

The Cascade Transformer: an Application for Efficient Answer Sentence Selection

Luca Soldaini

Amazon Alexa
Manhattan Beach, CA, USA
lssoldai@amazon.com

Alessandro Moschitti

Amazon Alexa
Manhattan Beach, CA, USA
amosch@amazon.com

Abstract

Large transformer-based language models have been shown to be very effective in many classification tasks. However, their computational complexity prevents their use in applications requiring the classification of a large set of candidates. While previous works have investigated approaches to reduce model size, relatively little attention has been paid to techniques to improve batch throughput during inference. In this paper, we introduce the Cascade Transformer, a simple yet effective technique to adapt transformer-based models into a cascade of rankers. Each ranker is used to prune a subset of candidates in a batch, thus dramatically increasing throughput at inference time. Partial encodings from the transformer model are shared among rerankers, providing further speed-up. When compared to a state-of-the-art transformer model, our approach reduces computation by 37% with almost no impact on accuracy, as measured on two English Question Answering datasets.

1 Introduction

Recent research has shown that transformer-based neural networks can greatly advance the state of the art over many natural language processing tasks. Efforts such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019c), XLNet (Dai et al., 2019), and others have led to major advancements in several NLP subfields. These models are able to approximate syntactic and semantic relations between words and their compounds by pre-training on copious amounts of unlabeled data (Clark et al., 2019; Jawahar et al., 2019). Then, they can easily be applied to different tasks by just fine-tuning them on training data from the target domain/task (Liu et al., 2019a; Peters et al., 2019). The impressive effectiveness of transformer-based neural networks can be partially attributed to their large number of parameters (ranging from 110 million

for “base” models to over 8 billion (Shoeybi et al., 2019)); however, this also makes them rather expensive in terms of computation time and resources. Being aware of this problem, the research community has been developing techniques to prune unnecessary network parameters (Lan et al., 2019; Sanh et al., 2019) or optimize the transformer architecture (Zhang et al., 2018; Xiao et al., 2019).

In this paper, we propose a completely different approach for increasing the efficiency of transformer models, which is orthogonal to previous work, and thus can be applied in addition to any of the methods described above. Its main idea is that a large class of NLP problems requires choosing one correct candidate among many. For some applications, this often entails running the model over hundreds or thousands of instances. However, it is well-known that, in many cases, some candidates can be more easily excluded from the optimal solution (Land and Doig, 1960), i.e., they may require less computation. In the case of hierarchical transformer models, this property can be exploited by using a subset of model layers to score a significant portion of candidates, i.e., those that can be *more easily* excluded from search. Additionally, the hierarchical structure of transformer models intuitively enables the re-use of the computation of lower blocks to feed the upper blocks.

Following the intuition above, this work aims at studying how transformer models can be cascaded to efficiently find the max scoring elements among a large set of candidates. More specifically, the contributions of this paper are:

First, we build a sequence of rerankers $SR_N = \{R_1, R_2, \dots, R_N\}$ of different complexity, which process the candidates in a pipeline. Each reranker at position i takes the set of candidates selected by $(i - 1)$ -th reranker and provides top k_i candidates to the reranker of position $i + 1$. By requiring that $k_i < k_{i-1} \forall i = 1, \dots, N - 1$, this approach

allows us to save computation time from the more expensive rerankers by progressively reducing the number of candidates at each step. We build R_i using transformer networks of 4, 6, 8, 10, and 12 blocks from RoBERTa pre-trained models.

Second, we introduce a further optimization on SR_N to increase its efficiency based on the observation that models R_i in SR_N process their input independently. In contrast, we propose the Cascade Transformer (CT), a sequence of rerankers built on top of a single transformer model. Rerankers R_1, \dots, R_N are obtained by adding small feed-forward classification networks at different transformer block positions; therefore, the partial encodings of the transformer blocks are used as both input to reranker R_i , as well as to subsequent transformer encoding blocks. This allows us to efficiently re-use partial results consumed by R_i for rankers R_{i+1}, \dots, R_N .

To enable this approach, the parameters of all rerankers must be compatible. Thus, we trained CT in a multi-task learning fashion, alternating the optimization for different i , i.e., the layers of R_i are affected by the back-propagation of its loss as well as by the loss of R_j , with $j \leq i$.

Finally, as a test case for CT, we target Answer Sentence Selection (AS2), a well-known task in the domain of Question Answering (QA). Given a question and a set of sentence candidates (e.g., retrieved by a search engine), this task consists in selecting sentences that correctly answer the question. We tested our approach on two different datasets: (i) ASNQ, recently made available by Garg et al. (2020); and (ii) a benchmark dataset built from a set of anonymized questions asked to Amazon Alexa. Our code, ASNQ split, and models trained on ASNQ are publicly available.¹

Our experimental results show that: (i) The selection of different k_i for SR_N determines different trade-off points between efficiency and accuracy. For example, it is possible to reduce the overall computation by 10% with just 1.9% decrease in accuracy. (ii) Most importantly, the CT approach largely improves over SR, reducing the cost by 37% with almost no loss in accuracy. (iii) The rerankers trained through our cascade approach achieve equivalent or better performance than transformer models trained independently. Finally, (iv) our results suggest that CT can be used with other

NLP tasks that require candidate ranking, e.g., parsing, summarization, and many other structured prediction tasks.

2 Related Work

In this section, we first summarize related work for sequential reranking of passages and documents, then we focus on the latest methods for AS2, and finally, we discuss the latest techniques for reducing transformer complexity.

Reranking in QA and IR The approach introduced in this paper is inspired by our previous work (Matsubara et al., 2020); there, we used a fast AS2 neural model to select a subset of instances to be input to a transformer model. This reduced the computation time of the latter up to four times, preserving most accuracy.

Before our paper, the main work on sequential rankers originated from document retrieval research. For example, Wang et al. (2011) formulated and developed a cascade ranking model that improved both top-k ranked effectiveness and retrieval efficiency. Dang et al. (2013) proposed two stage approaches using a limited set of textual features and a final model trained using a larger set of query- and document-dependent features. Wang et al. (2016) focused on quickly identifying a set of good candidate documents that should be passed to the second and further cascades. Gallagher et al. (2019) presented a new general framework for learning an end-to-end cascade of rankers using back-propagation. Asadi and Lin (2013) studied effectiveness/efficiency trade-offs with three candidate generation approaches. While these methods are aligned with our approach, they target document retrieval, which is a very different setting. Further, they only used linear models or simple neural models. Agarwal et al. (2012) focused on AS2, but just applied linear models.

Answer Sentence Selection (AS2) In the last few years, several approaches have been proposed for AS2. For example, Severyn and Moschitti (2015) applied CNN to create question and answer representations, while others proposed inter-weighted alignment networks (Shen et al., 2017; Tran et al., 2018; Tay et al., 2018). The use of compare and aggregate architectures has also been extensively evaluated (Wang and Jiang, 2016; Bian et al., 2017; Yoon et al., 2019). This family of approaches uses a shallow attention mechanism

¹<https://github.com/alexa/wqa-cascade-transformers>

over the question and answer sentence embeddings. Finally, [Tayyar Madabushi et al. \(2018\)](#) exploited fine-grained question classification to further improve answer selection.

Transformer models have been fine-tuned on several tasks that are closely related to AS2. For example, they were used for machine reading ([Devlin et al., 2019](#); [Yang et al., 2019a](#); [Wang et al., 2019](#)), ad-hoc document retrieval ([Yang et al., 2019b](#); [MacAvaney et al., 2019](#)), and semantic understanding ([Liu et al., 2019b](#)) tasks to obtain significant improvement over previous neural methods. Recently, [Garg et al. \(2020\)](#) applied transformer models, obtaining an impressive boost of the state of the art for AS2 tasks.

Reducing Transformer Complexity The high computational cost of transformer models prevents their use in many real-world applications. Some proposed solutions rely on leveraging knowledge distillation in the pre-training step, e.g., ([Sanh et al., 2019](#)), or used parameter reduction techniques ([Lan et al., 2019](#)) to reduce inference cost. However, the effectiveness of these approaches varies depending on the target task they have been applied to. Others have investigated methods to reduce inference latency by modifying how self-attention operates, either during encoding ([Child et al., 2019](#); [Guo et al., 2019b](#)), or decoding ([Xiao et al., 2019](#); [Zhang et al., 2018](#)). Overall, all these solutions are mostly orthogonal to our approach, as they change the architecture of transformer cells rather than efficiently re-using intermediate results.

With respect to the model architecture, our approach is similar to probing models² ([Adi et al., 2017](#); [Liu et al., 2019a](#); [Hupkes et al., 2018](#); [Belingov et al., 2017](#)), as we train classification layers based on partial encoding on the input sequence. However, (i) our intermediate classifiers are integral part of the model, rather than being trained on frozen partial encodings, and (ii) we use these classifiers not to inspect model properties, but rather to improve inference throughput.

Our approach also shares some similarities with student-teacher (ST) approaches for self-training ([Yarowsky, 1995](#); [McClosky et al., 2006](#)). Under this setting, a model is used both as a “teacher” (which makes predictions on unlabeled data to obtain automatic labels) and as a “student” (which learns both from gold standard and automatic labels). In recent years, many variants of ST have

²Also known as auxiliary or diagnostic classifiers.

been proposed, including treating teacher predictions as soft labels ([Hinton et al., 2015](#)), masking part of the label ([Clark et al., 2018](#)), or use multiple modules for the teacher ([Zhou and Li, 2005](#); [Ruder and Plank, 2018](#)). Unlike classic ST approaches, we do not aim at improving the teacher models or creating efficient students; instead, we trained models to be used as sequential ranking components. This may be seen as a generalization of the ST approach, where the student needs to learn a simpler task than the teacher. However, our approach is significantly different from the traditional ST setting, which our preliminary investigation showed to be not very effective.

3 Preliminaries and Task Definition

We first formalize the problem of selecting the most likely element in a set as a reranking problem; then, we define sequential reranking (SR); finally, we contextualize AS2 task in such framework.

3.1 Max Element Selection

In general, a large class of NLP (and other) problems can be formulated as a max element selection task: given a query q and a set of candidates $A = \{a_1, \dots, a_n\}$, select a_j that is an optimal element for q . We can model the task as a selector function $\pi : Q \times \mathcal{P}(A) \rightarrow A$, defined as $\pi(q, A) = a_j$, where $\mathcal{P}(A)$ is the powerset of A , $j = \operatorname{argmax}_i p(q, a_i)$, and $p(q, a_i)$ is the probability of a_i to be the required element. $p(q, a_i)$ can be estimated using a neural network model. In the case of transformers, said model can be optimized using a point-wise loss, i.e., we only use the target candidate to generate the selection probability. Pair-wise or list-wise approaches can still be used ([Bian et al., 2017](#)), but (i) they would not change the findings of our study, and (ii) point-wise methods have been shown to achieve competitive performance in the case of transformer models.

3.2 Search with Sequential Reranking (SR)

Assuming that no heuristics are available to pre-select a subset of most-likely candidates, max element selection requires evaluating each sample using a relevance estimator. Instead of a single estimator, it is often more efficient to use a sequence of rerankers to progressively reduce the number of candidates.

We define a reranker as a function $R : Q \times \mathcal{P}(A) \rightarrow \mathcal{P}(A)$, which takes a subset $\Sigma \subseteq$

A , and returns a set of elements, $R(q, \Sigma) = \{a_{i1}, \dots, a_{ik}\} \subset \Sigma$ of size k , with the highest probability to be relevant to the query. That is, $p(q, a) > p(q, b) \quad \forall a \in \Sigma, \quad \forall b \in A - \Sigma$.

Given a sequence of rerankers sorted in terms of computational efficiency, (R_1, R_2, \dots, R_N) , we assume that the ranking accuracy, \mathcal{A} (e.g., in terms of MAP and MRR), increases in reverse order of the efficiency, i.e., $\mathcal{A}(R_j) > \mathcal{A}(R_i)$ iff $j > i$. Then, we define a Sequential Reranker of order N as the composition of N rerankers: $SR_N(A) = R_N \circ R_{N-1} \circ \dots \circ R_1(A)$, where R_N can also be the element selector $\pi(q, \cdot)$. Each R_i is associated with a different $k_i = |R_i(\cdot)|$, i.e., the number of elements the reranker returns. Depending on the values of k_i , SR models with different trade-offs between accuracy and efficiency can be obtained.³

3.3 AS2 Definition

The definition of AS2 directly follows from the definition of element selection of Section 3.1, where the query is a natural language question and the elements are answer sentence candidates retrieved with any approach, e.g., using a search engine.

4 SR with transformers

In this section, we explain how to exploit the hierarchical architecture of a traditional transformer model to build an SR model. First, we briefly recap how traditional transformer models (we refer to them as “*monolithic*”) are used for sequence classification, and how to derive a set of sequential rerankers from a pre-trained transformer model (Section 4.1). Then, we introduce our Cascade Transformer (CT) model, a SR model that efficiently uses partial encodings of its input to build a set of sequential rerankers R_i (Section 4.3). Finally, we explain how such model is trained and used for inference in sections 4.3.1 and 4.3.2, respectively.

4.1 Monolithic Transformer Models

We first briefly describe the use of transformer models for sequence classification. We call them *monolithic* as, for all input samples, the computation flows from the first until the last of their layers.

Let $\mathcal{T} = \{E; L_1, L_2, \dots, L_n\}$ be a standard stacked transformer model (Vaswani et al., 2017), where E is the embedding layer, and L_i are the

³The design of an end-to-end algorithm to learn the optimal parameter set for a given target trade-off is left as future work.

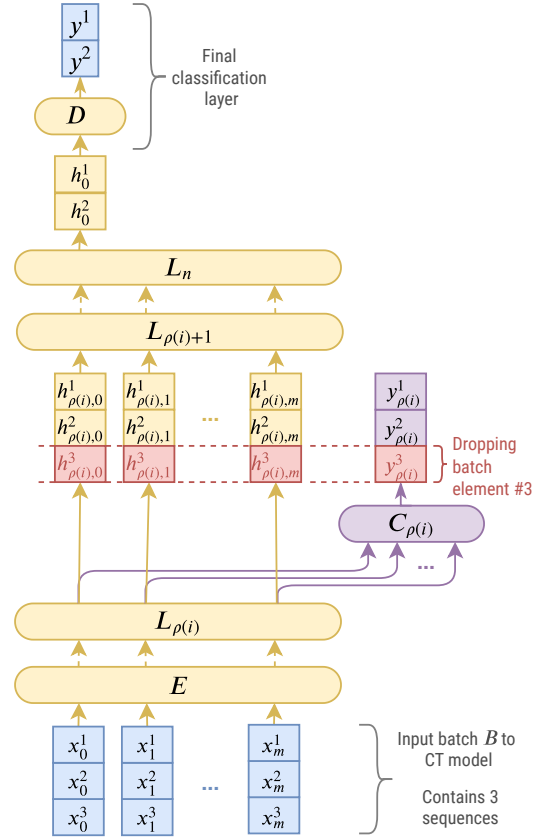


Figure 1: A visual representation of the Cascade Transformer (CT) model proposed in this paper. Components in **yellow** represent layers of a traditional transformer model, while elements in **purple** are unique to CT; input and outputs of the model are shown in **blue**. In this example, **drop rate** $\alpha = 0.4$ causes sample X^3 to be removed by partial classifier $C_{\rho(i)}$.

transformer layers⁴ generating contextualized representations for an input sequence; n is typically referred to as the depth of the encoder, i.e., the number of layers. Typical values for n range from 12 to 24, although more recent works have experimented with up to 72 layers (Shoeybi et al., 2019). \mathcal{T} can be pre-trained on large amounts of unlabeled text using a masked (Devlin et al., 2019; Liu et al., 2019c) or autoregressive (Yang et al., 2019c; Radford et al., 2019) language modeling objective.

Pre-trained language models are fine-tuned for the target tasks using additional layers and data, e.g., a fully connected layer is typically stacked on top of \mathcal{T} to obtain a sentence classifier. Formally, given a sequence of input symbols⁵, $X = \{x_0, x_1, \dots, x_m\}$, an encoding $H =$

⁴That is, an entire transformer block, constituted by layers for multi-head attention, normalization, feed forward processing and positional embeddings.

⁵For ranking tasks, the sequence of input symbols is typically a concatenation of the query q and a candidate a_j . In

$\{h_0, h_1, \dots, h_m\}$ is first obtained by recursively applying H_i to the input:

$H_0 = E(X), H_i = L_i(H_{i-1}) \quad \forall i = 1, \dots, n,$
 where $H = H_n$. Then, the first symbol of the input sequence⁶ is fed into a sequence of dense feed-forward layers D to obtain a final output score, i.e., $y = D(h_0)$. D is fine-tuned together with the entire model on a task-specific dataset (a set of question and candidate answer pairs, in our case).

4.2 Transformer-based Sequential Reranker (SR) Models

Monolithic transformers can be easily modified or combined to build a sequence of rerankers as described in Section 3.2. In our case, we adapt an existing monolithic \mathcal{T} to obtain a sequence of N rerankers R_i . Each R_i consists of encoders from \mathcal{T} up to layer $\rho(i)$, followed by a classification layer D_i , i.e., $R_i = \{E; L_1, \dots, L_{\rho(i)}, D_i\}$. For a sequence of input symbols X , all rerankers in the sequence are designed to predict $p(q, a)$, which we indicate as $R_i(X) = y_{\rho(i)}$. All rerankers in SR_N are trained independently on the target data.

In our experiments, we obtained the best performance by setting $N = 5$ and using the following formula to determine the architecture of each reranker R_i :

$$\rho(i) = 4 + 2 \cdot (i - 1) \quad \forall i = \{1, \dots, 5\}$$

In other words, we assemble sequential reranker SR_5 using five rerankers built with transformer models of 4, 6, 8, 10 and 12 layers, respectively. This choice is due to the fact that our experimental results seem to indicate that the information in layers 1 to 3 is not structured enough to achieve satisfactory classification performance for our task. This observation is in line with recent works on the effectiveness of partial encoders for semantic tasks similar to AS2 (Peters et al., 2019).

4.3 Cascade Transformer (CT) Models

During inference, monolithic transformer models evaluate a sequence X through the entire computation graph to obtain the classification scores Y .

order for the model to distinguish between the two, a special token such as “[SEP]” or “</s>” is used. Some models also use a second embedding layer to represent which sequence each symbol comes from.

⁶Before being processed by a transformer model, sequences are typically prefixed by a start symbol, such as “[CLS]” or “<s>”. This allows transformer models to accumulate knowledge about the entire sequence at this position without compromising token-specific representations (Devlin et al., 2019).

This means that when using SR_N , examples are processed multiple times by similar layers for different R_i , e.g., for $i = 1$, all R_i compute the same operations of the first $\rho(i)$ transformer layers, for $i = 2$, $N - 1$ rerankers compute the same $\rho(i) - \rho(i + 1)$, layers and so on. A more computationally-efficient approach is to share all the common transformer blocks between the different rerankers in SR_N .

We speed up this computation by using one transformer encoder to implement all required R_i . This can be easily obtained by adding a classification layer $C_{\rho(i)}$ after each $\rho(i)$ layers (see Figure 1). Consequently, given a sample X , the classifiers $C_{\rho(i)}$ produces scores $y_{\rho(i)}$ only using a partial encoding. To build a CT model, we use each $C_{\rho(i)}$ to build rerankers R_i , and select the top k_i candidates to score with the subsequent rerankers R_{i+1} . We use the same setting choices of N and $\rho(i)$ described in Section 4.2.

Finally, we observed the best performance when all encodings in $H_{\rho(i)}$ are used as input to partial classifier $C_{\rho(i)}$, rather than just the partial encoding of the classification token $h_{\rho(i),0}$. Therefore, we use their average to obtain score $y_{\rho(i)} = C_{\rho(i)}(\frac{1}{m} \sum_{l=1, \dots, m} h_{\rho(i),l})$. In line with Kovaleva et al. (2019), we hypothesize that, at lower encoding layers, long dependencies might not be properly accounted in $h_{\rho(i),0}$. However, in our experiments, we found no benefits in further parametrizing this operation, e.g., by either using more complex networks or weighting the average operation.

4.3.1 Training CT

The training of the proposed model is conducted in a multi-task fashion. For every mini-batch, we randomly sample one of the rankers R_i (including the final output reranker), calculate its loss against the target labels, and back-propagate its loss throughout the entire model down to the embedding layers. We experimented with several more complex sampling strategies, including a round-robin selection process and a parametrized bias towards early rankers for the first few epochs, but we ultimately found that uniform sampling works best. We also empirically determined that, for all classifiers $C_{\rho(i)}$, backpropagating the loss to the input embeddings, as opposed to stopping it at layer $\rho(i - 1)$, is crucial to ensure convergence. A possible explanation could be: enabling each classifier to influence the input representation during backpropagation ensures that later rerankers are more robust against

variance in partial encodings, induced by early classifiers. We experimentally found that if the gradient does not flow throughout the different blocks, the development set performance for later classifiers drops when early classifiers start converging.

4.3.2 Inference

Recall that we are interested in speeding up inference for classification tasks such as answer selection, where hundreds of candidates are associated with each question. Therefore, we can assume without loss of generality that each batch of samples $B = \{X^1, \dots, X^b\}$ contains candidate answers for the same question. We use our partial classifiers to throw away a fraction α of candidates, to increase throughput. That is, we discard $k_i = \lfloor \alpha \cdot k_{i-1} \rfloor$ candidates, where $\lfloor \cdot \rfloor$ rounds $\alpha \cdot k_{i-1}$ down to the closest integer.

For instance, let $\alpha = 0.3$, batch size $b = 128$; further, recall that, in our experiments, a CT consists of 5 cascade rerankers. Then, after layer 4, the size of the batch gets reduced to 90 ($\lfloor 0.3 \cdot 128 \rfloor = 38$ candidates are discarded by the first classifier). After the second classifier (layer 6), $\lfloor 0.3 \cdot 90 \rfloor = 27$ examples are further removed, for an effective batch size of 63. By layer 12, only 31 samples are left, i.e., the instance number scored by the final classifier is reduced by more than 4 times.

Our approach has the effect of improving the throughput of a transformer model by reducing the average batch size during inference: the throughput of any neural model is capped by the maximum number of examples it can process in parallel (i.e., the size of each batch), and said number is usually ceiled by the amount of memory available to the model (e.g., RAM on GPU). The monolithic models have a constant batch size at inference; however, because the batch size for a cascade model varies while processing a batch, we can size our network with respect to its average batch size, thus increasing the number of samples we initially have in a batch. In the example above, suppose that the hardware requirement dictates a maximum batch size of 84 for the monolithic model. As the average batch size for the cascading model is $(4 \cdot 128 + 2 \cdot 90 + 2 \cdot 63 + 2 \cdot 44 + 2 \cdot 28) / 12 = 80.2 < 84$, we can process a batch of 128 instances without violating memory constraints, increasing throughput by 52%.

We remark that using a fixed α is crucial to obtain the performance gains we described: if we were to employ a score-based thresholding ap-

| | | ASNQ | GPD | TRECQA | WikiQA |
|-------|-----------|--------|-------|--------|--------|
| TRAIN | Questions | 57,242 | 1,000 | 1,227 | 873 |
| | Avg cand. | 413.3 | 99.8 | 39.2 | 9.9 |
| | Avg corr. | 1.2 | 4.4 | 4.8 | 1.2 |
| DEV | Questions | 1,336 | 340 | 65 | 126 |
| | Avg cand. | 403.6 | 99.7 | 15.9 | 9.0 |
| | Avg corr. | 3.2 | 2.85 | 2.9 | 1.1 |
| TEST | Questions | 1,336 | 440 | 68 | 243 |
| | Avg cand. | 400.5 | 101.1 | 20.0 | 9.7 |
| | Avg corr. | 3.2 | 8.13 | 3.4 | 1.2 |

Table 1: Datasets statistics: ASNQ and GPD have more sentence candidates than TRECQA and WikiQA.

proach (that is, discard all candidates with score below a given threshold), we could not determine the size of batches throughout the cascade, thus making it impossible to efficiently scale our system. On the other hand, we note that nothing in our implementations prevents potentially correct candidates from being dropped when using CT. However, as we will show in Section 5, an opportune choice of a threshold and good accuracy of early classifiers ensure high probability of having at least one positive example in the candidate set for the last classifier of the cascade.

5 Experiments

We present three sets of experiments designed to evaluate CT. In the first (Section 5.3), we show that our proposed approach without any selection produces comparable or superior results with respect to the state of the art of AS2, thanks to its stability properties; in the second (Section 5.4), we compare our Cascade Transformer with a vanilla transformer, as well as a sequence of transformer models trained independently; finally, in the third (Section 5.5), we explore the tuning of the drop ratio, α .

5.1 Datasets

TRECQA & WikiQA Traditional benchmarks used for AS2, such as TRECQA (Wang et al., 2007) and WikiQA (Yang et al., 2015), typically contain a limited number of candidates for each question. Therefore, while they are very useful to compare accuracy of AS2 systems with the state of the art, they do not enable testing large scale passage reranking, i.e., inference on hundreds or thousand of answer candidates. Therefore, we evaluated our approach (Sec. 4.3) on two datasets: ASNQ, which is publicly available, and our GPD dataset. We still leverage TRECQA and WikiQA to show that that our

cascade system has comparable performance to state-of-the-art transformer models when no filtering is applied.

ASNQ The Answer Sentence Natural Questions dataset (Garg et al., 2020) is a large collection (23M samples) of question-answer pairs, which is two orders of magnitude larger than most public AS2 datasets. It was obtained by extracting sentence candidates from the Google Natural Question (NQ) benchmark (Kwiatkowski et al., 2019). Samples in NQ consists of tuples $\langle \text{question}, \text{answer}_{\text{long}}, \text{answer}_{\text{short}}, \text{label} \rangle$, where $\text{answer}_{\text{long}}$ contains multiple sentences, $\text{answer}_{\text{short}}$ is fragment of a sentence, and label is a binary value indicating whether $\text{answer}_{\text{long}}$ is correct. The positive samples were obtained by extracting sentences from $\text{answer}_{\text{long}}$ that contain $\text{answer}_{\text{short}}$; all other sentences are labeled as negative. The original release of ASNQ⁷ only contains train and development splits; we further split the dev. set to both have dev. and test sets.

GPD The General Purpose Dataset is part of our efforts to study large scale web QA and evaluate performance of AS2 systems. We built GPD using a search engine to retrieve up to 100 candidate documents for a set of given questions. Then, we extracted all candidate sentences from such documents, and rank them using a vanilla transformer model, such as the one described in Sec. 4.1. Finally, the top 100 ranked sentences were manually annotated as correct or incorrect answers.

We measure the accuracy of our approach on ASNQ and GPD using four metrics: Mean Average Precision (MAP), Mean Reciprocal Rank (MRR), Precision at 1 of ranked candidates (P@1), and Normalized Discounted Cumulative Gain at 10 of retrieved candidates (nDCG@10). While the first two metrics capture the overall system performance, the latter two are better suited to evaluate systems with many candidates, as they focus more on Precision. For WikiQA and TRECQA, we use MAP and MRR.

5.2 Models and Training

Our models are fine-tuned starting from a pre-trained RoBERTa encoder (Liu et al., 2019c). We chose this transformer model over others due to its strong performance on answer selection tasks (Garg et al., 2020). Specifically, we use the BASE

⁷https://github.com/alexa/wqa_tanda

| Model | WikiQA | | TRECQA | |
|---|-------------|-------------|-------------|-------------|
| | MAP | MRR | MAP | MRR |
| CA1 (Wang and Jiang, 2016) | 74.3 | 75.4 | – | – |
| CA2 (Yoon et al., 2019) | 83.4 | 84.8 | 87.5 | 94.0 |
| TANDA _{BASE} (Garg et al., 2020) | 88.9 | 90.1 | 91.4 | 95.2 |
| 4 layers TANDA | 80.5 | 80.9 | 77.2 | 83.1 |
| 6 layers TANDA | 82.1 | 82.9 | 78.5 | 88.4 |
| 8 layers TANDA | 85.7 | 86.7 | 88.2 | 94.7 |
| 10 layers TANDA | 89.0 | 90.0 | 90.5 | 95.9 |
| Our TANDA _{BASE} | 89.1 | 90.1 | 91.6 | 96.0 |
| CT (4 layers, $\alpha = 0.0$) | 60.1 | 60.2 | 67.9 | 74.7 |
| CT (6 layers, $\alpha = 0.0$) | 79.8 | 80.3 | 89.7 | 95.0 |
| CT (8 layers, $\alpha = 0.0$) | 84.8 | 85.4 | 92.3 | 95.3 |
| CT (10 layers, $\alpha = 0.0$) | 89.7 | 89.8 | 92.3 | 95.6 |
| CT (12 layers, $\alpha = 0.0$) | 89.9 | 91.0 | 92.4 | 96.7 |

Table 2: Comparison on two AS2 academic datasets. With the exception of a 4-layer transformer, both the partial and final classifiers from CT achieve comparable or better performance than state of the art models.

variant (768-dimensional embeddings, 12 layers, 12 heads, and 3072 hidden units), as it is more appropriate for efficient classification. When applicable⁸, we fine-tune our models using the two-step “transfer and adapt” (TANDA) technique introduced by Garg et al. (2020).

As mentioned in Section 4.3, we optimize our model in a multi-task setting; that is, for each mini-batch, we randomly sample one of the output layers of the CT classifiers to backpropagate its loss to all layers below.

While we evaluated different sampling techniques, we found that a simple uniform distribution is sufficient and allows the model to converge quickly.

Our models are optimized using Adam (Kingma and Ba, 2014) using triangular learning rate (Smith, 2017) with a 4,000 updates ramp-up⁹, and a peak learning rate $l_r = 1e^{-6}$. Batch size was set to up to 2,000 tokens per mini-batch for CT models. For the partial and final classifiers, we use 3-layers feed-forward modules with with 768 hidden units and *tanh* activation function. Like the original BERT implementation, we use dropout value of 0.1 on all dense and attention layers. We implemented our system using MxNet 1.5 (Chen et al., 2015) and GluonNLP 0.8.1 (Guo et al., 2019a) on a machine with 8 NVIDIA Tesla V100 GPUs, each with 16GB of memory.

⁸When fine-tuning on GPD, TRECQA, and WikiQA, we perform a “transfer” step on ASNQ before “adapting” to our target dataset; for ASNQ, we directly fine-tune on it.

⁹On ASNQ, it is roughly equivalent to 950k samples or about 4% of the training set.

| Method | Model | α | ASNQ | | | | GPD | | | | Cost reduction per batch |
|-----------------------------|--|----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------------------|
| | | | MAP | nDCG@10 | P@1 | MRR | MAP | nDCG@10 | P@1 | MRR | |
| Monolithic transformer (MT) | 4 layers TANDA | – | 31.5 | 30.8 | 25.9 | 30.8 | 38.9 | 50.1 | 40.8 | 54.0 | –67% |
| | 6 layers TANDA | – | 60.2 | 58.7 | 47.2 | 59.2 | 51.4 | 64.1 | 56.1 | 67.6 | –50% |
| | 8 layers TANDA | – | 63.9 | 62.2 | 49.2 | 62.4 | 56.3 | 68.7 | 61.2 | 70.4 | –33% |
| | 10 layers TANDA | – | 65.3 | 64.5 | 52.0 | 64.1 | 57.2 | 71.3 | 64.9 | 72.7 | –20% |
| | TANDA _{BASE} | – | 65.5 | 65.1 | 52.1 | 64.7 | 58.0 | 72.2 | 67.5 | 76.8 | baseline |
| Sequential Ranker (SR) | MT models, 4 to 12 layers, in sequence | 0.3 | 65.4 | 65.1 | 52.1 | 64.8 | 55.8 | 70.2 | 66.2 | 74.3 | +53% |
| | | 0.4 | 64.9 | 64.2 | 51.6 | 64.2 | 53.8 | 69.6 | 65.6 | 73.0 | +18% |
| | | 0.5 | 64.6 | 63.4 | 50.8 | 63.5 | 52.2 | 68.4 | 63.0 | 72.3 | –10% |
| Cascade transformer (CT) | 4 layers CT | 0.0 | 22.0 | 19.3 | 10.2 | 18.3 | 32.7 | 38.9 | 35.2 | 42.6 | –67% |
| | 6 layers CT | 0.0 | 49.1 | 47.2 | 32.7 | 47.7 | 44.8 | 56.0 | 47.3 | 58.5 | –50% |
| | 8 layers CT | 0.0 | 62.8 | 61.5 | 48.7 | 61.9 | 53.8 | 71.7 | 61.2 | 69.1 | –33% |
| | 10 layers CT | 0.0 | 65.6 | 65.1 | 53.0 | 65.2 | 55.8 | 72.0 | 63.1 | 72.1 | –20% |
| | | 0.0 | 66.3 | 66.1 | 53.2 | 65.4 | 57.8 | 71.9 | 67.5 | 76.9 | –0% |
| | Full CT | 0.3 | 65.3 | 65.3 | 52.9 | 65.3 | 55.7 | 69.8 | 66.2 | 75.1 | –37% |
| | (12 layers) | 0.4 | 64.8 | 65.0 | 52.5 | 64.8 | 52.8 | 68.6 | 65.6 | 74.3 | –45% |
| | 0.5 | 64.1 | 65.0 | 52.4 | 64.5 | 50.2 | 66.1 | 62.4 | 72.9 | –51% | |

Table 3: Comparison of Cascade Transformers with other models on the ASNQ and GPD datasets. “Monolithic transformer” refers to a single transformer model trained independently; “sequential ranker” (SR) is a sequence of monolithic transformer models of size 4, 6, . . . , 12 trained independently; and “Cascade Transformer” (CT) is the approach we propose. This can train models that equal or outperform the state of the art when no drop is applied (i.e., $\alpha = 0.0$); with drop, they obtain the same performance with 37% to 51% fewer operations.

5.3 Stability Results of Cascade Training

In order to better assess how our training strategy for CT models compare with a monolithic transformer, we evaluated the performance of our system on two well known AS2 datasets, WikiQA and TRECQA. The results of these experiments are presented in Table 2. Note how, in this case, we are not applying any drop to our cascade classifier, as it is not necessary on this dataset: all sentences fit comfortably in one mini batch (see dataset statistics in Table 1), so we would not observe any advantage in pruning candidates. Instead, we focus on evaluating how our training strategy affects performance of partial and final classifiers of a CT model.

Our experiment shows that classifiers in a CT model achieve competitive performance with respect to the state of the art: our 12-layer transformer model trained in cascade outperforms TANDA_{BASE} by 0.8 and 0.9 absolute points in MAP (0.9 and 0.7 in MRR). 10, 8, and 6 layer models are equally comparable, differing at most by 2.3 absolute MAP points on WikiQA, and outscoring TANDA by up to 11.2 absolute MAP points on TRECQA. However, we observed meaningful differences between the performance of the 4-layers cascade model and its monolithic counterparts. We hypothesize that this is due to the fact that lower layers are not typically well suited for classification when used as part of a larger model (Peters et al., 2019); this observation is reinforced by the fact that the 4 layers TANDA model shown in Table 2

takes four times the number of the iterations of any other model to converge to a local optimum.

Overall, these experiments show that our training strategy is not only effective for CT models, but can also produce smaller transformer models with good accuracy without separate fine-tuning.

5.4 Results on Effectiveness of Cascading

The main results for our CT approach are presented in Table 3: we compared it with (i) a state-of-the-art monolithic transformer (TANDA_{BASE}), (ii) smaller, monolithic transformer models with 4-10 layers, and (iii) a sequential ranker (SR) consisting of 5 monolithic transformer models with 4, 6, 8, 10 and 12 layers trained independently. For CT, we report performance of each classifier individually (layers 4 up to 12, which is equivalent to a full transformer model). We test SR and CT with drop ratio 30%, 40%, 50%. Finally, for each model, we report the relative cost per batch compared to a *base* transformer model with 12 layers.

Overall, we observed that our cascade models are competitive with monolithic transformers on both ASNQ and GPD datasets. In particular, when no selection is applied ($\alpha = 0.0$), a 12 layer cascade model performs equal or better to TANDA_{BASE}: on ASNQ, we improve P@1 by 2.1% (53.2 vs 52.1), and MAP by 1.2% (66.3 vs 65.5); on GPD, we achieve the same P@1 (67.5), and a slightly lower MAP (57.8 vs 58.0). This indicates that, despite the multitasking setup, our method is competitive with the state of the art.

A drop rate $\alpha > 0.0$ produces a small degradation in accuracy, at most, while significantly reducing the number of operations per batch (-37%). In particular, when $\alpha = 0.3$, we achieve less than 2% drop in P@1 on GPD, when compared to TANDA_{BASE}; on ANSQ, we slightly improve over it (52.9 vs 52.1). We observe a more pronounced drop in performance for MAP, this is to be expected, as intermediate classification layers are designed to drop a significant number of candidates.

For larger values of α , such as 0.5, we note that we achieve significantly better performance than monolithic transformer of similar computational cost. For example, CT achieves an 11.2% improvement in P@1 over a 6-layers TANDA model (62.4 vs 56.1) on GPD; a similar improvement is obtained on ANSQ (+11.0%, 52.4 vs 47.2).

Finally, our model is also competitive with respect to a sequential transformer with equivalent drop rates, while being between 1.9 to 2.4 times more efficient. This is because an SR model made of independent TANDA models cannot re-use encodings generated by smaller models as CT does.

5.5 Results on Tuning of Drop Ratio α

Finally, we examined how different values for drop ratio α affect the performance of CT models. In particular, we performed an exhaustive grid-search on a CT model trained on the GPD dataset for drop ratio values $\{\alpha_{p_1}, \alpha_{p_2}, \alpha_{p_3}, \alpha_{p_4}\}$, with $\alpha_{p_k} \in \{0.1, 0.2, \dots, 0.6\}$. The performance is reported in Figure 2 with respect to the relative computational cost per batch of a configuration when compared with a TANDA_{BASE} model.

Overall, we found that CT models are robust with respect to the choice of $\{\alpha_{p_k}\}_{k=1}^4$. We observe moderate degradation for higher drop ratio values (e.g., P@1 varies from 85.6 to 80.0). Further, as expected, performance increases for models with higher computational cost per batch, although they taper off for CT models with relative cost $\geq 70\%$. On the other hand, the grid search results do not seem to suggest an effective strategy to pick optimal values for $\{\alpha_{p_k}\}_{k=1}^4$, and, in our experiments, we ended up choosing the same values for all drop rates. In the future, we would be like to learn such values while training the cascade model itself.

6 Conclusions and Future Work

This work introduces CT, a variant of the traditional transformer model designed to improve inference

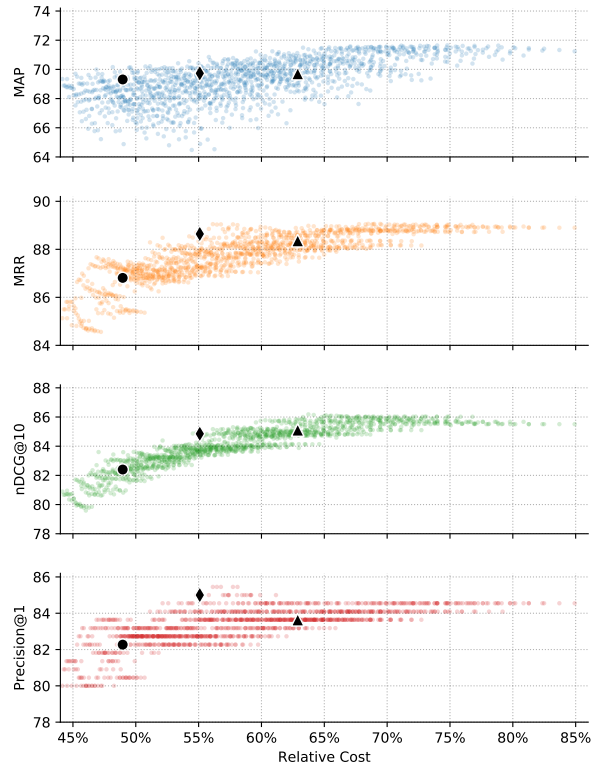


Figure 2: Grid search plot on the GPD validation set. Each point corresponds to a configuration of drop ratios $\{\alpha_{p_1}, \dots, \alpha_{p_4}\}$ with $\alpha_{p_k} \in \{0.1, 0.2, \dots, 0.6\}$; values on the x -axis represent the relative computational cost per batch of a configuration compared to TANDA_{BASE}. The three runs reported in Table 3 correspond to ▲ ($\alpha = 0.3$), ◆ ($\alpha = 0.4$), and ● ($\alpha = 0.5$).

throughput. Compared to a traditional monolithic stacked transformer model, our approach leverages classifiers placed at different encoding stages to prune candidates in a batch and improve model throughput. Our experiments show that a CT model not only achieves comparable performance to a traditional transformer model while reducing computational cost per batch by over 37%, but also that our training strategy is stable and jointly produces smaller transformer models that are suitable for classification when higher throughput and lower latency goals must be met. In future work, we plan to explore techniques to automatically learn where to place intermediate classifiers, and what drop ratio to use for each one of them.

References

Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *ICLR*.

- Arvind Agarwal, Hema Raghavan, Karthik Subbian, Prem Melville, Richard D. Lawrence, David C. Gondek, and James Fan. 2012. [Learning to rank for robust question answering](#). In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management, CIKM '12*, pages 833–842, New York, NY, USA. ACM.
- Nima Asadi and Jimmy Lin. 2013. [Effectiveness/efficiency tradeoffs for candidate generation in multi-stage retrieval architectures](#). In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, pages 997–1000, New York, NY, USA. ACM.
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017. [Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–10, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Weijie Bian, Si Li, Zhao Yang, Guang Chen, and Zhiqing Lin. 2017. [A compare-aggregate model with dynamic-clip attention for answer selection](#). In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 1987–1990, New York, NY, USA. ACM.
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. [Generating long sequences with sparse transformers](#). *arXiv preprint arXiv:1904.10509*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Christopher D Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). *CoRR*, abs/1901.02860.
- Van Dang, Michael Bendersky, and W. Bruce Croft. 2013. [Two-stage learning to rank for information retrieval](#). In *Proceedings of the 35th European Conference on Advances in Information Retrieval, ECIR'13*, pages 423–434, Berlin, Heidelberg. Springer-Verlag.
- 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.
- Luke Gallagher, Ruey-Cheng Chen, Roi Blanco, and J. Shane Culpepper. 2019. [Joint optimization of cascade ranking models](#). In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM '19*, pages 15–23, New York, NY, USA. ACM.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2020. [Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection](#). In *AAAI*.
- Jian Guo, He He, Tong He, Leonard Lausen, Mu Li, Haibin Lin, Xingjian Shi, Chenguang Wang, Junyuan Xie, Sheng Zha, et al. 2019a. [GluonCV and GluonNLP: Deep learning in computer vision and natural language processing](#). *arXiv preprint arXiv:1907.04433*.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Yunfan Shao, Xiangyang Xue, and Zheng Zhang. 2019b. [Star-transformer](#). 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 1315–1325.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*.
- Dieuwke Hupkes, Sara Veldhoen, and Willem Zuidema. 2018. [Visualisation and diagnostic classifiers’ reveal how recurrent and recursive neural networks process hierarchical structure](#). *Journal of Artificial Intelligence Research*, 61:907–926.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.

- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of bert. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4356–4365.
- 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. 2019. [Natural questions: a benchmark for question answering research](#). *Transactions of the Association of Computational Linguistics*.
- 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](#).
- A.H. Land and A.G. Doig. 1960. An automatic method for solving discrete programming problems. *Econometrica*, 28:497–520.
- Nelson F Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A Smith. 2019a. Linguistic knowledge and transferability of contextual representations. 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 1073–1094.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. Multi-task deep neural networks for natural language understanding. In *ACL*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. 2019. [CEDR: Contextualized embeddings for document ranking](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Yoshitomo Matsubara, Thuy Vu, and Alessandro Moschitti. 2020. Reranking for efficient transformer-based answer selection. In *To appear in Proceedings of the 43th international ACM SIGIR conference on research and development in information retrieval*. ACM.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Reranking and self-training for parser adaptation. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 337–344. Association for Computational Linguistics.
- Matthew Peters, Sebastian Ruder, and Noah A Smith. 2019. To tune or not to tune? adapting pretrained representations to diverse tasks. *arXiv preprint arXiv:1903.05987*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- Sebastian Ruder and Barbara Plank. 2018. Strong baselines for neural semi-supervised learning under domain shift. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1044–1054.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#).
- Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 373–382. ACM.
- Gehui Shen, Yunlun Yang, and Zhi-Hong Deng. 2017. [Inter-weighted alignment network for sentence pair modeling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1179–1189, Copenhagen, Denmark. Association for Computational Linguistics.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-LM: Training multi-billion parameter language models using gpu model parallelism. *arXiv preprint arXiv:1909.08053*.
- Leslie N Smith. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472. IEEE.
- Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. 2018. [Multi-cast attention networks for retrieval-based question answering and response prediction](#). *CoRR*, abs/1806.00778.
- Harish Tayyar Madabushi, Mark Lee, and John Barnard. 2018. [Integrating question classification and deep learning for improved answer selection](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3283–3294, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Quan Hung Tran, Tuan Lai, Gholamreza Haffari, Ingrid Zukerman, Trung Bui, and Hung Bui. 2018. [The context-dependent additive recurrent neural net](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1274–1283, New Orleans, Louisiana. Association for Computational Linguistics.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Lidan Wang, Jimmy Lin, and Donald Metzler. 2011. A cascade ranking model for efficient ranked retrieval. In *sigir*, pages 105–114.
- Mengqiu Wang, Noah A. Smith, and Teruko Mitamura. 2007. [What is the Jeopardy model? a quasi-synchronous grammar for QA](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 22–32, Prague, Czech Republic. Association for Computational Linguistics.
- Qi Wang, Constantinos Dimopoulos, and Torsten Suel. 2016. [Fast first-phase candidate generation for cascading rankers](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, pages 295–304, New York, NY, USA. ACM.
- Ran Wang, Haibo Su, Chunye Wang, Kailin Ji, and Jupeng Ding. 2019. To tune or not to tune? how about the best of both worlds? *ArXiv*, abs/1907.05338.
- Shuohang Wang and Jing Jiang. 2016. A compare-aggregate model for matching text sequences. *arXiv preprint arXiv:1611.01747*.
- Tong Xiao, Yinqiao Li, Jingbo Zhu, Zhengtao Yu, and Tongran Liu. 2019. Sharing attention weights for fast transformer. In *International Joint Conferences on Artificial Intelligence (IJCAI)*.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019a. [End-to-end open-domain question answering with BERTserini](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wei Yang, Haotian Zhang, and Jimmy Lin. 2019b. [Simple applications of BERT for ad hoc document retrieval](#). *CoRR*, abs/1903.10972.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. [WikiQA: A challenge dataset for open-domain question answering](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018, Lisbon, Portugal. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019c. [Xlnet: Generalized autoregressive pretraining for language understanding](#). *arXiv preprint arXiv:1906.08237*.
- David Yarowsky. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pages 189–196.
- Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, and Kyomin Jung. 2019. [A compare-aggregate model with latent clustering for answer selection](#). *CoRR*, abs/1905.12897.
- Biao Zhang, Deyi Xiong, and Jinsong Su. 2018. Accelerating neural transformer via an average attention network. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1789–1798.
- Zhi-Hua Zhou and Ming Li. 2005. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on Knowledge and Data Engineering*, 17(11):1529–1541.