

Question-Context Alignment and Answer-Context Dependencies for Effective Answer Sentence Selection

Minh Van Nguyen^{1*}, Kishan KC², Toan Nguyen², Thien Nguyen¹, Ankit Chadha², and Thuy Vu^{2**}

¹Department of Computer Science, University of Oregon, OR, USA

²Amazon Alexa AI, CA, USA

{minhmv, thien}@cs.uoregon.edu

{ckshan, amztoan, ankitrc, thuyvu}@amazon.com

Abstract

Answer sentence selection (AS2) in open-domain question answering finds answer for a question by ranking candidate sentences extracted from web documents. Recent work exploits answer context, i.e., sentences around a candidate, by incorporating them as additional input string to the Transformer models to improve the correctness scoring. In this paper, we propose to improve the candidate scoring by explicitly incorporating the dependencies between question-context and answer-context into the final representation of a candidate. Specifically, we use Optimal Transport to compute the question-based dependencies among sentences in the passage where the answer is extracted from. We then represent these dependencies as edges in a graph and use Graph Convolutional Network to derive the representation of a candidate, a node in the graph. Our proposed model achieves significant improvements on popular AS2 benchmarks, i.e., WikiQA and WDRASS, obtaining new state-of-the-art on all benchmarks.

Index Terms: question answering, human-computer interaction, large language model

1. Introduction

Voice-based virtual assistants powered by open-domain question answering (ODQA) [1, 2] have gained significant commercial market in recent years, e.g., Google Assistant, Siri, or Alexa, thanks to the progress in question answering using large pre-trained language models (LLMs) such as BERT [3], RoBERTa [4], and GPT-3 [5]. Recent works address the task via generative models [6, 7]. However, a well-known issue that has been shown to occur with the generative models is hallucination [8, 9, 10], where the models generate statements that are plausible looking but *factually* incorrect. Additionally, if the answers are composed by a pretrained LLM without external knowledge, the information contained in the answers might be outdated and no longer valid, e.g., the answer for the question “Which country is the reigning World Cup champion?” will change through time. To avoid such problems, this work follows a typical ODQA pipeline involving two main stages: web retrieval and answer sentence selection (AS2); the latter selects the most relevant *answer sentence* from the retrieved documents to return to the user. This stage is typically implemented as a point-wise model that scores each sentence individually either without any additional information [2, 11], jointly with other top candidates [12], or with other contextual information [13, 14]. We focus on the latter in this paper.

* This work was completed while the author was an intern at Amazon Alexa AI.

** Corresponding Author. Email: thuyvu@amazon.com

<i>q</i> :	What award did Lionel Messi win after the World Cup ?
<i>a</i> :	Lionel Messi has been crowned The Best FIFA Men’s Player for the second time.
<i>p_a</i> :	[prev] The Best FIFA Football Awards annually honour the most outstanding members of the world’s most popular sport. [cand] Lionel Messi has been crowned The Best FIFA Men’s Player for the second time. [next] The award , which is voted for by national team coaches and captains, journalists and also fans, recognises a year in which the former Barcelona star crowned his glorious career by leading Argentina to victory at the World Cup .

Table 1: The example shows contextual sentences, i.e., *prev* and *next* in a passage *p* for an answer *a*. Their relation with the question *q* is critical to select the **correct answer a**.

Previous work on AS2 exploits contextual information for better performance, however, is limited to the concatenation approach. In particular, Lauriola et al. [13] concatenated contextual sentences, i.e., previous (*prev*) and next sentences (*next*), with an answer candidate (*cand*) for a given question (*q*) as an input sequence to a Transformer architecture, i.e., using the structure “[CLS] *q* [SEP] *prev* [SEP] *cand* [SEP] *next* [EOS]” instead of “[CLS] *q* [SEP] *cand* [EOS]”. Han et al. [14], similarly, uses a concatenation approach for contextual sentences and document titles. While both models outperform the vanilla baseline, i.e., using only *cand*, we hypothesize that the improvement results from using a longer sequence for higher predictive power.

From a modeling viewpoint, the concatenation approach is effective largely due to the self-attention mechanism in the Transformer architecture [15]. In particular, the mechanism allows weighing the relevance of each token with other tokens from the input, i.e., question, candidate, and contextual sentences. The approach, however, could be sub-optimal for AS2 for two reasons. First, the concatenation approach fails to ignore irrelevant information in the context sentences, which could introduce noise to the prediction of the model. Second, the approach fails to explicitly capture the alignment between the question and the answer/context sentences, which can be done to better reveal the correctness for the answer candidate. Table 1 shows an example where a correct answer can be optimally selected by capturing the relevance between tokens from the question *q* and sentences in a paragraph *p* containing an answer candidate *a*.

We address these shortcomings in this paper. First, we propose to align the tokens of the question and the paragraph’s sentences to enhance the relevance computing. Specifically, we propose to employ Optimal Transport (OT) [16, 17] to solve the question-context alignment problem. Given two probability

distributions over two point sets and a cost function that measures the distance between any two points, the goal of OT is to find a mapping that moves probability from one distribution to another such that the total cost of transporting the probability mass between two point sets is minimized. We consider the question and a paragraph sentence (i.e., the answer candidate or context sentence) as two point sets, each word being a point. To measure the distance between two points, we employ the Euclidean distance between their semantic representations, which can be obtained from a pre-trained language model (PLM), e.g., RoBERTa [4]. A probability distribution is also defined over each point set via the frequencies of the words in training data. Intuitively, the optimal alignment between the question and the candidate sentence maps the words that are both statistically and semantically similar between the two sentences. In the end, the relevant context formed by the set of the candidate words aligned with the question is utilized to compute the representation for the paragraph sentence.

Second, we address the dependencies among sentences in the paragraph, i.e., the answer candidate and its contextual sentences, to further bolster the modeling of the answer candidate representation. In particular, we consider the paragraph sentences as nodes in a fully-connected graph and aim to learn a dependency weight between nodes. To compute the dependency weight for two sentences, we propose to employ their semantic representations and transportation costs with respect to the question, which are already obtained from the question-context alignment step. A feed-forward network with a sigmoid output function is then used to consume such information to estimate the dependency weight. Afterwards, the dependency weights are utilized to enhance the representations for the sentences via a Graph Convolutional Network (GCN) [18]. The output representation from the GCN for the answer candidate sentence can be directly sent to a binary classifier to obtain its correctness probability score. We finally exploit the mutual information (MI) between the GCN representations of the sentences to further encourage information sharing between them, i.e., by maximizing and minimizing the mutual information between the sentence pairs. To this end, we treat the GCN representations of the sentences as continuous random variables. The MI between the variables can then be optimized via the mutual information neural estimation (MINE) method [19, 20], which approximately estimates the lower bound of the MI via the binary cross entropy of a variable discriminator for optimization.

To demonstrate the effectiveness of our proposed model for AS2, we conduct experiments on a widely-used AS2 benchmark, i.e., WikiQA [21], used in previous work [13] and a recent large-scale dataset WDRASS [22]. Experimental results across the datasets show that our model achieves significant improvements compared to the previous work, obtaining new state-of-the-art performance for AS2.

2. Related work

AS2 is an important task in ODQA and is often solved by point-wise methods that rank answer sentences extracted from retrieved Web documents [23, 24, 25, 11]. Contextual information, e.g., neighboring sentences or document title, was recently incorporated to improve AS2 [13, 14]. Previous work, however, exploited the additional context by string concatenating [13], i.e., appending the input sequence. By contrast, we propose to explicitly model the dependencies between the answer sentence and its context. We propose to align words between the question and the answer/context sentences via Optimal Transport (OT),

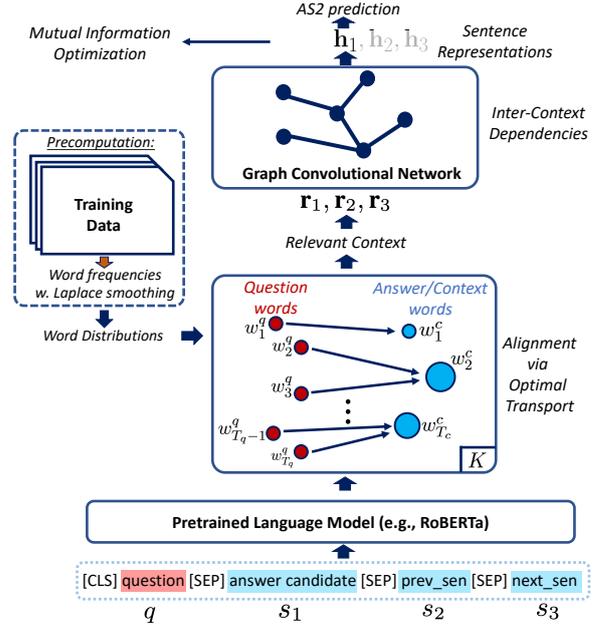


Figure 1: Illustration for our proposed model.

a popular alignment method used in machine translation [26], event argument extraction [27], and event coreference resolution [28]. Our work is the first to apply OT to perform the question-context alignment for AS2.

3. Architecture

Given a question q and a set of N answer sentence candidates $C = \{c_1, c_2, \dots, c_N\}$, the task of AS2 seeks to find correct answers $A \subset C$ via learning a reranking function $r : Q \times \phi(C) \rightarrow \phi(C)$, where Q is the question set and $\phi(C)$ is the set of all permutations of C , such that the answer sentences A are on top of the ranking produced by r . The reranker r is often implemented as a pointwise network $f(q, c_i)$, e.g., TANDA [11], which learns a correctness score $p_i \in (0, 1)$ for each answer candidate c_i for ranking. Our work focuses on contextual AS2 [13], where additional context such as surrounding sentences is considered to better determine the correctness score for an answer candidate.

Our proposed model, called “CASSIE”, for contextual AS2 has four main components: i) Encoding, ii) Question-Context Alignment with OT, iii) Answer-Context Dependencies, and iv) Mutual Information Optimization.

3.1. Encoding

We are given a question $q = [w_1^q, w_2^q, \dots, w_{T_q}^q]$ with T_q words and a set of N answer candidates $C = \{c_1, c_2, \dots, c_N\}$ (retrieved by a search engine), where each candidate is a sentence $c_i = [w_1^c, w_2^c, \dots, w_{T_c}^c]$ with T_c words. Following the previous work [13], we consider previous and next sentences s_{prev} , s_{next} as additional context for each answer candidate $c \in C$ ¹. The input for our model is then formed by concatenating the question, the answer candidate, and the context sentences to obtain a single input sequence: $[q; c; s_{prev}; s_{next}]$. The resulting sequence is fed into a pre-trained language model (PLM),

¹We employ padding sentences if any context sentence is missing.

e.g., RoBERTa [4], to obtain contextualized embeddings for the words. In addition, we employ different segment embeddings for the words in the question, the answer candidate, and the context sentences. These segment embeddings, which are randomly initialized and learnable during training, are added to initial embeddings for the words in the first layer of the PLM. For convenience, let $[\mathbf{w}_1^q, \mathbf{w}_2^q, \dots, \mathbf{w}_{T_q}^q]$ and $[\mathbf{w}_1^c, \mathbf{w}_2^c, \dots, \mathbf{w}_{T_c}^c]$ be the sequences of word representations obtained from the last layer of the PLM for the question q and the answer candidate $c \in C$ respectively.

3.2. Question-Context Alignment with OT

Optimal Transport (OT) [16, 17] is an established method to move probability from one distribution to another by finding an alignment between two point sets. In a discrete setting, we are given two probability distributions p_X and p_Y over two point sets $X = \{x_i\}_{i=1}^n$ and $Y = \{y_j\}_{j=1}^m$ respectively ($\sum_i p_{x_i} = 1$ and $\sum_j p_{y_j} = 1$). A function $D(x, y) : X \times Y \rightarrow \mathbb{R}^+$ is also provided to measure the distance between two points x and y . OT aims to find a mapping that moves probability mass from the points $\{x_i\}_{i=1}^n$ to the points $\{y_j\}_{j=1}^m$ such that the total cost of transporting the probability mass between the two point sets is minimized. Formally, the goal of OT is to find the transportation matrix $\pi_{XY} \in \mathbb{R}^{+n \times m}$ that minimizes the following transportation cost: $d_{XY} = \sum_{\substack{1 \leq i \leq n \\ 1 \leq j \leq m}} D(x_i, y_j) \pi_{XY}$

such that $\pi_{XY} \mathbf{1}_m = p_X$ and $\pi_{XY}^T \mathbf{1}_n = p_Y$. The transportation matrix π_{XY} represents the optimal alignment between the point sets X and Y , where the i -th row in the matrix provides the optimal alignment from a point $x_i \in X$ to each point $y_j \in Y$.

In our question-context alignment problem, we consider the question q and a candidate/context sentence c as two point sets: $\{w_i^q\}_{i=1}^{T_q}$ and $\{w_i^c\}_{i=1}^{T_c}$ respectively (each word is a point)². To obtain the probability distributions for the sets, we propose to measure the frequencies of the words and perform a sum normalization. In particular, the probability distribution for the question is computed as follows: $p_{w_i^q} = \frac{\text{freq}(w_i^q)}{\sum_{i'=1}^{T_q} \text{freq}(w_{i'}^q)}$, where $\text{freq}(w_i^q)$ is the number of questions that the word w_i^q appears in training data. Next, to estimate the distance between two words (points) $w_i^q \in q$ and $w_j^c \in c$, we measure their semantic divergence by computing the Euclidean distance of their contextualized representations obtained from the PLM: $D(w_i^q, w_j^c) = \|\mathbf{w}_i^q - \mathbf{w}_j^c\|$. The optimal transportation matrix π_{XY} , i.e., π_{qc} for the question q and the sentence c can then be solved efficiently using the Sinkhorn-Knopp algorithm [29, 17]. Finally, we obtain the relevant context r_c for the sentence c as: $r_c = \bigcup_{i=1}^{T_q} \{w_j^c | j = \arg\max_{1 \leq j' \leq T_c} \pi_{qc_{ij'}}\}$. In the end, we compute the representation for the sentence c as the average sum over the word representations for the relevant context:

$$\mathbf{r}_c = \frac{1}{|r_c|} \sum_{j | w_j^c \in r_c} \mathbf{w}_j^c \quad (1)$$

3.3. Answer-Context Dependencies

For convenience, let $[\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3]$ be the representations obtained from Equation (1) for the answer candidate $s_1 \equiv c$, the previous sentence $s_2 \equiv s_{prev}$, and the next sentence $s_3 \equiv s_{next}$. To learn the dependencies among the sentences, we consider each sentence as a node in a fully-connected graph $G = (V, E)$,

²We exclude stopwords and punctuations from the two point sets before performing the alignment.

where $V = \{s_i\}$ ($1 \leq i \leq 3$) is the node set and $E = \{(s_i, s_j)\}$ ($1 \leq i, j \leq 3$) is the edge set. Our goal is to learn a weight $\alpha_{ij} \in (0, 1)$ for each edge (s_i, s_j) to represent the dependence of s_i on s_j . To this end, we propose to employ their semantic representations $\mathbf{r}_i, \mathbf{r}_j$, and transportation costs to the question d_{qs_i}, d_{qs_j} to measure the dependency weight α_{ij} between the sentences s_i and s_j . Particularly, we first compute the score: $u_{ij} = FFN_{DEP}([\mathbf{r}_i \odot \mathbf{r}_j; d_{qs_i}; d_{qs_j}])$, where \odot is the element-wise product, $[\cdot]$ represents the concatenation operation, and FFN_{DEP} is a feed-forward network. Afterwards, the weight α_{ij} for the edge (s_i, s_j) is obtained via a softmax function: $\alpha_{ij} = \frac{\exp(u_{ij})}{\sum_{j'=1}^K \exp(u_{ij'})}$. The induced weights $\{\alpha_{ij}\}$ are then used to enhance the representations for the sentences via L layers of a Graph Convolutional Network (GCN) [18]:

$$\mathbf{h}_i^l = \text{ReLU}\left(\sum_{j=1}^K \alpha_{ij} \mathbf{W}^l \mathbf{h}_j^{l-1} + \mathbf{b}^l\right) \quad (2)$$

where $\mathbf{W}^l, \mathbf{b}^l$ are learnable weight matrix and bias for the layer l of the GCN ($1 \leq l \leq L$), and $\mathbf{h}_i^0 \equiv \mathbf{r}_i$ is the input representation for the sentence s_i . The output vectors $\mathbf{h}_i^L \equiv \mathbf{h}_i$ at the last layer of the GCN serve as the final representations for the sentences s_i . The representation \mathbf{h}_1 for the answer candidate $s_1 \equiv c$ is finally sent to a feed-forward network with a sigmoid output function to estimate the correctness score $p_c \in (0, 1)$ for the answer candidate c : $p_c = FFN_{AS2}(\mathbf{h}_1)$. For training, we minimize the binary cross-entropy loss L_{AS2} with the correctness scores p_c .

3.4. Mutual Information Optimization

In information theory, MI is defined as the KL divergence between the joint distribution and the product of the marginal distributions of two random variables. As a result, two random variables would be more dependent if they have large mutual information. In our case, some of the context sentences might also be correct/incorrect answers for the question. Therefore, we expect answer sentence pairs to share more semantic information and (answer, non-answer) sentence pairs to share less semantic information relevant to the question. This can be done by considering sentence representation vectors \mathbf{h}_i as random variables and maximizing/minimizing the mutual information between the variables, respectively. However, the sentence vectors \mathbf{h}_i are very high dimensional variables, making the exact calculation of the MI between the vectors impossibly expensive. To overcome this, we followed the mutual information neural estimation (MINE) method [19, 20] to estimate and optimize the lower bound of the MI between the variables via the binary cross entropy of a variable discriminator U , which is a feed-forward network with a sigmoid output function:

$$L_{MI} = - \sum_{(i,j) \in I_+} \log(U([\mathbf{h}_i; \mathbf{h}_j])) - \sum_{(i',j') \in I_-} \log(1 - U([\mathbf{h}_{i'}; \mathbf{h}_{j'}])) \quad (3)$$

where $I_+ = \{(i, j) | s_i, s_j \in A\}$ ($1 \leq i, j \leq 3$) is the index set for answer sentence pairs, and $I_- = \{(i', j') | s_{i'} \in A, s_{j'} \notin A\}$ ($1 \leq i', j' \leq 3$) is the index set for (answer, non-answer) sentence pairs among the three sentences.

3.5. Training and Inference

All the components in our proposed model are jointly trained via minimizing the loss function:

$$L = L_{AS2} + \gamma L_{MI} \quad (4)$$

where γ is a hyper-parameter to balance the contributions of each component to the training of the model. Following the previous work [11, 13], we consider all answer candidates for each question for training and inference.

Datasets	Train	Dev	Test	#A/Q
WikiQA	2,118	126	243	1.18
WDRASS	53,419	5,416	5,395	4.96

Table 2: The number of questions in training, development, and test data of the two datasets. The last column presents the average number of answers to a question.

4. Results

4.1. Experimental Setup

Datasets Following the previous works [11, 13], we use the same train/dev/test splits for the standard AS2 dataset, i.e., *WikiQA* [21]. In addition, we also experimented with a large-scale AS2 dataset called *WDRASS* [30] to investigate the models’ performance further.

- **WikiQA** is a QA dataset created by [21]. The dataset contains questions and answer candidates, manually annotated on Bing query logs over Wikipedia. Following the previous works, we conduct experiments with the *clean* version of the dataset, and combine the development and test sets to obtain a larger and more reliable set for model comparison.
- **WDRASS** is recently created by [30]. WDRASS is a large-scale dataset focusing on non-factoid questions requiring entire sentences to answer.

Statistics for all the datasets are shown in Table 2.

Hyper-parameters and Tools Following the previous work [13], we use a small portion of the WikiQA training data to tune hyper-parameters for our model and select the best hyper-parameters for all the datasets. We employ Adam optimizer to train the model with a learning rate of $1e - 5$ and a batch size of 64. We set 400 for the hidden vector sizes for all the feed-forward networks, $L = 2$ for the number of the GCN layers, and 0.3 for the trade-off weights γ . To implement the models, we use Pytorch version 1.7.1 and Huggingface Transformers version 3.5.1. We use the NLTK library version 3.5 [31] to preprocess the data and remove stopwords. The model performance is obtained over three runs with different random seeds.

Evaluation Metrics Following the previous works, we measure the model performance using the entire set of answer candidate sentences for each question, using the three metrics: Precision-at-1 (P@1), Mean Average Precision (MAP), and Mean Reciprocal Rank (MRR) scores.

4.2. Performance Comparison

We compare our proposed model with TANDA [11] and LOCT [13], which are the current state-of-the-art models for AS2. Table 3 shows the performance comparison between the models on two settings: i) using a non-finetuned RoBERTa base encoder, and ii) using a finetuned RoBERTa base encoder. The

Models	Non-finetuned RoBERTa Base		ASNQ-finetuned RoBERTa Base	
	P@1	MAP	P@1	MAP
TANDA	63.24*	75.00*	78.67*	86.74*
LOCT	68.09*	79.00*	81.31*	88.00*
CASSIE	74.16	83.29	83.77	89.28

Table 3: Performance comparison of the models on combination of development and test sets of WikiQA. * indicates the performance officially reported by the previous work [13].

Models	ASNQ-finetuned RoBERTa Base		
	P@1	MAP	MRR
TANDA	54.6	63.5	64.3
CASSIE	55.9	61.8	69.7
CASSIE (joint)	55.9	64.2	65.0

Table 4: Performance comparison on WDRASS test set.

non-finetuned RoBERTa Base is obtained from [4] while the other is produced by finetuning TANDA on the ASNQ dataset [11]. As can be seen from the table, all the models benefit from using the finetuned RoBERTa Base encoder. Across the two settings, our model outperforms the previous models by large margins, demonstrating its effectiveness for AS2.

Table 4 shows the performance of our proposed model compared to TANDA on the WDRASS test set. CASSIE significantly improves the performance for P@1 and MRR, however, decreases the performance for MAP. We attribute this to the fact that our model ranks the answer candidates individually. To deal with this, we explore another use case of our model where it can produce a joint reranking for multiple answer candidates ranked by TANDA, leading to significantly better performance for all the three metrics.

5. Conclusions and Future Work

In this work, we propose CASSIE, a novel LLM-based model that (i) efficiently learns the answer-context dependencies to improve representation learning for AS2 by (ii) leveraging relevant context in answer/context sentences captured via question-context alignments using Optimal Transport. Experimental results demonstrate the efficacy of our proposed model, resulting in significant improvements and new state-of-the-art performance across several widely-used AS2 benchmark WikiQA [21] and a recent large-scale AS2 dataset (WDRASS) [30].

Naturally, our proposed model CASSIE can approximate human judgement for the correctness of an answer for a given question. As a result, the model can be used to provide signals for assessing quality of answers produced by open-domain question answering (ODQA) systems. As such, we are exploring AS2 datasets and architecture designs for building better reward models for reinforcement learning from human feedback [7], which has shown impressive improvements for ODQA via training LLMs to generate answers for open-domain questions.

6. Acknowledgments

We would like to thank Zeyu Zhang for sharing the WDRASS dataset [22] and his helpful comments on the paper regarding the experiment setup and data preparation.

7. References

- [1] E. M. Voorhees, “The TREC-8 question answering track report,” in *Proceedings of The Eighth Text REtrieval Conference, TREC 1999, Gaithersburg, Maryland, USA, November 17-19, 1999*, ser. NIST Special Publication, vol. 500-246. National Institute of Standards and Technology (NIST), 1999.
- [2] M. Wang, N. A. Smith, and T. Mitamura, “What is the Jeopardy model? a quasi-synchronous grammar for QA,” in *Proceedings of EMNLP-CoNLL*, 2007.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL*, 2019.
- [4] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [5] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [6] C.-C. Hsu, E. Lind, L. Soldaini, and A. Moschitti, “Answer generation for retrieval-based question answering systems,” in *Findings of ACL-IJCNLP*, 2021.
- [7] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders *et al.*, “Webgpt: Browser-assisted question-answering with human feedback,” *arXiv preprint arXiv:2112.09332*, 2021.
- [8] J. Maynez, S. Narayan, B. Bohnet, and R. McDonald, “On faithfulness and factuality in abstractive summarization,” in *Proceedings of ACL*, 2020.
- [9] S. Roller, E. Dinan, N. Goyal, D. Ju, M. Williamson, Y. Liu, J. Xu, M. Ott, E. M. Smith, Y.-L. Boureau, and J. Weston, “Recipes for building an open-domain chatbot,” in *Proceedings of EACL*, 2021.
- [10] K. Shuster, S. Poff, M. Chen, D. Kiela, and J. Weston, “Retrieval augmentation reduces hallucination in conversation,” in *Findings of EMNLP*, 2021.
- [11] S. Garg, T. Vu, and A. Moschitti, “Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, p. 7780–7788, Apr 2020. [Online]. Available: <http://dx.doi.org/10.1609/AAAI.V34I05.6282>
- [12] Z. Zhang, T. Vu, and A. Moschitti, “Joint models for answer verification in question answering systems,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 3252–3262. [Online]. Available: <https://aclanthology.org/2021.acl-long.252>
- [13] I. Lauriola and A. Moschitti, “Answer sentence selection using local and global context in transformer models,” in *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021*. Springer, 2021.
- [14] R. Han, L. Soldaini, and A. Moschitti, “Modeling context in answer sentence selection systems on a latency budget,” in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. Online: Association for Computational Linguistics, Apr. 2021, pp. 3005–3010. [Online]. Available: <https://aclanthology.org/2021.eacl-main.261>
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [16] G. Monge, “Mémoire sur la théorie des déblais et des remblais,” *Mem. Math. Phys. Acad. Royale Sci.*, pp. 666–704, 1781.
- [17] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” *Advances in neural information processing systems*, vol. 26, 2013.
- [18] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [19] M. I. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, A. Courville, and D. Hjelm, “Mutual information neural estimation,” in *International conference on machine learning*. PMLR, 2018, pp. 531–540.
- [20] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio, “Learning deep representations by mutual information estimation and maximization,” in *ICLR*, 2019.
- [21] Y. Yang, W.-t. Yih, and C. Meeck, “WikiQA: A challenge dataset for open-domain question answering,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 2013–2018.
- [22] Z. Zhang, T. Vu, S. Gandhi, A. Chadha, and A. Moschitti, “Wdrass: A web-scale dataset for document retrieval and answer sentence selection,” in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022, pp. 4707–4711.
- [23] A. Severyn and A. Moschitti, “Learning to rank short text pairs with convolutional deep neural networks,” in *Proceedings of ACM SIGIR*, 2015.
- [24] G. Shen, Y. Yang, and Z.-H. Deng, “Inter-weighted alignment network for sentence pair modeling,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics, Sep. 2017, pp. 1179–1189.
- [25] S. Yoon, F. Deroncourt, D. S. Kim, T. Bui, and K. Jung, “A compare-aggregate model with latent clustering for answer selection,” in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019.
- [26] Z.-Y. Dou and G. Neubig, “Word alignment by fine-tuning embeddings on parallel corpora,” in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. Online: Association for Computational Linguistics, Apr. 2021, pp. 2112–2128. [Online]. Available: <https://aclanthology.org/2021.eacl-main.181>
- [27] A. Pouran Ben Veyseh, M. V. Nguyen, F. Deroncourt, B. Min, and T. Nguyen, “Document-level event argument extraction via optimal transport,” in *Findings of the Association for Computational Linguistics: ACL 2022*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 1648–1658. [Online]. Available: <https://aclanthology.org/2022.findings-acl.130>
- [28] D. Phung, H. Minh Tran, M. V. Nguyen, and T. H. Nguyen, “Learning cross-lingual representations for event coreference resolution with multi-view alignment and optimal transport,” in *Proceedings of the 1st Workshop on Multilingual Representation Learning*. Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 62–73.
- [29] R. Sinkhorn and P. Knopp, “Concerning nonnegative matrices and doubly stochastic matrices,” *Pacific Journal of Mathematics*, vol. 21, no. 2, pp. 343–348, 1967.
- [30] Z. Zhang, T. Vu, S. Gandhi, A. Chadha, and A. Moschitti, “Wdrass: A web-scale dataset for document retrieval and answer sentence selection,” in *Proceedings of the 31st ACM International Conference on Information and Knowledge Management*, ser. CIKM ’22. Association for Computing Machinery, 2022, p. 4707–4711.
- [31] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc., 2009.