

CK-Transformer: Commonsense Knowledge Enhanced Transformers for Referring Expression Comprehension

Zhi Zhang

ILLC, University of Amsterdam
z.zhang@uva.nl

Helen Yannakoudakis

Dept. of Informatics, King’s College London
helen.yannakoudakis@kcl.ac.uk

Xiantong Zhen

United Imaging
zhenxt@gmail.com

Ekaterina Shutova

ILLC, University of Amsterdam
e.shutova@uva.nl

Abstract

The task of multimodal referring expression comprehension (REC), aiming at localizing an image region described by a natural language expression, has recently received increasing attention within the research community. In this paper, we specifically focus on referring expression comprehension with commonsense knowledge (KB-Ref), a task which typically requires reasoning beyond spatial, visual or semantic information. We propose a novel framework for Commonsense Knowledge Enhanced Transformers (CK-Transformer) which effectively integrates commonsense knowledge into the representations of objects in an image, facilitating identification of the target objects referred to by the expressions. We conduct extensive experiments on several benchmarks for the task of KB-Ref. Our results show that the proposed CK-Transformer achieves a new state of the art, with an absolute improvement of 3.14% accuracy over the existing state of the art¹.

1 Introduction

Referring expression comprehension (REC) aims at locating a target object/region in an image given a natural language expression as input. The nature of the task requires multi-modal reasoning and joint visual and language understanding. In the past few years, several REC tasks and datasets have been proposed, such as RefCOCO (Yu et al., 2016), RefCOCog (Mao et al., 2016) and RefCOCO+ (Yu et al., 2016) (RefCOCOs). These ‘conventional’ REC tasks typically focus on identifying an object based on visual or spatial information of the object, such as its colour, shape, location, etc.; therefore primarily evaluating a model’s reasoning abilities over visual attributes and spatial relationships.

In practice, however, people often describe an object using non-visual or spatial information – consider, for example, the sentence (expression) “Give

¹The code will be available in <https://github.com/FightingFighting/CK-Transformer>

me something soft but rich in starch to eat” (Wang et al., 2020). Such instances require reasoning beyond spatial and visual attributes, and need to be interpreted with respect to the common sense knowledge (fact) embedded in the expressions, such as knowledge about which kind of objects are edible, soft and rich in starch in the given image. Recently, Wang et al. (2020) proposed a new dataset, KB-Ref, to evaluate the reasoning ability of a model over not only visual and spatial features but also commonsense knowledge. The dataset is devised such that at least one piece of fact from a knowledge base (KB) is required for a target object (referred to by an expression) to be identified.

Therefore, searching for appropriate facts from a KB is also crucial part in KB-Ref. In contrast to the only existing work (Wang et al., 2020), in which for each object candidate, the top-K facts with the highest cosine similarity between the averaged Word2Vec (Mikolov et al., 2013) embedding of the fact and the given expression are maintained, our framework focuses on multi-modal embedding and reasoning simultaneously over both the expression and the image to identify the top-K facts. Multi-modal features encode richer information helping to improve reasoning over varying (semantic) contexts and identification of relevant facts; for example, the above example of expression can be answered with the object “banana” in an image (or, equivalently, with the object “mashed potato” in another image).

In this paper, we propose a novel multi-modal framework for KB-Ref – Commonsense Knowledge Enhanced Transformers (CK-Transformer, CK-T for short) – that integrates (top-K) facts into all object candidates in an image for better identification of the target object. Specifically, our contributions are four-fold: 1) We propose the CK-T (see Figure 1) that effectively integrates diverse input from different modalities: vision, referring expressions and facts; 2) To the best of our knowledge,

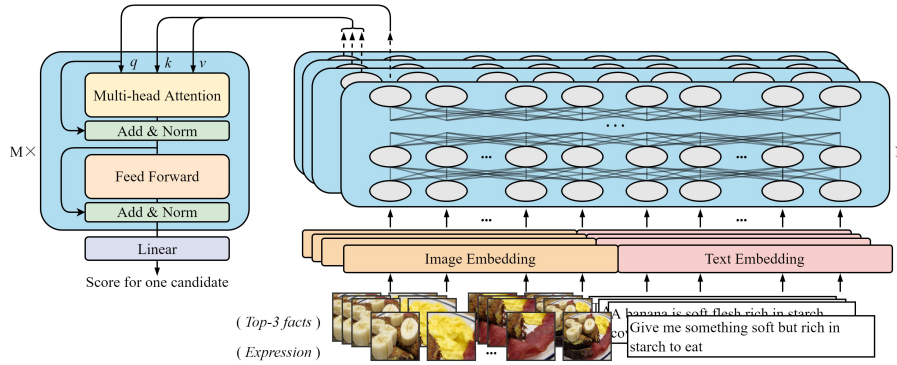


Figure 1: CK-Transformer. For each candidate (the first one in the figure), given an expression, a set of visual region candidates and top-K facts ($K=3$ in the figure), the model first encodes the expression and all top-K facts into corresponding multi-modal features, then fuses these features and maps them into a matching score for the candidate.

our approach is the first that introduces visual information into the identification of (top-K) relevant facts; 3) Our approach achieves a new state of the art using only top-3 facts per (candidate) object, which is furthermore substantially more efficient compared to existing work utilizing as much as top-50 facts; 4) We introduce facts into ‘conventional’ REC tasks, leading to improved performance.

2 Related Work

Referring expression comprehension with commonsense knowledge Different from conventional REC tasks (see Appendix A for details), KB-Ref focuses on querying objects given an expression that requires commonsense knowledge reasoning. The authors benchmarked a baseline model, ECIFA, for integration of facts, expression and image, and selects the target object by comparing the match scores between the image features and corresponding top-K fact features for all object candidates in the image (Wang et al., 2020). In our framework, we select top-K facts for each candidate by comparing the cosine similarity between the fact and expression embedding, where the embeddings are generated from a multi-modal encoder rather than a text encoder used in the ECIFA model.

Pre-trained vision–language encoders Several pre-trained multi-modal encoders (Su et al., 2019; Li et al., 2019; Chen et al., 2020; Tan and Bansal, 2019) have been proposed, achieving state-of-the-art results on vision–language tasks. Currently, UNITER (Chen et al., 2020) as one of powerful pre-trained encoders achieves the best performance on REC tasks (RefCOCOs). In this paper, we adapt UNITER such that it is used as a multi-modal en-

coder in fact search and as part of the CK-T.

3 Methodology

We formulate KB-Ref as a classification problem based on an image I consisting of a set of candidates (image regions) $I = \{c_j\}_{j=1}^n$ obtained from either ground-truth labels or predictions of a pre-trained object detector. Specifically, given an expression e , an image I and a KB, we first search for top-K facts $F_i^K = \{f_j\}_{j=1}^k$ from the KB for each candidate c_i , and then feed e , I , and F_i^K (the selected facts over I) into our CK-T simultaneously to predict the target object over all candidates.

3.1 Image-based fact search

For each candidate c_i in a given image, we retrieve all the facts from the KB (see Appendix D for details on the KB used in our framework) according to its category (e.g., a candidate object may belong to category ‘car’). Then, we calculate the cosine similarity between the facts and the given expression, where the similarities are obtained from a similarity extractor which we train by adapting UNITER. Specifically, given image–expression and image–fact pairs as input, we extract expression and fact features respectively from the position of the cross-modality output of UNITER (corresponding to the input of [CLS] token, see Appendix B for details), and then calculate the cosine similarity between the two. During training, inspired by Devlin et al. (2018), we replace 50% of ground-truth facts with random facts from the KB (with a similarity of 0), to help the model better distinguish useful facts from non-useful ones. Finally, we maintain top-K facts F_i^K with higher similarities to the expression

for each candidate c_i .

3.2 Commonsense Knowledge Enhanced Transformer

The CK-T consists of a bi-modal encoder (see 3.2.1) and a fact-aware classifier (see 3.2.2).

3.2.1 Bi-modal encoder

The bi-modal encoder (initialized by UNITER-base with N=12 layers (Chen et al., 2020)) integrates two modalities: image and text (e or f_i). Specifically, after generating the input embedding E_{Inp} consisting of image and text embedding (same with UNITER, see Appendix C for details), for each candidate c_i , we extract the expression-aware and fact-aware object features respectively (f_i^e and f_i^f) from the position of the visual output corresponding to c_i in the same encoder, based on the input of all candidates I , and e or f_i .

3.2.2 Fact-aware classifier

The fact-aware classifier is composed of multi-head attention layers and fully connected layers. For each candidate c_i , f_i^e and F_i^f (all K fact-aware object features for c_i) are fed into the integrator simultaneously (*Key* and *Value* are from F_i^f , and *Query* is from f_i^e), and fused into one three-source object features f_i^t (image, expression and top-K facts).

Finally, f_i^t is mapped into a match score s_i for c_i by a linear layer, and the optimization objective is to minimize the cross-entropy loss over all scores $\{s_j\}_{j=1}^n$ corresponding to all candidates I .

4 Results

We compare our CK-T to existing approaches on KB-Ref task without and with facts. Then we explore the importance of introducing visual information for fact search. Furthermore, we introduce facts into the traditional RefCOCO dataset, which was collected from MSCOCO (Lin et al., 2014) but differs in the types of expressions and object candidate settings. We extract image region features using an off-the-shelf detector, Faster R-CNN with ResNet-101 (Ren et al., 2015), based on bounding boxes (bbxes) (ground-truth labels or predicted results from the detector). See Appendix D and E for details about these datasets and experiment setting². Through parameter search on K and M (see Figure 3 and 4 in Appendix F), we keep M = 2 Fact-aware classifier blocks and top-3 facts for each candidate.

Model	Accuracy (%)	
	Val	Test
CMN (Hu et al., 2017)	41.28	40.03
SLR (Yu et al., 2017)	44.03	42.92
VC (Niu et al., 2019)	44.63	43.59
LGARNs (Wang et al., 2019)	45.11	44.27
MAttNet (Yu et al., 2018)	46.86	46.03
ECIFA-nf (Wang et al., 2020)	37.95	35.16
CK-T-nf (Ours)	58.02	57.53
ECIFA (Wang et al., 2020)	59.45	58.97
MAtt+E (Wang et al., 2020)	64.08	63.57
CK-T-Word2Vec	60.40	61.39
CK-T-Uw/oImage	64.44	64.78
CK-T (Ours)	65.62	66.71
Human	–	90.31

Table 1: Accuracy on KB-Ref dataset without and with facts (top and bottom part, respectively) using ground-truth bounding boxes and object categories.

Ground-truth bounding boxes and categories

By following Wang et al. (2020), we report our results on KB-Ref without and with facts. As can be seen in Table 1 (top), CK-T-nf, a version of CK-T without facts³, achieves an accuracy of 57.53% on the test set, outperforming existing approaches that do not utilize facts by approximately 11% – 22%. At the bottom part of the table we can see that our fact-enhanced CK-T model achieves the highest accuracy (66.71%) on the test set, which is 7.74% higher than that of ECIFA (a baseline model proposed by Wang et al. (2020)), and 3.14% higher than MAtt+E⁴. It is worth noting that both ECIFA and MAtt+E incorporate the top-50 facts for each candidate, which is considerably higher compared to top-3 facts in our CK-T. We surmise this is due to the fact that our fact search approach utilizes multi-modal fact and expression embeddings.

Predicted bounding boxes and categories To facilitate a fair comparison with ECIFA-d (Wang et al., 2020), we also use the maximum 10 detected bboxes for each image (CK-T-m10). As can be seen in Table 2, CK-T-m10 achieves an accuracy which is $\approx 5\%$ higher than that of ECIFA-d on the test set. CK-T-m100, a variant using at most 100 detected bboxes achieves a substantial improvement

²Including the efficiency discussion about our model

³all word tokens in fact sentences are replaced with only one [MASK] token.

⁴Wang et al. (2020) introduces their facts fusion module –Episodic Memory Module (E)–into MAttNet model (MAtt) (Yu et al., 2018) widely used for conventional REC.

Model	Accuracy (%)	
	Val	Test
ECIFA-d (Wang et al., 2020)	24.11	23.82
CK-T-m10 (Ours)	28.33	28.71
CK-T-m100 (Ours)	35.66	35.96

Table 2: Accuracy on KB-Ref using predicted bbxes and object categories.

with $\approx 7\%$, compared with CK-T-m10. This difference is primarily due to the increase in the number of correctly detected bbxes and predicted categories. Specifically, we find that with the top-100 bbxes, the number of samples containing the target bbxes rises from 18,901 to 31,653, while among these target bbxes, the number of correctly predicted categories grows from 11,324 to 15,928, out of a total of 43,284 samples in the KB-Ref dataset. This can also explain the dramatic decline on the accuracy between CK-T and CK-T-m10.

Incorporating image features into fact search

We experiment with various approaches to fact search and evaluate their effectiveness on KB-Ref (Table 1). We first utilize top-k facts searched in (Wang et al., 2020), where they use a pre-trained Word2Vec (Skip-Gram) (Mikolov et al., 2013) for searching facts (CK-T-Word2Vec). Then, we also selected facts from similarity predictors based on only text as input (CK-T-Uw/oImage)⁵, instead of image-text pairs in CK-T. As shown in Table 1, both CK-T-Uw/oImage and CK-T achieve better accuracy on the test set compared to CK-T-Word2Vec. Compared to CK-T-Uw/oImage, CK-T achieves around 2% higher accuracy. This is primarily due to the additional visual information used during fact search (see Appendix I for the examples of the selected facts by these fact search methods).

Introducing facts in traditional REC tasks

We incorporate facts from the KB used in KB-Ref into the tasks of RefCOCOs using CK-T. Table 3 shows the results comparison based on the ground-truth bbxes and categories (discussion about the detected results can be seen in Appendix G). Compared with U_{REC} ⁶, the model introducing facts achieves better or equal accuracy on all RefCOCOs tasks, where RefCOCOg is improved more than RefCOCO and

⁵Inspired by Frank et al. (2021), we replace all object candidate feature with the average of all image region features.

⁶Chen et al. (2020) achieve state-of-the-art results on RefCOCOs by finetuning UNITER. (We re-finetune the model for fair comparison and conduct McNemar Test)

Task	Accuracy (%)		
	U_{REC}	Intro Facts	
Ref-COCO	Val	90.98	91.43
	Test A	91.50	92.09
	Test B	90.89	90.95
Ref-COCO+	Val	83.23	83.45
	Test A	85.09	85.49
	Test B	79.08	79.08
Ref-COCOg	Val	86.23	87.21
	Test	85.79	87.59

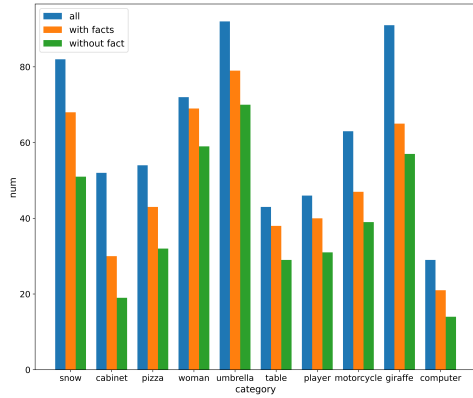
Table 3: Introducing facts into RefCOCO, RefCOCO+ and RefCOCOg. RefCOCO and RefCOCO+ have two different test sets, Test A and Test B, containing multiple persons and multiple objects in images respectively.

RefCOCO+. This is because RefCOCOg has less same-category object candidates in an image compared to RefCOCO and RefCOCO+ (an average of 1.63 and 3.9 per image, respectively) (Yu et al., 2016), and thus the retrieved facts integrated into different candidates are diversified (we first retrieve facts using the category), which contributes to distinguishing between candidates. This difference can also be proved in McNemar Test, where we find the change in the proportion of errors is statistically significant after introducing facts as compared to before on RefCOCOg ($p\text{-value} = 1.19e-08 < \alpha = 0.05$), while the similar proportions are found on RefCOCO and RefCOCO+ (see Appendix H for details about McNemar Test). The overall impact of commonsense knowledge in traditional REC is, however, not substantial. This is primarily due to much smaller number (78) of categories among the candidates in RefCOCOs, compared to 1805 in the KB-Ref (Wang et al., 2020). This limits the variety of selected facts, therefore impacting the extent to which commonsense knowledge is useful.

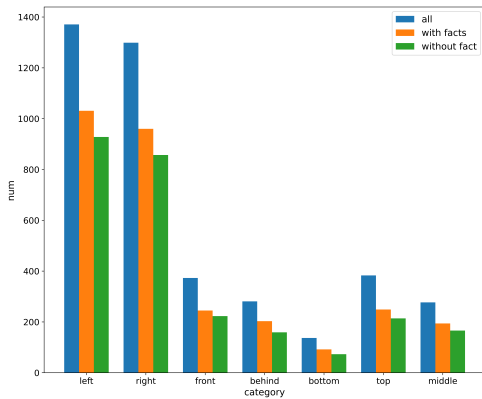
5 Analysis

To investigate in what cases commonsense knowledge helps, we conduct a fine-grained analysis of model performance on the test set of KB-Ref. Specifically, we compare the samples predicted by model with and without facts (CK-T and CK-T-nf) on three aspects: object categories, spatial relationships and the size of the bounding box.

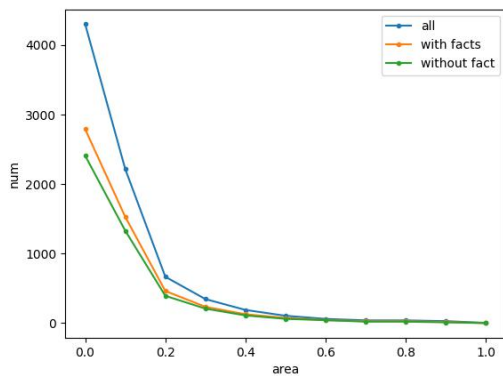
Object categories The test set contains 1502 categories and CK-T outperforms CK-T-nf on 1347 categories. Top 10 categories for which most improvement is observed are shown in Figure 2(a).



(a) Top 10 categories showing most improvement after introducing facts.



(b) The analysis of spatial relationships.



(c) The analysis of different bounding box sizes.

Figure 2: Fine-grained analysis. *all*: the total number of samples in the test set; *with fact*: the number of test samples that CK-T predicts correctly; *without fact*: the number of test samples that CK-T-nf predicts correctly.

In case of the 155 categories that do not show improvement, we find that the average number of samples per category is 6.68, making the results less reliable.

Spatial relationships We then investigate to what extent solving the REC task with and without facts relies on spatial reasoning, and whether there are particular spatial relationships between objects for which the use of facts is most crucial. Similar to the works of (Kazemzadeh et al., 2014; Johnson et al., 2017), we focus on the following spatial relationships: *left*, *right*, *front*, *behind*, *bottom*, *top*, *middle*. As shown in Figure 2(b), the model with facts (CK-T) outperforms that without facts (CK-T-nf) on all spatial relationships.

The size of the bounding box We then investigate the role of facts when identifying objects of different sizes, using the size of their bounding box as a proxy. We use the normalized area of the bounding box as the metric of bboxes size. As shown in Figure 2(c), the facts improve model performance on all bounding box sizes.

6 Conclusion

In this paper, we proposed CK-Transformer, which effectively integrates commonsense knowledge and the expression into the representations of the corresponding visual objects for multi-modal reasoning on KB-Ref. Our CK-Transformer achieves a new state-of-the-art performance on KB-Ref using only top-3 most relevant facts. We also demonstrated that visual information is beneficial for fact search. Finally, we show that commonsense knowledge improves conventional REC tasks across three different datasets.

7 Limitations

The computational requirements of our model are affected by the number of facts. Specifically, we train our CK-Transformer for 10000 steps with a batch size of 64 on one Titan RTX GPU, which takes 2.5, 3, 3.5, 7 days with the number of facts: 3, 5, 10, 20 respectively. The CK-Transformer processes 3.8, 2.8, 2.1, 0.7 samples on average per second at training time and 8.3, 7.3, 6.6, 1.1 samples per second at test time, with these amounts of facts. The computational requirements of our models are thus substantial, and future work should consider improving computational efficiency and thus reducing environmental impact.

References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Stella Frank, Emanuele Bugliarello, and Desmond Elliott. 2021. Vision-and-language or vision-for-language? on cross-modal influence in multimodal transformers. *arXiv preprint arXiv:2109.04448*.
- Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2017. Modeling relationships in referential expressions with compositional modular networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1115–1124.
- Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. 2016. Natural language object retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4555–4564.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2016. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Yulei Niu, Hanwang Zhang, Zhiwu Lu, and Shih-Fu Chang. 2019. Variational context: Exploiting visual and textual context for grounding referring expressions. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):347–359.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24:1143–1151.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypertexted, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- Niket Tandon, Gerard De Melo, and Gerhard Weikum. 2017. Webchild 2.0: Fine-grained commonsense knowledge distillation. In *Proceedings of ACL 2017, System Demonstrations*, pages 115–120.

- Peng Wang, Dongyang Liu, Hui Li, and Qi Wu. 2020. Give me something to eat: Referring expression comprehension with commonsense knowledge. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 28–36.
- Peng Wang, Qi Wu, Jiewei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. 2019. Neighbourhood watch: Referring expression comprehension via language-guided graph attention networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1960–1968.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. 2018. Mattnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer.
- Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L Berg. 2017. A joint speaker-listener-reinforcer model for referring expressions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7282–7290.
- Hanwang Zhang, Yulei Niu, and Shih-Fu Chang. 2018. Grounding referring expressions in images by variational context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4158–4166.

A Referring expression comprehension

Early approaches to REC use joint embedding of image and language by combination of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), and predict the target object that has the maximum probability given an input expression and an image (Mao et al., 2016; Hu et al., 2016; Zhang et al., 2018). In order to model different types of information encoded in input expression (subject appearance, location, and relationship to other objects), subsequent work used modular (attention) networks, to “match” the input to corresponding regions in the image, predicting as the target the region with the highest matched score (Hu et al., 2017; Yu et al., 2018).

B UNITER

UNITER is trained using four pre-training tasks, Masked Language Modeling (MLM), Masked Region Modeling (MRM), Image–Text Matching (ITM), and Word–Region Alignment (WRA), on four large-scale image–text datasets, COCO (Lin et al., 2014), Visual Genome (Krishna et al., 2016), Conceptual Captions (Sharma et al., 2018), and SBU Captions (Ordonez et al., 2011). This enables UNITER to capture fine-grained alignments between images and language. The architecture of UNITER is similar to BERT (Devlin et al., 2018) apart from the input and the output. Specifically, the input consists of an image (a set of visual region candidates), a sentence and [CLS] token, and they respectively lead to different outputs, i.e. vision output, language output and cross-modality output on the top of UNITER.

C Input embedding

Same with UNITER, we extract the input embeddings E_{Inp} consisting of an image and a text embedding corresponding to the object candidate I and text (an expression e or a fact f_i) respectively.

Image embedding The image embedding E_I is computed by summing three types of embeddings: visual feature embedding, visual geometry embedding and modality segment embedding. We first extract the visual features $V = \{v_1, v_2, \dots, v_n\}$ for all candidates using Faster R-CNN (pooled RoI features), and build a geometry feature $G = \{g_1, g_2, \dots, g_n\}$ for all candidates, where g_i is a 7-dimensional vector consisting of the geometry information of the bounding box corresponding to

candidate c_i , namely normalized top, left, bottom, right coordinates, width, height, and area, denoted by $g_i = [x1, y1, x2, y2, w, h, w * h]$. Visual feature embeddings and visual geometry embeddings are generated by mapping the visual features and the geometry features into the same vector space through a fully connection layer fc :

$$E_I = LN(fc(V) + fc(G) + M_I) \quad (1)$$

where LN is the layer normalization layer and M_I is the modality segment embedding for the image input (like segment embedding for two sentence in BERT model).

Text embedding Similarly, the text embedding E_T is computed based on three different types of embeddings: token embedding, position embedding and modality embedding (Normally there is a fourth embedding, sentence segment embedding similarly to BERT, but, in our task, both expressions and facts consist of one sentence only and so only the first sentence segment embedding is used). Similar to BERT (Devlin et al., 2018), the text $W = \{w_1, w_2, \dots, w_u\}$ is first tokenized by WordPieces (Wu et al., 2016), which are then built into token embeddings $T = \{t_1, t_2, \dots, t_v\}$ and position embeddings $P = \{p_1, p_2, \dots, p_v\}$ according to their position in the text sequence.

$$E_T = LN(T + P + M_T) \quad (2)$$

where M_T is the modality segment embedding for the text input.

Input embedding The final input embedding E_{Inp} is computed by concatenating image embedding E_I and text embedding E_T :

$$E_{Inp} = [E_I, E_T] \quad (3)$$

D Datasets

We use the KB-Ref dataset (Wang et al., 2020) aiming at evaluating the task of referring expression comprehension based on commonsense knowledge. KB-Ref consists of 43,284 expressions for 1,805 object categories on 16,917 images, as well as a knowledge base of key–value (category–fact) pairs collected from three common knowledge resources: Wikipedia, ConceptNet (Speer et al., 2017) and WebChild (Tandon et al., 2017)). KB-Ref is split into a training set (31,284 expressions with 9,925 images), a validation set (4,000 expressions with

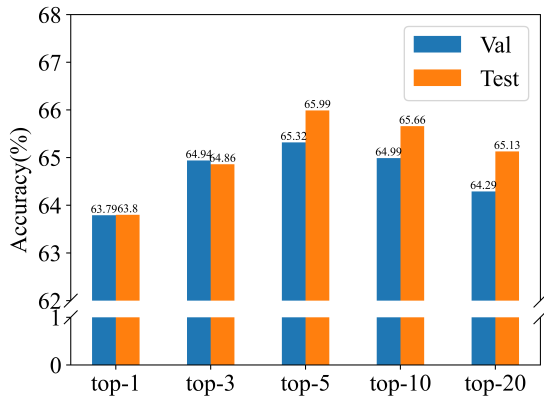


Figure 3: Accuracy across a varying number of facts (top-K).

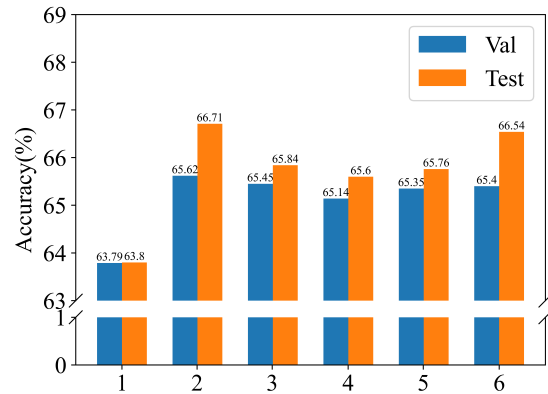


Figure 4: Accuracy across a varying number of fact-aware classifier block (M).

2,290 images) and a test set (8,000 expressions with 4,702 images).

We furthermore introduce commonsense knowledge into traditional tasks/datasets of referring expression comprehension, namely RefCOCO, RefCOCOG and RefCOCO+⁷. The datasets are devised from the MSCOCO image dataset (Lin et al., 2014) but differ in the types of expressions and object candidate settings. Specifically, RefCOCO+ does not allow the use of absolute location words in the expressions, and most expressions focus on the appearance of the objects. The expressions in RefCOCOG are longer and contain more descriptive words. RefCOCO and RefCOCO+ contain more objects of the same category within an image.

E Experimental settings

We extract image region features using Faster R-CNN with ResNet-101 (Ren et al., 2015) which was pre-trained on Visual Genome (Krishna et al., 2016) using object and attribute annotations (Anderson et al., 2018). For bounding box detection, we keep the bounding boxes with at least 0.2 confidence score indicating the extent of detection. In the CK-T, the hidden layer dimension is 768 and the number of multi-head attention heads is 12. The models are trained using Adamw (Loshchilov and Hutter, 2017) with a learning rate of $6e^{-5}$ and a batch size of 64 on Titan RTX GPUs. Our CK-Transformer has 120M parameters in total where fact-aware classifier has 34M and bi-modal encoder has 86M. As for UNITER model, we use same setting with *UNITER-base*, except for using Nvidia

⁷following Apache License 2.0

Apex⁸ for speeding up training. The efficiency of our model is effected by the number of facts. Specifically, we train our CK-Transformer 10000 steps and a batch per step, which takes 2.5, 3, 3.5, 7 days with the number of facts: 3, 5, 10, 20 respectively. The CK-Transformer trains 3.8, 2.8, 2.1, 0.7 sample in average per second and tests 8.3, 7.3, 6.6, 1.1 sample per second.

F Impact of CK-T structure

We explore the impact in performance on KB-Ref as we vary the number of top-K facts (K) and fact-aware classifier block (M) on the development set. We first keep the number of the fact-aware classifier block constant and set it to 1 to experiment with different values for K from 1 to 20. As shown in Figure 3, as K increases, performance starts to improve with a peak at K=5 before starting to gradually decrease performance.

In the second experiment, we keep K constant and set it to 3 and explore the effect of varying values for M. We observe that the highest accuracy is achieved with with top-3 facts and 2 integrator layers as shown in Figure 4.

G Introducing facts in traditional REC tasks based on detection

The results of introducing facts in traditional REC tasks based on detected bbxes and categories are shown in Table 4. Compared to result based on ground-truth bbxes and categories (Table 3), the improvement on models based on detection is less or even worse than the models without facts.

⁸<https://github.com/NVIDIA/apex>

Task		Accuracy (%)	
		U_{REC}	Intro Facts
Ref-COCO	Val ^d	81.15	81.06
	Test A ^d	86.85	86.87
	Test B ^d	74.48	73.97
Ref-COCO+	Val ^d	74.74	74.68
	Test A ^d	81.05	80.70
	Test B ^d	65.88	66.07
Ref-COCOG	Val ^d	74.49	74.69
	Test ^d	75.24	74.86

Table 4: Introducing facts into RefCOCO, RefCOCO+ and RefCOCOG based on detection (d).

Task		McNemar Test
		(p -value)
RefCOCO	Test A	0.049
	Test B	0.905
RefCOCO+	Test A	0.297
	Test B	0.966
RefCOCOG	Test	$1.19e-08$

Table 5: The McNemar Test between models before and after introducing facts on the tasks of RefCOCOs.

H McNemar Test

We also report the statistical significance for accuracy (shown in Table 3) on the tasks of RefCOCOs. Specifically, we conduct the McNemar Test between models before and after introducing facts, on the test set of RefCOCO, RefCOCO+ and RefCOCOG, respectively. As shown in Table 5, as for Test set on RefCOCOG and Test A on RefCOCO p -value = $1.19e-08$ and p -value = 0.049 (< 0.05) respectively, which means the proportion of errors is statistically significantly different after introducing facts as compared to before. However, the change in the proportion of errors after introducing facts on other tasks (Test B on RefCOCO, Test A and Test B on RefCOCO+) is not statistically significant. This is reasonable, as the error from detection will affect the fact search (we first retrieve facts using the category) and thus more error information is introduced into CK-Transformer, which make the performance worse.

I Example searched fact using different methods

As shown in Figure 5, there are several facts which are selected from three different fact search methods: CK-Transformer, CK-T-Uw/oImage and CK-

T-Word2Vec. As we can see in the Table, normally the facts of CK-Transformer model (green) is the best relevant with the referring expression (blue) and the facts in CK-T-Word2Vec model is the worst relevant with the expression.



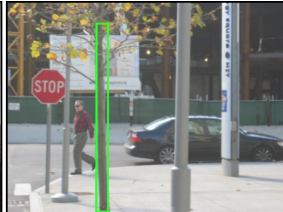

			
Exp: The tool under the mouse can improve the usability of the mouse	Exp: The tall building has instruments on the upper exterior walls that used as a reference to find out the time.	Exp: The most important part of the tree with branches and leaves	Exp: An access point as an underground public utility.
Fact: A mousepad enhances the usability of the mouse compared to using a mouse directly on a table.	Fact: Clock towers are a specific type of building which houses a turret clock and has one or more clock faces on the upper exterior walls.	Fact: The trunk is the most important part of the tree for timber production.	Fact: Manholes are often used as an access point for an underground public utility, allowing inspection, maintenance, and system upgrades.
Fact (Uw/oImage): A mousepad is a surface for placing and moving a computer mouse.	Fact (Uw/oImage): The tower has four clock faces, two of which are in diameter, at about high.	Fact (Uw/oImage): An automobile has a trunk.	Fact (Uw/oImage): A manhole is an opening to a confined space such as a shaft, utility vault, or large vessel.
Fact (Word2Vec): Mousepad on the mouse.	Fact (Word2Vec): Before the middle of the twentieth century, most people did not have watches, and prior to the 18th century even home clocks were rare.	Fact (Word2Vec): A trunk is an example of a box.	Fact (Word2Vec): These covers are traditionally made of metal, but may be constructed from precast concrete, glass reinforced plastic or other composite materials.

Figure 5: Example fact search process (using the top-1 fact) for different search methods: CK-T (green), CK-T-Uw/oImage (orange) and CK-T-Word2Vec (yellow).