Decoupled Context Processing for Context Augmented Language Modeling

Zonglin Li *
Google Research, New York
lizonglin@google.com

Ruiqi Guo *
Google Research, New York
guorq@google.com

Sanjiv Kumar
Google Research, New York
sanjivk@google.com

Abstract

Language models can be augmented with a context retriever to incorporate knowledge from large external databases. By leveraging retrieved context, the neural network does not have to memorize the massive amount of world knowledge within its internal parameters, leading to better parameter efficiency, interpretability and modularity. In this paper we examined a simple yet effective architecture for incorporating external context into language models based on decoupled `Encoder-Decoder` architecture. We showed that such a simple architecture achieves competitive results on auto-regressive language modeling and open domain question answering tasks. We also analyzed the behavior of the proposed model which performs grounded context transfer. Finally we discussed the computational implications of such retrieval augmented models.

1 Introduction

Transformers have proven to be powerful language models that capture an impressive amount of world knowledge in its internal parameters and generalize to a variety of downstream tasks [39, 31]. Recently, there has been a lot of success in improving language model quality by increasing the number of parameters in transformers, often on the order of hundreds of billion [10, 30, 38, 6]. However, the scaling of model size also contributes to the exponential rise of the computation costs, both in terms of the number of accelerators needed and energy consumption [29].

To overcome the exponential increase in the number of parameters, one natural idea is to utilize information retrieved from an external source such as a massive external database, therefore freeing the neural network from having to memorize world knowledge. To this end, researchers proposed multiple context augmented language model architectures [21, 15, 4, 44, 26]. Such architecture typically has two components: a retriever that embeds the input sequence and retrieves relevant context from external source through vector similarity search; a neural network that integrates both the input and retrieved external context into the prediction of target sequence, formally:

$$P(\mathbf{y}|\mathbf{x}, \mathbf{C} = Retrieve(\mathbf{x}, \mathcal{D}); \theta) \geq P(\mathbf{y}|\mathbf{x}, \theta') \tag{1}$$

Here, $\mathbf{C} = \{\mathbf{c}\}$ is a set of context retrieved from the external database $\mathcal{D}$. θ' is a self-contained language model which predicts target sequence $\mathbf{y}$ based solely on the input x whereas θ corresponds to the context augmented language model which incorporates both the input $\mathbf{x}$ and the retrieved context $\mathbf{C}$.

One of the challenges for such context augmented language model is the computational cost of context retrieval and incorporation, especially when multiple pieces of context is present or the context sequence is long. In this paper, we propose a computationally efficient architecture for

*Equal contribution

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

incorporating context based on vanilla `Encoder-Decoder`, which decouples the encoding of context and the prediction of target sequence. We show that the model with such a simple architecture is competitive when compared with customized mechanisms such as Chunked-Cross-Attention [4] on language modeling score (as measured by bits-per-byte, BPB), while being more efficient in terms of parameter count and computation cost. Then, we define metrics to measure the utility of the retrieved context and use it to guide the training of the retriever. We further show competitive results on downstream tasks of question answering, and demonstrate that the model takes advantage of the retrieved context without memorizing facts within its internal parameters. Finally, we study the implication of context retrieval in terms of retrieval latency, accuracy and computation cost.

To summarize the main contributions of this article:

- Proposed a novel Encoder-Decoder based architecture for incorporating retrieved external context, which decouples context encoding from language model inference.
- Demonstrated the competitive results of the proposed model on both the auto-regressive language modeling task and the open domain question answering task.
- Analyzed model behavior by understanding how context improves language modeling prediction on tokens with different linguistic properties and how the model performs grounded context transfer.
- Discussed computational cost and retrieval efficiency in context augmentation.

2 Related Works

Large language models, typically in the form of big neural networks, are trained with a huge amount of training data rich in unstructured knowledge. Researchers have found that after model training, the neural networks often end up storing a surprisingly large amount of memorized information within its weights [2, 7] which are then leveraged as a knowledge base. Multiple hypotheses have been developed on how components such as fully-connected layers [13] and attention layers [5] may be responsible for such memorization behavior. While the capability of storing world knowledge is desirable, memorization also contributes to huge model sizes and the lack of explicit control over knowledge base, such as performing selection or updates.

An alternative strategy is to enable language models to incorporate world knowledge in the form of retrieved context from external sources, instead of having to memorize them. Multiple works have proposed architectures that support external retrieval, usually composed of a context retriever that searches a large external key value store and a method of integrating retrieved information.

There are various ways to construct the key value store. The keys are primarily used for similarity matching, and they can be sparse vectors such as BM25 [34], or dense embeddings extracted from part of the model [21, 25, 40], or from pretrained embedders [12, 28], or embedders trained for specific downstream tasks [15, 20, 41, 35, 24, 11]. The values also have various different forms. For example, TOME [11] stores a dense embedding about the contextual information of an entity mention, while Realm [15], RAG [26], FID [18], Retro [4], MARGE [25], DenSPI [35] and DensePhrases [24] store the raw text as the value. Works such as kNN-LM [21] and Spalm [44] store one token as a value. Finally the key value store is searched over using vector similarity techniques, typically with some off-the-shelf nearest neighbor search implementations such as FAISS [19], ScaNN [14], HNSW [27] or SPTAG [9].

There are also many different ways to integrate the retrieved context. A popular approach is to concatenate the retrieval results with the original input and jointly process them. It has been adopted by works such as Realm [15], RAG [26], and FiD [18]. Other works utilize some forms of cross attention for the context integration, such as the Chunked-Cross-Attention with input conditioning in Retro [4], kNN Attention in Memorizing Transformer [40], Memory Attention in TOME [11] and Cross-Attention in MARGE [25]. For token level integration, kNN-LM [21] uses simple linear interpolation while Spalm [44] uses a learned gate based on the last layer embedding. There are also works that directly utilize the retrieval results without any integration, such as DenSPI [35] and DensePhrases [24]. Most of the works use retrieval as a way to augment tasks such as language modeling or question answering, with the exception of MARGE [25] where retrieval is treated as an autoencoder bottleneck for multilingual pretraining, and is not strictly necessary for inference. We compare representative previous works and contrast with our proposal in Table 1.

Method	Retrieval Granularity	Retrieval Encoding	Context Integration	Decoupled Context Encoding	Tasks
kNN-LM [21]	Token	Last layer	Interpolation	Yes	LM
Spalm [44]	Token	Last layer	Gating	Yes	LM
Realm [15], RAG [26], FID [18]	Input	Trained	Concat	No	OpenQA
Retro [4]	Chunk	Frozen	Chunked-Cross-Attention	No	LM, OpenQA
Proposed	Chunk	Frozen / trained	Encoder-Decoder Cross-Attention	Yes	LM, OpenQA

Table 1: Architectural differences between previous retrieval augmented model and ours in (i) context retrieval, (ii) context integration and (iii) targeted applications.

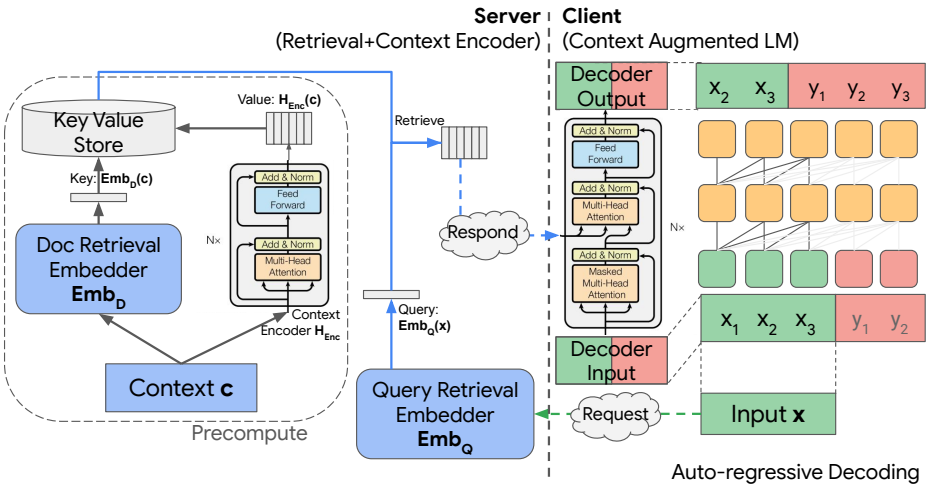

Figure 1: **Architecture of context augmented language model.** We opt to use the standard `Encoder-Decoder` cross attention mechanism for context incorporation which allows us to decouple context encoding from LM inference. $\mathbf{c}$, $\mathbf{x}$, $\mathbf{y}$ serve as context, decoder input and decoder target respectively. In other words, the client sends over input $\mathbf{x}$, and the server conducts retrieval to find relevant context and returns the encoded representation $\mathbf{H}_{\texttt{Enc}}(\mathbf{c}; \theta_{\texttt{Enc}})$. The encoded representation $\mathbf{H}_{\texttt{Enc}}$ is pre-computed offline and returned as "metadata" of the retrieval. Note that training of the `Encoder` and `Decoder` is joint, while they are decoupled at inference time: the client does not need to store parameters or the run inference on the `Encoder` component.

3 Architecture

We use `Encoder-Decoder` Transformer architecture [39] to integrate language model input and retrieved context. We denote the context encoder and LM decoder as `Enc` and `Dec` respectively. Given an input token sequence $\mathbf{x} = (x_1, x_2, ..., x_n)$, the task is to predict the next tokens $\mathbf{y} = (y_1, y_2, ..., y_s)$.

Without external context, we concatenate $\mathbf{x}$ before $\mathbf{y}$ and the task becomes a traditional auto-regressive language modeling to predicts targets $\mathbf{y}$ following input (or "promopt") $\mathbf{x}$. In this setting, only the decoder is involved (denoted as "No-retrieval"). To incorporate external context, we use $\mathbf{c}$, $\mathbf{x}$, $\mathbf{y}$ to serve as encoder input, decoder input and decoder target respectively. We first use a retriever to identify the context $\mathbf{c} = (c_1, c_2, \ldots, c_m)$ given input $\mathbf{x}$, then fetch the the pre-computed encoder output of the corresponding context tokens $\mathbf{H}_{\texttt{Enc}}(\mathbf{c}; \theta_{\texttt{Enc}})$ as output. In this setting, encoder output $\mathbf{H}_{\texttt{Enc}}$ is directly used by the decoder through Encoder-Decoder cross-attention to influence the final prediction. The decoder does not have to know the exact tokens of $\mathbf{c}$ that are retrieved.

$$\texttt{No-retrieval}: P(y_i|y_{<i}, x_1, x_2, \ldots x_n; \theta'_{\texttt{Dec}})$$
$$\texttt{Retrieval}: P(y_i|y_{<i}, x_1, x_2, \ldots x_n, \{\mathbf{H}_{\texttt{Enc}}(\mathbf{c})\}; \theta_{\texttt{Dec}})$$

Under this formulation, only decoder parameters $\theta_{\texttt{Dec}}$ are required at inference time. The retriever retrieves indices of the relevant context and looks up their encodings. The context encodings are processed ahead of the time, and are completely decoupled from online operations. This is in contrast to previous works of Realm [15], Rag [26] or Retro [4] where the interaction between input $\mathbf{x}$ and context $\mathbf{c}$ is bi-directional, which necessitates context encoding at inference time. In our model, information flows uni-directionally from $\mathbf{c}$ to $\mathbf{x}$ and $\mathbf{y}$, and that the encoding of each context $\mathbf{c}$ is processed independently. On one hand, this is more restrictive than bi-directional interaction; on the other hand, such a design ensures complete decoupling of context processing and the online language model inference. The exact mechanism is detailed in Figure. 1.

Conceptually the retriever can be an arbitrary blackbox. In practice, we use a dual encoder formulation [8, 16], which first represents $\mathbf{x}$ as a query embedding $Emb_Q(\mathbf{x})$ and performs vector similarity search over a database of $\mathcal{D}$ to find the indices of documents whose document embedding has the highest inner products with the query embedding. We then look up the context encoder outputs that correspond to retrieved indices and return them as the retriever output.

$$l^* = \arg\max_{\mathbf{v}\in\mathcal{D}} \langle Emb_Q(\mathbf{x}), \mathbf{v})\rangle; \mathcal{D} = \{Emb_D(\mathbf{c}); \mathbf{c} \in \mathcal{C}\}$$

$$\mathbf{H}_{\texttt{Enc}}[l^*] = \texttt{Enc}(\mathbf{c}_{l^*}; \theta_{\texttt{Enc}})$$

In the case of multiple supporting context, k-$\arg\max$ is used instead of $\arg\max$. The encoder outputs of each supporting context are then concatenated:

$$P(y_i|y_{<i}, x_1, x_2, \ldots x_n, Concat(\mathbf{H}_{\texttt{Enc}}[l_1], \mathbf{H}_{\texttt{Enc}}[l_2], \cdots \mathbf{H}_{\texttt{Enc}}[l_k]); \theta_{\texttt{Dec}});$$

Where $Concat$ is simply vector concatenation:

$$Concat((\mathbf{h}_1, \mathbf{h}_2, \cdots, \mathbf{h}_n], [\mathbf{g}_1, \mathbf{g}_2, ..., \mathbf{g}_m], ...) = [\mathbf{h}_1, \mathbf{h}_2, ..., \mathbf{h}_n, \mathbf{g}_1, \mathbf{g}_2, ..., \mathbf{g}_m, ...]$$

At training time, the encoder and decoder are jointly trained. We first perform offline retrieval to form triplets of $(\mathbf{x}, \mathbf{y}, \mathbf{c})$, where $\mathbf{c}$ is retrieved by some predefined retriever. The loss is masked and only defined on the targets $\mathbf{y}$. Because the encoding of each context is independent and there is no interaction between context, the attention matrix of encoder is block diagonal and we process them in a linear loop over each diagonal block. Thus, the computation cost of both encoder and decoder at each step is linear in the number of context.

For online language model inference, only the retriever and the decoder are involved. The retrieval embedding $Emb_D(\mathbf{c})$ and the encoder output $H_{\texttt{Enc}(c)}$ of the context are both offline pre-processed and prepared into the retrieval database of such associated key-value pairs. When a new input sequence arrives, the retriever is only responsible for the approximate nearest neighbor search and lookup of the associated value that is the pre-computed encoder output. The decoder then takes in the input sequence and cross attends on the concatenation of pre-computed encoder output to generate the targeted tokens. Thanks to the decoupling, neither the retriever nor the decoder needs to store encoder parameters. Hence, such an approach is more parameter efficient compared to similar works such as Retro [4] by saving the storage and computation budget on encoder, which is helpful in "client-server" scenario where the capacity of the "client" can be limited. When accounting the parameter count in comparison with other models, we only need to count the decoder and cross attention parameters. We also followed Retro's [4] approach of excluding the embedding matrices from the parameter count.

4 Auto-regression Language Modeling

We experimented with the same "encoder-decoder" context incorporation mechanism for both auto-regressive language modeling and open domain question answering. The only difference is that auto-regressive language modeling processes input sequences in a sliding window fashion, while question answering task receives the full input sequence (the question) at once.

4.1 Experimental Setup

For auto-regressive language modeling, we use English C4 [32] version `2.2.1`, the same as Retro. We train the language model and prepare the retrieval database using the `train` split and evaluate the results using `validation` split. The language model target sequence is a sliding window (chunk)

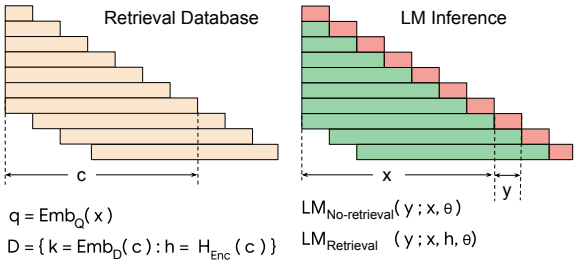

(a)

Database Split	# Articles	# Entries
`C4 Train`	364.6M	3382M
`C4 Val, unfiltered`	0.3645M	3.369M
`C4 Val, filtered`	-	2.868M
`NQ-open Train`	79k	-
`NQ-open Dev`	8.8k	-
`NQ-open Test`	3.6k	-
`Wiki Database`	-	21M

(b)

Figure 2: (a) Chunking scheme for language modeling training with C4. For each article in C4, we divide the tokenized text into non-overlapping blocks of at most 64 tokens and use them as targets $\mathbf{y}$. We use the preceding tokens of at most 448 in length as $\mathbf{x}$. (b) The number of entries in the retrieval database of C4 auto-regressive LM and Natural Question QA task.

of $s = 64$ tokens, with at most $n = 448$ preceding tokens are used as input sequence. This setup is similar to XLNet [43] and Retro [4]. The target and input sequences that are smaller than the given window size (64 and 448, respectively) are padded with zeros.

To construct the retrieval database, the same sliding window processing is also used for the context sequences. The database is formed as associated pairs of retrieval embedding and encoder output: $\{Emb_D(\mathbf{c}) : H_{\texttt{Enc}}(\mathbf{c})\}$, where context $\mathbf{c}$ are the sliding window of 512 tokens with a stride of 64 tokens. We choose our hyper-parameters to be comparable to Retro: chunk size $s = 64$ and input window size $n = 448$ (smaller than 2048 that of Retro). This also implies that the number of entries in the database is larger than the number of articles, but smaller than the number of tokens.

Our training corpus is in the form of triplets $(\mathbf{x}, \mathbf{y}, \mathbf{c})$. $\mathbf{x}$ and $\mathbf{y}$ are acquired directly by applying sliding window on the `train` split of C4. Then BM25 [34] is used as a bootstrapping retriever to mine relevant context $\mathbf{c}$ to from the database. The first retrieval results with no more than 8 consecutive token overlap with the target is used as context. Figure 2a illustrate the sliding window construction of database as well as sequence served as input and targets. Table 2b gives the exact number of entries in the resulting database used as external context.

We use mT5 [42] as the backbone architecture for our context augmented `Encoder-Decoder`, and train our models from scratch. `Train` split is used both for training and retrieval, while `validation` split is used for bits-per-byte (Bpb) evaluation. In auto-regressive language model evaluation, due to the fact that text are crawled from web sources, there can be a non-trivial overlap of tokens between the training and validation splits. In such cases, tokens are often "copy-pasted" from retrieved context into targets without changing. Such copying leads to near-zero perplexity on targets and has a big effect on final bits-per-byte measurement. Following the discussion on dataset leakage of Retro [4], we filtered any example whose targets and context sequences have more than 8 common consecutive tokens (correspond to 12.5% filtering of Retro). We found 14.87% of the validation chunks are removed by the filtering of longest common substring.

Figure 3a reports our auto-regressive language model results with different model sizes under the Bit-per-byte (bpb) values. Bpb is tokenizer agnostic and is often used to compare models with different vocabulary. We use mT5 `base`, `large` and `XL` respectively without modification, and the results compare favorably to Retro models with similar sizes that uses customized chunked-cross-attention which couples context encoding and LM decoding. Our experiments demonstrate that context incorporation can be achieved with simple `Encoder-Decoder` cross attention, with the additional benefit of decoupled encoder processing.

4.2 Retriever Training

The goal of context retrieval is to identify the context sequences $\mathbf{c}^*$ that maximize the improvement of some utility function, such as the log-likelihood improvement on target prediction. i.e.

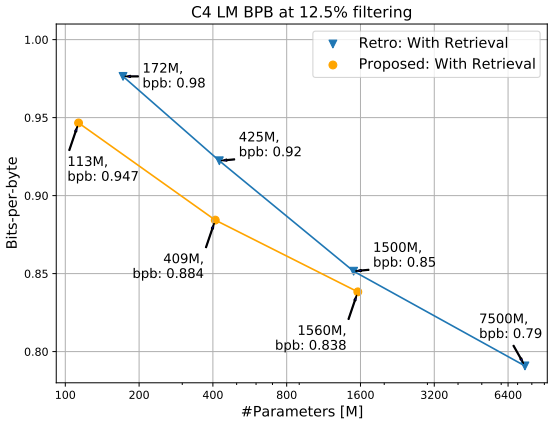

(a)

Model	Model Size	Exact Match Accuracy
`Realm` [15]	110M	40.4
`DPR` [20]	110M	41.5
`Ours (Large)`	409M	44.35
`RAG` [26]	400M	44.5
`Retro` [4]	7.5B	45.5
`Ours (XL)`	1.56B	47.95
`FiD` [18]	770M	51.4
`EMDR` [37]	440M	52.5
`FiD + Distill` [17]	770M	54.4

(b)

Figure 3: Performance of decoupled encoder-decoder on auto-regressive language modeling (C4) and question answering tasks (Natural Question). (a) Comparing the proposed method and Retro on `c4-en-2.2.1` validation split with 12.5% token overlap filtering. The y-axis measures bits-per-byte (bpb, lower is better), which is the perplexity normalized by token length. The x-axis shows the number of non-embedding parameters in log scale. (b) End-to-end result on Natural Question test split. Following previous works, we measure exact match (EM) accuracy on "short answer type" with at most five tokens.

$$\mathbf{c}^* = \arg\max_{\mathbf{c}} U(\mathbf{x}, \mathbf{y}, \mathbf{c});$$

$$U(\mathbf{x}, \mathbf{y}, \mathbf{c}) = \sum_{y_i} \log P(y_i | y_{<i}, \mathbf{x}, \mathbf{c}, \theta) - \log P(y_i | y_{<i}, \mathbf{x}, \theta')$$

However, it is infeasible to evaluate the utility function $U(\mathbf{x}, \mathbf{y}, \mathbf{c})$ for all triplets $(\mathbf{x}, \mathbf{y}, \mathbf{c})$. Therefore, we introduce a proxy to approximate U by computing the expectation of U conditioned on the token y_i and whether it appeared in the input $\mathbf{x}$ and context sequence $\mathbf{c}$:

$$\hat{U}(\mathbf{x}, \mathbf{y}, \mathbf{c}) = \sum_{y_i \in \mathbf{y}, \mathbf{c}} \mathbb{1}_{y_i \in \mathbf{x}} \bar{U}(y_i | y_i \in \mathbf{x}, y_i \in \mathbf{c}) + \mathbb{1}_{y_i \notin \mathbf{x}} \bar{U}(y_i | y_i \notin \mathbf{x}, y_i \in \mathbf{c})$$

Intuitively, the "context utility" $\hat{U}$ is a weighted token overlap between context $\mathbf{c}$ and target $\mathbf{y}$. The weight is higher if the token y_i did not appear in input but is contained in context, and when the token is "sensitive" to context by showing larger loglikelihood change when the context is present. To train a retriever, we adopted the typical dual encoder formulation with in-batch softmax training similar to that of DPR [20, 16]. We use training data bootstrapped from BM25 retrieval, where (1) the valid retrieval with highest context utility $\hat{U}$ is the positive; (2) the other top retrieval from BM25 that has less than 80% of context utility of the highest one is used as hard negative; (3) The rest of the in-batch samples are regarded as random negatives. All retrievals are subject to the filtering criteria of no more than 8 consecutive tokens.

We trained our retriever using T5X retrieval framework [28] based on mT5-`small` with an embedding dimensionality of 128. The model is trained for 100,000 iterations on a batch size of 8,192 on 64 TPUv3 chips. Unfortunately, using the dense trained retriever does not lead to a visible improvement on the final C4 Bpb evaluation, possibly because both the dense and BM25 retriever end up retrieving similar context. We report results using trained dense retriever because of the better retrieval efficiency.

4.3 Ablation and Analysis

To analyze the effect of the context augmentation, we show first in Figure 4a that the context augmentation in language model consistently helps across different model sizes in the proposed

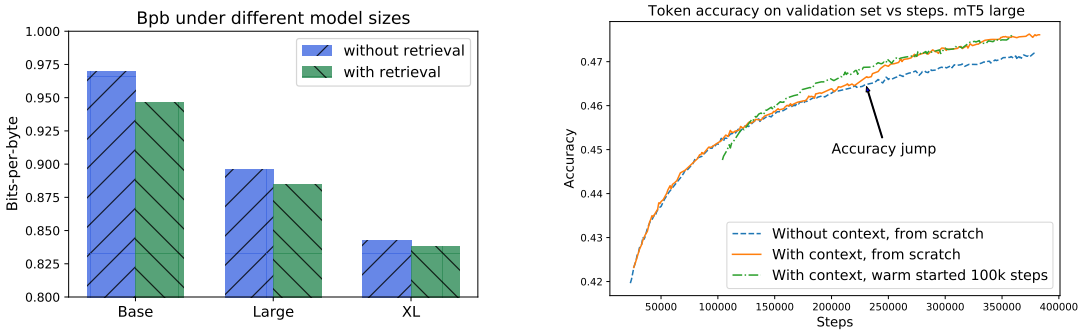

(a) (b)

Figure 4: (a) Bpb evaluation (lower is better) comparing models of the same sizes trained from scratch, with and without context retrieval, respectively. (b) Comparing the accuracy history of `Large` model training with and without retrieval. Notice that the curve "jumps" when training with retrieval but without warm-starting.

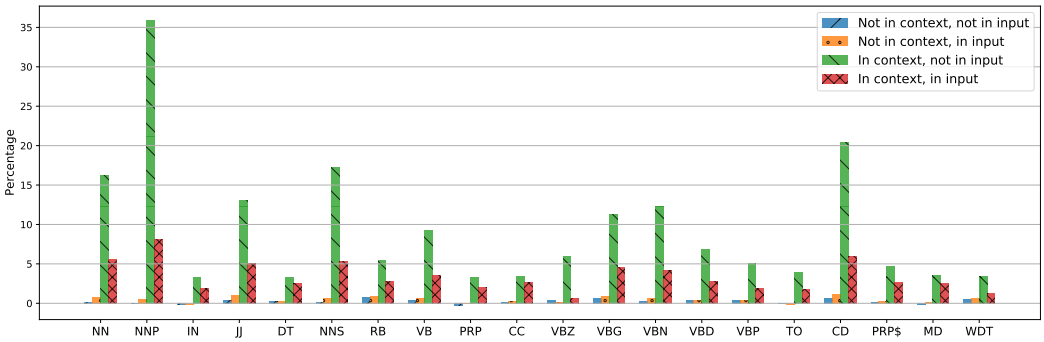

Figure 5: Percentage of improvement on target log-likelihood when there's context. We show the most frequent 20 POS tags, excluding punctuation. For the data on all the POS tags, see Appendix A.

method. While the absolute improvement on Bpb over "No-retrieval" may seem small when it is filtered to 12.5% in longest common sequence, we note that the improvements were made on a limited set of important tokens. Intuitively, tokens of functional words such as "the", "a", "to" should not benefit much from external information and can be predicted with its local context. Yet functional words make up a significant portion in perplexity because of their high frequencies in text.

To this end, we consider the difference in log-likelihood of two language models of same size, trained with and without retrieved context. We evaluate the models on validation split of C4 and obtain a breakdown using part-of-speech (POS) tagging. We used `NLTK` [3] to assign each target token to a POS tag. Note that the tokenization boundaries of `NLTK` and sentencepiece [22] could be different, and the assignment is done by majority voting in that case.

From Figure 5, it is immediately clear that the impact of context retrieval on different types of tokens is not uniform. Noun (`NN`, `NNP`, `NNS`, `NNPS`), and number (`CD`) benefited the most from the context, followed by adjective (`JJ`) and verb (`VB`, `VBZ`, `VBG`, `VBN`, `VBD`, `VBP`). Tokens such as preposition (`IN`), coordinating conjunction (`CC`), etc. are less helped by context. We believe it is because their prediction usually only depends on immediate local context. We also found that whether the target token appeared in the context sequence or input sequence has a large impact on the improvements, and included this in the breakdown. Intuitively, tokens that are grounded in the retrieved context $\mathbf{c}$ but did not appear in the input $\mathbf{x}$ benefit the most from context retrieval.

4.3.1 Training Details

We trained the Encoder-Decoder LM model for a total of $1,100,000$ steps with a batch size of 512 and a default learning rate schedule of square-root decay. This corresponds to $10,000$ warmup steps with a fixed learning rate of 0.01, followed by square-root decay for $990,000$ steps. We found training from scratch instead of fine-tuning on pre-trained models is important to achieve competitive results

on C4. We also found it beneficial to include additional $100,000$ steps with a lower learning rate of 2×10^{-4} after standard training schedule, which slightly improves the results. For `base` and `large` we used 64 TPUv3 chips whereas 128 TPUv3 chips for training `XL`. `XL` runs 1.5 training steps per second. `XXL` is left out due to the insufficient computation resources.

Comparing the training curve of model with and without retrieval, as depicted in Figure 4b, we noticed a phenomenon where both models have roughly the same performance initially but the one with retrieval suddenly increases at some point and continues to improve afterwards. This is especially true in `large` and `XL`. We suspect that decoders of larger models have higher memorization capacity, and in the initial phase they can improve training objective without delegating memorization to context retrieval. This is a challenge to training stability and one way to address this is by first "warm start" the model with $100,000$ steps of training on a subset of training data with highest 10% of the context utility before proceed on the full dataset.

5 Open Domain Question Answering

5.1 Natural Question

Large language models are useful because they generalize to downstream tasks, in addition to performing auto-regressive language generation. To demonstrate that the proposed context augmentation scheme is effective on downstream applications we evaluated our proposed model of decoupled `Encoder-Decoder` on the OpenQA task of Natural Question [23].

We use the same question and context processed by [18], where the context is retrieved with DPR retriever [20]. We construct each context sequence in the format of `"title: {title} source: {source}"` and pad or trim it to 256 sentencepiece tokens. The context sequences are then encoded by the encoder independently and simply concatenated, to ensure there is no interaction between them.

The decoder is trained to receive and respond on the sequence containing both the question and answer in format `"question: {question} \n answer: {answer}"` similar to [4], where `"question: {question} \n answer: "` is the input prompt (and there is no loss defined on these tokens) and `{answer}` is the target sequence the decoder is expected to predict. Again, the decoder cross-attends to the concatenated encoder output $\mathbf{H}_{\texttt{Enc}}(\mathbf{c})$ from all the retrieved context and the decoder is completely decoupled from encoder. It does not even have to know the exact tokens of context sequence being returned by the retriever.

The model weights are initialized from T5.1.1 checkpoints [32] We jointly fine-tuned the encoder and decoder for $40,000$ steps with 20 context passages for each input of the `train` split, and validated on the `dev` split every 1,000 steps. We used a batch size of 64, a fixed learning rate of 10^{-4} and Adafactor optimizer [36]. Finally, we selected the checkpoint with the best validation accuracy on the `dev` split and evaluated on the `test` split. Each model is trained on 64 TPUv3 chips, and the evaluation metrics is string exact match with SQuAD [33] normalization. We report our final results in Table 3b. Overall, the decoupled `Encoder-Decoder` produced competitive results on Natural Question task, albeit being a much simpler architecture and not specifically designed for QA tasks. We note that FiD [18] uses similar architecture. However, FiD concatenates the question with the retrieved context passages. Therefore the context encoding is dependent on the input sequence and needs to be computed at inference time.

5.2 Grounded Answer Generation

Context augmented models tend to generate answers by transferring tokens from context to the output. It is interesting to quantify how often the model output comes from grounded transfer of context tokens and how often this leads to the correct answer. We use the same model trained in 5.1 to (a) run language model decoding with the original set of retrieved context to obtain "original output": $\mathbf{y}^* = Decode(\mathbf{x}, \mathbf{C}; \theta)$; and (b) rerun model prediction but remove all context passages that contain the "original output": $\mathbf{y}' = Decode(\mathbf{x}, \mathbf{C}' = \mathbf{C} \setminus \{\mathbf{c}; \mathbf{y}^* \subseteq \mathbf{c}\}; \theta)$. In the case of grounded context transfer, we expect the "secondshot output" to change, but still grounded in the remaining context.

We show grounded-ness analysis in Table 2. Overall, we found 76% (2,765 out of 3,354) of the cases are likely due to grounded context transfer, because the "secondshot output" has changed to

Original output $\mathbf{y}^* = Decode(\mathbf{x}, \mathbf{C}; \theta)$	Original output in context				Original output out of context
	3354				256
Secondshot output $\mathbf{y}' = Decode(\mathbf{x}, \mathbf{C}'; \theta)$ $\mathbf{C}' = \mathbf{C} \setminus \{\mathbf{c}; \mathbf{y}^* \subseteq \mathbf{c}\}$	Output changed still in context $\mathbf{y}' \neq \mathbf{y}^*, \mathbf{y}' \subseteq \mathbf{C}'$	Output changed out of context $\mathbf{y}' \neq \mathbf{y}^*, \mathbf{y}' \nsubseteq \mathbf{C}'$	Output unchanged	Excluded	-
	2765	261	291	37	-

Table 2: Analysis of grounded context transfer on Natural Question. First row shows the number of questions broken down by whether the original model output appeared in the retrieved passages. To demonstrate the model follows grounded context transfer, we removed context passages containing original output and reran inference on the remaining context passages. 76% of output changed due to the removal of "original output" but the new prediction is still contained in remaining context.

something different from "original output", but is still grounded in the remaining context passages. In addition, we found that the accuracy is 51.5% when model output is from grounded transfer, higher than 35.4% for the rest. This implies that the model has higher accuracy when the output is grounded. We exclude 37 examples from the analysis because they have less than 20 neighbors after removing original output.

6 Computational Discussion

We evaluate the computation implications of our method around inference latency and space-time trade off. We estimate the latency based on benchmarks and publicly available packages and performance data. In the Natural Question experiment, we use the 1.56B XL model in Table 3b, whose encoder output have dimension of 2048 and the activation data type is `bfloat16` [1]. When profiled on a single TPUv3 core, it takes 200ms to encode all 20 retrieved context passages each containing 256 tokens (with paddings). This is equivalent to 10ms per context. Thanks to the decoupling of the context encoding, the context passages encoding is done offline, and simply looked up at inference time. Inferencing the DPR query embedder model takes around 8ms unbatched on one TPUv2 chip. We use ScaNN [14] to perform approximate similarity search on the DPR embeddings with a dimensionality of 768, which takes roughly 12ms to retrieve 20 neighbors with recall@20=0.97 on a single CPU core. For each passage, the uncompressed encoded embeddings have a total size of 1 MiB (`sizeof(bfloat16) * 2048 * 256`). When reading the encoding from an NVMe SSD with a read throughput of 660 MB/s, the lookup takes around 1.5ms and we use another 0.8ms to transmit of one encoding over network. End to end, it takes around 66ms (`((1.5 + 0.8) * 20 + 12 + 8)`) to retrieve 20 encoded passages. This is about a third of time for running inference of the encoder, and everything except query embedder is conducted with generic hardware without accelerators which is cheaper and consumes less energy.

However, one limitation of this approach is the amount of disk space and network bandwidth needed to store and transmit all the pre-computed encodings, which scales with the dimensionality of encoder output. We leave it for future research to reduce the encoder output size while maintaining similar quality, which may be possible through compression, dimensionality reduction, or adding projection layer to the encoder.

7 Conclusion

In this paper, we advocate the design of context augmented language model based on `Encoder-Decoder` architecture. `Encoder-Decoder` models are simple, proven and enjoy the unique computational advantage that the context encoding and language generation are decoupled. That is, the context encoding can be produced with offline pre-computation and caching, while the decoder is parameter efficient because it is agnostic to the encoder computation. We further demonstrated the effectiveness of this simple architecture by comparing with competitive baselines on common NLP tasks such as language modeling on C4 and question answering on Natural Questions. We also analyzed the model behavior and showed the context augmentation resulted in big improvements on content words and the model generates output grounded in retrieved context most of the time. Finally, we estimated the latency improvement from caching enabled by the decoupled computation.

References

[1] Ieee standard for floating-point arithmetic. *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pages 1–84, 2019.

[2] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017.

[3] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009.

[4] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. Improving language models by retrieving from trillions of tokens. *arXiv preprint arXiv:2112.04426*, 2021.

[5] Trenton Bricken and Cengiz Pehlevan. Attention approximates sparse distributed memory. *Advances in Neural Information Processing Systems*, 34, 2021.

[6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[7] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650, 2021.

[8] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*, 2018.

[9] Qi Chen, Haidong Wang, Mingqin Li, Gang Ren, Scarlett Li, Jeffery Zhu, Jason Li, Chuanjie Liu, Lintao Zhang, and Jingdong Wang. *SPTAG: A library for fast approximate nearest neighbor search*, 2018.

[10] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.

[11] Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, and William Cohen. Mention memory: incorporating textual knowledge into transformers through entity mention attention. *arXiv preprint arXiv:2110.06176*, 2021.

[12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[13] Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer feed-forward layers are key-value memories. *arXiv preprint arXiv:2012.14913*, 2020.

[14] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pages 3887–3896. PMLR, 2020.

[15] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*, 2020.

[16] Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. Efficient natural language response suggestion for smart reply. *arXiv preprint arXiv:1705.00652*, 2017.

[17] Gautier Izacard and Edouard Grave. Distilling knowledge from reader to retriever for question answering. *arXiv preprint arXiv:2012.04584*, 2020.

[18] Gautier Izacard and Edouard Grave. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*, 2020.

[19] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 2019.

[20] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*.

[21] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. *ICLR 2020*, 2019.

[22] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018.

[23] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*, 2019.

[24] Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. Learning dense representations of phrases at scale. *arXiv preprint arXiv:2012.12624*, 2020.

[25] Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. Pre-training via paraphrasing. *Advances in Neural Information Processing Systems*, 33:18470–18481, 2020.

[26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

[27] Yu A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836, 2018.

[28] Jianmo Ni, Gustavo Hernández Ábrego, Noah Constant, Ji Ma, Keith B Hall, Daniel Cer, and Yinfei Yang. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*, 2021.

[29] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021.

[30] Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.

[31] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.

[32] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.

[33] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.

[34] Stephen Robertson and Hugo Zaragoza. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc, 2009.

[35] Minjoon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur P Parikh, Ali Farhadi, and Hannaneh Hajishirzi. Real-time open-domain question answering with dense-sparse phrase index. *arXiv preprint arXiv:1906.05807*, 2019.

[36] Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR, 2018.

[37] Devendra Singh, Siva Reddy, Will Hamilton, Chris Dyer, and Dani Yogatama. End-to-end training of multi-document reader and retriever for open-domain question answering. *Advances in Neural Information Processing Systems*, 34, 2021.

[38] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*, 2022.

[39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[40] Yuhuai Wu, Markus N Rabe, DeLesley Hutchins, and Christian Szegedy. Memorizing transformers. *arXiv preprint arXiv:2203.08913*, 2022.

[41] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *International Conference on Learning Representations (ICLR)*, April 2021.

[42] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online, June 2021. Association for Computational Linguistics.

[43] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.

[44] Dani Yogatama, Cyprien de Masson d'Autume, and Lingpeng Kong. Adaptive semiparametric language models. *Transactions of the Association for Computational Linguistics*, 9:362–373, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] We mentioned the limitations in Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We still need to clean up the code before it's ready.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4 and Section 5.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] It would take too long and too much resources to run each experiment multiple times.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We reported the type of device we used to train the model as well as the steps per second for mT5 XL model. For more information on run time, see Appendix.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes] Will be mentioned in Appendix.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

Appendix A Full data on POS tags and percentage of improvement on the log likelihood score.

	Not in context, not in input	Not in context, in input	In context, not in input	In context, in input	Occurrence
NN	0.157771	0.698062	16.197878	5.578312	23639620
NNP	0.020258	0.527689	35.891013	8.141562	16840255
IN	-0.169446	-0.187711	3.308690	1.816983	13603903
JJ	0.400632	1.011774	13.077988	5.019096	12885333
DT	0.224373	0.196791	3.253798	2.453526	11128747
NNS	0.131760	0.594879	17.203110	5.242418	10088414
.	0.074150	-0.292448	7.026237	0.411417	6330229
RB	0.692010	0.897788	5.454908	2.711667	6174253
VB	0.384166	0.659601	9.219075	3.491015	5390427
,	0.731329	0.710795	3.420206	2.353697	5245713
PRP	-0.260621	-0.006360	3.305721	1.998650	5219071
CC	0.055603	0.277403	3.446472	2.634180	4463355
VBZ	0.355881	0.093804	5.872791	0.625349	3895160
VBG	0.588187	0.913552	11.268538	4.581329	3834853
VBN	0.260189	0.650080	12.332930	4.213668	3646850
VBD	0.312395	0.378780	6.878956	2.758350	3377739
VBP	0.333019	0.362408	5.038212	1.873432	3308882
TO	0.000127	-0.183880	3.921325	1.767642	3205519
CD	0.619997	1.148265	20.414486	5.950232	2723028
PRP$	0.040851	0.269473	4.617825	2.683837	1942225
MD	-0.115006	0.047148	3.572827	2.559091	1590480
:	-0.729198	-0.690925	5.812855	3.404822	650535
WDT	0.455384	0.571508	3.444940	1.234093	576220
WRB	0.638024	0.921696	5.063341	3.163921	572072
JJR	1.010684	1.612511	7.880883	4.808698	483001
(	0.226634	0.602507	3.785168	2.695342	450846
WP	1.296407	1.746595	5.018384	3.681695	437709
RP	0.347023	0.406551	5.269570	3.544124	416556
POS	-1.028733	0.472410	8.667241	2.427072	387925
)	0.200041	0.019679	4.084305	1.586972	369449
JJS	0.499507	0.790719	7.975327	4.348411	342354
NNPS	-0.206552	0.327515	30.412756	6.301600	233250
EX	0.793452	1.222533	4.670563	3.714079	201943
RBR	1.037655	2.046677	5.190104	4.050341	192250
PDT	0.604735	1.509266	3.752347	3.438381	84221
RBS	0.114155	0.759014	4.438339	4.140244	83746
FW	0.779163	0.560080	13.321369	4.776943	80379
$	-0.071967	0.671289	18.502472	5.415212	63854
UH	0.154138	0.207683	7.501694	3.188253	45710
”	0.967397	1.077652	17.579383	2.770341	30181
WP$	2.702986	3.293384	6.577276	3.215185	18004
#	0.767495	4.536762	31.027920	14.052459	14969
SYM	0.343766	-3.428490	20.566550	3.729796	5532
“	1.352541	4.927900	35.315082	16.644833	2998
LS	3.417315	-1.039910	57.893078	8.573192	836

Appendix B Visualization of context and change in token log likelihood

Context:
was humbly decorated but Grand Master Raphael Cotton er ordered the renovation of the interior. Mattia Preti the artist from Calabria Italy was commissioned to embellish the interior with detailed carved stone wall designs finished with gold leaf gilding including the painting of the vaulted ceiling and side altars. They depicting the various scenes from the life of Saint John. The Knights of St. John were very decisive to spend and dedicate their resources for the best ornamentals for the cathedral. During 1748, Grand Master António Manoel de Vilhena built new adjoining buildings to enhance the requirements of the time. When Sir Walter Scott came to Malta in 1831 he said that the cathedral was a "magn ificent church, the most striking interior ever seen." He spent three weeks in Malta during October 1831. The floor of the co-cathedral is covered with 400 marble carved tombstones with superbly intricate, marble-inlaid created by skilled tradesmen, tombstones create a formidable appearance to the large floors around all the church. All the buried knights are sons of Europe's noble families from the time of the 16th to the 18th century. Here lies the mastermind behind the building of Valletta, his name was Grand Master Jean Paris ot de la Vall ette among the Grand Masters buried within the crypt. St. John's Co-Cathedral always remained the central religious place for the Knights of Malta until the French occupied Malta in 1798. The French ransacked the church from its very rich belongings which were melted down into coins to pay their soldiers. The museum adjoining the cathedral contain several art objects. The most important are the Fle mish 29 priceless tapestries according to designs by Peter Paul Rubens at the Flemish atelier of Judecos de Vos, famous as the court Weaver for King Louis XIV. They were donated by Aragonese Grand Master Ramon Perellos y Roccaful in 1697. It was the traditional custom that when a newly elected Grandmaster he is to bestow a gift or gioi

Input:
nod to 21st century architecture – the new parliament designed by the eminent Italian architect Renzo Piano. The Knights of St. John's 268 year rule very much shaped Malta's built environment with miles of bastions, forts and palaces which still define the islands. But once inside the story is a very different one. It is here that one is likely to run out of superlatives in attempting to come to describe this magnificent building. The knights certainly went all out to create an interior quite unlike any other and no expense was spared in the embellishment of their main place of worship. The cathedral's walls are richly carved and the monuments to the various Grandmasters fill almost every available space. The church is the proud owner of two paintings by Caravaggio – including his largest and only signed canvas The Beheading of Saint John, commissioned for the church's Oratory where it hangs to this day. The Knights employed the Calabrian artist Mattia Preti for the most ambitious part of the church's décor – the magnificent cycle of ceiling paintings depicting the life of the saint. The renowned Maltese sculptor Melchiorre Gafa – at the time making a name for himself in Rome - was commissioned for the apse's sculpture of the Baptism of Christ. Sadly his early demise meant that the work was never completed and the project was eventually entrusted to Gafa's only pupil, the Italian Giuseppe Mazzu oli. Not to be overlooked is the cathedral's magnificent floor – made up of around 364 knights' tombstones in inlaid marble laid one next to the other. In the adjoining museum one finds a fine set of large Flemish tapestries which used to hang in the church on special feast days. The cathedral'

Targets:
~~s~~ small-crypt-contains-~~the~~ or nate-~~tombs of~~ some-~~of the~~ prominent-~~Grand~~ Masters—among them Jean Paris ot-~~de~~ Val ette—~~the~~ found er-~~of the~~ city- No visit-~~to the islands is~~ complete-without-~~a~~ visit-~~to this magnificent church—the~~ ape x-~~of the Knights'~~

(a) Sampled from context-input-target triplet with top 1% of context utility.

Context:
just brings your outfit all together. So I am listing some of my favorite bangles that I hope to one day add to my collection. This first one that I would love to be able to own one day is a Cartier gold bangle bracelet. I just think these are so stylish and can be worn with any outfit. Okay I won't lie, Kourtney Kardashian was a total inspiration for it... yes Kourtney not Kylie... but I just think they look so cute and can really bring an outfit together. Let's just ignore the price though because unless I become a millionaire there is a good chance I will never be able to afford one of these bad boys. That's okay though, a girl can dream. The next bracelet is one that I actually already own a couple of. These Kendra Scott bracelets can make an outfit pop a little bit more. Even though these bracelets aren't technical ly bangles, they're cuffs, my wrist is so tiny they close all the way around and look like a bangle. Anyways, these K endra Scott Erica Bracelets with stones can help an outfit stand out because of the different colored stones. They have so many different options and they match the stone-to-metal colors very nicely. Some of my go-to bracelets in my everyday outfits for sure. Now okay, I know you shouldn't own one of these bracelets until you actually "tie the knot," but I can't help it! I'm obsessed with the simplicity and such special meaning behind them. The Kate Spade knot bangles are definitely a gift that I want for when I get married. This next bangle is actually something that has been on my Christmas list for about three years now. The Coordinate bangle. I just think all of their jewelry is pretty flawless and can have such a special meaning behind it. For one of Omar's Christmas gifts last year (2014) I got us both matching leather bracelets. His bracelet has my home coordinates and my bracelet had his base in Japan's coordinates

Input:
We have our first big family vacation coming up this fall and it is a serious road trip...all the way to Texas. As each day brings us closer to our departure date I am already planning the packing in my head. This includes all of the making that will need to take place beforehand. There is nothing like a trip to speed up my progress on a garment, or gifty for Desmond, in the works. I saw these lush, silk wrapped bracelets in the latest Real Simple and they begged me to bring them along on our deep south adventure. Let the making begin! Many of you know that I am no strangers to w rapped bangles. Wrapped in wool, wrapped in cotton, I love a wrapped bangle, so this project wasn't new territory. I even had some bangles lying around. These were hand-me-downs from my Mom. I've worn them for years and their shape made them great candidates for a silk-wrapped update. Did I mention I happened to have some hand -

Targets:
dy~~ed silk~~-ribbon-~~in the~~ studio? Inconceivable~~, I know. I~~-actually won-~~these~~ (I NEVER win-anything)-~~in a~~-drawing-~~from the~~-Richmond Craft Mafia-~~a~~ couple-~~of years~~ ago-~~I'd~~ been-saving-~~them for~~ something glamorous-~~and it was~~-

(b) Randomly sampled.

Figure 6: Visualize the difference in log-likelihood for each target token with and without context.

Visualize the delta log-likelihood for each target token with and without context. The color of each target token is determined by $\log P(y_i|y_{<i}, \mathbf{x}, \mathbf{c}, \theta) - \log P(y_i|y_{<i}, \mathbf{x}, \theta')$. The color is green if the difference is positive (context helped). The color is red if the difference is negative (context harmed). We put a box around a target token if it only appeared in the context, but not input. We put a strike-

through on a token if it appeared in the input. We show two examples here. Figure 6a is sampled from the context-input-target triplet with top 1% of context utility. Figure 6b is randomly sampled from the whole C4 train split. We provide more examples in the supplementary materials. They are interactive HTML files that you can mouse-over the tokens to see the difference in log-likelihood.

Appendix C Model training steps per second

Model Size	Steps/Sec	Training Devices	Number of Steps	Batch size
C4 Language Modeling				
`113M`	3.3	64 TPUv3	1,100,000	512
`409M`	2.2	64 TPUv3	1,100,000	512
`1560M`	1.5	128 TPUv3	1,100,000	512
Natural Question				
`409M`	212.5	64 TPUv3	40,000	64
`1560M`	74.6	64 TPUv3	40,000	64

Appendix D License for assets used in this work

Name	License	Link
C4 English Dataset	ODC-By-1.0	https://commoncrawl.org
Natural Question	CC BY-SA 3.0	https://github.com/google-research-datasets/natural-questions/tree/master/nq_open
T5X	Apache-2.0	https://github.com/google-research/t5x
mT5	Apache-2.0	https://github.com/google-research/multilingual-t5
T5 Retrieval	Apache-2.0	https://github.com/google-research/t5x_retrieval