

Unsupervised Question Decomposition for Question Answering

Ethan Perez^{1,2} Patrick Lewis^{1,3} Wen-tau Yih¹ Kyunghyun Cho^{1,2,4} Douwe Kiela¹

Abstract

We aim to improve question answering (QA) by decomposing hard questions into easier sub-questions that existing QA systems can answer. Since collecting labeled decompositions is cumbersome, we propose an unsupervised approach to produce sub-questions. Specifically, by leveraging $>10M$ questions from Common Crawl, we learn to map from the distribution of multi-hop questions to the distribution of single-hop sub-questions. We answer sub-questions with an off-the-shelf QA model and incorporate the resulting answers in a downstream, multi-hop QA system. On a popular multi-hop QA dataset, HOTPOTQA, we show large improvements over a strong baseline, especially on adversarial and out-of-domain questions. Our method is generally applicable and automatically learns to decompose questions of different classes, while matching the performance of decomposition methods that rely heavily on hand-engineering and annotation.

1. Introduction

Question answering (QA) systems have become remarkably good at answering simple, single-hop questions but still struggle with compositional, multi-hop questions (Yang et al., 2018; Hudson & Manning, 2019). In this work, we examine if we can answer hard questions by leveraging our ability to answer simple questions. Specifically, we approach QA by breaking a hard question into a series of sub-questions that can be answered by a simple, single-hop QA system. The system’s answers can then be given as input to a downstream QA system to answer the hard question, as shown in Fig. 1. Our approach thus answers the hard question in multiple, smaller steps, which can be easier than answering the hard question all at once. For example, it may be easier to answer “What profession do H. L. Mencken and

¹Facebook AI Research ²New York University ³University College London ⁴CIFAR Azrieli Global Scholar. Correspondence to: Ethan Perez <perez@nyu.edu>.

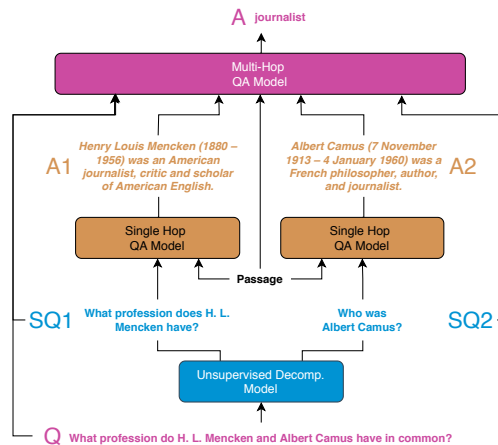


Figure 1. Overview: Using unsupervised learning, we decompose a multi-hop question into single-hop sub-questions, whose predicted answers are given to a downstream question answering model.

Albert Camus have in common?” when given the answers to the sub-questions “What profession does H. L. Mencken have?” and “Who was Albert Camus?”

Prior work in learning to decompose questions into sub-questions has relied on extractive heuristics, which generalizes poorly to different domains and question types, and requires human annotation (Talmor & Berant, 2018; Min et al., 2019b). In order to scale to any arbitrary question, we would require sophisticated natural language generation capabilities, which often relies on large quantities of high-quality supervised data. Instead, we find that it is possible to learn to decompose questions *without supervision*.

Specifically, we learn to map from the distribution of hard questions to the distribution of simpler questions. First, we automatically construct a noisy, “pseudo-decomposition” for each hard question by retrieving relevant sub-question candidates based on their similarity to the given hard question. We retrieve candidates from a corpus of 10M simple questions that we extracted from Common Crawl. Second, we train neural text generation models on that data with (1) standard sequence-to-sequence learning and (2) unsupervised sequence-to-sequence learning. The latter has the advantage that it can go beyond the noisy pairing between questions and pseudo-decompositions. Fig. 2 overviews our decomposition approach.

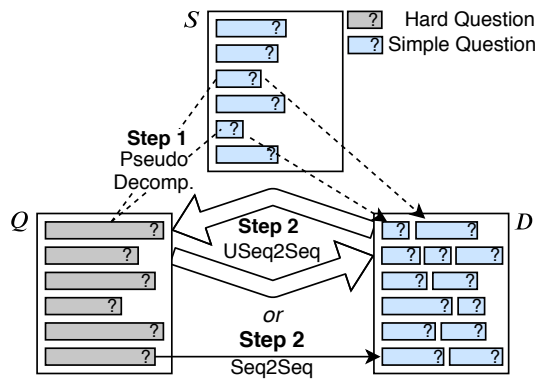


Figure 2. *Unsupervised Decomposition*: **Step 1**: We create a corpus of pseudo-decompositions D by finding candidate sub-questions from a simple question corpus S which are similar to a multi-hop question in Q . **Step 2**: We learn to map multi-hop questions to decompositions using Q and D as training data, via either standard or unsupervised sequence-to-sequence learning.

We use decompositions to improve multi-hop QA. We first use an off-the-shelf single-hop QA model to answer decomposed sub-questions. We then give each sub-question and its answer as additional input to a multi-hop QA model. We test our method on HOTPOTQA (Yang et al., 2018), a popular multi-hop QA benchmark.

Our contributions are as follows. First, QA models relying on decompositions improve accuracy over a strong baseline by 3.1 F1 on the original dev set, 11 F1 on the multi-hop dev set from Jiang & Bansal (2019a), and 10 F1 on the out-of-domain dev set from Min et al. (2019b). Our most effective decomposition model is a 12-block transformer encoder-decoder (Vaswani et al., 2017) trained using unsupervised sequence-to-sequence learning, involving masked language modeling, denoising, and back-translation objectives (Lample & Conneau, 2019). Second, our method is competitive with state-of-the-art methods SAE (Tu et al., 2020) and HGN (Fang et al., 2019) which leverage strong supervision. Third, we show that our approach automatically learns to generate useful decompositions for all 4 question types in HOTPOTQA, highlighting the general nature of our approach. In our analysis, we explore how sub-questions improve multi-hop QA, and we provide qualitative examples that highlight how question decomposition adds a form of interpretability to black-box QA models. Our ablations show that each component of our pipeline contributes to QA performance. Overall, we find that it is possible to successfully decompose questions without any supervision and that doing so improves QA.

2. Method

We now formulate the problem and overview our high-level approach, with details in the following section. We

aim to leverage a QA model that is accurate on simple questions to answer hard questions, without using supervised question decompositions. Here, we consider simple questions to be “single-hop” questions that require reasoning over one paragraph or piece of evidence, and we consider hard questions to be “multi-hop.” Our aim is then to train a multi-hop QA model M to provide the correct answer a to a multi-hop question q about a given a context c (e.g., several paragraphs). Normally, we would train M to maximize $\log p_M(a|c, q)$. To help M , we leverage a single-hop QA model that may be queried with sub-questions s_1, \dots, s_N , whose “sub-answers” to each sub-question a_1, \dots, a_N may be provided to the multi-hop QA model. M may then instead maximize the (potentially easier) objective $\log p_M(a|c, q, [s_1, a_1], \dots, [a_N, s_N])$.

Supervised decomposition models learn to map each question $q \in Q$ to a decomposition $d = [s_1; \dots; s_N]$ of N sub-questions $s_n \in S$ using annotated (q, d) examples. In this work, we do not assume access to strong (q, d) supervision. To leverage the single-hop QA model without supervision, we follow a three-stage approach: 1) map a question q into sub-questions s_1, \dots, s_N via unsupervised techniques, 2) find sub-answers a_1, \dots, a_N with the single-hop QA model, and 3) provide s_1, \dots, s_N and a_1, \dots, a_N to help predict a .

2.1. Unsupervised Question Decomposition

To train a decomposition model, we need appropriate training data. We assume access to a hard question corpus Q and a simple question corpus S . Instead of using supervised (q, d) training examples, we design an algorithm that constructs pseudo-decompositions d' to form (q, d') pairs from Q and S using an unsupervised approach (§2.1.1). We then train a model to map q to a decomposition. We explore learning to decompose with standard and unsupervised sequence-to-sequence learning (§2.1.2).

2.1.1. CREATING PSEUDO-DECOMPOSITIONS

For each $q \in Q$, we construct a pseudo-decomposition set $d' = \{s_1; \dots; s_N\}$ by retrieving simple question s from S . We concatenate all N simple questions in d' to form the pseudo-decomposition used downstream. N may be chosen based on the task or vary based on q . To retrieve useful simple questions for answering q , we face a joint optimization problem. We want sub-questions that are both (i) similar to q according to some metric f and (ii) maximally diverse:

$$d'^* = \operatorname{argmax}_{d' \subset S} \sum_{s_i \in d'} f(q, s_i) - \sum_{\substack{s_i, s_j \in d' \\ s_i \neq s_j}} f(s_i, s_j) \quad (1)$$

2.1.2. LEARNING TO DECOMPOSE

Having now retrieved relevant pseudo-decompositions, we examine different ways to learn to decompose (with implementation details in the following section):

No Learning We use pseudo-decompositions directly, employing retrieved sub-questions in downstream QA.

Sequence-to-Sequence (Seq2Seq) We train a Seq2Seq model with parameters θ to maximize $\log p_\theta(d'|q)$.

Unsupervised Sequence-to-Sequence (USeq2Seq) We start with paired (q, d') examples but do not learn from the pairing, because the pairing is noisy. We use unsupervised sequence-to-sequence learning to learn a $q \rightarrow d$ mapping instead of training directly on the noisy pairing.

2.2. Answering Sub-Questions

To answer the generated sub-questions, we use an off-the-shelf QA model. The QA model may answer sub-questions using any free-form text (i.e., a word, phrase, sentence, etc.). Any QA model is suitable, so long as it can accurately answer simple questions in S . We thus leverage good accuracy on questions in S to help QA models on questions in Q .

2.3. QA using Decompositions

Downstream QA systems may use sub-questions and sub-answers in various ways. We add sub-questions and sub-answers as auxiliary input for a downstream QA model to incorporate in its processing. We now describe the implementation details of our approach outlined above.

3. Experimental Setup

3.1. Question Answering Task

We test unsupervised decompositions on HOTPOTQA (Yang et al., 2018), a standard benchmark for multi-hop QA. We use HOTPOTQA’s “Distractor Setting,” which provides 10 context paragraphs from Wikipedia. Two (or more) paragraphs contain question-relevant sentences called “supporting facts,” and the remaining paragraphs are irrelevant, “distractor paragraphs.” Answers in HOTPOTQA are either `yes`, `no`, or a span of text in an input paragraph. Accuracy is measured with F1 and Exact Match (EM) scores between the predicted and gold spans.

3.2. Unsupervised Decomposition

3.2.1. QUESTION DATA

We use HOTPOTQA questions as our initial multi-hop, hard question corpus Q . We use SQUAD 2 questions as our

initial single-hop, simple question corpus S . However, our pseudo-decomposition corpus should be large, as the corpus will be used to train neural Seq2Seq models, which are data hungry. A larger $|S|$ will also improve the relevance of retrieved simple questions to the hard question. Thus, we take inspiration from work in machine translation on parallel corpus mining (Xu & Koehn, 2017; Artetxe & Schwenk, 2019) and in unsupervised QA (Lewis et al., 2019). We augment Q and S by mining more questions from Common Crawl. We choose sentences which start with common “wh”-words and end with “?” Next, we train a FastText classifier (Joulin et al., 2017) to classify between 60K questions sampled from Common Crawl, SQUAD 2, and HOTPOTQA. Then, we classify Common Crawl questions, adding questions classified as SQUAD 2 questions to S and questions classified as HOTPOTQA questions to Q . Question mining greatly increases the number of single-hop questions (130K \rightarrow 10.1M) and multi-hop questions (90K \rightarrow 2.4M). Thus, our unsupervised approach allows us to make use of far more data than supervised counterparts.

3.2.2. CREATING PSEUDO-DECOMPOSITIONS

To create pseudo-decompositions, we set the number N of sub-questions per question to 2, as questions in HOTPOTQA usually involve two reasoning hops. In the appendix, we discuss how our approach works when N varies per question.

Similarity-based Retrieval To retrieve question-relevant sub-questions, we embed any text t into a vector \mathbf{v}_t by summing the FastText vectors (Bojanowski et al., 2017)¹ for words in t .² We use cosine similarity as our similarity metric f . Let q be a multi-hop question used to retrieve pseudo-decomposition (s_1^*, s_2^*) , and let $\hat{\mathbf{v}}$ be the unit vector of \mathbf{v} . Since $N = 2$, Eq. 1 reduces to:

$$(s_1^*, s_2^*) = \operatorname{argmax}_{\{s_1, s_2\} \in S} [\hat{\mathbf{v}}_q^\top \hat{\mathbf{v}}_{s_1} + \hat{\mathbf{v}}_q^\top \hat{\mathbf{v}}_{s_2} - \hat{\mathbf{v}}_{s_1}^\top \hat{\mathbf{v}}_{s_2}] \quad (2)$$

The last term requires $O(|S|^2)$ comparisons, which is expensive as $|S|$ is large ($>10M$). Instead of solving Eq. (2) exactly, we find an approximate pseudo-decomposition (s_1', s_2') by computing Eq. (2) over $S' = \operatorname{topK}_{\{s \in S\}} [\hat{\mathbf{v}}_q^\top \hat{\mathbf{v}}_s]$, using $K = 1000$. We use FAISS (Johnson et al., 2017a) to efficiently build S' .

Random Retrieval For comparison, we test random pseudo-decompositions, where we randomly retrieve s_1, \dots, s_N by sampling from S . USeq2Seq trained on random $d' = [s_1; \dots; s_N]$ should at minimum learn to map q to multiple simple questions.

¹We use 300-dim. vectors trained on English Common Crawl: fasttext.cc/docs/en/english-vectors.html

²We also tried TFIDF and BERT representations but did not see significant improvements over FastText (see the appendix).

Editing Pseudo-Decompositions Since the sub-questions are retrieval-based, the sub-questions are often not about the same entities as q . As a post-processing step, we replace entities in (s'_1, s'_2) with entities from q . We find all entities in (s'_1, s'_2) that do not appear in q using spaCy (Honnibal & Montani, 2017). We replace these entities with a random entity from q with the same type (e.g., “Date” or “Location”) if and only if one exists. We use entity replacement on pseudo-decompositions from both random and similarity-based retrieval.

3.2.3. UNSUPERVISED DECOMPOSITION MODELS

Pre-training Pre-training is a key ingredient for unsupervised Seq2Seq methods (Artetxe et al., 2018; Lample et al., 2018), so we initialize all decomposition models with the same pre-trained weights, regardless of training method (Seq2Seq or USeq2Seq). We warm-start our pre-training with the pre-trained, English Masked Language Model (MLM) from Lample & Conneau (2019), a 12-block decoder-only transformer model (Vaswani et al., 2017) trained to predict masked-out words on Toronto Books Corpus (Zhu et al., 2015) and Wikipedia. We train the model with the MLM objective for one epoch on the augmented corpus Q (2.4 M questions), while also training on decompositions D formed via random retrieval from S . For our pre-trained encoder-decoder, we initialize a 6-block encoder with the first 6 MLM blocks, and we initialize a 6-block decoder with the last 6 MLM blocks, randomly initializing the remaining weights as in Lample & Conneau (2019).

Seq2Seq We fine-tune the pre-trained encoder-decoder using maximum likelihood. We stop training based on validation BLEU (Papineni et al., 2002) between generated decompositions and pseudo-decompositions.

USeq2Seq We follow the approach by Lample & Conneau (2019) in unsupervised translation.³ Training follows two stages: (1) MLM pre-training on the training corpora (described above), followed by (2) training simultaneously with denoising and back-translation objectives. For denoising, we produce a noisy input \hat{d} by randomly masking, dropping, and locally shuffling tokens in $d \sim D$, and we train a model with parameters θ to maximize $\log p_\theta(d|\hat{d})$. We likewise maximize $\log p_\theta(q|\hat{q})$. For back-translation, we generate a multi-hop question \hat{q} for a decomposition $d \sim D$, and we maximize $\log p_\theta(d|\hat{q})$. Similarly, we maximize $\log p_\theta(q|\hat{d})$. To stop training without supervision, we use a modified version of round-trip BLEU (Lample et al., 2018) (see the appendix for details). We train with denoising and back-translation on smaller corpora of HOTPOTQA questions (Q) and their pseudo-decompositions (D).⁴

³We use github.com/facebookresearch/XLM

⁴Using the augmented corpora here did not improve QA.

3.3. Single-hop Question Answering Model

We train our single-hop QA model following prior work from Min et al. (2019b) on HOTPOTQA.⁵

Model Architecture We fine-tune a pre-trained model to take a question and several paragraphs and predicts the answer, similar to the single-hop QA model from Min et al. (2019a). The model computes a separate forward pass on each paragraph (with the question). For each paragraph, the model learns to predict the answer span if the paragraph contains the answer and to predict “no answer” otherwise. We treat `yes` and `no` predictions as spans within the passage (prepended to each paragraph), as in Nie et al. (2019) on HOTPOTQA. During inference, for the final softmax, we consider all paragraphs as a single chunk. Similar to Clark & Gardner (2018), we subtract a paragraph’s “no answer” logit from the logits of all spans in that paragraph, to reduce or increase span probabilities accordingly. In other words, we compute the probability $p(s_p)$ of each span s_p in a paragraph $p \in \{1, \dots, P\}$ using the predicted span logit $l(s_p)$ and “no answer” paragraph logit $n(p)$ as follows:

$$p(s_p) = \frac{e^{l(s_p) - n(p)}}{\sum_{p'=1}^P \sum_{s'_p} e^{l(s'_p) - n(p')}} \quad (3)$$

We use ROBERTA_{LARGE} (Liu et al., 2019) as our pre-trained initialization. Later, we also experiment with using the BERT_{BASE} ensemble from Min et al. (2019b).

Training Data and Ensembling Similar to Min et al. (2019b), we train an ensemble of 2 single-hop QA models using data from SQUAD 2 and HOTPOTQA questions labeled as “easy” (single-hop). To ensemble, we average the logits of the two models before predicting the answer. SQUAD is a single-paragraph QA task, so we adapt SQUAD to the multi-paragraph setting by retrieving distractor paragraphs from Wikipedia for each question. We use the TFIDF retriever from DrQA (Chen et al., 2017) to retrieve 2 distractor paragraphs, which we add to the input for one model in the ensemble. We drop words from the question with a 5% probability to help the model handle any ill-formed sub-questions. We use the single-hop QA ensemble as a black-box model once trained, never training the model on multi-hop questions.

Returned Text We have the single-hop QA model return the sentence containing the model’s predicted answer span, alongside the sub-questions. Later, we compare against alternatives, i.e., returning the predicted answer span without its context or not returning sub-questions.

⁵Our code is based on `transformers` (Wolf et al., 2019)

Unsupervised Question Decomposition for Question Answering

Decomp. Method	Pseudo-Decomps.	HOTPOTQA F1		
		Orig	MultiHop	OOD
\times	\times (1hop)	66.7	63.7	66.5
\times	\times (Baseline)	77.0 \pm .2	65.2 \pm .2	67.1 \pm .5
No Learn	Random	78.4 \pm .2	70.9 \pm .2	70.7 \pm .4
	FastText	78.9 \pm .2	72.4 \pm .1	72.0 \pm .1
Seq2Seq	Random	77.7 \pm .2	69.4 \pm .3	70.0 \pm .7
	FastText	78.9 \pm .2	73.1 \pm .2	73.0 \pm .3
USeq2Seq	Random	79.8 \pm .1	76.0 \pm .2	76.5 \pm .2
	FastText	80.1\pm.2	76.2\pm.1	77.1\pm.1
DecompRC*		79.8 \pm .2	76.3 \pm .4	77.7 \pm .2
SAE (Tu et al., 2020) †		80.2	61.1	62.6
HGN (Fang et al., 2019) †		82.2	78.9 \ddagger	76.1 \ddagger
		Ours	SAE \ddagger	HGN \ddagger
Test (EM/F1)		66.33/79.34	66.92/79.62	69.22/82.19

Table 1. Unsupervised decompositions significantly improve the F1 on HOTPOTQA over the baseline. We achieve comparable F1 to methods which use supporting fact supervision (†). (*) We use supervised and heuristic decompositions from Min et al. (2019b). (‡) Scores are approximate due to mismatched Wikipedia dumps.

3.4. Multi-hop Question Answering Model

Our multi-hop QA architecture is identical to the single-hop QA model, but the multi-hop QA model also uses sub-questions and sub-answers as input. We append each (sub-question, sub-answer) pair in order to the multi-hop question along with separator tokens. We train one multi-hop QA model on all of HOTPOTQA, also including SQUAD 2 examples used to train the single-hop QA model. Later, we experiment with using BERT_{LARGE} and BERT_{BASE} instead of ROBERTA_{LARGE} as the multi-hop QA model. All reported error margins show the mean and std. dev. across 5 multi-hop QA training runs using the same decompositions.

4. Results on Question Answering

We compare variants of our approach that use different learning methods and different pseudo-aligned training sets. As a baseline, we compare ROBERTA with decompositions to a ROBERTA model that does not use decompositions but is identical in all other respects. We train the baseline for 2 epochs, sweeping over batch size $\in \{64, 128\}$, learning rate $\in \{1 \times 10^{-5}, 1.5 \times 10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}\}$, and weight decay $\in \{0, 0.1, 0.01, 0.001\}$; we choose the hyperparameters that perform best on our dev set. We then use the best hyperparameters for the baseline to train our ROBERTA models with decompositions.

We report results on 3 versions of the dev set: (1) the original version,⁶ (2) the multi-hop version from Jiang & Bansal (2019a) which created some distractor paragraphs adversari-

⁶The test set is private, so we randomly halve the dev set to form validation and held-out dev sets. We will release our splits.

ally to test multi-hop reasoning, and (3) the out-of-domain version from Min et al. (2019b) which retrieved distractor paragraphs using the same procedure as the original version, but excluded paragraphs in the original version.

Main Results Table 1 shows how unsupervised decompositions affect QA. Our ROBERTA baseline performs quite well on HOTPOTQA (77.0 F1), despite processing each paragraph separately, which prohibits inter-paragraph reasoning. The result is in line with prior work which found that a version of our baseline QA model using BERT (Devlin et al., 2019) does well on HOTPOTQA by exploiting single-hop reasoning shortcuts (Min et al., 2019a). We achieve significant gains over our strong baseline by leveraging decompositions from our best decomposition model, trained with USeq2Seq on FastText pseudo-decompositions; we find a 3.1 F1 gain on the original dev set, 11 F1 gain on the multi-hop dev set, and 10 F1 gain on the out-of-domain dev set. Unsupervised decompositions even match the performance of using (within our pipeline) supervised and heuristic decompositions from DECOMP RC (i.e., 80.1 vs. 79.8 F1 on the original dev set).

More generally, all decomposition methods improve QA over the baseline by leveraging the single-hop QA model (“1hop” in Table 1). Using FastText pseudo-decompositions as sub-questions directly improves QA over using random sub-questions on the multi-hop set (72.4 vs. 70.9 F1) and out-of-domain set (72.0 vs. 70.7 F1). USeq2Seq on random pseudo-decompositions also improves over the random sub-question baseline (e.g., 79.8 vs. 78.4 F1 on HOTPOTQA). However, we only find small improvements when training USeq2Seq on FastText vs. Random pseudo-decompositions (e.g., 77.1 vs. 76.5 F1 on the out-of-domain dev set).

The best decomposition methods learn with USeq2Seq. Using Seq2Seq to generate decompositions gives similar QA accuracy as the “No Learning” setup, e.g. both approaches achieve 78.9 F1 on the original dev set for FastText pseudo-decompositions. The results are similar perhaps since supervised learning is directly trained to place high probability on pseudo-decompositions.⁷ USeq2Seq may improve over Seq2Seq by learning to align hard questions and pseudo-decompositions while ignoring the noisy pairing.

After our experimentation, we chose USeq2Seq trained on FastText pseudo-decompositions as the final model, and we submitted the model for hidden test evaluation. Our approach achieved a test F1 of 79.34 and Exact Match (EM) of 66.33. Our approach is competitive with concurrent, state-of-the-art systems SAE (Tu et al., 2020) and HGN (Fang

⁷We also tried using the Seq2Seq model to initialize USeq2Seq. Seq2Seq initialization resulted in comparable or worse downstream QA accuracy, suggesting that pre-training on noisy decompositions did not help bootstrap USeq2Seq (see appendix for details).

Decomps.	Bridge	Comp.	Intersec.	Single-hop
✗	80.7±.2	73.8±.4	78.1±.6	73.8±.6
✓	82.3±.4	80.1±.3	81.2±.4	76.7±.6

Table 2. F1 scores on 4 types of questions in HOTPOTQA. Unsupervised decompositions improves QA for all types.

SubQs	SubAs	QA F1
✗	✗	77.0±.2
✓	Sentence	80.1±.2
✓	Span	77.8±.3
✓	Random Entity	76.9±.2
✓	✗	76.9±.2
✗	Sentence	80.2±.1

Table 3. **Ablation Study:** QA model F1 when trained with different sub-answers: the sentence containing the predicted sub-answer, the predicted sub-answer span, and a random entity from the context. We also train QA models with (✓) or without (✗) sub-questions and sub-answers.

et al., 2019), which both (unlike our approach) learn from additional, strong supervision about which sentences are necessary to answer the question.

4.1. Question Type Breakdown

To understand where decompositions help, we break down QA performance across 4 question types from Min et al. (2019b). “Bridge” questions ask about an entity not explicitly mentioned in the question (“When was Erik Watts’ father born?”). “Intersection” questions ask to find an entity that satisfies multiple separate conditions (“Who was on CNBC and Fox News?”). “Comparison” questions ask to compare a property of two entities (“Which is taller, Momhil Sar or K2?”). “Single-hop” questions are likely answerable using single-hop shortcuts or single-paragraph reasoning (“Where is Electric Six from?”). We split the original dev set into the 4 types using the supervised type classifier from Min et al. (2019b). Table 2 shows F1 scores for ROBERTA with and without decompositions across the 4 types.

Unsupervised decompositions improve QA across all question types. Our single decomposition model generates useful sub-questions for all question types without special case handling, unlike earlier work from Min et al. (2019b) which handled each question type separately. For single-hop questions, our QA approach does not require falling back to a single-hop QA model and instead learns to leverage decompositions to better answer questions with single-hop shortcuts (76.7 vs. 73.8 F1 without decompositions).

4.2. Answers to Sub-Questions are Crucial

To measure the usefulness of sub-questions and sub-answers, we train the multi-hop QA model with various, ablated in-




Figure 3. Multi-hop QA is better when the single-hop QA model answers with the ground truth “supporting fact” sentences. We plot mean and std. across 5 random QA training runs.

puts, as shown in Table 3. Sub-answers are crucial to improving QA, as sub-questions with no answers or random answers do not help (76.9 vs. 77.0 F1 for the baseline). Only when sub-answers are provided do we see improved QA, with or without sub-questions (80.1 and 80.2 F1, respectively). It is important to provide the sentence containing the predicted answer span instead of the answer span alone (80.1 vs. 77.8 F1, respectively), though the answer span alone still improves over the baseline (77.0 F1).

4.3. How Do Decompositions Help?

Decompositions help to answer questions by retrieving important supporting evidence to answer questions. Fig. 3 shows that multi-hop QA accuracy increases when the sub-answer sentences are the “supporting facts” or sentences needed to answer the question, as annotated by HOTPOTQA. We retrieve supporting facts without learning to predict them with strong supervision, unlike many state-of-the-art models (Tu et al., 2020; Fang et al., 2019; Nie et al., 2019).

4.4. Example Decompositions

To illustrate how decompositions help QA, Table 4 shows example sub-questions from our best decomposition model with predicted sub-answers. Sub-questions are single-hop questions relevant to the multi-hop question. The single-hop QA model returns relevant sub-answers, sometimes in spite of grammatical errors (Q_1, SQ_1) or under-specified questions (Q_2, SQ_1). The multi-hop QA model then returns an answer consistent with the predicted sub-answers. The decomposition model is largely extractive, copying from the multi-hop question rather than hallucinating new entities, which helps generate relevant sub-questions. To better understand our system, we analyze the model for each stage: decomposition, single-hop QA, and multi-hop QA.

Q1: Are both Coldplay and Pierre Bouvier from the same country?
SQ₁: Where are Coldplay and Coldplay from?
└ Coldplay are a <u>British</u> rock band formed in 1996 by lead vocalist and keyboardist Chris Martin and lead guitarist Jonny Buckland at University College London (UCL).
SQ₂: What country is Pierre Bouvier from?
└ Pierre Charles Bouvier (born 9 May 1979) is a <u>Canadian</u> singer, songwriter, musician, composer and actor who is best known as the lead singer and guitarist of the rock band Simple Plan.
Â: No
Q2: How many copies of Roald Dahl’s variation on a popular anecdote sold?
SQ₁: How many copies of Roald Dahl’s?
└ His books have sold more than <u>250 million</u> copies worldwide.
SQ₂: What is the name of the variation on a popular anecdote?
└ “Mrs. Bixby and the Colonel’s Coat” is a short story by <u>Roald Dahl</u> that first appeared in the 1959 issue of Nugget.
Â: more than 250 million
Q3: Who is older, Annie Morton or Terry Richardson?
SQ₁: Who is Annie Morton?
└ Annie Morton (born October 8, 1970) is an <u>American model</u> born in Pennsylvania.
SQ₂: When was Terry Richardson born?
└ Kenton Terry Richardson (born 26 July 1999) is an English professional footballer who plays as a defender for League Two side Hartlepool United.
Â: Annie Morton

Table 4. Example sub-questions generated by our model, along with predicted sub-answer sentences (answer span underlined) and final predicted answer.

5. Analysis

5.1. Unsupervised Decomposition Model

Intrinsic Evaluation of Decompositions We evaluate the quality of decompositions on other metrics aside from downstream QA. To measure the fluency of decompositions, we compute the likelihood of decompositions using the pre-trained GPT-2 language model (Radford et al., 2019). We train a classifier on the question-wellformedness dataset of Faruqi & Das (2018), and we use the classifier to estimate the proportion of sub-questions that are well-formed. We measure how abstractive decompositions are by computing (i) the token Levenstein distance between the multi-hop question and its generated decomposition and (ii) the ratio between the length of the decomposition and the length of the multi-hop question. We compare our best decomposition model against the supervised+heuristic decompositions from DECOMPRC (Min et al., 2019b) in Table 5.

Unsupervised decompositions are both more natural and well-formed than decompositions from DECOMPRC. Unsupervised decompositions are also closer in edit distance and length to the multi-hop question, consistent with our obser-

Decomp. Method	GPT2 NLL	% Well-Formed	Edit Dist.	Length Ratio
USeq2Seq	5.56	60.9	5.96	1.08
DecompRC	6.04	32.6	7.08	1.22

Table 5. Analysis of sub-questions produced by our method vs. the supervised+heuristic method of Min et al. (2019a). From left-to-right: Negative Log-Likelihood (NLL) according to GPT2 (lower is better), % Well-Formed according to a classifier, Edit Distance between decomposition and multi-hop question, and token-wise Length Ratio between decomposition and multi-hop question.




Figure 4. **Left:** We decode from the decomposition model with beam search and use n^{th} -ranked hypothesis as a question decomposition. We plot the F1 of a multi-hop QA model trained to use the n^{th} -ranked decomposition. **Right:** Multi-hop QA is better when the single-hop QA model places high probability on its sub-answer.

vation that our decomposition model is largely extractive.

Quality of Decomposition Model Another way to test the quality of the decomposition model is to test if the model places higher probability on decompositions that are more helpful for downstream QA. We generate $N = 5$ hypotheses from our best decomposition model using beam search, and we train a multi-hop QA model to use the n^{th} -ranked hypothesis as a question decomposition (Fig. 4, left). QA accuracy decreases as we use lower probability decompositions, but accuracy remains relatively robust, at most decreasing from 80.1 to 79.3 F1. The limited drop suggests that decompositions are still useful if they are among the model’s top hypotheses, another indication that our model is trained well for decomposition.

5.2. Single-hop Question Answering Model

Sub-Answer Confidence Figure 4 (right) shows that the model’s sub-answer confidence correlates with downstream multi-hop QA performance for all HOTPOTQA dev sets. A low confidence sub-answer may be indicative of (i) an unanswerable or ill-formed sub-question or (ii) a sub-answer that is more likely to be incorrect. In both cases, the single-hop QA model is less likely to retrieve the useful supporting evidence to answer the multi-hop question.

Changing the Single-hop QA Model We find that our approach is robust to the single-hop QA model that answers sub-questions. We use the BERT_{BASE} ensemble from Min et al. (2019b) as the single-hop QA model. The model performs much worse compared to our ROBERTA_{LARGE} single-hop ensemble when used directly on HOTPOTQA (56.3 vs. 66.7 F1). However, the model results in comparable QA when used to answer single-hop sub-questions within our larger system (79.9 vs. 80.1 F1 for our ROBERTA_{LARGE} ensemble).

5.3. Multi-hop Question Answering Model

Varying the Base Model To understand how decompositions impact performance as the multi-hop QA model gets stronger, we vary the base pre-trained model. Table 6 shows the impact of adding decompositions to BERT_{BASE}, BERT_{LARGE}, and finally ROBERTA_{LARGE} (see the appendix for hyperparameters). The gain from using decompositions grows with strength of the multi-hop QA model. Decompositions improve QA by 1.2 F1 for a BERT_{BASE} model, by 2.6 F1 for the stronger BERT_{LARGE} model, and by 3.1 F1 for our best ROBERTA_{LARGE} model.

6. Related Work

Answering complicated questions has been a long-standing challenge in natural language processing. To this end, prior work has explored decomposing questions with supervision or heuristic algorithms. IBM Watson (Ferrucci et al., 2010) decomposes questions into sub-questions in multiple ways or not at all. DECOMPRC (Min et al., 2019b) largely frames sub-questions as extractive spans of a multi-hop question, learning to predict span-based sub-questions via supervised learning on human annotations. In other cases, DECOMPRC decomposes a multi-hop question using a heuristic algorithm, or DecomPRC does not decompose at all. Watson and DECOMPRC use special case handling to decompose different questions, while our algorithm is fully automated and requires minimal hand-engineering.

More traditional, semantic parsing methods map questions to compositional programs, whose sub-programs can be viewed as question decompositions in a formal language (Talmor & Berant, 2018; Wolfson et al., 2020). Examples include classical QA systems like SHRDLU (Winograd, 1972) and LUNAR (Woods et al., 1974), as well as neural Seq2Seq semantic parsers (Dong & Lapata, 2016) and neural module networks (Andreas et al., 2015; 2016). Such methods usually require strong, program-level supervision to generate programs, as in visual QA (Johnson et al., 2017b) and on HOTPOTQA (Jiang & Bansal, 2019b). Some models use other forms of strong supervision, e.g. predicting the “supporting evidence” to answer a question annotated by

Multi-hop QA Model	QA F1 (w/o → w/ Decomps.)
BERT _{BASE}	71.8±.4 → 73.0±.4
BERT _{LARGE}	76.4±.2 → 79.0±.1
ROBERTA _{LARGE}	77.0±.3 → 80.1±.2

Table 6. Stronger QA models benefit more from decompositions.

HOTPOTQA. Such an approach is taken by SAE (Tu et al., 2020) and HGN (Fang et al., 2019), whose methods may be combined with our approach.

Unsupervised decomposition complements strongly and weakly supervised decomposition approaches. Our unsupervised approach enables methods to leverage millions of otherwise unusable questions, similar to work on unsupervised QA (Lewis et al., 2019). When decomposition supervision exists, supervised and unsupervised learning can be used in tandem to learn from both labeled and unlabeled examples. Such semi-supervised methods outperform supervised learning for tasks like machine translation (Sennrich et al., 2016). Other work on weakly supervised question generation uses a downstream QA model’s accuracy as a signal for learning to generate useful questions. Weakly supervised question generation often uses reinforcement learning (Nogueira & Cho, 2017; Wang & Lake, 2019; Strub et al., 2017; Das et al., 2017; Liang et al., 2018), where an unsupervised initialization can greatly mitigate the issues of exploring from scratch (Jaderberg et al., 2017).

7. Conclusion

We proposed an algorithm that decomposes questions without supervision, using 3 stages: (1) learning to decompose using pseudo-decompositions without supervision, (2) answering sub-questions with an off-the-shelf QA system, and (3) answering hard questions more accurately using sub-questions and their answers as additional input. When evaluated on HOTPOTQA, a standard benchmark for multi-hop QA, our approach significantly improved accuracy over an equivalent model that did not use decompositions. Our approach relies only on the final answer as supervision but works as effectively as state-of-the-art methods that rely on strong supervision, such as supporting fact labels or example decompositions. Qualitatively, we found that unsupervised decomposition resulted in fluent sub-questions whose answers often match the annotated supporting facts in HOTPOTQA. Our unsupervised decompositions are largely extractive, which is effective for compositional, multi-hop questions but not all complex questions, showing room for future work. Overall, this work opens up exciting avenues for leveraging methods in unsupervised learning and natural language generation to improve the interpretability and generalization of machine learning systems.

References

- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. Neural module networks. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 39–48, 2015. URL <https://arxiv.org/abs/1511.02799>.
- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. Learning to compose neural networks for question answering. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1545–1554, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1181. URL <https://www.aclweb.org/anthology/N16-1181>.
- Artetxe, M. and Schwenk, H. Margin-based parallel corpus mining with multilingual sentence embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3197–3203, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1309. URL <https://www.aclweb.org/anthology/P19-1309>.
- Artetxe, M., Labaka, G., Agirre, E., and Cho, K. Unsupervised neural machine translation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Sy2ogebAW>.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5: 135–146, 2017. doi: 10.1162/tacl.a.00051. URL <https://www.aclweb.org/anthology/Q17-1010>.
- Chen, D., Fisch, A., Weston, J., and Bordes, A. Reading Wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1870–1879, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1171. URL <https://www.aclweb.org/anthology/P17-1171>.
- Clark, C. and Gardner, M. Simple and effective multi-paragraph reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 845–855, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1078. URL <https://www.aclweb.org/anthology/P18-1078>.
- Das, A., Kottur, S., Moura, J. M., Lee, S., and Batra, D. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. URL <https://arxiv.org/abs/1511.02799>.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pp. 4171–4186, June 2019. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- Dong, L. and Lapata, M. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 33–43, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1004. URL <https://www.aclweb.org/anthology/P16-1004>.
- Fang, Y., Sun, S., Gan, Z., Pillai, R., Wang, S., and Liu, J. Hierarchical graph network for multi-hop question answering, 2019. URL <https://arxiv.org/abs/1911.03631>.
- Faruqui, M. and Das, D. Identifying well-formed natural language questions. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 798–803, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1091. URL <https://www.aclweb.org/anthology/D18-1091>.
- Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A. A., Lally, A., Murdock, J. W., Nyberg, E., Prager, J., Schlaefter, N., and Welty, C. Building watson: An overview of the deepqa project. *AI Magazine*, 31(3):59–79, Jul. 2010. doi: 10.1609/aimag.v31i3.2303. URL <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2303>.
- Honnibal, M. and Montani, I. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear, 2017. URL <https://github.com/explosion/spaCy>.
- Hudson, D. A. and Manning, C. D. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. URL <https://arxiv.org/abs/1902.09506>.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=SJ6yPD5xg>.

- Jiang, Y. and Bansal, M. Avoiding reasoning shortcuts: Adversarial evaluation, training, and model development for multi-hop QA. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2726–2736, Florence, Italy, July 2019a. Association for Computational Linguistics. doi: 10.18653/v1/P19-1262. URL <https://www.aclweb.org/anthology/P19-1262>.
- Jiang, Y. and Bansal, M. Self-assembling modular networks for interpretable multi-hop reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, Hong Kong, China, nov 2019b. Association for Computational Linguistics. URL <https://arxiv.org/abs/1909.05803>.
- Johnson, J., Douze, M., and Jégou, H. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017a.
- Johnson, J., Hariharan, B., van der Maaten, L., Hoffman, J., Fei-Fei, L., Zitnick, C. L., and Girshick, R. Inferring and executing programs for visual reasoning. In *ICCV*, 2017b. URL <https://arxiv.org/abs/1705.03633>.
- Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/E17-2068>.
- Lample, G. and Conneau, A. Cross-lingual language model pretraining. In *NeurIPS*, 2019. URL <https://arxiv.org/abs/1901.07291>.
- Lample, G., Conneau, A., Denoyer, L., and Ranzato, M. Unsupervised machine translation using monolingual corpora only. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkYTTf-AZ>.
- Lewis, P., Denoyer, L., and Riedel, S. Unsupervised question answering by cloze translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4896–4910, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1484. URL <https://www.aclweb.org/anthology/P19-1484>.
- Liang, C., Norouzi, M., Berant, J., Le, Q. V., and Lao, N. Memory augmented policy optimization for program synthesis and semantic parsing. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 9994–10006. Curran Associates, Inc., 2018. URL <https://arxiv.org/abs/1807.02322>.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. URL <https://arxiv.org/abs/1907.11692>.
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., and Wu, H. Mixed precision training. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1gs9JgRZ>.
- Min, S., Wallace, E., Singh, S., Gardner, M., Hajishirzi, H., and Zettlemoyer, L. Compositional questions do not necessitate multi-hop reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4249–4257, Florence, Italy, July 2019a. Association for Computational Linguistics. doi: 10.18653/v1/P19-1416. URL <https://www.aclweb.org/anthology/P19-1416>.
- Min, S., Zhong, V., Zettlemoyer, L., and Hajishirzi, H. Multi-hop reading comprehension through question decomposition and rescoring. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6097–6109, Florence, Italy, July 2019b. Association for Computational Linguistics. doi: 10.18653/v1/P19-1613. URL <https://www.aclweb.org/anthology/P19-1613>.
- Nie, Y., Wang, S., and Bansal, M. Revealing the importance of semantic retrieval for machine reading at scale. In *2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019. URL <https://www.aclweb.org/anthology/D19-1258/>.
- Nogueira, R. and Cho, K. Task-oriented query reformulation with reinforcement learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 574–583, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1061. URL <https://www.aclweb.org/anthology/D17-1061>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/

- 1073083.1073135. URL <https://www.aclweb.org/anthology/P02-1040>.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. 2019. URL https://d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Sennrich, R., Haddow, B., and Birch, A. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 86–96, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1009. URL <https://www.aclweb.org/anthology/P16-1009>.
- Strub, F., de Vries, H., Mary, J., Piot, B., Courville, A., and Pietquin, O. End-to-end optimization of goal-driven and visually grounded dialogue systems. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pp. 2765–2771, 2017. doi: 10.24963/ijcai.2017/385. URL <https://doi.org/10.24963/ijcai.2017/385>.
- Talmor, A. and Berant, J. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 641–651, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1059. URL <https://www.aclweb.org/anthology/N18-1059>.
- Tu, M., Huang, K., Wang, G., Huang, J., He, X., and Zhou, B. Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents. In *AAAI 2020 (accepted)*, 2020. URL <https://arxiv.org/abs/1911.00484>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *NeurIPS*, pp. 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Wang, Z. and Lake, B. M. Modeling question asking using neural program generation. *CoRR*, abs/1907.09899, 2019. URL <http://arxiv.org/abs/1907.09899>.
- Winograd, T. *Understanding Natural Language*. Academic Press, Inc., USA, 1972. ISBN 0127597506. URL <https://www.sciencedirect.com/science/article/pii/0010028572900023>.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv*, abs/1910.03771, 2019. URL <https://arxiv.org/abs/1910.03771>.
- Wolfson, T., Geva, M., Gupta, A., Gardner, M., Goldberg, Y., Deutch, D., and Berant, J. Break it down: A question understanding benchmark. *Transactions of the Association for Computational Linguistics*, 2020. URL <https://arxiv.org/abs/2001.11770>.
- Woods, W., Kaplan, R., and Nash-Webber, B. The lunar sciences natural language information system. Final Report 2378, Bolt, Beranek and Newman, Inc., Cambridge, MA, 1974. URL <https://dl.acm.org/doi/10.1145/1499586.1499695>.
- Xu, H. and Koehn, P. Zipporah: a fast and scalable data cleaning system for noisy web-crawled parallel corpora. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2945–2950, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1319. URL <https://www.aclweb.org/anthology/D17-1319>.
- Yang, Z., Qi, P., Zhang, S., Bengio, Y., Cohen, W., Salakhutdinov, R., and Manning, C. D. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL <https://www.aclweb.org/anthology/D18-1259>.
- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., and Fidler, S. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV 15, pp. 1927, USA, 2015. IEEE Computer Society. ISBN 9781467383912. doi: 10.1109/ICCV.2015.11. URL <https://doi.org/10.1109/ICCV.2015.11>.

A. Pseudo-Decompositions

Tables 8-13 show examples of pseudo-decompositions and learned decompositions from various models.

A.1. Variable Length Pseudo-Decompositions

In §3.2.2, we leveraged domain knowledge about the task to fix the pseudo-decomposition length $N = 2$. A general algorithm for creating pseudo-decompositions should find a suitable N for each question. We find that Eq. 1 in 2.1.1 always results in decompositions of length $N = 2$, as the regularization term grows quickly with N . Thus, we test another formulation based on Euclidean distance:

$$D'^* = \operatorname{argmin}_{D' \subset S} \left\| \mathbf{v}_q - \sum_{s \in D'} \mathbf{v}_s \right\|_2 \quad (4)$$

We create pseudo-decompositions in an similar way as before, first finding a set of candidate sub-questions $S' \in S$ with high cosine similarity to \mathbf{v}_q , then performing beam search up to a maximum value of N . We test pseudo-decomposition formulations by creating synthetic compositional questions by combining 2-3 single-hop questions with “and.” We then measure the ranking of the correct decomposition (a concatenation of the single-hop questions). For $N = 2$, both methods perform well, but Eq. 1 does not work for decompositions where $N = 3$, whereas Eq. 4 does, achieving a mean reciprocal rank of 30%. However, Eq. 1 outperforms Eq. 4 on HOTPOTQA, and Eq. 1 is faster to compute and easier to scale. Moreover, Eq. 4 requires an embedding space where summing sub-question representations is meaningful, whereas Eq. 1 only requires embeddings that encode semantic similarity. Thus, we adopt Eq. 1 for our main experiments. Table 8 contains an example where the variable length decomposition method mentioned above produces a three-subquestion decomposition whereas the other methods are fixed to two subquestions.

A.2. Impact of Question Corpus Size

In addition to our previous results on FastText vs. Random pseudo-decompositions, we found it important to use a large question corpus to create pseudo-decompositions. QA F1 increased from 79.2 to 80.1 when we trained decomposition models on pseudo-decompositions comprised of questions retrieved from Common Crawl (>10M questions) rather than only SQUAD 2 (~130K questions), using an appropriately larger beam size (100 → 1000).

A.3. Pseudo-Decomposition Retrieval Method

Table 7 shows QA results with pseudo-decompositions retrieved using sum-bag-of-word representations from FastText, TFIDF, BERT_{LARGE} first layer hidden states. We also

Decomp. Method	Pseudo-Decomps.	HOTPOTQA F1		
		Dev	Advers.	OOD
\times	\times (1hop)	66.7	63.7	66.5
\times	\times (Baseline)	77.0±.2	65.2±.2	67.1±.5
No Learn	Random	78.4±.2	70.9±.2	70.7±.4
	BERT	78.9±.4	71.5±.3	71.5±.2
	TFIDF	79.2±.3	72.2±.3	72.0±.5
	FastText	78.9±.2	72.4±.1	72.0±.1
Seq2Seq	Random	77.7±.2	69.4±.3	70.0±.7
	BERT	79.1±.3	72.6±.3	73.1±.3
	TFIDF	79.2±.1	73.0±.3	72.9±.3
	FastText	78.9±.2	73.1±.2	73.0±.3
CSeq2Seq	Random	79.4±.2	75.1±.2	75.2±.4
	BERT	78.9±.2	74.9±.1	75.2±.2
	TFIDF	78.6±.3	72.4±.4	72.8±.2
	FastText	79.9±.2	76.0±.1	76.9±.1
USeq2Seq	Random	79.8±.1	76.0±.2	76.5±.2
	BERT	79.8±.3	76.2±.3	76.7±.3
	TFIDF	79.6±.2	75.5±.2	76.0±.2
	FastText	80.1±.2	76.2±.1	77.1±.1
DecompRC		79.8±.2	76.3±.4	77.7±.2
SAE (Tu et al., 2020)		80.2	61.1	62.6
HGN (Fang et al., 2019)		82.2	78.9	76.1

Table 7. QA F1 scores for all combinations of learning methods and pseudo-decomposition retrieval methods that we tried.

vary the learning method and include results Curriculum Seq2Seq (CSeq2Seq), where we initialize the USeq2Seq approach with the Seq2Seq model trained on the same data.

B. Unsupervised Decomposition Model

B.1. Unsupervised Stopping Criterion

To stop USeq2Seq training, we use an unsupervised stopping criterion to avoid relying on a supervised validation set of decompositions. We generate a decomposition \hat{d} for a multi-hop question q , and we measure BLEU between q and the model-generated question \hat{q} for \hat{d} , similar to round-trip BLEU in unsupervised translation (Lample et al., 2018). We scale round-trip BLEU score by the fraction of “good” decompositions, where a good decomposition has (1) 2 sub-questions (question marks), (2) no sub-question which contains all words in the multi-hop question, and (3) no sub-question longer than the multi-hop question. Without scaling, decomposition models achieve perfect round-trip BLEU by copying the multi-hop question as the decomposition. We measure scaled BLEU across multi-hop questions in HOTPOTQA dev, and we stop training when the metric does not increase for 3 consecutive epochs.

It is possible to stop training the decomposition model based on downstream QA accuracy. However, training a QA model on each decomposition model checkpoint (1) is computationally expensive and (2) ties decompositions to a specific, downstream QA model. In Figure 5, we show

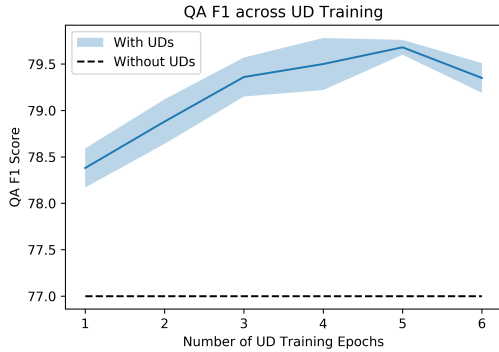


Figure 5. How multi-hop QA accuracy varies over the course of decomposition model training, for one training run of USeq2Seq on FastText pseudo-decompositions. Our unsupervised stopping criterion selects the epoch 3 checkpoint, which performs roughly as well as the best checkpoint (epoch 5).

downstream QA results across various USeq2Seq checkpoints when using the BERT_{BASE} single-hop QA ensemble from (Min et al., 2019b). The unsupervised stopping criterion does not significantly hurt downstream QA compared to using a weakly-supervised stopping criterion.

B.2. Training Hyperparameters

MLM Pre-training We pre-train our encoder-decoder distributed across 8 DGX-1 machines, each with 8, 32GB NVIDIA V100 GPUs interconnected by Infiniband. We pre-train using the largest possible batch size (1536), and we choose the best learning rate (3×10^{-5}) based on training loss after a small number of iterations. We chose a maximum sequence length of 128. We keep other hyperparameters identical to those from Lample & Conneau (2019) used in unsupervised translation.

USeq2Seq We train each decomposition model with distributed training across 8, 32GB NVIDIA V100 GPUs. We chose the largest possible batch size (256) and then the largest learning rate which resulted in stable training (3×10^{-5}). Other hyperparameters are the same as Lample & Conneau (2019).

Seq2Seq We use a large batch size (1024) and chose the largest learning rate which resulted in stable training across many pseudo-decomposition training corpora (1×10^{-4}). We keep other training settings and hyperparameters the same as for USeq2Seq.




Figure 6. Performance of downstream, multi-hop QA model, with and without decompositions, when varying the amount of training data. We also assess the impact of removing single-hop training data (SQUAD 2.0 and HOTPOTQA “easy” questions).

C. Multi-hop QA Model

C.1. Varying the Number of Training Examples

To understand how decompositions impact performance given different amounts of training data, we vary the number of multi-hop training examples. We use the “medium” and “hard” level labels in HOTPOTQA to determine which examples are multi-hop. We consider training setups where the multi-hop QA model does or does not use data augmentation via training on hotpot “easy”/single-hop questions and SQUAD 2 questions. Fig. 6 shows the results. Decompositions improve QA, so long as the multi-hop QA model has enough training data, either via single-hop QA examples or enough multi-hop QA examples.

C.2. Training Hyperparameters

To train ROBERTA_{LARGE}, we fix the number of training epochs to 2, as training longer did not help. We sweep over batch size $\in \{64, 128\}$, learning rate $\in \{1 \times 10^{-5}, 1.5 \times 10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}\}$, and weight decay $\in \{0, 0.1, 0.01, 0.001\}$, similar to the ranges used in the original paper (Liu et al., 2019). We chose the hyperparameters that did best for the baseline QA model (without decompositions) on our validation set: batch size 64, learning rate 1.5×10^{-5} , and weight decay 0.01. Similarly, for the experiments with BERT, we fix the number of epochs to 2 and choose hyperparameters by sweeping over the recommended ranges from Devlin et al. (2019) for learning rate ($\{2 \times 10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}\}$) and batch size ($\{16, 32\}$). For BERT_{BASE}, we thus choose learning rate 2×10^{-5} and batch size 16, and for BERT_{LARGE}, we use the whole-word masking model with learning rate 2×10^{-5} and batch size 32. We train all QA models with mixed precision floating point arithmetic (Micikevicius et al., 2018), distributing training across 8, 32GB NVIDIA V100 GPUs.

Unsupervised Question Decomposition for Question Answering




Figure 7. Performance difference for various answer entity types when the QA model does vs. does not use decompositions. We see the largest, consistent gains for entity-centric answers.




Figure 9. Performance difference for yes/no and span answer types for comparison questions when the QA model does vs. does not use decompositions. Decompositions are roughly as helpful for yes/no questions as for span-based questions.




Figure 8. Performance difference for bridge and comparison questions when the QA model does vs. does not use decompositions. Here, we use the original bridge/comparison splits from HOTPOTQA, which does not have a one-hop category and categorizes intersection questions as bridge. For the original dev set, the improvement with decompositions is greater for comparison questions than bridge questions. The multi-hop set does not alter comparison questions from the original version, so these scores do not change much.

C.3. Improvements across Detailed Question Types

To better understand where decompositions improve QA, we show the improvement across various fine-grained splits of the evaluation sets in Figures 5-11.




Figure 10. Performance difference for various multi-hop “wh”-words when the QA model does vs. does not use decompositions. Improvements by question word vary across dev sets.




Figure 11. Performance difference between a model when the QA model does vs. does not use decompositions, stratified by whether the gold final answer is in a sub-answer sentence. We find a larger improvement over the baseline when the gold answer contained in a sub-answer sentence.

Unsupervised Question Decomposition for Question Answering

Q: What is the name of the singer who's song was released as the lead single from the album "Confessions", and that had popular song stuck behind for eight consecutive weeks?	
USeq2Seq + Random	<p>Sub-Q1 what is the name of the singer who 's song was released as the lead single from the album " confessions?</p> <p>Sub-A1 is a song by American singer Usher. He co-wrote the song with Sean Garrett, Patrick J. Que Smith, Robert McDowell, LRoc, Ludacris, Da'Ron, and Lil Jon.</p> <p>Sub-Q2 why did popular song stuck behind for eight consecutive weeks?</p> <p>Sub-A2 It reached number 1 in the United Kingdom and in Netherlands, and number 2 in the United States, stuck for eight consecutive weeks behind the Usher singles "Yeah" and "Burn", respectively.</p>
USeq2Seq + FastText	<p>Sub-Q1 what is the name of the singer?</p> <p>Sub-A1 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p> <p>Sub-Q2 what song was released as the lead single from the album " confessions , " and that had popular song stuck behind for eight consecutive weeks?</p> <p>Sub-A2 "Yeah!" is a song by American singer Usher.</p>
Seq2Seq + Random	<p>Sub-Q1 what is the maximum number of students allowed to take part in the program?</p> <p>Sub-A1 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p> <p>Sub-Q2 what is the maximum number of students allowed per year?</p> <p>Sub-A2 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p>
Seq2Seq + FastText	<p>Sub-Q1 what is the name of the lead singer of the group?</p> <p>Sub-A1 Billy Ray Cyrus is an American country music singer, songwriter, actor and philanthropist, who helped make country music a worldwide phenomenon.</p> <p>Sub-Q2 when was the last time you heard a song called " on and on " by a female vocalist?</p> <p>Sub-A2 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007). It was written and produced by Keys, Kerry Brothers Jr., and George M. Harry, and released as the album's lead single on September 11, 2007.</p>
No Learn + Random	<p>Sub-Q1 what was sold to foreign pts?</p> <p>Sub-A1 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p> <p>Sub-Q2 what tends to lead to less money?</p> <p>Sub-A2 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p>
No Learn + FastText	<p>Sub-Q1 which single was released as the album 's lead single?</p> <p>Sub-A1 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p> <p>Sub-Q2 who did the song writing , what are some of the stories behind the songs?</p> <p>Sub-A2 It was written and produced by Keys, Kerry Brothers Jr., and George M. Harry, and released as the album's lead single on September 11, 2007.</p>
DecompRC	<p>Sub-Q1 which singer?</p> <p>Sub-A1 "No One" is a song by American singer and songwriter Alicia Keys from her third studio album, "As I Am" (2007).</p> <p>Sub-Q2 what is the name of alicia keys 's song was released as the lead single from the album "confessions", and that had popular song stuck behind for eight consecutive weeks?</p> <p>Sub-A2 It reached number 1 in the United Kingdom and in Netherlands, and number 2 in the United States, stuck for eight consecutive weeks behind the Usher singles "Yeah" and "Burn", respectively.</p>
Variable USeq2Seq + FastText	<p>Sub-Q1 what is the name of the singer?</p> <p>Sub-A1 is a song by american singer usher. he co-wrote the song with sean garrett, patrick j. que smith, robert mcdowell, iroc, ludacris, da'ron, and lil jon.</p> <p>Sub-Q2 who ' s the song was released as the lead single from the album?</p> <p>Sub-A2 "no one" is a song by american singer and songwriter alicia keys from her third studio album, "as i am" (2007).</p> <p>Sub-Q3 what popular song was stuck behind for eight consecutive weeks?</p> <p>Sub-A3 it reached number 1 in the united kingdom and in netherlands, and number 2 in the united states, stuck for eight consecutive weeks behind the usher singles "yeah" and "burn", respectively.</p>

Table 8. Various decomposition methods for the question "What is the name of the singer who's song was released as the lead single from the album "Confessions", and that had popular song stuck behind for eight consecutive weeks?" Here, the Variable USeq2Seq model has decomposed the question into three subquestions rather than two.

Unsupervised Question Decomposition for Question Answering

Q: Who is older, Annie Morton or Terry Richardson?	
USeq2Seq + Random	<p>Sub-Q1 who is older , annie morton?</p> <p>Sub-A1 Annie Morton (born October 8, 1970) is an American model born in Pennsylvania.</p> <p>Sub-Q2 who is terry richardson?</p> <p>Sub-A2 Terrence "Uncle Terry" Richardson (born August 14, 1965) is an American fashion and portrait photographer who has shot advertising campaigns for Marc Jacobs, Aldo, Supreme, Sisley, Tom Ford, and Yves Saint Laurent among others.</p>
USeq2Seq + FastText	<p>Sub-Q1 who is annie morton?</p> <p>Sub-A1 Annie Morton (born October 8, 1970) is an American model born in Pennsylvania.</p> <p>Sub-Q2 when was terry richardson born?</p> <p>Sub-A2 Kenton Terry Richardson (born 26 July 1999) is an English professional footballer who plays as a defender for League Two side Hartlepool United.</p>
Seq2Seq + Random	<p>Sub-Q1 what is the maximum number of students allowed to take part in the program?</p> <p>Sub-A1 Kenton Terry Richardson (born 26 July 1999) is an English professional footballer who plays as a defender for League Two side Hartlepool United.</p> <p>Sub-Q2 what is the maximum number of students allowed to take part in the program?</p> <p>Sub-A2 Kenton Terry Richardson (born 26 July 1999) is an English professional footballer who plays as a defender for League Two side Hartlepool United.</p>
Seq2Seq + FastText	<p>Sub-Q1 who is terry morton?</p> <p>Sub-A1 Terrence "Uncle Terry" Richardson (born August 14, 1965) is an American fashion and portrait photographer who has shot advertising campaigns for Marc Jacobs, Aldo, Supreme, Sisley, Tom Ford, and Yves Saint Laurent among others.</p> <p>Sub-Q2 who is terry morton?</p> <p>Sub-A2 Terrence "Uncle Terry" Richardson (born August 14, 1965) is an American fashion and portrait photographer who has shot advertising campaigns for Marc Jacobs, Aldo, Supreme, Sisley, Tom Ford, and Yves Saint Laurent among others.</p>
No Learn + Random	<p>Sub-Q1 what did deenet phase i become?</p> <p>Sub-A1 Snoecks is a Belgian magazine. The huge, 550-plus-page magazine appears once a year in October and focuses on the most interesting new international developments in the arts, photography and literature.</p> <p>Sub-Q2 what group can amend the victorian constitution?</p> <p>Sub-A2 Kenton Terry Richardson (born 26 July 1999) is an English professional footballer who plays as a defender for League Two side Hartlepool United.</p>
No Learn + FastText	<p>Sub-Q1 who was terry richardson?</p> <p>Sub-A1 Terrence "Uncle Terry" Richardson (born August 14, 1965) is an American fashion and portrait photographer who has shot advertising campaigns for Marc Jacobs, Aldo, Supreme, Sisley, Tom Ford, and Yves Saint Laurent among others.</p> <p>Sub-Q2 who changes positions more , annie morton , cotton hill , or terry richardson?</p> <p>Sub-A2 Annie Morton (born October 8, 1970) is an American model born in Pennsylvania.</p>
DecompRC	<p>Sub-Q1 annie morton is born when?</p> <p>Sub-A1 Annie Morton (born October 8, 1970) is an American model born in Pennsylvania.</p> <p>Sub-Q2 terry richardson is born when?</p> <p>Sub-A2 Terrence "Uncle Terry" Richardson (born August 14, 1965) is an American fashion and portrait photographer who has shot advertising campaigns for Marc Jacobs, Aldo, Supreme, Sisley, Tom Ford, and Yves Saint Laurent among others.</p>
Variable USeq2Seq + FastText	<p>Sub-Q1 who is annie morton?</p> <p>Sub-A1 annie morton (born october 8, 1970) is an american model born in pennsylvania.</p> <p>Sub-Q2 who was terry morton?</p> <p>Sub-A2 madonna is a biography by english author andrew morton, chronicling the life of american recording artist madonna.</p>

Table 10. Various decomposition methods for the question "Who is older, Annie Morton or Terry Richardson?"

Unsupervised Question Decomposition for Question Answering

Q: Since 2 June 2017, The Leader of Fine Gael had been held by which Irish Fine Gael politician who has served as Taoiseach and Minister for Defence?	
USeq2Seq + Random	<p>Sub-Q1 how long since 2 june 2017 , the leader of fine gael had been in office?</p> <p>Sub-A1 Enda Patrick Kenny (born 24 April 1951) is an Irish Fine Gael politician who served as Taoiseach from 2011 to 2017, Leader of Fine Gael from 2002 to 2017, Minister for Defence from May 2014 to July 2014 and 2016 to 2017, Leader of the Opposition from 2002 to 2011, Minister for Tourism and Trade from 1994 to 1997 and Minister of State for Youth Affairs from 1986 to 1987.</p> <p>Sub-Q2 which irish fine gael politician has served as taoiseach and minister for defence?</p> <p>Sub-A2 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p>
USeq2Seq + FastText	<p>Sub-Q1 since 2 june 2017 , the leader of fine gael had been?</p> <p>Sub-A1 Since 2 June 2017, the office had been held by Leo Varadkar following the resignation of Enda Kenny.</p> <p>Sub-Q2 which irish fine gael politician has served as taoiseach and minister for defence?</p> <p>Sub-A2 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p>
Seq2Seq + Random	<p>Sub-Q1 what is the maximum number of students allowed to take part in the program?</p> <p>Sub-A1 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p> <p>Sub-Q2 what is the maximum number of students allowed per year?</p> <p>Sub-A2 The Leader of Fine Gael is the most senior politician within the Fine Gael political party in Ireland. Since 2 June 2017, the office had been held by Leo Varadkar following the resignation of Enda Kenny.</p>
Seq2Seq + FastText	<p>Sub-Q1 who has been appointed as the new deputy leader of fine gael since 2 june 2017?</p> <p>Sub-A1 Simon Anthony Coveney (born 16 June 1972) is an Irish Fine Gael politician who has served as Minister for Foreign Affairs and Trade and Deputy Leader of Fine Gael since June 2017.</p> <p>Sub-Q2 the fine gael fine gael , the fine gael of fine gael?</p> <p>Sub-A2 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p>
No Learn + Random	<p>Sub-Q1 what was considered to be a major milestone?</p> <p>Sub-A1 The 2017 Fine Gael leadership election was triggered in May 2017, when Enda Kenny resigned as party leader.</p> <p>Sub-Q2 what was the air force not interested in for their message system?</p> <p>Sub-A2 The 2017 Fine Gael leadership election was triggered in May 2017, when Enda Kenny resigned as party leader.</p>
No Learn + FastText	<p>Sub-Q1 what if fine gael did support fine gael after the next election?</p> <p>Sub-A1 With Fine Gael being the governing party at the time, this election effectively appointed a new Taoiseach for Ireland.</p> <p>Sub-Q2 who has been appointed as defence minister of india?</p> <p>Sub-A2 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p>
DecompRC	<p>Sub-Q1 which leader of fine gael?</p> <p>Sub-A1 Since 2 June 2017, the office had been held by Leo Varadkar following the resignation of Enda Kenny.</p> <p>Sub-Q2 since 2 june 2017 enda patrick kenny had been held by which irish fine gael politician who has served as taoiseach and minister for defence?</p> <p>Sub-A2 Leo Eric Varadkar (; Irish: "Leo de Varad" ;born 18 January 1979) is an Irish Fine Gael politician who has served as Taoiseach, Minister for Defence and Leader of Fine Gael since June 2017.</p>
Variable USeq2Seq + FastText	<p>Sub-Q1 since 2 june 2017 , the leader of fine gael had been held by?</p> <p>Sub-A1 since 2 june 2017, the office had been held by leo varadkar following the resignation of enda kenny.</p> <p>Sub-Q2 which irish fine gael politician has served as taoiseach and minister for defence?</p> <p>Sub-A2 enda patrick kenny (born 24 april 1951) is an irish fine gael politician who served as taoiseach from 2011 to 2017, leader of fine gael from 2002 to 2017, minister for defence from may 2014 to july 2014 and 2016 to 2017, leader of the opposition from 2002 to 2011, minister for tourism and trade from 1994 to 1997 and minister of state for youth affairs from 1986 to 1987. he has been a teachta dla (td) since 1975, currently for the mayo constituency.</p>

Table 13. Various decomposition methods for the question "Since 2 June 2017, The Leader of Fine Gael had been held by which Irish Fine Gael politician who has served as Taoiseach and Minister for Defence?"