



Boosting Neural Language Inference via Cascaded Interactive Reasoning

Min Li^{1 †} and Chun Yuan^{2 †}

[†] Tsinghua Shenzhen International Graduate School, Tsinghua University, Beijing, China

¹limin21@mails.tsinghua.org.cn

²yuanc@sz.tsinghua.edu.cn

Abstract. Natural Language Inference (NLI) focuses on ascertaining the logical relationship (entailment, contradiction, or neutral) between a given premise and hypothesis. This task presents significant challenges due to inherent linguistic features such as diverse phrasing, semantic complexity, and contextual nuances. While Pre-trained Language Models (PLMs) built upon the Transformer architecture have yielded substantial advancements in NLI, prevailing methods predominantly utilize representations from the terminal layer. This reliance on final-layer outputs may overlook valuable information encoded in intermediate layers, potentially limiting the capacity to model intricate semantic interactions effectively. Addressing this gap, we introduce the Cascaded Interactive Reasoning Network (CIRN), a novel architecture designed for deeper semantic comprehension in NLI. CIRN implements a hierarchical feature extraction strategy across multiple network depths, operating within an interactive space where cross-sentence information is continuously integrated. This mechanism aims to mimic a process of progressive reasoning, transitioning from surface-level feature matching to uncovering more profound logical and semantic connections between the premise and hypothesis. By systematically mining latent semantic relationships at various representational levels, CIRN facilitates a more thorough understanding of the input pair. Comprehensive evaluations conducted on several standard NLI benchmark datasets reveal consistent performance gains achieved by CIRN over competitive baseline approaches, demonstrating the efficacy of leveraging multi-level interactive features for complex relational reasoning.

Keywords: Neural Language Inference, Deep Learning, Neural Language Process.

1 Introduction

Natural Language Inference (NLI), often framed as Recognizing Textual Entailment (RTE), seeks to discern the logical connection between a pair of sentences, classifying it as "entailment," "contradiction," or "neutral." As a foundational task in natural language understanding, NLI presents significant hurdles. Its difficulty arises not just from recognizing diverse linguistic constructions but also from the necessity of capturing underlying commonsense knowledge and performing nuanced reasoning.

The attention mechanism has proven highly effective across numerous NLP domains. It enhances cross-lingual semantic mapping in machine translation, aids in identifying crucial content for text summarization [3], and directs models towards relevant passages for accurate question answering. As formalized by [27], attention calculates weighted context representations based on query-key similarities, adeptly modeling dependencies irrespective of distance in the sequence. Building on this success, we posit that developing more sophisticated attention-derived representations, which capture richer interaction patterns, can further elevate performance on complex reasoning tasks like NLI.

Conventional attention mechanisms typically compute an alignment matrix representing word-level correspondences between sentences. While multi-head attention [27] enriches this by offering multiple alignment perspectives, these often focus on pairwise token interactions. In this work, we aim to model more complex, higher-order dependencies that span beyond individual words to encompass phrase-level and structural relationships. We propose a method to generate intermediate representations encoding these intricate interactions, moving beyond simple lexical matching to capture structural configurations and combinatorial semantics between the sentence pair. Our experimental findings suggest that leveraging these enriched semantic features markedly improves logical reasoning accuracy and enhances model robustness in challenging scenarios involving paraphrases, antonyms, and lexical ambiguities.

To realize this, we introduce a novel architecture termed the Cascaded Interactive Reasoning Network (CIRN). CIRN is specifically designed to model intricate semantic interactions hierarchically. It operates by progressively extracting and integrating features across multiple layers of representation within an interactive space, simulating a gradual refinement of understanding from surface-level cues to deeper logical connections. We validate CIRN extensively on standard NLI benchmarks, including SNLI and MultiNLI [8]. Furthermore, to assess its versatility, we adapt CIRN for paraphrase identification, treating matched pairs as entailment and non-matching as neutral, and test it on the large-scale Quora Question Pair (QQP) dataset and seven other datasets. Across these diverse evaluations, CIRN consistently surpasses strong baseline models, underscoring its effectiveness and broad applicability.

The primary contributions of this work can be outlined as follows:

- Introduction of the Cascaded Interactive Reasoning Network (CIRN), a general framework designed for hierarchical aggregation and interactive reasoning over multi-level semantic information pertinent to NLI.
- A method embedded within CIRN to explicitly model complex, higher-order semantic dependencies and interactions between sentence pairs, extending beyond traditional token-level alignment approaches.
- Comprehensive empirical validation on multiple public datasets (including NLI and paraphrase identification), demonstrating consistent and significant performance improvements achieved by CIRN over competitive baseline methods.

2 Related Work

Traditional Neural Network-based Methods. Early approaches to Natural Language Inference (NLI) primarily relied on handcrafted syntactic features, transformation rules, and relation extraction techniques, which performed well on small-scale datasets but struggled with generalization and scalability [7]. The introduction of large-scale datasets [8] and the development of deep learning frameworks triggered a major shift in the field. Attention mechanisms, in particular, revolutionized NLI by enabling models to explicitly capture alignment and dependency relationships for enhanced semantic reasoning [9,10,29]. Sentence-encoding architectures [9] as well as joint models leveraging cross-attention mechanisms improved performance by learning hierarchical word- and phrase-level interactions [11]. Innovations like residual networks further enhanced model depth while retaining lower-layer information [16,33]. Neural network architectures have been widely applied for comprehensive sentence-level semantic modeling in NLI settings. Sequential dependencies and contextual semantic relationships are better modeled using Recurrent Neural Networks (RNNs) [12,21,30,32], while convolutional frameworks capture critical local features through sliding-window operations [13]. Attention-based mechanisms further improve upon this by identifying salient components, aligning key phrases, and distilling semantic relevance [14,31,39]. Techniques such as bidirectional Long Short-Term Memory (Bi-LSTM) networks offer a deeper representation of feature differences by incorporating contextual information at multiple levels [15,38]. Similarly, other architectures, including CNNs and RNNs, focus on either local sensitivity [17] or sequential semantics [17]. Recently, Graph Neural Networks (GNNs) have gained traction for leveraging sentence structural information and modeling global connections [18]. Additionally, some models adhere to the principle of encoding independent sentence representations without cross-interaction during feature extraction, relying on classifier layers for downstream tasks [10].

Pre-trained Language Model-based Approaches. The emergence of pre-trained language models such as BERT [3] fundamentally transformed NLI tasks by introducing versatile representations obtained through pretext self-supervised objectives. These representations, fine-tuned on downstream tasks, led to significant gains in accuracy and robustness. Follow-up models such as XLNet [19], RoBERTa [4], and CharBERT [20] further refined pretraining methodologies, thereby shrinking the performance gap between pretraining and fine-tuning. Cross-feature attention modules within these models allow them to explicitly focus on granular word- and phrase-level alignments, improving semantic reasoning [12,36,41]. Additionally, sentence interaction mechanisms, such as co-attentive DenseNets, further enhance multi-sentence relationships by integrating multiple layers of dependencies [20,37]. A complementary avenue of progress involves incorporating linguistic knowledge to refine sentence representations. Methods introduce explicit syntactic and semantic attributes, such as leveraging part-of-speech (POS) tags and named entities [22,40], syntactic dependency parsing [34], semantic roles [5,42], word synonym relations [6,35], and syntax tree structures [23].

These augmentations help capture task-specific nuances, particularly for subtle distinctions that are not evident in global sentence semantics. However, current methods tend to prioritize global semantic similarity, often neglecting fine-grained distinctions at the token or phrase level, limiting performance on tasks requiring precise semantic matching.

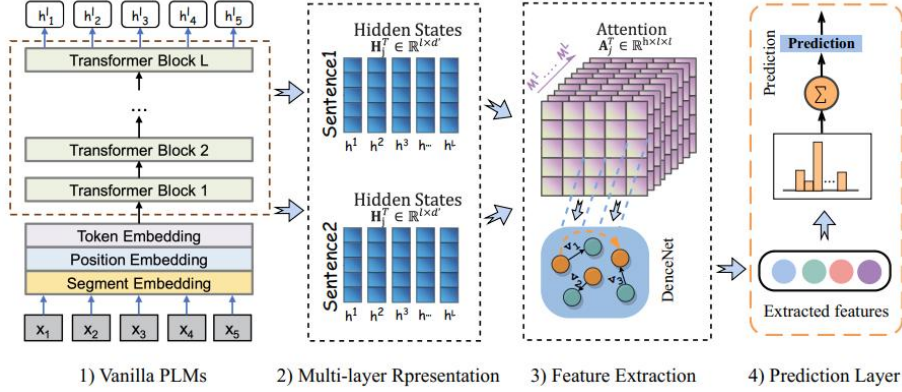


Fig. 1. The overall architecture of the Cascaded Interactive Reasoning Network (CIRN): 1) **PLM Encoding**: concatenates sentence pair (S_1, S_2) as input to a pre-trained language model (e.g., BERT, RoBERTa), obtaining representations $H^{(l)}$ from each Transformer layer; 2) **Multi-layer Representation**: separates $H^{(l)}$ into sentence representations $H_1^{(l)}, H_2^{(l)}$, and computes interaction matrices $M^{(l)} \in \mathbb{R}^{n \times m \times d}$ by element-wise multiplication; 3) **Feature Extraction**: stacks interaction matrices across layers into an interaction tensor $\bar{M} \in \mathbb{R}^{n \times m \times d \times L}$, and applies DenseNet to extract deep interactive features; 4) **Prediction Layer**: aggregates extracted features through a fully connected layer followed by softmax to output classification probabilities.

3 Model

We frame the natural language inference (NLI) task as a multi-class classification problem aimed at predicting the relationship $y \in \mathcal{Y}$ between a sentence pair, where $\mathcal{Y} = \{\textit{entailment}, \textit{contradiction}, \textit{neutral}\}$. To tackle this, we propose the **Cascaded Interactive Reasoning Network (CIRN)**. The overall model architecture is depicted in the left panel of Figure 1.

3.1 Input Preprocessing

For a given sentence pair $S_1 = \{x_1^1, x_2^1, \dots, x_n^1\}$ and $S_2 = \{x_1^2, x_2^2, \dots, x_m^2\}$, we introduce special tokens [CLS] and [SEP] to conform to BERT's input format. For BERT, the input embeddings are constructed as the sum of three components:

$$X_i = T_i + S_i + P_i, \quad (1)$$

where T_i , T_i , and P_i correspond to the Token Embedding, Segment Embedding, and Positional Embedding for the i -th token, respectively. Here, Segment Embedding differentiates the first text (Sentence A) from the second text (Sentence B).

3.2 Multi-layer Transformer Representations

The backbone of BERT is composed of L stacked Transformer layers, each comprising Multi-Head Self-Attention and a Feed-Forward Network. Let $H^{(0)} = [X_1, X_2, \dots, X_{N'}]$ denote the input embedding sequence, where $N' = n + m + 3$ (including [CLS] and two [SEP] tokens). The output of the l -th Transformer layer can be expressed as:

$$H^{(l)} = \text{TransformerLayer}_l(H^{(l-1)}), \quad l = 1, 2, \dots, L. \quad (2)$$

From $H^{(l)}$, we extract the token embeddings for the first and second sentences, denoted as $H_1^{(l)} \in \mathbb{R}^{n \times d}$ and $H_2^{(l)} \in \mathbb{R}^{m \times d}$, where d is the hidden dimension of the model.

3.3 Interaction Matrices Across Layers

Element-wise Interaction Representation. To explicitly model interactions between tokens from the two sentences, we compute an element-wise product between their embeddings. For the l -th Transformer layer, given token embeddings $h_1^{(l,i)} \in \mathbb{R}^d$ (from the first sentence) and $h_2^{(l,j)} \in \mathbb{R}^d$ (from the second sentence), the interaction tensor is defined as:

$$I_{i,j}^{(l)} = h_1^{(l,i)} \odot h_2^{(l,j)}, \quad \forall i \in [1, n], j \in [1, m], \quad (3)$$

where \odot represents element-wise multiplication. This operation produces a layer-specific interaction tensor $I^{(l)} \in \mathbb{R}^{n \times m \times d}$, encoding token-level semantic dependencies for both sentences.

Stacking Multi-layer Interactions. The layer-wise interaction matrices are stacked along a new dimension to form a comprehensive multi-layer interaction representation:

$$I_{\text{stack}} = [I^{(1)}; I^{(2)}; \dots; I^{(L)}] \in \mathbb{R}^{n \times m \times d \times L}, \quad (4)$$

where $[\cdot; \cdot]$ denotes concatenation along the layer dimension. The aggregated tensor I_{stack} captures interactions at all layers, offering a detailed and hierarchical view of semantic relationships between the two sentences.

3.4 DenseNet for Feature Extraction

DenseNet, originally designed for image processing, effectively enhances gradient flow and feature reuse through dense connections across layers. In our model, we adapt DenseNet to extract features from the stacked interaction tensor I_{stack} . Convolutional or

pooling operations are first applied to reduce dimensionality, followed by cascading dense blocks.

Let the input of the first dense block be Z_0 . Each layer in a dense block computes its output as:

$$z_k = f_k(\text{Concat}(Z_0, z_1, \dots, z_{k-1})), \quad (5)$$

where f_k represents the transformation function (e.g., convolution or fully connected operation) of layer k . The final output of the Dense Block is obtained as:

$$Z_m = \text{Concat}(Z_0, z_1, \dots, z_{m-1}) \quad (6)$$

where Z_m aggregates features across all layers within the block.

DenseNet Output Representation. After passing I_{stack} through multiple Dense Blocks and Transition Layers, the final feature tensor F is derived:

$$F = \text{DenseNet}(I_{\text{stack}}) \quad (7)$$

where $F \in \mathbb{R}^{d'}$ encapsulates high-level semantic interactions and refined sentence pair representations.

3.5 Classification Layer

The extracted features F are fed into a fully-connected layer, followed by a softmax classifier to predict the probability distribution over the relation categories $\mathcal{Y} = \{\textit{entailment}, \textit{contradiction}, \textit{neutral}\}$. Formally:

$$p = \text{softmax}(WF + b), \quad (8)$$

where $W \in \mathbb{R}^{3 \times d'}$, $b \in \mathbb{R}^3$ are trainable weights. The predicted label is selected as:

$$\hat{y} = \arg \max_i p_i. \quad (9)$$

Training Objective. The training objective minimizes the cross-entropy loss between the true label y (one-hot encoded) and the predicted probability distribution p :

$$\mathcal{L} = -\sum_{i=1}^3 y_i \log(p_i), \quad (10)$$

where y represents the ground-truth label, and $y_i \in \{0,1\}$ denotes the label for class i .

4 Experimental Configuration

This section details the experimental methodology used to evaluate our proposed Cascaded Interactive Reasoning Network (CIRN). We describe the datasets employed, the

baseline models selected for comparison, and the specific implementation settings and hyperparameters used during training and evaluation.

4.1 Evaluation Benchmarks

To thoroughly assess the performance of CIRN, our experiments utilize ten large-scale, publicly accessible benchmark datasets commonly employed in natural language understanding research. A primary evaluation suite is the General Language Understanding Evaluation (GLUE) benchmark [24], a standard testbed featuring diverse NLP tasks including sentence similarity and textual entailment assessment¹. We focus on six sentence-pair classification tasks within GLUE: MRPC, QQP, STS-B, MNLI, RTE, and QNLI. Complementing the GLUE tasks, we incorporate four additional widely-used datasets to gauge model generalizability across different domains and task nuances: SNLI [8], SICK [24], TwitterURL [25], and Scitail [26].

4.2 Comparative Baselines

We benchmark CIRN against a representative set of existing models to demonstrate its relative effectiveness. Our main points of comparison are Pre-trained Language Models (PLMs), including the foundational BERT [3] architecture. We also compare against several PLM variants enhanced with external knowledge sources, namely SemBERT [5], SyntaxBERT [23], and UERBERT [6], allowing us to contextualize the gains from our specific interactive reasoning approach versus other knowledge integration techniques. Furthermore, we include competitive models that do not rely on large-scale pre-training. This category features architectures like ESIM [2], a standard Transformer encoder [27] trained specifically for the task, and other relevant models focusing on bilateral multi-perspective matching or compare-aggregate strategies [11,28]. These baselines help isolate the benefits derived from pre-training versus architectural innovations like CIRN. For robustness evaluations (details potentially in another section referencing TextFlint), performance comparisons are primarily made against the standard BERT baseline under various input perturbations. Detailed architectural descriptions of baseline models are omitted for conciseness and can be found in their respective publications.

4.3 Implementation and Training Parameters

Our implementation utilizes the PyTorch deep learning framework. Optimization of all trainable model parameters is performed using the Adadelata optimizer, configured with $\rho = 0.95$ and $\epsilon = 1 \times 10^{-8}$. The initial learning rate is set to 0.5, and training proceeds with a batch size of 70. A learning rate adaptation mechanism is employed: if validation performance on the target task does not improve for 30,000 consecutive steps, the optimizer is switched to Stochastic Gradient Descent (SGD) with a reduced learning rate of 3×10^{-4} to potentially refine the solution towards a better local minimum. Dropout

¹ <https://huggingface.co/datasets/glue>

layers are inserted before all linear transformation layers and following the word embedding layer to mitigate overfitting. L2 weight decay is applied to all model weights for regularization. The effective L2 decay coefficient at training step t , denoted $\mathcal{R}_{L2}(t)$, is scheduled dynamically:

$$\mathcal{R}_{L2}(t) = \sigma\left(\frac{(t-T_{\text{ramp}}/2) \times 8}{T_{\text{ramp}}/2}\right) \times R_{L2}^{\text{max}} \quad (11)$$

where R_{L2}^{max} represents the maximum L2 decay ratio, set to 0.9×10^{-5} , and T_{ramp} defines the number of steps over which the decay ramps up to its maximum, set to 100,000. $\sigma(\cdot)$ signifies the sigmoid function. Additionally, a specific L2 penalty with a coefficient of 1×10^{-3} is applied to the squared difference between weights of corresponding encoder layers (if a dual-encoder setup is relevant). Regarding the DenseNet-based feature extraction component (as described in Section 3.4), architectural choices include setting the number of layers n within each dense block to 8 and the growth rate g to 20. The initial feature map reduction ratio η is 0.3, and the transition layer compression factor θ is 0.5. To manage computational resources, maximum sequence lengths are enforced via truncation: 48 tokens for MultiNLI, 32 for SNLI, and 24 for QQP. Following established practice [21], the MultiNLI training incorporates 15% of the SNLI dataset.

5 Results and Analysis

5.1 Model Performance Evaluation

Table 1 presents a performance comparison between **CIRN** and various competing approaches across 10 benchmark datasets. Consistent with expectations, methods based on pre-trained language models significantly surpass non-pretrained counterparts, owing to the extensive knowledge acquired during pre-training and sophisticated contextual representation capabilities. Our proposed **CIRN** demonstrates notable improvements over strong baseline models. When integrated with BERT-base, it achieves an average accuracy gain of 0.8%, and with BERT-large, the improvement is 0.7%. These results highlight the efficacy of **CIRN** for tasks requiring nuanced language understanding, such as neural language inference. Furthermore, **CIRN** surpasses the performance of RoBERTa-base [4] by 1.5% and RoBERTa-large by 0.5% on average across the evaluated datasets. We attribute these performance enhancements to **CIRN**'s effective modeling strategy, which concurrently analyzes sentence relationships from both fine-grained local interactions and broader global contexts. This multi-perspective approach enables the model to discern more subtle and intricate semantic connections that might be missed by methods focusing on only one level of analysis. The empirical evidence strongly suggests that incorporating mechanisms to model both local difference information and global semantic context, as **CIRN** does, is advantageous for semantic analysis tasks. Overall, the experimental results position **CIRN** as a highly competitive method for evaluating semantic similarity and related tasks, empirically validating the effectiveness of our proposed architecture and approach across diverse benchmarks.

Table 1. The performance comparison of CIRN with other methods.

Model	Pre- train	m/mm	MRPC	QQP	STS-B	QNLI	RTE	SNLI	Sci	SICK	Tw	MRPC
BiMPM	×	72.3/72.1	79.6	85	-	81.4	56.4	-	-	-	-	-
CAFE	×	78.7/77.9	82.4	88	-	81.5	56.8	88.5	83.3	72.3	-	-
ESIM	×	75.8/75.6	80.3	88.2	-	80.5	-	88	70.6	71.8	-	-
Transformer	×	72.3/71.4	81.7	84.4	73.6	80.3	58	84.6	72.9	70.3	68.8	74.4
BiL+ELMo+Attn	✓	76.4/76.1	84.6	86.7	73.3	79.8	56.8	89	85.8	78.9	81.4	78.9
OpenAI GPT	✓	82.1/81.4	82.3	81.3	80	87.4	56	88.4	84.8	79.5	81.9	80.4
UERBERT	✓	84.2/83.5	88.3	90.5	85.1	90.6	67.1	90.8	92.2	87.8	86.2	86
SemBERT	✓	84.4/84.0	88.2	90.2	87.3	90.9	69.3	90.9	92.5	87.9	86.8	86.5
SyntaxBERT	✓	84.9/84.6	89.2	89.6	88.1	91.1	68.9	91	92.7	88.7	87.3	86.3
DABERT	✓	84.9/84.7	89.1	91.3	88.2	91.4	69.5	91.3	93.6	88.6	87.5	86.7
BERT-Base	✓	84.3/83.7	87.2	89.1	86.8	90.4	67.2	90.7	91.8	87.2	84.8	85.8
BERT-Base-CIRN	✓	85.1/84.9	89.3	89.1	87.1	91.2	68.5	91.2	92.1	87.8	86.6	86.8
BERT-Large	✓	86.8/86.3	88.9	89.3	87.6	92.7	70.1	91	94.4	91.1	91.5	88
BERT-Large-CIRN	✓	86.8/86.7	89.9	90.2	88.1	93	72	91.1	94.3	91.2	92.4	88.8

5.2 Ablation Study

To isolate and evaluate the contribution of individual architectural components within our proposed model, we conducted an ablation study on the MultiNLI dataset [8]. The results of this analysis are presented in Table 2. First, we assessed the **impact of the interaction mechanism**. Removing the component responsible for element-wise interactions (specifically, the Hadamard product) resulted in a noticeable decrease in performance, with accuracy dropping from 85.1% to 84.6% on MNLI. This finding underscores the utility of this interaction feature for capturing fine-grained semantic relationships between the input sequences. Second, we examined the **role of the DenseNet module**. Excising DenseNet from the architecture led to a more significant performance decline, reducing accuracy by 1.2% on MNLI. This highlights DenseNet's crucial function in effectively extracting and consolidating high-level features and complex semantic patterns derived from the text interactions. Third, the **advantages of multi-layer representation integration** were evaluated. When the model was restricted to using only the semantic information from the final layer of the underlying BERT encoder, performance degraded considerably compared to the full model which leverages representations from multiple layers. This result confirms the importance of integrating hierarchical semantic information captured across different depths of the Transformer architecture. Collectively, these ablation experiments validate that the superior performance of our model is not attributable to a single component but arises from the synergistic contribution of its key elements. Both the detailed interaction modeling (via the interaction matrix) and the deep feature extraction capabilities facilitated by DenseNet, built upon multi-layer semantic inputs, are integral to the model's effectiveness.

Table 2. Results of ablation experiment of various composition functions.

Ablation Experiments	Dev Accuracy	
	Matched	Mismatched
1. CIRN _{bert base}	85.1	84.9
2. CIRN _{bert base} - Remove first 11 layer	84.6	83.9
3. CIRN _{bert base} - Remove interaction matrix	84.7	84.1
4. CIRN _{bert base} - Remove feature extraction	83.9	83.4

6 Conclusion

In this paper, we delve into multi-level semantic feature extraction and propose a Cascaded Interactive Reasoning Network (CIRN). By aggregating multi-level semantic information, CIRN improves the accuracy of natural language reasoning, while CIRN breaks away from the limitations of traditional alignment. It encodes high-order complex relationships between sentence pairs, considering not only simple word-to-word connections, but also deeper lexical and semantic associations. This expands the receptive field of the model and enhances its ability to capture complex semantics. The experimental results on 10 public datasets show consistent improvements, especially the significant improvement in robustness tests, highlighting the effectiveness of CIRN.

Acknowledgments. This work was supported by the National Key R&D Program of China (2022YFB4701400/4701402, SSTIC Grant (KJZD20230923115106012, KJZD20230923114916032, GJHZ20240218113604008) and Beijing Key Lab of Networked Multimedia.

References

1. D. Liang, F. Zhang, Q. Zhang, and X.-J. Huang, “Asynchronous deep interaction network for natural language inference,” in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, 2019, pp. 2692–2700.
2. Q. Chen, X. Zhu, Z. Ling, S. Wei, H. Jiang, and D. Inkpen, “Enhanced lstm for natural language inference,” arXiv preprint arXiv:1609.06038, 2016.
3. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” arXiv preprint.
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. Z. Zhang, Y. Wu, H. Zhao, Z. Li, S. Zhang, X. Zhou, and X. Zhou, “Semantics-aware bert for language understanding,” in AAAI, vol. 34, 2020, pp. 9628–9635.
6. T. Xia, Y. Wang, Y. Tian, and Y. Chang, “Using prior knowledge to guide bert’s attention in semantic textual matching tasks,” in WWW 2021.
7. M. Wang, N. A. Smith, and T. Mitamura, “What is the jeopardy model? a quasi-synchronous grammar for qa,” in EMNLP 2007.



8. S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, "A large annotated corpus for learning natural language inference," arXiv preprint arXiv:1508.05326, 2015
9. A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, "Supervised learning of universal sentence representations from natural language inference data," arXiv preprint arXiv:1705.02364, 2017.
10. J. Choi, K. M. Yoo, and S.-g. Lee, "Learning to compose task-specific tree structures," in Thirty-Second AAAI Conference on Artificial Intelligence, 2018.
11. Z. Wang, W. Hamza, and R. Florian, "Bilateral multi-perspective matching for natural language sentences," arXiv preprint arXiv:1702.03814, 2017.
12. S. Peng, H. Cui, N. Xie, S. Li, J. Zhang, and X. Li, "Enhanced-rnn: an efficient method for learning sentence similarity," in WWW 2020.
13. S. Xu, E. Shijia, and Y. Xiang, "Enhanced attentive convolutional neural networks for sentence pair modeling," Expert Systems with Applications.
14. K. Cho, A. Courville, and Y. Bengio, "Describing multimedia content using attention-based encoder-decoder networks," IEEE Transactions on Multimedia.
15. Y. Liu, M. L. Di Liang, F. Giunchiglia, X. Li, S. Wang, W. Wu, L. Huang, X. Feng, and R. Guan, "Local and global: temporal question answering via information fusion," in Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, 2023, pp. 5141–5149.
16. He K, Zhang X, Ren S, et al. Deep residual learning for image recognition[C]//Proceedings of the IEEE conference on computer vision and pattern recognition. 2016: 770-778.
17. R. Yang, J. Zhang, X. Gao, F. Ji, and H. Chen, "Simple and effective text matching with richer alignment features," arXiv preprint arXiv:1908.00300, 2019.
18. J. Dong, M.-A. Rondeau, and W. L. Hamilton, "Distilling structured knowledge for text-based relational reasoning," in EMNLP 2020.
19. Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," Advances in neural information processing systems, vol. 32, 2019.
20. C. Xue, D. Liang, P. Wang, and J. Zhang, "Question calibration and multi-hop modeling for temporal question answering," in AAAI.
21. Y. Liu, C. Sun, L. Lin, and X. Wang, "Learning natural language inference using bidirectional lstm model and inner-attention," arXiv preprint arXiv:1605.09090.
22. L. Li, Q. Liao, M. Lai, D. Liang, and S. Liang, "Local and global: Text matching via syntax graph calibration," in ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).
23. J. Bai, Y. Wang, Y. Chen, Y. Yang, J. Bai, J. Yu, and Y. Tong, "Syntax-bert: Improving pre-trained transformers with syntax trees," arXiv preprint, 2021.
24. A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," arXiv preprint arXiv:1804.07461, 2018.
25. W. Lan, S. Qiu, H. He, and W. Xu, "A continuously growing dataset of sentential paraphrases," arXiv preprint arXiv:1708.00391, 2017.
26. T. Khot, A. Sabharwal, and P. Clark, "Scitail: A textual entailment dataset from science question answering," in AAAI 2018.
27. Vaswani, Ashish, et al. "Attention is all you need." Advances in neural information processing systems 30 (2017).
28. Y. Tay, L. A. Tuan, and S. C. Hui, "A compare-propagate architecture with alignment factorization for natural language inference," arXiv preprint arXiv:1801.00102

29. D. Liang, F. Zhang, W. Zhang, Q. Zhang, J. Fu, M. Peng, T. Gui, and X. Huang, "Adaptive multi-attention network incorporating answer information for duplicate question detection," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 95–104.
30. S. Wang, D. Liang, J. Song, Y. Li, and W. Wu, "DABERT: Dual attention enhanced BERT for semantic matching," in *Proceedings of the 29th International Conference on Computational Linguistics*, Oct. 2022.
31. C. Xue, D. Liang, S. Wang, J. Zhang, and W. Wu, "Dual path modeling for semantic matching by perceiving subtle conflicts," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023.
32. Y. Liu, D. Liang, F. Fang, S. Wang, W. Wu, and R. Jiang, "Time-aware multiway adaptive fusion network for temporal knowledge graph question answering," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
33. Z. Fei, Q. Zhang, T. Gui, D. Liang, S. Wang, W. Wu, and X.-J. Huang, "Cqg: A simple and effective controlled generation framework for multi-hop question generation," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 6896–6906.
34. Y. Liu, M. Li, D. Liang, X. Li, F. Giunchiglia, L. Huang, X. Feng, and R. Guan, "Resolving word vagueness with scenario-guided adapter for natural language inference," *arXiv preprint arXiv:2405.12434*, 2024.
35. R. Zheng, R. Bao, Y. Zhou, D. Liang, S. Wang, W. Wu, T. Gui, Q. Zhang, and X. Huang, "Robust lottery tickets for pre-trained language models," *arXiv preprint arXiv:2211.03013*, 2022.
36. R. Ma, Y. Tan, X. Zhou, X. Chen, D. Liang, S. Wang, W. Wu, T. Gui, and Q. Zhang, "Searching for optimal subword tokenization in cross-domain ner," *arXiv preprint arXiv:2206.03352*, 2022.
37. J. Song, D. Liang, R. Li, Y. Li, S. Wang, M. Peng, W. Wu, and Y. Yu, "Improving semantic matching through dependency-enhanced pre-trained model with adaptive fusion," in *Findings of the Association for Computational Linguistics: EMNLP 2022*. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, Dec. 2022, pp. 45–57. [Online].
38. T. Gui, Q. Zhang, J. Gong, M. Peng, D. Liang, K. Ding, and X.-J. Huang, "Transferring from formal newswire domain with hypernet for twitter pos tagging," in *Proceedings of the 2018 conference on empirical methods in natural language processing*, 2018, pp. 2540–2549.
39. X. Wu, J. Yang, T. Li, S. Zhang, Y. Du, L. Chai, D. Liang, and Z. Li, "Unleashing potential of evidence in knowledge-intensive dialogue generation," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
40. X. Wu, J. Yang, L. Chai, G. Zhang, J. Liu, X. Du, D. Liang, D. Shu, X. Cheng, T. Sun et al., "Tablebench: A comprehensive and complex benchmark for table question answering," *arXiv preprint arXiv:2408.09174*, 2024.
41. B. Li, D. Liang, and Z. Zhang, "Comateformer: Combined attention transformer for semantic sentence matching," *arXiv preprint arXiv:2412.07220*, 2024.
42. J. Chen and J. Liu, "S3prompt: Instructing the model with self-calibration, self-recall and self-aggregation to improve in-context learning," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024, pp. 14259–14271.