Durham Research Online

Deposited in DRO:

20 July 2021

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Gajbhiye, Amit and Al Moubayed, Noura and Bradley, Steven (2021) 'ExBERT: An External Knowledge Enhanced BERT for Natural Language Inference.', in Artificial Neural Networks and Machine Learning – ICANN 2021 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part V. , pp. 460-472. Theoretical Computer Science and General Issues., 12895

Further information on publisher's website:

https://doi.org/10.1007/978-3-030-86383-8_37

Publisher's copyright statement:

The final authenticated version is available online at https://doi.org/10.1007/978-3-030-86383-8_37

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

ExBERT: An External Knowledge Enhanced BERT for Natural Language Inference

Amit Gajbhiye¹, Noura Al Moubayed², and Steven Bradley²

¹ University of Sheffield, Sheffield, UK
a.gajbhiye@sheffield.ac.uk

² University of Durham, Durham, UK
{noura.al-moubayed, s.p.bradley}@durham.ac.uk

Abstract. Neural language representation models such as BERT, pre-trained on large-scale unstructured corpora lack explicit grounding to real-world commonsense knowledge and are often unable to remember facts required for reasoning and inference. Natural Language Inference (NLI) is a challenging reasoning task that relies on common human understanding of language and real-world commonsense knowledge. We introduce a new model for NLI called External Knowledge Enhanced BERT (ExBERT), to enrich the contextual representation with real-world commonsense knowledge from external knowledge sources and enhance BERT's language understanding and reasoning capabilities. ExBERT takes full advantage of contextual word representations obtained from BERT and employs them to retrieve relevant external knowledge from knowledge graphs and to encode the retrieved external knowledge. Our model adaptively incorporates the external knowledge context required for reasoning over the inputs. Extensive experiments on the challenging SciTail and SNLI benchmarks demonstrate the effectiveness of ExBERT: in comparison to the previous state-of-the-art, we obtain an accuracy of 95.9% on SciTail and 91.5% on SNLI.

Keywords: Natural Language Inference · Contextual Representations.

1 Introduction

Natural Language Inference (NLI), also known as Recognising Textual Entailment, is formulated as a - 'directional relationship between pairs of text expressions, denoted by T (the entailing "Text") and H (the entailed "Hypothesis"). Text T, entails hypothesis H, if humans reading T would typically infer that H is most likely true.' [4]. The NLI task definition relies on common human understanding of language and real-world commonsense knowledge. NLI is a complex reasoning task, and NLI models can not rely solely on training data to acquire all the real-world commonsense knowledge required for reasoning and inference [3]. For example, consider the premise-hypothesis pair in Table 1, if the training data do not provide the common knowledge that, (*wave RelatedTo crash*) and (*crash IsA hit*), it will be hard for the NLI model to correctly recognise that the premise entails the hypothesis.

2 A. Gajbhiye et al.

Table 1: SNLI premise (**P**) & hypothesis (**H**) and commonsense triples (red) from ConceptNet KG. Commonsense knowledge enrich premise-hypothesis context and helps the NLI model in reasoning.

P:	Four boys are about to be hit by an approaching wave. (wave RelatedTo crash)
H:	A giant wave is about to crash on some boys. (crash IsA hit)

Recently, deep pre-trained language representations models (PTLMs) such as BERT [6] achieved impressive performance improvements on a wide range of NLP tasks. These models are trained on large amounts of raw text using a self-supervised language modelling objective. Although pre-trained language representations have significantly improved the state-of-the-art on many complex natural language understanding tasks, they lack grounding to real-world knowledge and are often unable to remember real-world facts when required [14]. Investigations into the learning capabilities of PTLMs reveal that the models fail to recall facts learned at training time, and do not generalise to rare/unseen entities [17]. Knowledge probing tests [14] on the commonsense knowledge of ConceptNet [20] reveal that PTLMs such as BERT have critical limitations when solving problems involving commonsense knowledge. Hence, infusing external real-world commonsense knowledge can enhance the language understanding capabilities of PTLMs and the performance on the complex reasoning tasks such as NLI.

Incorporating external commonsense knowledge into pre-trained NLI models is challenging. The main challenges are (i) **Structured Knowledge Retrieval:** Given a premise-hypothesis pair how to effectively retrieve specific and relevant external knowledge from the massive amounts of data in Knowledge Graphs (KGs). Existing models [3], use heuristics and word surface forms of premises-hypothesis which may be biased and the retrieved knowledge may not be contextually relevant for reasoning over premise-hypothesis pair. (ii) **Encoding Retrieved Knowledge:** Learning the representations of the retrieved external knowledge amenable to be fused with deep contextual representations of premise-hypothesis is challenging. Various KG embedding techniques [22] are employed to learn these representations. However, while learning, the KG embeddings are required to be valid within the individual KG fact and hence might not be predictive enough for the downstream tasks [22] (iii) **Feature Fusion:** How to fuse the learned external knowledge features with the premise-hypothesis contextual embeddings. This feature fusion requires substantial NLI model adaptations with marginal performance gains [3].

To overcome the aforementioned challenges, we propose, **ExBERT** - an External knowledge enhanced BERT model which enhances the contextual representations of BERT model with external commonsense knowledge to improve BERT's real-world grounding and reinforce its reasoning and inference capabilities. ExBERT utilises BERT for learning the contextual representation of premise-hypothesis as well as the representations of retrieved external knowledge. The aim here is to take full advantage of contextual word representations obtained from pre-trained language models and the real-world commonsense knowledge from Knowledge Graphs (KGs).

The main contributions of this paper are: (i) we devise a new approach, ExBERT, to incorporate external knowledge in contextual word representations. (ii) we investigate and demonstrate the feasibility of using contextual word representation for encoding external knowledge obviating learning specialised KG embeddings. To the best of our knowledge, this is the first study of its kind, indicating a potential future research direction. (iii) we introduce a new external knowledge retrieval mechanism for NLI that is capable of retrieving fine-grained contextually relevant external knowledge from KGs.

2 Related Work

Traditional Attention-Based Models do not utilise contextual representations from PTLMs [7]. KG-Augmented Entailment System (KES) [11] augments the NLI model with external knowledge encoded using graph convolutional networks. Knowledge-based Inference Model (KIM) [3] incorporates lexical-level knowledge (such as synonym and antonym) into its attention and composition components. ConSeqNet [23], a system of a text-based model and a graph-based model, concatenates the output of the two models, to be fed to a classifier. The AdvEntuRe [10] framework train the decomposable attention model with adversarial training examples generated by incorporating knowledge from linguistic resources, and with a sequence-to-sequence neural generator. BiCAM models improve the performance of NLI baselines via the incorporation of knowledge from the ConcepNet and Aristo Tuple KGs by factorised bilinear pooling [9].

PTLM-Based Models PTLM-based models such as OpenAI GPT [19] and BERT [6] leverage transfer learning from a large textual corpus and are fine-tuned on NLI datasets. Specifically, OpenAI GPT pre-trains the Transformer [21] model in an unsupervised manner with the standard language modelling objective and fine-tunes the model in a supervised manner for NLI. Semantics-aware BERT (SemBERT) [25] demonstrated the benefit of enriching the BERT's contextual representation with the semantic roles. BERT model has shown to be robust to adversarial examples when external knowledge is incorporated to the attention mechanism using simple transformations [15]. The KES [11] model highlighted above further evaluates their system with BERT contextual embeddings in the framework.

3 Methodology

ExBERT architecture is depicted in Fig.1. In this section, we describe the key components of ExBERT and their detailed implementation including the model architecture in Section 3.2. We start by describing the contextual representation based external knowledge retrieval procedure in Section 3.1.

3.1 External Knowledge Retrieval: Selection and Ranking

Retrieval and preparation of contextually specific and relevant information from knowledge graphs are complex and challenging tasks. The crucial challenge is to

4 A. Gajbhiye et al.

Fig. 1: A high-level view of ExBERT architecture.

identify the knowledge specific and relevant to the task at hand from the massive amounts of noisy KG data [1]. Different from the previous approaches that use word surface forms to retrieve external knowledge, we use the cosine similarity between the contextual representations of premise-hypothesis tokens and the external knowledge. The external knowledge for the premise and hypothesis is retrieved individually. Below we explain the procedure for the premise. The same procedure is applied to the hypothesis. The output of external knowledge retrieval is the set of contextually relevant KG triples for the premise and hypothesis. We divide the external knowledge retrieval process into two parts: Selection and Ranking.

Selection We first filter the stop words from the premise. Then we retrieve all the KB triples that contain the tokens of the premise as one of the words in the head entity of KG triples. For example, for the token “*speaking*” one of the retrieved KG fact is “*public_speaking IsA speaking*”. The retrieved triples are converted to a sentence. For example, the previous triple is transformed into “*public speaking is a speaking*”. The selection process retrieves a large number of KG triples which are not all relevant to the context of the premise. We filter the selected triples in the ranking step.

Ranking step ranks the selected KG triples according to the contextual similarity to the fine-grained context of the premise. Specifically, given the BERT generated context-aware representation of the premise tokens, we group all the bigrams of the representations. Each group of the bigram representation is averaged, and the cosine similarity is calculated with the average of the BERT representation of each of the selected KG triple sentence (created in selection step). For each bigram, we choose the KG triple sentence with the highest cosine similarity score. To capture the fine-grained context of the premise, we repeat the ranking step with the trigrams, fourgrams, and the average of the whole premise BERT representations and retrieve the KG record with highest cosine

similarity for each of the grams. Duplicate KG triple sentences are removed from the final set of retrieved knowledge.

The advantages of our knowledge retrieval mechanism are that it is free from heuristic biases, requires no feature engineering, and the retrieved external knowledge is contextually relevant to the fine-grained context of the premise and hypothesis.

3.2 Model Architecture

BERT Encoding Layer This layer uses the BERT encoder to learn the context-aware representations of premise-hypothesis and the set of retrieved external knowledge.

Specifically, given premise $P = \{p_i\}_{i=1}^n$, hypothesis $H = \{h_j\}_{j=1}^m$, and the set of external knowledge $EXT = \{\{e_r\}_{r=1}^l\}^k$, where r is the number of tokens in the external knowledge and k is the number of retrieved external knowledge. For encoding the premise and hypothesis, we input P and H to BERT in the following form

$$S = [\langle \text{CLS} \rangle, P, \langle \text{SEP} \rangle, H, \langle \text{SEP} \rangle] \quad (1)$$

$$H = \text{BERT}(S) \in \mathbb{R}^{(n+m+3) \times h} \quad (2)$$

where $\langle \text{SEP} \rangle$ is the token separating P and H , $\langle \text{CLS} \rangle$ is the classification token, and h is the dimension of the hidden state. When the BERT model is fine-tuned for the downstream task, the fine-tuned hidden state vector ($\mathbf{h}_{cls}$) corresponding to the classification token is used as the aggregate representation for the sequence. For each of the external knowledge in the set EXT , we generate the context-aware representations using the same BERT encoder as used for premise-hypothesis above as follows

$$EXT^k = [\langle \text{CLS} \rangle, e_1, \dots, e_l, \langle \text{SEP} \rangle] \quad (3)$$

$$E^k = \text{BERT}(EXT^k) \in \mathbb{R}^{l+2 \times h} \quad (4)$$

$$\mathbf{e}^k = \text{MeanPooling}(E^k) \in \mathbb{R}^{1 \times h} \quad (5)$$

The averaged context-aware vector representation ($\mathbf{e}$) generated for each of the (k) retrieved external knowledge are stacked to create the context-aware matrix, $E \in \mathbb{R}^{k \times h}$.

Knowledge Integration Layer This layer integrates external knowledge into the premise-hypothesis contextual representations by means of multi-head dot product attention. The layer uses a mixture model to allow a better trade-off between the context from external knowledge and the premise-hypothesis context [24]. The mixture model learns two parameter matrices, A and B , weighing the importance of premise-hypothesis context and the external knowledge context.

Multi-head Attentions To measure the importance of external knowledge to each context-aware premise-hypothesis representation, we apply multi-head dot product attention [21] between the context-aware representations of external knowledge and that of premise-hypothesis.

6 A. Gajbhiye et al.

In multi-head dot product attention, the context-aware representations are projected linearly to generate the queries, keys and values. As we use the multi-head attention to highlight the external knowledge important to premise-hypothesis context, premise-hypothesis representation (H) generates the query matrix (H^q) via linear projection and the two linear projections of external knowledge representation (E) generate the keys and values. The attention function is defined as

$$\text{Attention}(H^q, E^k, E^v) = \text{softmax}\left(\frac{H^q E^{kT}}{\sqrt{h_k}}\right) E^v \quad (6)$$

Then the multi-head attention is

$$\begin{aligned} C_{ph}^{ext} &= \text{MH}(H^q, E^k, E^v) \\ &= \text{Concat}(\text{head}_1, \dots, \text{head}_h) W^o \end{aligned} \quad (7)$$

where $\text{head}_i = \text{Attention}(H^q W_i^q, E^k W_i^k, E^v W_i^v)$ and W_i^q, W_i^k, W_i^v , and W^o are projection matrices and i is the number of attention heads (12 in our case). The output of multi-head attention, $C_{ph}^{ext} \in \mathbb{R}^{(m+n+3) \times h}$ is an attention-weighted context matrix measuring the importance of the external knowledge context to each of the context-aware premise-hypothesis representation.

Similarly, to measure the importance of each premise-hypothesis BERT representation (H) to the aggregate premise-hypothesis representation (hidden representation $\mathbf{h}_{cls}$ corresponding to CLS token), we apply the multi-head attention between $\mathbf{h}_{cls}$ token representation and H as

$$\text{Attention}(C_{cls}^q, H^k, H^v) = \text{softmax}\left(\frac{C_{cls}^q H^{kT}}{\sqrt{h_k}}\right) H^v \quad (8)$$

where $C_{cls}^q \in \mathbb{R}^{(m+n+3) \times h}$ is a matrix obtained by repeating $\mathbf{h}_{cls}$ hidden state $(n + m + 3)$ number of times. The multi-head attention is calculated similar to (Eq. 7) that outputs a context matrix $C_{ph}^{cls} \in \mathbb{R}^{(m+n+3) \times h}$. The output matrix C_{ph}^{cls} is an attention-weighted context matrix measuring the importance of each of the premise-hypothesis representation to the aggregate premise-hypothesis representation.

Mixture Model The mixture model learns a trade-off between the premise-hypothesis context and the context from external knowledge and is defined as

$$M = A C_{ph}^{ext} + B C_{ph}^{cls} \quad (9)$$

where A and B are the parameter metrics, learned with a single layer neural network and $A + B = J \in \mathbb{R}^{(n+m+3) \times 1}$, J is a matrix of all ones. The parameters A and B learn to balance the proportion of incorporating the premise-hypothesis context and the context from external knowledge. Each of the representations in $M \in \mathbb{R}^{(m+n+3) \times h}$ can be regarded as a knowledge aware state representation that encodes external knowledge context information with respect to the context of each of the premise-hypothesis representation.

Composition Layer We compose the knowledge state representation (M) to the corresponding premise-hypothesis representation to obtain knowledge-aware matrix $\hat{H}$ as

$$\hat{H} = H + M \quad (10)$$

Pooling Layer The pooling layer creates a fixed-length representation from premise-hypothesis representations H and the knowledge-aware representations $\hat{H}$. We apply the standard mean and max pooling mechanisms as

$$\mathbf{h}_{mean} = \text{MeanPooling}(H) \quad \mathbf{h}_{max} = \text{MaxPooling}(H) \quad (11)$$

$$\hat{\mathbf{h}}_{mean} = \text{MeanPooling}(\hat{H}) \quad \hat{\mathbf{h}}_{max} = \text{MaxPooling}(\hat{H}) \quad (12)$$

Classification Layer We classify the relationship between premise and hypothesis using a Multilayer Perceptron (MLP) classifier. The input to the MLP is the concatenation of pooled representations as

$$\mathbf{f} = [\mathbf{h}_{mean}; \hat{\mathbf{h}}_{mean}; \mathbf{h}_{max}; \hat{\mathbf{h}}_{max}] \quad (13)$$

The MLP consists of two hidden layers with tanh activation and a softmax output layer to obtain the probability distribution for each class. The network is trained in an end-to-end manner using multi-class cross-entropy loss.

4 Experiments

4.1 Datasets

NLI & KGs The key contribution of this paper is the unique method to incorporate external knowledge into the pre-trained BERT representations. ExBERT is capable of incorporating knowledge from any external knowledge source that allows the knowledge to be retrieved, given an entity as input. This includes KBs with (*head, relation, tail*) graph structure, KBs that contain only entity metadata without a graph structure and those that combine both a graph and entity metadata.

In this work, we retrieve external commonsense knowledge from ConceptNet [20] for evaluating ExBERT on SNLI (570,000 examples) [2] and SciTail (27,000 examples)[12] benchmarks, and from the science domain-targeted KG, Aristo Tuple [5] for evaluation on science domain SciTail dataset.

ConceptNet is a multilingual KG comprising of 83 languages. We pre-process the ConceptNet data to retrieve the facts with head and tail entities in the English language. The final pre-processed ConceptNet that we retrieve the external knowledge from contains 3,098,816 ($\approx 3M$) commonsense facts connected by 47 relations. Aristo Tuple is an English language KG that contains 294,000 science domain facts connected with 955 unique relations. We search the whole Aristo Tuple KG to retrieve relevant external knowledge.

8 A. Gajbhiye et al.

4.2 Experimental Setup

Following our external knowledge retrieval mechanism discussed in Section 3.1, we first retrieve the external knowledge from ConceptNet and Aristo Tuple for SNLI and SciTail datasets via selection and ranking steps. In the ranking step, the English uncased BERT_{BASE} [6] model is employed in feature extraction mode (i.e. without fine-tuning) to learn the contextual representations of the premise, the hypothesis and to each of the selected KG triple sentences. We then use the retrieved external knowledge to train the following three versions of ExBERT.

Models We used the English uncased BERT_{BASE} to train three versions of ExBERT: Two ExBERT+ConceptNet models on SNLI and SciTail respectively and one ExBERT+AristoTuple model on SciTail. The models utilise the external knowledge from the KG their name is suffixed.

Training Details ExBERT is implemented in PyTorch using the base implementation of BERT³. The underlying BERT is initialised with the pre-trained BERT parameters and follows the same fine-tuning procedure as the original BERT. During training, the pre-trained BERT parameters are fine-tuned with the other ExBERT parameters. We use the Adam optimiser [13] with the initial learning rate fine-tuned from {8e-6, 2e-5, 3e-5, 5e-5} and warm-up rate of 0.1. The batch size is selected from {16, 24, 32}. The maximum number of epochs is chosen from {2, 3, 4, 5}. Dropout ratio of 0.5 is used at the classification layer [8]. Texts are tokenised using word pieces, with a maximum length of 40 for SNLI, 60 for SciTail, and 15 for external knowledge. The hyper-parameters are fine-tuned on the dev set of each NLI dataset.

5 Results

The results of top-performing models on the SNLI⁴ and SciTail⁵ leaderboards are summarised in Table 2. For fairness of comparison, we compare ExBERT with only the PTLMs based NLI models that leverage external knowledge. On **SNLI**, the performance of the state-of-the-art models is highly competitive.

We observe that ExBERT outperforms all the existing baselines on the SNLI dataset and pushing the benchmark to 91.5% within the models using BERT_{BASE} as the base model. ExBERT achieves a maximum performance improvement of 1.9% over the previous state-of-the-art BERT_{BASE} + SRL [26] baseline.

Among the models built on BERT_{LARGE} with more than 340M million parameters [6], our ExBERT⁶ with BERT_{BASE} (110M parameter) remarkably outperforms the BERT_{LARGE} and BERT_{LARGE} + SRL [26] models with the absolute improvements of 0.5% and 0.2% respectively, and is able to match the performance of SemBERT_{LARGE} [25] and MT – DNN_{LARGE} [16] models.

³ <https://github.com/huggingface/transformers>

⁴ <https://nlp.stanford.edu/projects/snli/>

⁵ <https://leaderboard.allenai.org/scitail/submissions/public>

⁶ We expect further improvements in ExBERT’s performance with BERT_{LARGE}, however we left the evaluation for future work due to the limited computing resources.

Table 2: Results on SNLI and SciTail dataset: For SNLI, ExBERT uses ConceptNet KG. For SciTail ExBERT uses ConceptNet KG and AristoTuple KGs.

SNLI Dataset		SciTail Dataset	
BERT _{BASE} as Base Model		BERT _{BASE} as Base Model	
NLI Model	Test Acc(%)	NLI Model	Test Acc(%)
BERT BERT _{BASE} + SRL [26]	89.6	OpenAI GPT [19]	88.3
OpenAI GPT [19]	89.9	BERT _{BASE} [16]	92.5
BERT _{BASE} [15]	90.5	BERT+LF [18]	92.8
BERT _{BASE} [16]	90.8	MT – DNN _{BASE} [16]	94.1
BERT+LF [18]	90.5	MT-DNN+LF[18]	94.3
SemBERT _{BASE} [25]	91.0	BERT _{LARGE} as Base Model	
MT – DNN _{BASE} [16]	91.1	BERT _{LARGE} [16]	94.4
MT-DNN+LF [18]	91.1	MT – DNN _{LARGE} [16]	95.0
BERT _{LARGE} as Base Model		ExBERT+ConceptNet (Ours)	95.2
BERT _{LARGE} [16]	91.0	ExBERT+AristoTuple (Ours)	95.9
BERT _{LARGE} + SRL [26]	91.3		
SemBERT _{LARGE} [25]	91.6		
MT – DNN _{LARGE} [16]	91.6		
ExBERT+ConceptNet (Ours)	91.5		

On **SciTail** (Table 2), ExBERT outperforms all the existing models including the models built on BERT_{LARGE} model. Our best performing model, ExBERT+AristoTuple demonstrates an absolute improvement of 7.6% over the established baseline of OpenAI GPT [19]. Moreover, using only BERT_{BASE} as the underlying model, our ExBERT+AristoTuple outperforms BERT_{LARGE} based MT – DNN_{LARGE} [16] model by 0.9%.

We observe higher performance improvements on the smaller SciTail dataset which demonstrates that incorporating external knowledge helps the model with small training data. Further, we observe that ExBERT attains higher accuracy when external knowledge is incorporated from the science domain-specific KG, Aristo Tuple as compared to when external knowledge is added from the commonsense KG, ConceptNet. The specialised scientific knowledge in Aristo Tuple is more beneficial to SciTail.

6 Analysis

6.1 Number of External Features

To investigate the effect of incorporating various numbers of external knowledge features, we vary the number of KG triple sentences input to ExBERT. Particularly, we are interested in answering the question: How many commonsense features are required for the optimal model performance? Figure 2 illustrates the results of the experiment. **For SNLI**, ExBERT achieves the highest accuracy (91.5%) using 11 external knowledge sentences. We observe a decrease in accuracy when increasing the number of external knowledge sentences after 11. The fewer number of external knowledge sentences required, compared to SciTail dataset, to achieve the maximum accuracy on SNLI dataset, is attributed to the limited linguistic and semantic variation and the short average length of stop-word filtered premise (7.35 for entailment and neutral class) and hypothesis

10 A. Gajbhiye et al.

(3.61 for entailment and 4.45 for neutral class) [12] of the SNLI dataset, which limits its ability to fully extract and exploit external KG knowledge.

Fig. 2: ExBERT accuracy with varying amount of external knowledge.

For SciTail, ExBERT when evaluated using the general commonsense knowledge source ConceptNet, requires a relatively high number of external knowledge sentences (15) to achieve the maximum accuracy. This is due to the higher syntactic and semantic complexity of SciTail, which needs more knowledge to reason. However, when evaluated with the domain-specific Aristo Tuple KG, the model achieves the highest accuracy with fewer (7) external knowledge sentences. To reiterate, domain specific knowledge in Aristo Tuple improves the model performance with less external knowledge.

6.2 Qualitative Analysis

Case Study This section provides the case study of different premise-hypothesis pairs and the corresponding external knowledge, to vividly show the effectiveness of ExBERT in adaptively identifying the relevant features from the supplied external knowledge. Recall that given a context-aware representation of premise-hypothesis token, the relevance of the retrieved external knowledge in E is measured by the multi-head attention defined in Eq.(6). We average the attention weights of all heads and plot a heat map.

Fig.3 presents the heat map showing the attention of premise-hypothesis tokens to the retrieved external knowledge sentences from ConceptNet. In Fig.3(a), we can see, these attention distribution is quite meaningful, with the “*speaking*” and “*talking*” attending mainly to the retrieved external knowledge “*speaking is talking*”. Similarly, the tokens “*speaking*” “*talking*” and “*man*” attends to “*talking is a human activity*”. In Fig.3(b) among the other attentions, the most prominent can be observed between the tokens “*performing for cash*” and the external knowledge sentence “*performing used for earning*”.

Attending to the relevant external knowledge demonstrates the ExBERT’s ability to effectively utilise the retrieved external knowledge based on the context from the premise and hypothesis.

Fig. 3: Case Study. Visualisation of ExBERT's attention between external knowledge from ConceptNet (y axis) and SNLI premise-hypothesis pair tokens (x axis).

7 Conclusion

We introduced ExBERT to enrich the contextual representation of BERT with real-world commonsense knowledge from external knowledge sources and to enhance its language understanding and reasoning capabilities. ExBERT can incorporate external knowledge from any external knowledge source that allows the knowledge to be retrieved, given an entity. We devised a novel external knowledge retrieval mechanism utilising contextual representations to retrieve relevant external knowledge. Experimental results on SNLI and SciTail NLI benchmarks in conjunction with two KGs, ConceptNet and Aristo Tuple, shows that ExBERT achieves significant performance improvements over the previous state-of-the-art methods, including those which are enhanced by BERT_{LARGE}. Further, we demonstrated the feasibility of utilising contextual representations for encoding external knowledge from KGs, which indicates a potential direction for future research.

References

1. Bast, H., Björn, B., Haussmann, E.: Semantic search on text and knowledge bases. *Found. Trends Inf. Retr.* **10**(2–3), 119–271 (Jun 2016)
2. Bowman, S.R., Angeli, G., Potts, C., Manning, C.D.: A large annotated corpus for learning natural language inference. In: *ACL* (2015)
3. Chen, Q., Zhu, X., Ling, Z.H., Inkpen, D., Wei, S.: Neural natural language inference models enhanced with external knowledge. In: *ACL* (2018)
4. Dagan, I., Glickman, O., Magnini, B.: The pascal recognising textual entailment challenge. pp. 177–190. Springer Berlin Heidelberg (2006)
5. Dalvi, M.B., Tandon, N., Clark, P.: Domain-targeted, high precision knowledge extraction. *Transactions of the Association for Computational Linguistics* **5**, 233–246 (2017), <https://www.transacl.org/ojs/index.php/tacl/article/view/1064>
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the NAACL-HLT 2019, Volume 1 (Long and Short Papers)*. pp. 4171–4186 (2019)
7. Gajbhiye, A., Jaf, S., Moubayed, N.A., Bradley, S., McGough, A.S.: Cam: A combined attention model for natural language inference. In: *2018 IEEE International Conference on Big Data (Big Data)*. pp. 1009–1014 (Dec 2018)

12 A. Gajbhiye et al.

8. Gajbhiye, A., Jaf, S., Moubayed, N.A., McGough, A.S., Bradley, S.: An exploration of dropout with rnns for natural language inference. In: ICANN 2018. pp. 157–167. Springer, Cham (2018)
9. Gajbhiye, A., Winterbottom, T., Al Moubayed, N., Bradley, S.: Bilinear fusion of commonsense knowledge with attention-based nli models. In: ICANN 2020. pp. 633–646. Springer, Cham (2020)
10. Kang, D., Khot, T., Sabharwal, A., Hovy, E.: AdvEntuRe: Adversarial training for textual entailment with knowledge-guided examples. In: ACL. Melbourne, Australia (Jul 2018)
11. Kapanipathi, P., Thost, V., Patel, S.S., Whitehead, S., Abdelaziz, I., Balakrishnan, A., Chang, M., Fadnis, K., Gunasekara, C., Makni, B., et al.: Infusing knowledge into the textual entailment task using graph convolutional networks. arXiv preprint arXiv:1911.02060 (2019)
12. Khot, T., Sabharwal, A., Clark, P.: Scitail: A textual entailment dataset from science question answering. In: AAAI, New Orleans, Louisiana, USA, February 2–7, 2018 (2018)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) ICLR 2015
14. Kwon, S., Kang, C., Han, J., Choi, J.: Why do masked neural language models still need common sense knowledge? CoRR **abs/1911.03024** (2019)
15. Li, A.H., Sethy, A.: Knowledge enhanced attention for robust natural language inference. arXiv preprint arXiv:1909.00102 (2019)
16. Liu, X., He, P., Chen, W., Gao, J.: Multi-task deep neural networks for natural language understanding. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 4487–4496. Florence, Italy (Jul 2019)
17. Logan, R., Liu, N.F., Peters, M.E., Gardner, M., Singh, S.: Barack’s wife hillary: Using knowledge graphs for fact-aware language modeling. In: Proceedings of the 57th ACL. pp. 5962–5971. Florence, Italy (Jul 2019)
18. Pang, D., Lin, L.H., Smith, N.A.: Improving natural language inference with a pretrained parser. arXiv preprint arXiv:1909.08217 (2019)
19. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training (2018)
20. Speer, R., Chin, J., Havasi, C.: Conceptnet 5.5: An open multilingual graph of general knowledge (2017)
21. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. pp. 6000–6010 (2017)
22. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: A survey of approaches and applications. IEEE Transactions on Knowledge and Data Engineering **29**(12), 2724–2743 (Dec 2017). <https://doi.org/10.1109/TKDE.2017.2754499>
23. Wang, X., Kapanipathi, P., Musa, R., Yu, M., Talamadupula, K., Abdelaziz, I., Chang, M., Fokoue, A., Makni, B., Mattei, N., et al.: Improving natural language inference using external knowledge in the science questions domain. In: Proceedings of the AAAI. vol. 33, pp. 7208–7215 (2019)
24. Yang, B., Mitchell, T.: Leveraging knowledge bases in LSTMs for improving machine reading. In: ACL. pp. 1436–1446. Vancouver, Canada (Jul 2017)
25. Zhang, Z., Wu, Y.W., Hai, Z., Li, Z., Zhang, S., Zhou, X., Zhou, X.: Semantics-aware bert for language understanding. ArXiv **abs/1909.02209** (2020)
26. Zhang, Z., Wu, Y., Li, Z., Zhao, H.: Explicit contextual semantics for text comprehension. CoRR **abs/1809.02794** (2018), <http://arxiv.org/abs/1809.02794>