

VAULT: VArIable Unified Long Text Representation for Machine Reading Comprehension

Haoyang Wen^{2†*}, Anthony Ferritto^{1†}, Heng Ji²
Radu Florian¹, Avirup Sil¹

¹ IBM Research AI, ² University of Illinois Urbana-Champaign
wen17@illinois.edu, aferritto@ibm.com
hengji@illinois.edu, {raduf, avi}@us.ibm.com

Abstract

Existing models on Machine Reading Comprehension (MRC) require complex model architecture for effectively modeling long texts with paragraph representation and classification, thereby making inference computationally inefficient for production use. In this work, we propose VAULT: a light-weight and parallel-efficient paragraph representation for MRC based on contextualized representation from long document input, trained using a new Gaussian distribution-based objective that pays close attention to the partially correct instances that are close to the ground-truth. We validate our VAULT architecture showing experimental results on two benchmark MRC datasets that require long context modeling; one Wikipedia-based (Natural Questions (NQ)) and the other on TechNotes (TechQA). VAULT can achieve comparable performance on NQ with a state-of-the-art (SOTA) complex document modeling approach while being 16 times faster, demonstrating the efficiency of our proposed model. We also demonstrate that our model can also be effectively adapted to a completely different domain – TechQA – with large improvement over a model fine-tuned on a previously published large PLM.

1 Introduction

Machine Reading Comprehension (MRC) has seen great advances in recent years with the rise of pre-trained language models (PLM) (Devlin et al., 2019; Liu et al., 2019; Lan et al., 2019) and public leaderboards (Rajpurkar et al., 2016, 2018; Yang et al., 2018; Joshi et al., 2017; Welbl et al., 2018; Kwiatkowski et al., 2019). While some challenges (Rajpurkar et al., 2016, 2018) focus on reading comprehension with shorter contexts, many others

(Welbl et al., 2018; Joshi et al., 2017; Kwiatkowski et al., 2019; Tanaka et al., 2021) focus on longer contexts that cannot fit into a typical 512 sub-token transformer window. Motivated by this, we focus on reading comprehension with long contexts.

One newer approach to this task (Zheng et al., 2020) focuses on modeling document hierarchy to represent multi-grained information for answer extraction. Although this approach creates a strong representation of the text, it suffers from a significant drawback. The graph-based methods (Veličković et al., 2018) are inefficient on parallel hardware, such as GPUs, resulting in slow inference speed (Zhou et al., 2018; Zheng et al., 2020). Motivated by this, in this paper, we propose a reading comprehension model that addresses the above issue and uses a more light-weight, parallel-efficient (i.e. efficient on parallel hardware) paragraph representation based on long contextual representations for providing paragraph answers to questions. Instead of modeling document hierarchy from tokens to document pieces, we first introduce a base model that builds on top of a large “long-context” PLM (we use Longformer, Beltagy et al., 2020) to model longer contexts with lightweight representations of each paragraph. We note that while our approach could work with any PLM, we expect it to perform better with models that can support long contexts and therefore see more paragraph representations at once (Gong et al., 2020). To provide our model a notion of paragraph position relative to a text we also introduce position-aware paragraph representations (PAPR) utilizing special markup tokens and provide them as input for efficient paragraph classification. This approach allows us to encode paragraph-level position in the text and teach the model to impute information on each paragraph into the hidden outputs for these tokens that we can exploit to determine in which paragraph the answer resides. We then predict the

* Work done during an internship at IBM Research AI.

† Equal contributions.

answer span from this identified paragraph.

While previous MRC methods (Chen et al., 2017; Devlin et al., 2019) use ground-truth start and end span positions exclusively as training objectives when extracting answer spans from the context and consider all other positions as incorrect instances equally. However, spans that overlap with the ground-truth should be considered as partially correct. Motivated by Li et al. (2020) which proposes a new optimization criteria based on constructing prior distribution over synonyms for machine translation, we further propose to improve the above base model by considering the start and end positions of ground-truth answer spans as Gaussian-like distributions, instead of single points, and optimize our model using statistical distance.

We call this final model, VAULT (VAriable Unified Long Text representation) as it can handle a variable number and lengths of paragraphs at any position with the same unified model structure to handle long texts.

To evaluate the performance of VAULT, we select the new Natural Questions (NQ, Kwiatkowski et al., 2019) and TechQA (Castelli et al., 2020) datasets. NQ attempts to make Machine Reading Comprehension (MRC) more realistic by providing longer Wikipedia documents as contexts and real user search-engine queries as questions, and aims at avoiding *observation bias*: high lexical overlap between the question and the answer context which can happen frequently if the question is created after the user sees the paragraph (Rajpurkar et al., 2016, 2018; Yang et al., 2018; Chakravarti et al., 2020; Karpukhin et al., 2020; Lee et al., 2019; Murdock et al., 2018). The task introduces the extraction of long answers (henceforth LA; typically paragraphs) besides also requiring short answers (henceforth SA) similar to SQuAD (Rajpurkar et al., 2016). In Figure 1 we examine an example from NQ along with the answers of VAULT and (Zheng et al., 2020). We see that while VAULT can extract answers from the very bottom of a page – if relevant – the existing system suffers from positional bias. It often predicts answers from the first paragraph of Wikipedia (a region which often contains the most relevant information). We evaluate our model for domain adaptation on TechQA, a recently introduced challenging dataset for QA on technical support articles where answers are typically 3-5 times longer than standard MRC datasets (Rajpurkar et al., 2016, 2018).

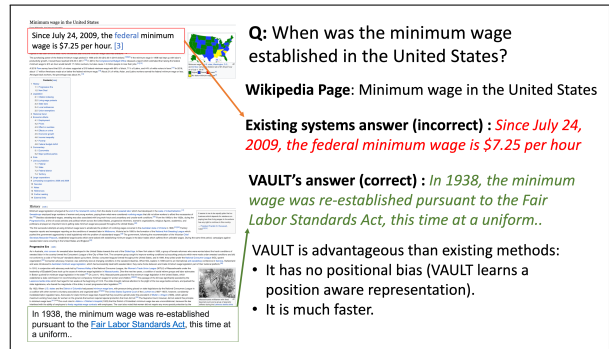


Figure 1: Example from the NQ dataset with answers from VAULT and (Zheng et al., 2020).

Empirically we first show that VAULT achieves comparable performance on NQ with (Zheng et al., 2020)'s document modeling architecture based on graph neural networks while being 16 times faster, demonstrating the efficiency of our proposed model. Secondly, we show the generalization of our model architecture for domain adaptation on TechQA. Our experiments show that our model pre-trained on NQ can be effectively adapted to TechQA outperforming a standard fine-tuned model trained on a large PLM such as RoBERTa. To summarize, our contributions include:

1. We introduce a novel and effective yet simple paragraph representation.
2. We introduce soft labels to leverage information from local contexts near ground-truth during training which is novel for MRC.
3. Our model provides similar performance to a SOTA system on NQ while being 16 times faster and also effectively adapts to a new domain: TechQA.

2 Related Work

Machine reading comprehension has been widely modeled as cloze-type span extraction (Chen et al., 2017; Cui et al., 2017; Devlin et al., 2019). In NQ, we need to identify answers in two levels, long and short answers. (Alberti et al., 2019a) adapt a span extraction model for short answer extraction. (Zheng et al., 2020; Liu et al., 2020) construct complex networks for paragraph-level representation to enhance long answer classification along with span extraction for short answers. In this work, we propose a more light-weight and parallel-efficient way for constructing paragraph-level representation and classification by using longer context and

modeling the negative instance through Gaussian prior optimization.

Using the hierarchical nature of a long document for question answering has been previously studied by (Choi et al., 2017), where they use a hierarchical approach to select candidate sentences and extract answers in those candidates. However, due to the limit of input length for large PLMs, existing methods (Alberti et al., 2019b; Zheng et al., 2020; Chakravarti et al., 2020) slice long documents into document pieces and perform prediction for each piece separately. In our work, we show that by modeling longer input with position-aware paragraph representation coupled with Gaussian prior optimization (which is novel for MRC), we can achieve comparable performance using much simpler architecture compared to previous models, which coincide with recent new PLM for long inputs on question answering (Ainslie et al., 2020)¹.

3 Model Architecture

In this section, we introduce VAULT, our proposed model that uses a simple yet effective paragraph representation based on a longer context. VAULT starts from a base classifier that utilizes position-aware paragraph representations trained on top of a large PLM: Longformer (Beltagy et al., 2020). Next, we further introduce our Gaussian Prior-based training objective that considers partial credits for positions near the ground-truth, instead of only focusing on one ground-truth position. We show an overview of VAULT on the example from Figure 1 in Figure 2.

3.1 A Base “Paragraph” Predictor Model

SOTA methods for paragraph prediction (Zheng et al., 2020; Liu et al., 2020) represent paragraphs through expensive graph modeling, making it inefficient for “large-scale” production MRC systems. On the other hand, simply selecting the first paragraph performs poorly (Kwiatkowski et al., 2019). We hypothesize that by modeling a much longer context even simple paragraph representation can be effective for paragraph (i.e., long answers) classification. For this purpose, we employ a large-window PLM: Longformer (Beltagy et al., 2020), which has shown effectiveness in modeling long contexts for QA (Yang et al., 2018; Welbl et al., 2018; Joshi et al., 2017). Compared to conven-

¹The code and model weights of ETC has not been released at the time of writing of the paper for us to have an accurate comparison.

tional Transformer-based PLMs e.g. RoBERTa (Liu et al., 2019) that can only take up to 512 subword tokens, Longformer provides a much larger maximum input length of 4,096.

Position-aware Paragraph Representation (PAPR): To address the fact that many popular *unstructured* texts such as Wikipedia pages have relatively standard ways of displaying certain relevant information (e.g. birthdays are usually in the first paragraph vs. spouse names are in the “Personal Life” paragraph), we provide the base model with a representation of which part of the text it is reading by marking the paragraphs with special atomic markup tokens (`[paragraph=i]`) at the beginning of each paragraph, indicating the position of the paragraph within the text². With this input representation, we then directly perform long answer classification using the special paragraph token output embedding. Formally, for every paragraph $l_i \in P$, where P are all paragraphs in a text and the representation for the corresponding markup token h_i^p , the logit of a paragraph answer a it computes is as $a_i^p = \mathbf{W}h_i^p + \mathbf{b}$.

We obtain additional document-piece representation from the standard `[CLS]` (Devlin et al., 2019) token to model document pieces that do not contain paragraph answers. The probability of choosing the paragraph given context c , is computed as the softmax over paragraph candidate (with an answer span) logits and not containing answer logit:

$$p_l(l_i | c) = \text{softmax}(a_i^p).$$

We pad the paragraph representations to ensure a rectangular tensor in a batch. Our final prediction strategy is similar to Zheng et al. (2020) as we first choose the paragraph candidate with the highest logit among all candidates. We then extract span answers within the selected paragraph answer candidate using a standard pointer network.

3.2 Gaussian Prior Optimization (GPO)

Conventional span extraction models (Chen et al., 2017; Glass et al., 2020; Liu et al., 2020) optimize the probability of predicted start and end positions of the answer spans with ground-truth spans via maximum likelihood estimation (MLE, Wilks et al., 1938). MLE methods promote the probability for the ground-truth positions while suppressing the probability for all other positions. However we hypothesize that, for all those negative instances, the positions that are near the ground-truth should

²Similar tags are added for lists and tables.




Figure 2: Overview of VAULT answering the example from Figure 1. The 10th paragraph containing the correct answer is underlined. The span linear layer receives hidden state outputs from all 4096 tokens in the window to create the start and end logits. The paragraph linear layer receives the orange-highlighted [CLS] and markup tokens (e.g. [Para=10]) to predict in which paragraph the answer resides. These logits are then used together to first select the best paragraph (LA) and finally select the best answer within said paragraph (SA).

be given higher credit than farther distant positions, since the extracted answers will be partially overlapping with the ground-truth.

To tackle this problem, we follow the intuition from Li et al. (2020) which proposes to promote the probability of generating synonyms using a Gaussian-like distribution for machine translation. We construct the distribution where it has the highest probability at ground-truth positions, and drop the probability exponentially as computed by the distance to the corresponding ground-truth positions. Specifically, for a groundtruth start or end position at y_s , where $s \in \{\text{start}, \text{end}\}$, we use a Gaussian distribution $\mathcal{N}(y_s, \sigma)$, where the mean is the position y_s and variance σ is a hyperparameter. We consider the probability density $\varphi(y | y_s, \sigma)$ of the Gaussian distribution at each position y as the logit for the corresponding position. We then use the softmax function with temperature T to rescale the logits to get the Gaussian-like distribution $q(y | \hat{y}_s)$ for ground-truth distribution at position y_s ,

$$q(y | y_s) = \text{softmax}(\varphi(y | y_s, \sigma)/T).$$

We augment our MLE objective with an additional KL divergence (Kullback and Leibler, 1951) term

between constructed distribution $q(y | y_s)$ and model prediction $p_s(y | c)$, $s \in \{\text{start}, \text{end}\}$, so that we can guide our model to follow the Gaussian-like distribution for partial credit.

$$\begin{aligned} L_D &= KL(q(y | y_s) || p_s(y | c)) \\ &= \sum_y q(y | y_s) \log p_s(y | c) \\ &\quad - \sum_y q(y | y_s) \log q(y | y_s). \end{aligned}$$

We refer to this final model as VAULT.

4 Experiments

Datasets: We experiment with two challenging “natural” MRC datasets: NQ (Kwiatkowski et al., 2019) and TechQA (Castelli et al., 2020). We provide a brief summary of the datasets and direct interested readers to the corresponding papers. NQ consists of crowdsourced-annotated *full* Wikipedia pages which appear in Google search logs with two tasks: the start and end offsets for the short answer (SA) and long answer (LA, eg. paragraph) – if they exist. TechQA is developed from real user questions in the customer support domain where each question is accompanied by 50 documents –

at most one of which has an answer – with answers significantly longer ($\sim 3\text{-}5\times$) than standard MRC datasets like SQuAD. We report official F1 scores for each dataset.

Results on NQ: We train VAULT on NQ – predicting the paragraph and span answers as NQ’s LA and SA respectively – and compare against ROBERTA_{DM}: a RoBERTa (Liu et al., 2019) variant of the SOTA document model (DM) (Zheng et al., 2020) using the base variants for a more systematic comparison. Although it may seem fair to include a Longformer DM baseline in our table, doing so would be infeasible (and unwise) due to production resource constraints. We further show the impact of VAULT by providing ablation experiments where its components (GPO and PAPR) are removed. The base LM (Longformer in our experiments) without GPO and PAPR, is implemented in the style of (Alberti et al., 2019b; Chakravarti et al., 2020) where we first predict the SA and then select the enclosing LA. We aim to show that our proposed method provides comparable results to ROBERTA_{DM} while being considerably faster while decoding and displaying improved performance over experiments just using the language model. To do this we consider development set SA and LA F1 (the F1 metrics with respect to the span and paragraph answers respectively) as well as decoding time t_{decode} (on a V100) as metrics.

Table 1 shows the results on the NQ dev set. We see VAULT and ROBERTA_{DM} provide comparable F1 performance (precision and recall are shown in the Appendix). However, when it comes to decoding time, we can find VAULT decodes over 16 times faster than ROBERTA_{DM}. We additionally see in the ablation experiments that our enhancements increase both F1 metrics by multiple points, at the expense of some decoding time. In particular we note that the F1 performance of Longformer is not competitive with VAULT. We conclude that VAULT provides the best balance of F1 and decoding time as it is effectively tied on F1 (with ROBERTA_{DM}) and is only around 20 minutes slower to decode than the quickest model.

Domain Adaptation: Results on TechQA: Since VAULT has shown to be effective on NQ, we evaluate it on a new domain, TechQA. We compare it against a RoBERTa base model trained with the same hyper-parameters as (Castelli et al., 2020) – except we use 11 epochs instead of 20. We chose base instead of large (as is used for the TechQA

Model	SA F1	LA F1	t_{decode}
ROBERTA _{DM}	52.2	70.1	11h
VAULT	51.6	70.4	40m
- GPO	49.1	67.6	41m
- PAPR (Longformer)	49.5	65.6	22m

Table 1: Comparison of VAULT vs. ROBERTA_{DM} on NQ. We achieve comparable performance while being 16 times faster.

baseline) to give a fair comparison since we are using a base PLM for our experiments with VAULT. Similarly, we use RoBERTa rather than BERT as it is closer to Longformer. Having already established the run-time effectiveness of VAULT on NQ, we focus on F1 metrics here, including “has answer” (HA) F1. We consider HA F1 our primary metric as we are exploring paragraph answer extraction in this work and (as previously mentioned) answers in TechQA are much longer than other datasets. We believe that the improvements in HA F1, at least partially, come from GPO.

Model	F1	HA F1
RoBERTa	48.6	7.6
VAULT	49.3	16.1

Table 2: Results on TechQA dev set. VAULT clearly outperforms RoBERTa on both F1 and Has Answer F1.

Results on TechQA are reported in Table 2. We see that our VAULT model provides an improvement of 0.7 F1 and 8.5 HA F1 (denotes Has Answer); thus showing the effectiveness of our approach. In particular, we see that this approach of imputing a paragraph structure to classify provides a large boost to performance when a non-null answer exists (HA F1).

5 Conclusions

In this work we introduce and examine a powerful yet simple model for reading comprehension on long texts which we call VAULT, based on the hypothesis that with a large sequence length long answers can be classified effectively without computationally heavy graph-based models. We validate our approach by showing it yields F1 scores competitive with heavier methods at a fraction of the decoding cost on two very different domain benchmark datasets that require reading long texts.

References

Joshua Ainslie, Santiago Ontañón, Chris Alberti, Philip Pham, Anirudh Ravula, and Sumit Sanghai. 2020.

- ETC: encoding long and structured data in transformers. *CoRR*, abs/2004.08483.
- Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. 2019a. [Synthetic QA corpora generation with roundtrip consistency](#). *CoRR*, abs/1906.05416.
- Chris Alberti, Kenton Lee, and Michael Collins. 2019b. [A BERT baseline for the natural questions](#). *CoRR*, abs/1901.08634.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.
- Vittorio Castelli, Rishav Chakravarti, Saswati Dana, Anthony Ferritto, Radu Florian, Martin Franz, Dinesh Garg, Dinesh Khandelwal, J. Scott McCauley, Mike McCawley, Mohamed Nasr, Lin Pan, Cezar Pendus, John F. Pitrelli, Saurabh Pujar, Salim Roukos, Andrzej Sakrajda, Avirup Sil, Rosario Uceda-Sosa, Todd Ward, and Rong Zhang. 2020. [The techqa dataset](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1269–1278. Association for Computational Linguistics.
- Rishav Chakravarti, Anthony Ferritto, Bhavani Iyer, Lin Pan, Radu Florian, Salim Roukos, and Avi Sil. 2020. [Towards building a robust industry-scale question answering system](#). In *Proceedings of the 28th International Conference on Computational Linguistics: Industry Track*, pages 90–101, Online. International Committee on Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Eunsol Choi, Daniel Hewlett, Jakob Uszkoreit, Illia Polosukhin, Alexandre Lacoste, and Jonathan Berant. 2017. [Coarse-to-fine question answering for long documents](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 209–220, Vancouver, Canada. Association for Computational Linguistics.
- Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2017. [Attention-over-attention neural networks for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 593–602, Vancouver, Canada. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Michael Glass, Alfio Gliozzo, Rishav Chakravarti, Anthony Ferritto, Lin Pan, G P Shrivatsa Bhargav, Dinesh Garg, and Avi Sil. 2020. [Span selection pre-training for question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2782, Online. Association for Computational Linguistics.
- Hongyu Gong, Yelong Shen, Dian Yu, Jianshu Chen, and Dong Yu. 2020. [Recurrent chunking mechanisms for long-text machine reading comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6751–6761, Online. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [Albert: A lite bert for self-supervised learning of language representations](#).
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.

- Zuchao Li, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, Zhuosheng Zhang, and Hai Zhao. 2020. [Data-dependent gaussian prior objective for language generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Dayiheng Liu, Yeyun Gong, Jie Fu, Yu Yan, Jiusheng Chen, Daxin Jiang, Jiancheng Lv, and Nan Duan. 2020. [RikiNet: Reading Wikipedia pages for natural question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6762–6771, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- J. William Murdock, Lin Pan, Chung-Wei Hang, Mary Swift, Zhiguo Wang, Chris Nolan, Prathyusha Peddi, Nisarga Markandaiah, Eunyoung Ha, Kazi Hasan, and et al. 2018. [Engineered ai still matters for question answering](#). *Advances in Cognitive Systems*, 6:140–158.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). *EMNLP*.
- Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. 2021. [Visualmrc: Machine reading comprehension on document images](#). *CoRR*, abs/2101.11272.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *International Conference on Learning Representations*.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents](#). *Transactions of the Association for Computational Linguistics*, 6:287–302.
- SS Wilks et al. 1938. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Bo Zheng, Haoyang Wen, Yaobo Liang, Nan Duan, Wanxiang Che, Daxin Jiang, Ming Zhou, and Ting Liu. 2020. [Document modeling with graph attention networks for multi-grained machine reading comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6708–6718, Online. Association for Computational Linguistics.
- Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2018. [Graph neural networks: A review of methods and applications](#). *CoRR*, abs/1812.08434.

A Additional Experimental Results

For interested readers we further show precision and recall numbers for the NQ experiments in Table 3.

B Implementation Details

B.1 NQ

All models for this work are implemented in (Wolf et al., 2020). We use the following hyperparameters for VAULT when finetuning on NQ: sequence length 4096, doc stride 2048 (Ainslie et al., 2020), negative instance subsampling rates (has answer/no answer) 0.02/0.08, learning rate 5e-5, and 4 epochs of training.

B.2 TechQA

While TechQA does provide full HTML for its Technotes, the answers are annotated with respect to the cleaned plaintext. Therefore to determine paragraph breaks for VAULT we split on the "\n\n" token "ĈĈ" in the vocabulary. By imputing paragraph answers in this way, we are then able to predict the paragraph answer and then a contained span answer.

C Example Analysis

We examine additional examples in Figure 3 to provide insight on the improvements of VAULT. We

Model	SA F1	SA P	SA R	LA F1	LA P	LA R
ROBERTA _{DM}	52.2	57.2	48.0	70.1	69.4	70.9
VAULT	51.6	61.5	44.4	70.4	69.5	71.4
- GPO	49.1	57.6	42.7	67.6	67.0	68.1
- PAPER (Longformer)	49.5	56.4	44.2	65.6	62.4	69.3

Table 3: Comparison of VAULT vs. ROBERTA_{DM} on NQ with precision (P) and recall (R) statistics.

compare the correct answers produced by VAULT with the incorrect answers produced by the ablated model from the last row of Table 3 (NQ) and Roberta baseline from the first row of Table 2 (TechQA).

In the first example the gold SA is null, however there is a gold LA. This indicates that there is no short span which answers the question: the correct answer here is an entire paragraph. This does not confuse VAULT which is able to identify the correct answer directly. However the ablated model which attempts to predict SA first struggles here – predicting the incorrect LA – as there is no gold SA.

In the second example we see that in this Technote both the correct and incorrect answers are single sentence paragraphs surrounded by paragraph breaks. Our VAULT is able to identify the correct paragraph using our imputed structure and select the correct answer – whereas the Roberta baseline selects a nearby but incorrect answer.

Example A1 (NQ)
Question: why did government sponsored surveys and land acts encourage migration to the west
Wikipedia Page: Homestead Acts
Text: ...
An extension of the Homestead Principle in law, the Homestead Acts were an expression of the " Free Soil " policy of Northerners who wanted individual farmers to own and operate their own farms, as opposed to Southern slave-owners who wanted to buy up large tracts of land and use slave labor, thereby shutting out free white men.
The first of the acts, the Homestead Act of 1862 , opened up millions of acres. Any adult who had never taken up arms against the U.S. government could apply. Women and immigrants who had applied for citizenship were eligible. The 1866 Act explicitly included black Americans and encouraged them to participate, but rampant discrimination slowed black gains. Historian Michael Lanza argues that while the 1866 law pack was not as beneficial as it might have been, it was part of the reason that by 1900 one fourth of all Southern black farmers owned their own farms. [1]
 ...

Example A2 (TechQA)
Question: Are there any probes that can connecto the Nokia NSP EPC v17.9 and Nokia NSP RAN v17.3 using JMS/HTTP?
Text: release notice; downloads; nco-p-nokia-nfmp; Probe for Nokia Network Functions Manager for Packet NEWS

ABSTRACT
 This new probe will be ready for downloading on July 20, 2017.

CONTENT
This probe is written to support Nokia Network Functions Manager for Packet release 17.3.
You can download the package you require from the IBM Passport Advantage website:
 www-01.ibm.com...

Figure 3: Additional Examples of questions in the NQ and TechQA datasets. VAULT’s correct answer is shown in green, incorrect baseline in red. (A1) The correct answer is a paragraph LA; only VAULT identifies the correct LA directly even though the gold SA is null. (A2) VAULT identifies the correct "paragraph" answer.