

# Task Transfer and Domain Adaptation for Zero-Shot Question Answering

**Xiang Pan\***  
New York University  
xp2030@nyu.edu

**Alex Sheng\***  
New York University  
alexsheng4@gmail.com

**David Shimshoni\***  
New York University  
ds5396@nyu.edu

**Aditya Singhal\***  
New York University  
adis@nyu.edu

**Sara Rosenthal**  
IBM Research AI  
sjrosenthal@us.ibm.com

**Avirup Sil**  
IBM Research AI  
avi@us.ibm.com

## Abstract

Pretrained language models have shown success in various areas of natural language processing, including reading comprehension tasks. However, when applying machine learning methods to new domains, labeled data may not always be available. To address this, we use supervised pretraining on source-domain data to reduce sample complexity on domain-specific downstream tasks. We evaluate zero-shot performance on domain-specific reading comprehension tasks by combining task transfer with domain adaptation to fine-tune a pre-trained model with no labelled data from the target task. Our approach outperforms Domain-Adaptive Pretraining on downstream domain-specific reading comprehension tasks in 3 out of 4 domains.

## 1 Introduction

Pretrained language models (Liu et al., 2019; Wolf et al., 2020) require substantial quantities of labeled data to learn downstream tasks. For domains that are novel or where labeled data is in short supply, supervised learning methods may not be suitable (Zhang et al., 2020; Madasu and Rao, 2020; Rietzler et al., 2020). Collecting sufficient quantities of labeled data for each new application can be resource intensive, especially when aiming for both a specific task type and a specific data domain. By traditional transfer learning methods, it is prohibitively difficult to fine-tune a pretrained model on a domain-specific downstream task for which there is no existing training data. In light of this, we would like to use more readily available labeled in-domain data from unrelated tasks to domain-adapt our fine-tuned model.

In this paper, we consider a problem setting where we have a domain-specific target task (QA) for which we do not have any in-domain training

data (SQuAD). However, we assume that we have generic training data for the target task type, and in-domain training data for another task. To address this problem setting, we present Task and Domain Adaptive Pretraining (T+DAPT), a technique that combines domain adaptation and task adaptation to improve performance in downstream target tasks. We evaluate the effectiveness of T+DAPT in zero-shot domain-specific machine reading comprehension (MRC) (Hazen et al., 2019; Reddy et al., 2020; Wiese et al., 2017) by pretraining on in-domain NER data and fine-tuning for generic domain-agnostic MRC on SQuADv1 (Rajpurkar et al., 2018), combining knowledge from the two different tasks to achieve zero-shot learning on the target task. We test the language model’s performance on domain-specific reading comprehension data taken from 4 domains: News, Movies, Biomedical, and COVID-19. In our experiments, RoBERTa-Base models trained using our approach perform favorably on domain-specific reading comprehension tasks compared to baseline RoBERTa-Base models trained on SQuAD as well as Domain Adaptive Pretraining (DAPT). Our code is publicly available for reference.<sup>1</sup>

We summarize our contributions as follows:

- We propose Task and Domain Adaptive Pretraining (T+DAPT) combining domain adaptation and task adaptation to achieve zero-shot learning on domain-specific downstream tasks.
- We experimentally validate the performance of T+DAPT, showing our approach performs favorably compared to both a previous approach (DAPT) and a baseline RoBERTa fine-tuning approach.
- We analyze the adaptation performance on different domains, as well as the behavior of

\*Equal Contribution

<sup>1</sup><https://github.com/adityaarunsinghal/Domain-Adaptation>

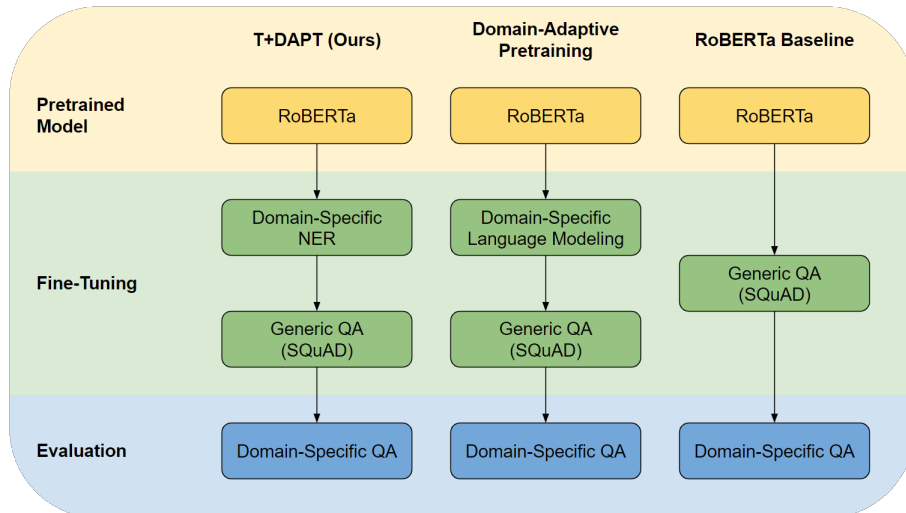


Figure 1: sequential transfer learning procedures of T+DAPT, DAPT, and a RoBERTa baseline for zero-shot question answering.

DAPT and T+DAPT under various experimental conditions.

## 2 Related Work

It has been shown that pretrained language models can be domain-adapted with further pretraining (Pruksachatkun et al., 2020) on unlabeled in-domain data to significantly improve the language model’s performance on downstream supervised tasks in-domain. This was originally demonstrated by BioBERT (Lee et al., 2019). Gururangan et al. (2020) further explores this method of domain adaptation via unsupervised pretraining, referred to as Domain-Adaptive Pretraining (DAPT), and demonstrates its effectiveness across several domains and data availability settings. This procedure has been shown to improve performance on specific domain reading comprehension tasks, in particular in the biomedical domain (Gu et al., 2021). In this paper, as a baseline for comparison, we evaluate the performance of DAPT-enhanced language models in their respective domains, both in isolation with SQuAD1.1 fine-tuning and in conjunction with our approach that incorporates the respective domain’s NER task. DAPT models for two of our domains, News and Biomedical, are initialized from pre-trained weights as provided by the authors of Gururangan et al. (2020). We train our own DAPT baselines on the Movies and COVID-19 domains. Xu et al. (2020) explore methods to reduce catastrophic forgetting during language model fine-tuning. They apply topic modeling on the MS MARCO dataset (Bajaj et al., 2018) to generate 6 narrow domain-

specific data sets, from which we use BioQA and MoviesQA as domain-specific reading comprehension benchmarks.

## 3 Experiments

We aim to achieve zero-shot learning for an unseen domain-specific MRC task by fine-tuning on both a domain transfer task and a generic MRC task. The model is initialized by pretrained RoBERTa weights (Liu et al., 2019), then fine-tuned using our approach with a domain-specific supervised task to augment domain knowledge, and finally trained on SQuAD to learn generic MRC capabilities to achieve zero-shot MRC in the target domain on an unseen domain-specific MRC task without explicitly training on the final task. This method is illustrated in Figure 1.

### 3.1 Datasets

We explore the performance of this approach in the Movies, News, Biomedical, and COVID-19 domains. Specifically, our target domain-specific MRC tasks are MoviesQA (Xu et al., 2020), NewsQA (Trischler et al., 2017), BioQA (Xu et al., 2020), and CovidQA (Möller et al., 2020), respectively. We choose to use named entity recognition (NER) as our supervised domain adaptation task for all four target domains, as labeled NER data is widely available across various domains. Furthermore, NER and MRC share functional similarities, as both rely on identifying key tokens in a text as entities or answers. The domain-specific NER tasks are performed using supervised training data

Dataset	Dev Set	Sample
MoviesQA	755	Q: After its re-opening, which types of movies did the Tower Theatre show? A: second and third run movies, along with classic films
NewsQA	934	Q: Who is the struggle between in Rwanda? A: The struggle pits ethnic Tutsis, supported by Rwanda, against ethnic Hutu, backed by Congo.
BioQA	4,790	Q: What is hemophilia? A: a bleeding disorder characterized by low levels of clotting factor proteins.
CovidQA	2,019	Q: What is the molecular structure of bovine coronavirus? A: single-stranded, linear, and nonsegmented RNA

Table 1: Overview of the domain-specific MRC datasets used in our experiments. The number of question-answer pairs in the train set and development set for each domain is shown, along with a sample question-answer pair from each domain. The datasets share the same format as SQuAD.

from the MIT Movie Corpus (Liu et al., 2013), CoNLL 2003 News NER (Tjong Kim Sang and De Meulder, 2003), NCBI-Disease (Doğan et al., 2014) and COVID-NER<sup>2</sup>. The domain-specific language modeling tasks for DAPT are performed using unsupervised text from IMDB (Maas et al., 2011), the RealNews Corpus (Zellers et al., 2020), the Semantic Scholar Open Research Corpus (Lo et al., 2020) and the Covid-19 Corpus<sup>3</sup>.

### 3.2 Methods

We compare our approach (T+DAPT) to a previous approach (DAPT) as well as a baseline model. For the baseline, the pretrained RoBERTa-Base model is fine-tuned on SQuAD and evaluated on domain-specific MRC without any domain adaptation. In the DAPT approach, RoBERTa-Base is first initialized with fine-tuned DAPT weights (NewsRoBERTa and BioRoBERTa) provided by Gururangan et al. (2020) or implemented ourselves using the methodology described in Gururangan et al. (2020) and different Movies and COVID-19 datasets (Maas et al., 2011; Danescu-Niculescu-Mizil and Lee, 2011; Pang et al., 2019). These models are initialized by DAPT weights—which have been fine-tuned beforehand on unsupervised text corpora for domain adaptation—from the HuggingFace model hub (Wolf et al., 2020), fine-tuned on SQuAD, and evaluated on domain-specific MRC.

### 3.3 Results

We compare the effectiveness of our approach, which uses NER instead of language modeling

<sup>2</sup><https://github.com/tsantosh7/COVID-19-Named-Entity-Recognition>

<sup>3</sup><https://github.com/davidcampos/covid19-corpus>

(as in DAPT) for the domain adaptation method in a sequential training regime. Our experiments cover every combination of domain (Movies, News, Biomedical, or COVID) and domain adaptation method (T+DAPT which uses named entity recognition vs. DAPT which uses language modeling vs. baseline with no domain adaptation at all).

Our results are presented in Table 2. We use F1 score to evaluate the QA performance of each model in its target domain. In our experiments, DAPT performs competitively with baseline models and outperforms in one domain (CovidQA). Our T+DAPT approach (RoBERTa + Domain NER + SQuAD) outperforms the baseline in three out of four domains (Movies, Biomedical, COVID) and outperforms DAPT in three out of four domains (Movies, News, Biomedical). We also test a combination of DAPT and T+DAPT by retraining DAPT models on domain NER then SQuAD, and find that this combined approach underperforms compared to either T+DAPT alone or DAPT alone in all four domains. We further discuss the possible reasons for these results in Section 4.

## 4 Analysis

**Specific domains learn from adaptation:** Our approach shows promising performance gains when used for zero-shot domain-specific question answering, particularly in the biomedical, movies, and COVID domains, where the MRC datasets were designed with the evaluation of domain-specific features in mind. Performance gains are less apparent in the News domain, where the NewsQA dataset was designed primarily to evaluate causal reasoning and inference abilities—which correlate strongly with SQuAD and base-

<b>RoBERTa Retraining Procedure</b>	<b>MoviesQA</b>	<b>NewsQA</b>	<b>BioQA</b>	<b>CovidQA</b>
SQuAD1.1	67.1	<b>57.0</b>	58.0	42.0
DAPT + SQuAD1.1	60.7	54.4	57.8	<b>47.2</b>
<i>T+DAPT</i> (ours)	<b>68.0</b>	56.0	<b>58.9</b>	42.7
DAPT + <i>T+DAPT</i>	66.4	54.2	55.1	43.1

Table 2: F1 score of pretrained RoBERTa-Base models on dev sets of MRC datasets for given domains with the stated retraining regimens

line RoBERTa pretraining—rather than domain-specific features and adaptation. The lack of performance gains from either T+DAPT or DAPT in the News domain could also possibly be attributed to the nature of the domain: Gururangan et al. (2020) found that the News domain had the highest vocabulary overlap of any domain (54.1%) with the RoBERTa pretraining corpus, so the baseline for this domain could have had an advantage in the News domain that would be lost due to catastrophic forgetting while little relevant knowledge is gained from domain adaptation. We perform follow-up experiments with varying amounts of epochs and training data in SQuAD fine-tuning to analyze the tradeoff between more thorough MRC fine-tuning and better preservation of source domain knowledge from DAPT and auxiliary domain adaptation tasks. The results from these runs are in the Appendix (Table 4).

**When does DAPT succeed or fail:** In zero-shot QA, DAPT performs competitively with the baseline in all domains and outperforms in the COVID domain. This builds upon the results of Gururangan et al. (2020), which reports superior performance on tasks like relation classification, sentiment analysis, and topic modeling, but does not address reading comprehension tasks, which DAPT may not have originally been optimized for. Unsupervised language modeling may not provide readily transferable features for reading comprehension, as opposed to NER which identifies key tokens and classifies those tokens into specific entities. These entities are also often answer tokens in reading comprehension, lending to transferable representations between NER and reading comprehension. Another possible factor is that RoBERTa was pretrained on the English Wikipedia corpus, the same source that the SQuAD questions were drawn from. Because of this, it is possible that pretrained RoBERTa already has relevant representations that would provide an intrinsic advantage for SQuAD-style reading comprehension which

would be lost due to catastrophic forgetting after retraining on another large language modeling corpus in DAPT.

In the COVID domain, we use the article dataset from Wang et al. (2020). These articles also make the basis for the CovidNER and CovidQA (Möller et al., 2020) datasets, which may explain the large performance improvement from DAPT in this domain. These results suggest that the performance of DAPT is sensitive to the similarity of its language modeling corpus to the target task dataset.<sup>1</sup>

## 5 Conclusion

We evaluate the performance of our T+DAPT approach with domain-specific NER, achieving positive results in a zero-shot reading comprehension setting in four different domain-specific QA datasets. These results indicate that our T+DAPT approach robustly improves performance of pre-training language models in zero-shot domain QA across several domains, showing that T+DAPT is a promising approach to domain adaptation in low-resource settings for pretrained language models, particularly when directly training on target task data is difficult.

In future work, we intend to explore various methods to improve the performance of T+DAPT by remedying catastrophic forgetting and maximizing knowledge transfer. For this we hope to emulate the regularization used by Xu et al. (2020) and implement multi-task learning and continual learning methods like AdapterNet (Hazan et al., 2018). In order to improve the transferability of learned features, we will explore different auxiliary tasks such as NLI and sentiment analysis in addition to few-shot learning approaches.

## 6 Ethical Considerations

Question answering systems are useful tools in complement to human experts, but the “word-of-

<sup>1</sup>Additional experiments in the COVID domain with different auxiliary tasks are presented in the Appendix A.1



<b>BioQA Samples</b>
Q: what sugar is found in rna DAPT: ribose, whereas the sugar in DNA is deoxyribose T+DAPT: ribose
Q: normal blood pressure range definition DAPT: 120 mm Hg1 T+DAPT: a blood pressure of 120 mm Hg1 when the heart beats (systolic) and a blood pressure of 80 mm Hg when the heart relaxes (diastolic)
<b>MoviesQA Samples</b>
Q: what is cyborgs real name DAPT: Victor Stone/Cyborg is a hero from DC comics most famous for being a member of the Teen Titans T+DAPT: Victor Stone
Q: who plays klaus baudelaire in the show DAPT: Liam Aiken played the role of Klaus Baudelaire in the 2004 movie A Series of Unfortunate Events. T+DAPT: Liam Aiken

Table 3: Samples from BioQA and MoviesQA where T+DAPT achieves exact match with the label answer, and DAPT produces a different answer. Answers from each approach are shown side-by-side for comparison.

machine effect” (Longoni and Cian, 2020) demonstrates the effects of a potentially dangerous over-trust in the results of such systems. While the methods proposed in this paper would allow more thorough usage of existing resources, they also bestow confidence and capabilities to models which may not have much domain expertise. T+DAPT models aim to mimic extensively domain-trained models, which are themselves approximations of real experts or source documents. Use of domain adaptation methods for low-data settings could propagate misinformation from a lack of source data. For example, while making an information-retrieval system for biomedical and COVID information could become quicker and less resource-intensive using our approach, people should not rely on such a system for medical advice without extensive counsel from a qualified medical professional.

## References

- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2018. [Ms marco: A human generated machine reading comprehension dataset](#). *arXiv:1611.09268 [cs]*.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. [Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs](#).
- Rezarta Islamaj Doğan, Robert Leaman, and Zhiyong Lu. 2014. [Ncbi disease corpus: A resource for disease name recognition and concept normalization](#). *Journal of Biomedical Informatics*, 47:1–10.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#).
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#).
- Alon Hazan, Yoel Shoshan, Daniel Khapun, Roy Aladjem, and Vadim Ratner. 2018. [Adaptnet - learning input transformation for domain adaptation](#).
- Timothy J. Hazen, Shehzaad Dhuliawala, and Daniel Boies. 2019. [Towards domain adaptation from limited data for question answering using deep neural networks](#). *arXiv:1911.02655 [cs]*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*.
- Jingjing Liu, Panupong Pasupat, Yining Wang, Scott Cyphers, and Jim Glass. 2013. [Query understanding enhanced by hierarchical parsing structures](#).
- 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](#).

- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Dan S. Weld. 2020. [S2orc: The semantic scholar open research corpus](#).
- Chiara Longoni and Luca Cian. 2020. [Artificial intelligence in utilitarian vs. hedonic contexts: The “word-of-machine” effect](#). *Journal of Marketing*.
- Andrew Maas, Raymond Daly, Peter Pham, Dan Huang, Andrew Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#).
- Avinash Madasu and Vijjini Anvesh Rao. 2020. [Sequential domain adaptation through elastic weight consolidation for sentiment analysis](#). *arXiv:2007.01189 [cs]*.
- Timo Möller, Anthony Reina, Raghavan Jayakumar, and Malte Pietsch. 2020. [COVID-QA: A question answering dataset for COVID-19](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*, Online. Association for Computational Linguistics.
- Timo Möller, Anthony Reina, Raghavan Jayakumar, and Malte Pietsch. 2020. [Covid-qa: A question answering dataset for covid-19](#).
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2019. [Thumbs up? sentiment classification using machine learning techniques](#).
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [Intermediate-task transfer learning with pretrained models for natural language understanding: When and why does it work?](#)
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2018. [Squad: 100,000+ questions for machine comprehension of text](#).
- Revanth Gangi Reddy, Bhavani Iyer, Md Arafat Sultan, Rong Zhang, Avi Sil, Vittorio Castelli, Radu Florian, and Salim Roukos. 2020. [End-to-end qa on covid-19: Domain adaptation with synthetic training](#).
- Alexander Rietzler, Sebastian Stabinger, Paul Opitz, and Stefan Engl. 2020. [Adapt or get left behind: Domain adaptation through BERT language model finetuning for aspect-target sentiment classification](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4933–4941, Marseille, France. European Language Resources Association.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the conll-2003 shared task](#). *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. [Newsqa: A machine comprehension dataset](#).
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Michael Kinney, et al. 2020. [Cord-19: The covid-19 open research dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*.
- Jerry Wei, Chengyu Huang, Soroush Vosoughi, and Jason Wei. 2020. [What are people asking about covid-19? a question classification dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*.
- Georg Wiese, Dirk Weissenborn, and Mariana Neves. 2017. [Neural domain adaptation for biomedical question answering](#).
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).
- Y. Xu, X. Zhong, A. J. J. Yepes, and J. H. Lau. 2020. [Forget me not: Reducing catastrophic forgetting for domain adaptation in reading comprehension](#). *arXiv:1911.00202 [cs]*.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, Yejin Choi, and Paul Allen. 2020. [Defending against neural fake news](#).
- Rong Zhang, Revanth Gangi Reddy, Md Arafat Sultan, Vittorio Castelli, Anthony Ferritto, Radu Florian, Efsun Sarioglu Kayi, Salim Roukos, Avirup Sil, and Todd Ward. 2020. [Multi-stage pre-training for low-resource domain adaptation](#). *arXiv:2010.05904 [cs]*.

## A Appendix

RoBERTa Adaptation Procedure	CovidQA
CovidQA (upper bound)	52.1416
SQuAD only	42.0485
DAPT	47.2190
CovidNER	42.6584
CovidQCLS	42.6300
DAPT+Covid-NER	43.0710
DAPT+Covid-QCLS	<b>45.8314</b>
DAPT+CovidNER+CovidQCLS	43.0854

Table 4: Zero-shot F1 performance of RoBERTa-Base models on dev sets of QA data for COVID domain with SQuAD1.1 following different intermediate pretraining regimens. The CovidQA upper bound score is attained by training directly on the CovidQA train set.

Model	NewsQA
RoBERTa-Base	
1 Epoch, 1000 Samples	19.9953
2 Epochs, 1000 Samples	35.2666
2 Epochs, 5000 Samples	47.0090
2 Epochs, All Samples	<b>56.9803</b>
2 Epochs, All Samples (Head)	05.5891
NewsRoBERTa (DAPT)	
1 Epoch, 1000 Samples	17.9025
2 Epochs, 1000 Samples	28.4453
2 Epochs, 5000 Samples	44.1206

Table 5: Zero-shot F1 performance of RoBERTa-Base models on NewsQA following different amounts of SQuAD fine-tuning. For comparison the score of our News model from the main paper with 2 epochs and all samples is included as an upper bound, alongside a head tuning baseline where all weights are frozen except the classifier layer.

### A.1 Experiment Details and Additional Experiments

**Freezing Layer** - We tried to freeze the bottom layer after NER training and only train the QA layer on SQuAD, the performance is worse than fine-tuning the whole RoBERTa and QA layer. NER and QA may not rely on the exact same features for the final task which may be the reason that freezing causes a performance decrease.

**Different Training Epoch and Training Examples** - When selecting the best performance model, we use a validation set in target domain to evaluate the performance. From Table 5, we show our trials

with different amounts of SQuAD training in the News Domain and how it affected performance in NewsQA.

**Different Training Order** - We tried to use different training order, for example, we train on SQuAD1.1 task first and then on NER, the F1 score is 42.15 in CovidQA, which has some improvement, but QA as the last task performs better.

**Another Auxiliary Task** - In the Covid domain, we also do experiments on a more QA-relevant task, question classification (QCLS) (Wei et al., 2020). We show the result in Table 4. The experiments show that QCLS task have more improvements than NER task. In addition, we test the model trained on CovidQA as the performance upper bound.