top-k sam-

<sup>4</sup><https://hotpotqa.github.io/>

Model	F1
<b>MLE</b>	81.1
+CE-AC ( $k = 1$ )	82.5
+CE-AC ( $k = 2$ )	83.3
+CE-AC ( $k = 3$ )	82.1
+CE-AC ( $k = 4$ )	81.8

Table 2: F1 on the full HotpotQA dev set with increasing number ( $k$ ) of top-k negative answer candidates

pling to create instance bundles, where the top-k answer candidates were sampled from the MLE model we use as a baseline. Table 2 shows the performance of the fine-tuned model as we vary the number of answers in  $\mathcal{B}_{\mathcal{A}}$  with CE-AC loss. The overall performance gets better with CE up to  $|\mathcal{B}_{\mathcal{A}}| = 3$ , but reduces after that. On a closer examination of the samples, we find that on average we get two distinct answer candidates and the rest of the candidates are ungrammatical variations of the two distinct candidates (including the oracle answer). These ungrammatical variations provide a noisy signal that hurts model performance.

**ROPES** Since ROPES already contains minimally different QA pairs, we use question mining to create instance bundles. We use as the most closely comparable prior work the multi-step model of (Liu et al., 2020), which adds a ROPES-specific architecture on top of RoBERTa-large (Liu et al., 2019)<sup>5</sup>, while our baseline MLE model is a generic T5-large model without any special architecture. Table 3 shows gains over MLE of around 12% absolute on the dev set in both EM and consistency, though these drop to just a few points on the test set. Liu et al. (2020) saw similar behavior in their experiments with ROPES, which they attribute to high distributional shift from train to test. They recommend splitting the training set into train and train-dev, and treating the original dev set as an in-domain test set. Following their protocol, we find that using CE gives an absolute improvement of more than 9% EM over an MLE model on this dev-test set, while UL gives only a few point gain.

## 4.2 Joint Inference

In cases where we can generate a bundle given only a question (that is, the answer candidates are clear and our heuristics can generate a contrasting ques-

<sup>5</sup>UnifiedQA (Khashabi et al., 2020) also evaluated on ROPES, but they used much more training data from many other datasets, and a much larger model than we experiment with, so their performance is not particularly comparable.

Model	Dev EM	Dev C	Test EM	Dev-Test EM
Multi-step (Liu et al., 2020)	71.4	-	57.4	63.4
<b>T5-large MLE + UL</b>	65.7	52.1	56.7	66.1
	68.3	55.6	58.6	69.0
+ Answer Cond.	74.5	62.1	59.5	71.3
<b>+ Question Cond.</b>	<b>76.6</b>	<b>64.7</b>	<b>60.5</b>	<b>75.5</b>
+ Two way	73.5	59.7	59.1	70.3
+ Joint	72.5	58.2	59.3	73.4
+ Full Partition	75.1	63.2	59.6	73.1

Table 3: Performance on ROPES dev and test set, where C measures consistency. Dev-test EM is computed by splitting train into train and dev and using the original dev set as test, because of severe distributional shift in the original test set.

tion), we can treat test time inference as a hard assignment problem between questions and answers in the generated bundle. We use the scoring function  $\psi(q, a)$  to align each question to an answer in the bundle by optimizing objective below:

$$\begin{aligned} \max \quad & \sum_{\substack{a_j \in \mathcal{B}_A, \\ q_i \in \mathcal{B}_Q}} \psi(q_i, a_j) x_{ij} \\ \text{s.t.} \quad & \sum_{j=0}^{|\mathcal{B}_A|} x_{ij} = 1, \sum_{i=0}^{|\mathcal{B}_Q|} x_{ij} = 1 \end{aligned}$$

We refer to this as joint inference. Intuitively, even if the model is only given a single question at test time, if it can reason jointly about two competing assignments it can potentially use the alternatives to arrive at a better response than if it only considered the single question it was given. As shown in Figure 3, when using joint inference the performance of a baseline MLE model on COMPARISON improves from 79 F1 to 85.5. The CE model manages to achieve this performance (85.8 F1) without enforcing these constraints at test time, but joint inference improves CE to 90.1 F1.

## 5 Discussion

Here we try to understand how CE compares to MLE and UL and when is it effective.

### 5.1 Relation between MLE, UL and CE

In this section we describe the relationship between CE, MLE and UL in the special case when the scoring function  $\psi$  comes from a locally-normalized generative model, as it does in this work. Let

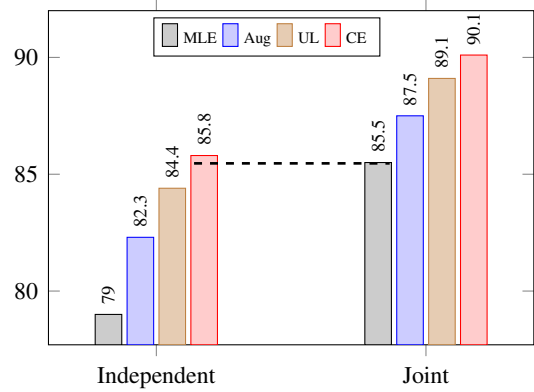


Figure 3: Performance (F1) on COMPARISON dev with independent prediction versus joint inference. Joint inference improves all models. The results are on the subset of COMPARISON for which we have paired instances (~93%).

$p_V$  be the locally-normalized probability of an answer candidate (i.e., combined likelihood of the tokens under a given generative model).  $\psi(q, a)$  then equals  $p_V(a|q)$ . The CE-AC loss with locally-normalized compatibility score can be written as

$$\mathcal{L}_{\text{CE-AC}}(q_g, a_g, \mathcal{B}) = \log \frac{p_V(a_g|q_g)}{\sum_{c \in \mathcal{B}_A} p_V(a_c|q_g)}. \quad (2)$$

We can decompose and rewrite Eq. 2 as

$$\begin{aligned} \mathcal{L}_{\text{CE-AC}}(q_g, a_g, \mathcal{B}) &= \log p_V(a_g|q_g) - \log \sum_{c \in \mathcal{B}_A} p_V(a_c|q_g) \\ &= \mathcal{L}_{\text{MLE}}(q_g, a_g) + \text{Reg}(\mathcal{B}_A, q_g). \end{aligned} \quad (3)$$

Eq. 3 shows that  $\mathcal{L}_{\text{CE-AC}}$  is a linear combination of MLE and a regularization term that decreases the probability of incorrect answers in the bundle.

On a closer look we can see an interesting connection between the regularization term and unlikelihood. The regularization term in CE-AC is essentially the log of an unlikelihood term, except the unlikelihood objective in Section 2.2 in practice gets applied at each timestep of decoding, while the regularization term in CE-AC is applied over the entire answer sequence.

Our formulation of CE is more general than the specific case we are analyzing here, but we make note of it as this is the function that we used in our experiments, and it significantly outperformed unlikelihood training. The theoretical connections shown here could benefit from further exploration.

Dataset	MLE	UL	CE	Entropy <sub>10</sub>	Top-2 ratio
COMP.	77.7	82.4	84.7	2.31	0.5
ROPES	65.7	68.3	77.6	0.97	2.3
QUOREF	84.8	83.9	85.0	0.06	3.1

Table 4: Comparison between QUOREF, COMPARISON, and ROPES datasets with Top-k bundling. MLE, UL, and CE results are on the corresponding development sets ( $F_1$  for COMPARISON and QUOREF, EM for ROPES) and Entropy<sub>10</sub> and Top-2 ratio are measured on random samples of the training sets. UL and CE columns show results after fine-tuning the baseline MLE model with the respective objectives.

## 5.2 Importance of close instance bundles

Experiments on ROPES and COMPARISON show strong improvements by using CE and UL when instances can be grouped into closely related bundles. But effective groupings may not be possible on all datasets. To analyze the applicability of our methods to a dataset without natural bundles, we looked at QUOREF (Dasigi et al., 2019). Table 4 shows a comparison between UL and CE across QUOREF, ROPES and COMPARISON with bundles created using top-k sampling. On QUOREF, UL does not improve on top of MLE, and CE shows only a very small improvement which is likely statistical noise. To understand why, we analyzed the  $p(a|q)$  distribution of the baseline MLE model, and computed the following two measures on a random sample of the training set.

- Entropy<sub>10</sub> =  $-\sum_{i=1}^{10} p(a_i|q) \log p(a_i|q)$
- Top-2 ratio =  $\log p(a_1|q)/p(a_2|q)$

As seen in Table 4, QUOREF has a lower Entropy<sub>10</sub>, and a higher Top-2 ratio than the other datasets, indicating that the baseline MLE model places a lot more weight on the top-1 answer in this task. Manual analysis additionally found that most of the top predictions were ungrammatical variations of the top-1 answer, similar to (but more extreme than) what was seen on the full HotpotQA dataset. This could explain why the top-k bundling heuristic is not as effective in the case of QUOREF as the other two datasets. More generally, these results indicate the importance of effective instance bundling heuristics, and future work could focus on identifying more general ways to create bundles.

## 6 Related Work

Learning with negative samples has been explored in many natural language tasks, such as

dialogue generation (Cai et al., 2020), word embeddings (Mikolov et al., 2013), language modeling (Noji and Takamura, 2020), etc., and computer vision tasks such as image captioning (Dai and Lin, 2017), unsupervised representation learning (Hadsell et al., 2006), etc. Similarly, mutual information minimization based learners in question answering (Yeh and Chen, 2019) and image classification (Hjelm et al., 2019) try to decrease the mutual information between positive and negative samples.

Natural language applications often sample negative examples either randomly from the data or based on likelihood (or unlikelihood) metrics from a reference model. However, the negative samples extracted in this manner are often unrelated. A growing body of literature is exploring ways to obtain closely-related examples, either manually (Kaushik et al., 2020; Gardner et al., 2020) or automatically (Ribeiro et al., 2020; Ross et al., 2021; Wu et al., 2021). This is complementary to our work, as we show how to make better use of these related examples during training. There is also work on consistent cluster assignments in coreference resolution (Chang et al., 2011); factually consistent summaries (Kryscinski et al., 2020) and language models (Elazar et al., 2021).

Another growing body of literature on training with closely related examples, to which we are contributing, includes methods that make use of logical or domain specific consistency rules, in natural language inference tasks (Minervini and Riedel, 2018), reading comprehension (Asai and Hajishirzi, 2020; Gupta et al., 2021), and visual question answering (Teney et al., 2019, 2020; Jacovi et al., 2021). In open domain QA, re-ranking extracted answer spans from a baseline model has shown promising improvements and shares connections with our answer conditional setup (Iyer et al., 2020). Instead of training just a ranking model (which is similar to answer conditional CE) on top of a baseline (MLE) model, we jointly train a single QA model with both objectives. This promotes better representation learning in the baseline QA model.

## 7 Conclusion

We have presented a way to use contrastive estimation in a supervised manner to learn from distinguishing cues between multiple related QA pairs, or *instance bundles*. Our experiments with multiple CE-based loss functions, defined over a joint neighborhood of questions and answers, have shown



that these models outperform existing methods on two datasets: ROPES and HotpotQA. Apart from presenting several ways to create instance bundles, we also explore theoretical connections between unlikelihood training and contrastive estimation, and initial exploration into when instance bundles are likely to be effective with these methods. We believe our results give strong motivation for further work in techniques to both create and use instance bundles in NLP datasets. The code is available at <https://github.com/dDua/contrastive-estimation>.

## Acknowledgements

We would like to thank Noah Smith, Nitish Gupta, Robert Logan, Anthony Chen, Dylan Slack, Sanjay Subramanian and the anonymous reviewers for the discussions and feedback on earlier versions. We would also like to thank Hasso Plattner Institute (HPI) for supporting the first author through UCI-HPI fellowship.

## References

- Jacob Andreas. 2020. [Good-enough compositional data augmentation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566, Online. Association for Computational Linguistics.
- Akari Asai and Hannaneh Hajishirzi. 2020. [Logic-guided data augmentation and regularization for consistent question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5642–5650, Online. Association for Computational Linguistics.
- Hengyi Cai, Hongshen Chen, Yonghao Song, Zhuoye Ding, Yongjun Bao, Weipeng Yan, and Xiaofang Zhao. 2020. [Group-wise contrastive learning for neural dialogue generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 793–802, Online. Association for Computational Linguistics.
- Kai-Wei Chang, Rajhans Samdani, Alla Rozovskaya, Nick Rizzolo, Mark Sammons, and Dan Roth. 2011. [Inference protocols for coreference resolution](#). In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task*, pages 40–44, Portland, Oregon, USA. Association for Computational Linguistics.
- Bo Dai and Dahua Lin. 2017. [Contrastive learning for image captioning](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 898–907.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasović, Noah A. Smith, and Matt Gardner. 2019. [Quoref: A reading comprehension dataset with questions requiring coreferential reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5925–5932, Hong Kong, China. Association for Computational Linguistics.
- Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. 1997. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *ArXiv preprint*, abs/2102.01017.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Nitish Gupta, S. Singh, Matt Gardner, and D. Roth. 2021. [Paired examples as indirect supervision in latent decision models](#). *ArXiv preprint*, abs/2104.01759.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, pages 1735–1742. IEEE.
- R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Srinivasan Iyer, Sewon Min, Yashar Mehdad, and Wentaoh Yih. 2020. [Reconsider: Re-ranking using span-focused cross-attention for open domain question answering](#). *ArXiv preprint*, abs/2010.10757.

- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Alon Jacovi, Swabha Swayamdipta, Shauli Ravfogel, Yanai Elazar, Yejin Choi, and Yoav Goldberg. 2021. [Contrastive explanations for model interpretability](#). *ArXiv preprint*, abs/2103.01378.
- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Divyansh Kaushik, Amrith Setlur, Eduard Hovy, and Zachary C Lipton. 2020. [Explaining the efficacy of counterfactually-augmented data](#). *ArXiv preprint*, abs/2010.02114.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. 2019. [Reasoning over paragraph effects in situations](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 58–62, Hong Kong, China. Association for Computational Linguistics.
- Jiangming Liu, Matt Gardner, Shay B. Cohen, and Mirella Lapata. 2020. [Multi-step inference for reasoning over paragraphs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3040–3050, Online. Association for Computational Linguistics.
- 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*, abs/1907.11692.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Pasquale Minervini and Sebastian Riedel. 2018. [Adversarially regularising neural NLI models to integrate logical background knowledge](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 65–74, Brussels, Belgium. Association for Computational Linguistics.
- Hiroshi Noji and Hiroya Takamura. 2020. [An analysis of the utility of explicit negative examples to improve the syntactic abilities of neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3375–3385, Online. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Carlos Guestrin, and Sameer Singh. 2019. [Are red roses red? evaluating consistency of question-answering models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6174–6184, Florence, Italy. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Alexis Ross, Ana Marasović, and Matthew Peters. 2021. [Explaining NLP models via minimal contrastive editing \(MiCE\)](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3840–3852, Online. Association for Computational Linguistics.
- Noah A. Smith and Jason Eisner. 2005. [Contrastive estimation: Training log-linear models on unlabeled data](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 354–362, Ann Arbor, Michigan. Association for Computational Linguistics.
- Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. 2020. [Learning what makes a difference from counterfactual examples and gradient supervision](#). *ArXiv preprint*, abs/2004.09034.
- Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. 2019. [On incorporating semantic prior knowledge in deep learning through embedding-space constraints](#). *ArXiv preprint*, abs/1909.13471.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Tongshuang Wu, Marco Túlio Ribeiro, J. Heer, and Daniel S. Weld. 2021. [Polyjuice: Automated, general-purpose counterfactual generation](#). *ArXiv preprint*, abs/2101.00288.

- 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.
- Yi-Ting Yeh and Yun-Nung Chen. 2019. [QAInfomax: Learning robust question answering system by mutual information maximization](#). 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 3370–3375, Hong Kong, China. Association for Computational Linguistics.
- Min-Ling Zhang and Zhi-Hua Zhou. 2007. Ml-knn: A lazy learning approach to multi-label learning. *Pattern recognition*, 40(7):2038–2048.
- Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach, and Ryan Cotterell. 2019. [Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, Florence, Italy. Association for Computational Linguistics.