

VLSP 2021 - VieCap4H Challenge: Automatic Image Caption Generation for Healthcare Domain in Vietnamese

Thao Minh Le
Deakin University
thao.le@deakin.edu.au

Long Hoang Dang
Deakin University
hldang@deakin.edu.au

Thanh-Son Nguyen
A*STAR, Singapore
nguyen_thanh_son@ihpc.a-star.edu.sg

Thi Minh Huyen Nguyen
Hanoi University of Science
huyenntm@hus.edu.vn

Xuan-Son Vu
Umeå University
sonvx@cs.umu.se

Abstract

This paper presents VieCap4H, a grand data challenge on automatic image caption generation for the healthcare domain in Vietnamese. VieCap4H is held as part of the eighth annual workshop on Vietnamese Language and Speech Processing (VLSP 2021). The task is considered as an image captioning task. Given a static image, mostly about healthcare-related scenarios, participants are asked to design machine learning methods to generate natural language captions in Vietnamese to describe the visual content of the image. We introduce VieCap4H, a novel human-annotated image captioning dataset in Vietnamese that contains over 10,000 image-caption pairs collected from real-world scenarios in the healthcare domain. All the models proposed by the challenge participants are evaluated using BLEU scores against groundtruths. The challenge was run on AIHUB.VN platform. Within less than two months, the challenge has attracted over 90 individual participants and recorded more than 900 valid submissions.

1 Introduction

Humans are unique in their capability to interpret and describe their visual perception in natural language. In the last decade, we have witnessed ground-breaking success of modern AI powered by deep learning in different tasks in visual understanding (He et al., 2016; Ren et al., 2015; Xu et al., 2017; Gordo et al., 2017) and natural language processing (Lample et al., 2016; Devlin et al., 2019; Vaswani et al., 2017; Zhang et al., 2019).

However, building a machine that learns to talk about what it sees remains very challenging. In this playground, *image captioning*, a machine learning task that automatically generates natural language descriptions of a given image, has emerged and attracted enormous attention in the AI research community (Hossain et al., 2019; You et al., 2016; Yao




Figure 1: An example of the image captioning task in VieCap4H data challenge. This image can be described in words in two different ways: (1) describing partial information of the image; and (1+2) describing overall information of the image.

et al., 2017; Vinyals et al., 2016; Aneja et al., 2018; Rennie et al., 2017; Anderson et al., 2018; Feng et al., 2019; Vu et al., 2020). The task is fascinating and yet challenging at the same time as it sits on the bridge between computer vision and natural language processing, the two most important fields of modern AI.

The COVID-19 pandemic has exacerbated the ongoing shortage of healthcare workers globally, posing an urgent need for smart assistants that can effectively cooperate with humans to fill the gap. More specifically, these agents require a capability to understand the healthcare domain and summarize the visual information captured by a digital camera to produce accurate descriptions and analyses to humans. While there have been multiple image captioning datasets available in prominent languages such as in English (Lin et al., 2014; Young et al., 2014), Chinese (Wu et al., 2017), datasets with Vietnamese captions are scarce. In this work, we introduce a so-called image captioning dataset VieCap4H to challenge machine learning models in their ability to generate descriptions in Vietnamese for visual content, mainly in the healthcare domain. Figure 1 shows an example of the image captioning task in VieCap4H data challenge - i.e., an image is associated with at least one caption in Vietnamese.

In some cases, the same image can be described in multiple ways to assess the generalizability of different solutions.

Along with the introduction of the dataset, we also run a challenge held as part of the eighth annual workshop on Vietnamese Language and Speech Processing (VLSP 2021). The challenge was hosted on *aihub.vn*¹ which consisted of two phases, including *public test* and *private test* during September 20 to October 25, 2021. The challenge recorded over 900 valid submitted entries in total by 14 teams of nearly 90 participants.

The contribution of this work is two folds: (1) we introduce a reliable dataset for image captioning in healthcare specifically for the Vietnamese research community as a result of a conscientious annotation process, providing a fair benchmark for interested researchers to train and evaluate their models (2) we run a challenge with an automatic evaluation framework to encourage participants to make use of the provided dataset and contribute their knowledge to advance the field, making potential applications of the task in either healthcare settings and general settings (e.g. virtual assistants for blind and visually impaired people, or visual content indexing and searching) accessible for the local community.

The remaining sections of this paper are organized as follows. The next section will detail the data collection and annotation process. Section 3 will discuss different approaches in solving the task during the challenge. We will discuss the results in detail in Section 4. Finally, we conclude the paper with some discussions on future directions and challenges of the topic.

2 The VieCap4H Dataset

2.1 Data collection

As we focus on images reported in news, we employed the following steps for constructing the dataset. First, we select a set of articles collected in the GDEL project². An article is selected if its url contains the substring */health/*. This is to ensure that the articles are within the healthcare domain. We then downloaded all the selected articles and extracted their images using the services provided by Diffbot³. Lastly, the images are downloaded and manually examined by two undergraduate stu-

dents. The images that are not truly related to the healthcare domain are discarded.

A group of undergraduate students were then recruited to write the caption for each of the remaining images. The students were given with detailed instructions. Finally, all their captions are revised by senior researchers. This last revision helps to correct all the linguistic annotation mistakes and also to improve the consistency among the captions generated by different students. Figure 2 summarizes our data collection and data annotation process.

2.2 Dataset splits

We carefully design a splitting method to form different sets (i.e., training, public test, and private test sets) from the collected data to provide a reliable evaluation of model performance. A good data splitting method should divide the dataset into splits in which each of them covers all usual and unusual patterns in the challenge domain. In addition, as one image can inherently associate with multiple captions and vice versa, the splitting method needs to avoid identical images and captions appear across the splits.

In VieCap4H, the K-mean clustering technique is applied to distribute the collected data into three different splits including training, public testing, and private testing splits. To be specific, we firstly use Fasttext-Sent2Vec (Pagliardini et al., 2018) with the pre-trained model provided in ETNLP (Vu et al., 2019) to extract the features of ground truth captions. Since an image can include multiple captions, we take mean of the extracted features of all the captions associated with the same image to return a single feature vector per image. These textual features are then served as the input to the K-mean algorithm with $K = 50$ to cluster the dataset into 50 distinguished non-overlapping groups. Figure 3 presents word clouds showing the most frequent words in representative clusters of the ground-truth captions.

For each cluster, we take 80%, 10% and 10% of the total images and their associated captions for the train, public test and private test split, respectively. Finally, we join the respective subsets over all the clusters together to obtain the train, public test and private test. The VieCap4H splits are made available for public use⁴. Interested researchers are required to sign up for an *aihub.vn* account to

¹<https://aihub.vn/competitions/40>

²<https://www.gdelproject.org/>

³<https://www.diffbot.com/>

⁴<https://aihub.vn/competitions/40#participate>

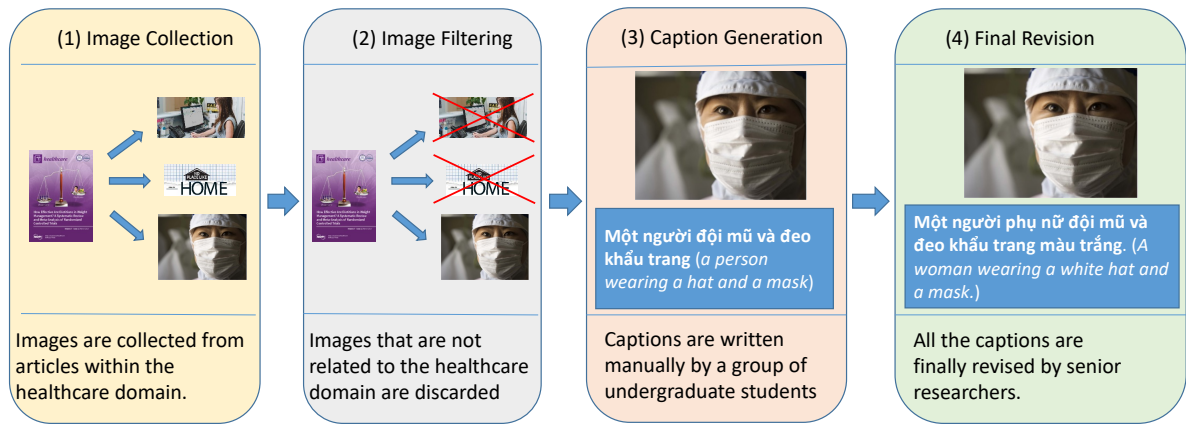


Figure 2: Overall pipeline for data collection and data annotation process. It consists of four steps: (1) Image Collection, (2) Image Filtering, (3) Caption Generation, and (4) Final Revision.



Figure 3: Word clouds showing the most frequent words in representative clusters of the ground-truth captions.

access the dataset.

During the two phases of the challenge, only the visual part of the test splits is accessible to participants while ground-truth captions are used at inference time for evaluation purposes only.

2.2.1 Data statistics

We provide in Table 1 the statistics of the VieCap4H dataset. The *train split* contains over 8,000 images with over 9,000 associated captions, while the *public test split* and *private test split* have approximately 1,000 images and ground-truth captions each. All splits have captions with similar average length of approximately 12 words per caption.

	Images	Captions	Avg. captions per image	Avg. caption's length (tokens)
Train	8032	9429	1.17	11.88
Public test	1002	1039	1.04	11.86
Private test	1034	1095	1.05	11.97
All	10068	11563	1.15	11.89

Table 1: Data statistics of the VieCap4H dataset.

Method Team	Encoder- decoder	Unified V-L Pretraining
tiendv		✓
gpt-team		✓
coder_phuho	✓	
caodoanhuit	✓	
vingovan	✓	
NguyenNghia	✓	

Table 2: Methods used by top teams in the final round of the VieCap4H challenge.

3 The VieCap4H Challenge

3.1 Methods

During the challenge, participants have utilized different network architectures to learn to generate textual captions for visual content. These methods can be broadly classified into two categories (See Table 2): *Encoder-decoder framework* based methods and *Unified vision-language pretraining* based methods. As deep neural models severely suffer from data hungry, all the participants leveraged transfer knowledge by adapting different pre-trained image captioning models and fine-tuning them on the provided data to deal with the problem of small training data provided by the challenge.

To be fair between all the participants, we asked all the teams to register with the organizers beforehand to use these existing pre-trained models. Table 3 summarizes all the pre-trained models that are used by the challenge’s participants.

3.1.1 Methods using encoder-decoder framework

A large body of works proposed for image captioning task utilizes an encoder-decoder framework. Specifically, the encoder is often a CNN based model in early works (Vinyals et al., 2015) or a transformer based model in more recent works (Herdade et al., 2019), which takes as input a given image and returns a vector representation of the image. The visual encoding is then served as input of the decoder, which is often a recurrent neural network, to generate the output caption. These models predict one word at a time using the visual encoding and a ground-truth caption during training, while inputs are the visual encoding and words that are generated in previous time steps during inference time. Figure 4 illustrates a typical encoder-decoder framework for image captioning.

Four out of six teams in the final round of the

VieCap4H challenge made use of the encoder-decoder framework for their solutions with different choices for network models for the encoder and decoder. These teams have tried different settings with different network architectures for visual encoding, ranging from RoI pooling region features extracted by Faster R-CNN (Ren et al., 2015) to more recent Transformer-based models such as Swin-Transformer (Liu et al., 2021) and Vision Transformer (Dosovitskiy et al., 2021). Regarding the decoder, all the teams either used LSTM coupled with attention mechanism or Transformer (Vaswani et al., 2017).

3.1.2 Fine-tuning unified vision-language pre-training models

A number of unified vision and language pre-training models (Zhou et al., 2020; Chen et al., 2019; Lu et al., 2019) have been proposed in the last few years to solve multiple vision and language tasks with a single model by a simple fine-tuning process. These models are usually trained with a large amount of image-text data using self-supervised learning techniques. They achieved superior performance on various tasks including Visual Question Answering, Image-Text Retrieval and Image Captioning.

In the scope of our challenge, the UIT AI team, who achieved the highest performance on both public and private leaderboards, made use of the VLP model proposed by Zhou et. al. (Zhou et al., 2020). Meanwhile, the GPT-team, who is the runner-up of the competition, fine-tuned a variant of the GPT model (Le, 2021) to solve both image captioning and machine translation at once. Thanks to the superb generalization capabilities of these pre-trained models and the access to larger training data via data augmentation and external data, these teams outperformed other teams in the challenge utilizing more traditional approaches by a large margin.

3.2 Training strategies

As the performance of deep neural networks is largely reliant on the amount of training data, multiple teams of the challenge have applied different augmentation techniques to diversify the training data in order to better facilitate the learning of their models.

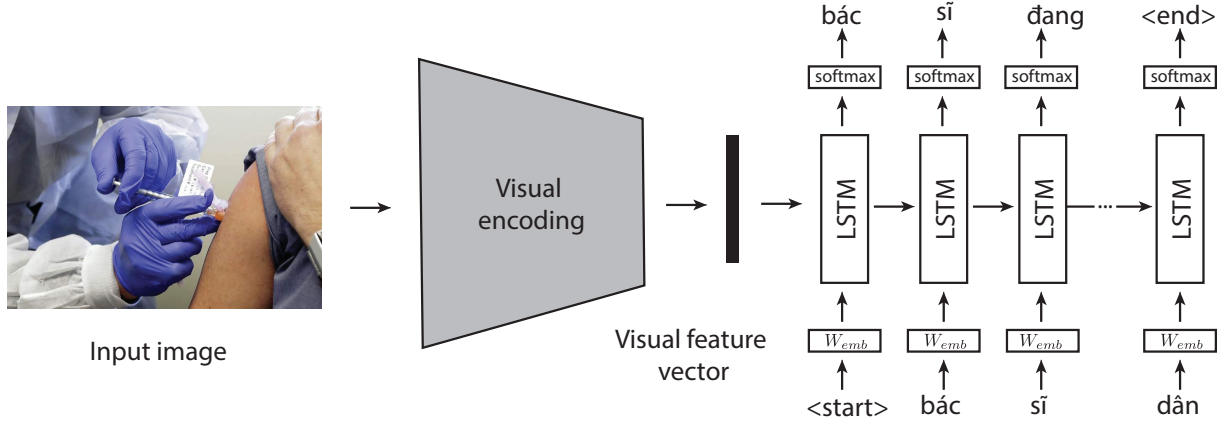


Figure 4: Illustration of a typical encoder-decoder framework for image captioning task. Example is taken from the VieCap4H dataset.

Model	Language	Vision	Description
Faster R-CNN (Ren et al., 2015)		✓	Trained on Visual Genome (Krishna et al., 2016)/COCO (Lin et al., 2014)
ResNet-152 (He et al., 2016)		✓	Trained on ImageNet (Deng et al., 2009)
Swin-Transformer (Liu et al., 2021)		✓	Trained on ImageNet-1K (Deng et al., 2009)
ViT Transformer (Dosovitskiy et al., 2021)		✓	Trained on ImageNet (Deng et al., 2009)
EfficientNet (Tan and Le, 2019)		✓	Trained on ImageNet (Deng et al., 2009)
GPT2News (Le, 2021)	✓		https://huggingface.co/imthanhlv/gpt2news
BERT (Devlin et al., 2019)	✓		Trained on large scale text dataset extracted from BooksCorpus (Zhu et al., 2015) and English Wikipedia
PhoBERT (Nguyen and Nguyen, 2020)	✓		Trained on 20GB texts obtained from Vietnamese news and Vietnamese Wikipedia

Table 3: Pre-trained models used by all the participants of VieCap4H challenge.

3.2.1 Data augmentation

The UIT AI team used NLP tools such as Underthesea⁵ and PyVi library⁶ to partly change the linguistic structure of the annotated captions while keeping the semantic meaning of the sentences unchanged. They used a pre-trained PhoBERT (Nguyen and Nguyen, 2020) to ensure the similarity of the augmented captions and the original captions. The Fruit AI team, who is ranked third in the challenge, instead injected noise to each of the annotated captions by replacing some characters in the captions with random characters and assigned the augmented caption-image pairs with *fake* labels.

Regarding image augmentation, most of the participants used standard image augmentation techniques such as *horizontal flip*, *random crop* to obtain more training data. Compared to text augmentation, image augmentation is more common and convenient to do. The UIT AI team additionally made use of external images crawled from

external sources by matching the semantic of the images with keywords in the VieCap4H annotated captions.

3.2.2 Network training

The Fruit AI team trained a network model that performs multi-task learning to predict the next word in a caption while being able to identify corrupted sentences with *fake* labels. The GPT team, on the other hand, opted for a multi-task vision-and-language model where they leveraged the Clip-Cap model (Mokady et al., 2021) to perform image captioning and language translation tasks at the same time. These two teams ranked 3rd and 2nd in the competition, respectively. All other teams formulated image captioning task as a classical classification task and trained their models with cross-entropy loss. By having access to a much larger amount of training data, the UIT AI team considerably outperformed all the other teams in the challenge in terms of quantitative results.

⁵<https://github.com/undertheseanlp/underthesea>

⁶<https://github.com/trungtv/pyvi>

4 Results

4.1 Data format

The VieCap4H dataset includes images and their annotated captions split into three subsets, including train split, public test split and private split. The textual annotations are provided in JSON files in the following format:

```
[{"id": "uuid_img1", "captions": "caption1",
  "id": "uuid_img2", "captions": "caption2"},
```

where “uuid_img1” refers to the unique index (ID) of a specific image in the dataset and “caption1” denotes the corresponding annotated caption by human annotators as described in Section 2.

4.2 Submission format for evaluation

Participants are required to submit a JSON file that includes the IDs of testing images and their corresponding predicted captions. Note that the image IDs in the submission and the provided sample file must be in the same order. The submission format is as follows:

```
[{"id": "uuid_img1", "captions": "caption1",
  "id": "uuid_img2", "captions": "caption2"}].
```

4.3 Evaluation metrics

Referring to (Chen et al., 2015), we use the BiLingual Evaluation Understudy (BLEU), which is a common metric for machine translation tasks, to compute the n -gram co-occurrence between ground truth captions and generated captions. Let $G = \left\{ \{g_{i,t}\}_{t=1}^T \right\}_{i=1}^D$ and $P = \{p_i\}_{i=1}^D$ denote all the ground truth captions $g_{i,t}$ and corresponding generated captions p_i of all the images in a test set, where the index i refers to the i -th image in the corpus, D is the size of the test set and T is the number of ground truth captions of the image i .

Given K_n is the number of selected n -gram for a specific value of n included in a caption, the n -gram precision scores at the corpus-level are calculated by:

$$\text{PS}_n(P, G) = \frac{\sum_{i=1}^D \sum_{j=1}^{K_n} \min \left(\max_{t \in T} (C_j(g_{i,t})), C_j(p_i) \right)}{\sum_{i=1}^D \sum_{j=1}^{K_n} C_j(p_i)}. \quad (1)$$

	Value
# individual participants	90
# valid teams	8
# selected teams	7
# submitted technical reports	6

Table 4: The competition participation summary.

Here $C_j(g_{i,t})$ and $C_j(p_i)$ are functions to count the number of n -gram $_j$ in the ground truth and generated captions. To overcome the weakness of the n -gram precision measure in which it can receive a high score if the generated caption is a substring of the ground truth caption, the PS_n score calculated above is penalized by the brevity penalty defined by:

$$\text{BP}(P, G) = \begin{cases} e^{1 - \frac{L_P}{L_G}} & L_P \leq L_G \\ 1 & L_P > L_G \end{cases}, \quad (2)$$

where L_P is the sum of the length of the generated captions and L_G is the sum of the effective length of the ground truth captions. In case multiple ground-truth captions of an image are available, a representative ground-truth caption is chosen as the one that has closest length to the generated caption. As a result, the BLEU score is computed as:

$$s = \exp \left(\sum_{n=1}^N w_n \log \text{PS}_n(P, G) \right), \quad (3)$$

$$\text{BLEU}(P, G) = s \cdot \text{BP}(P, G), \quad (4)$$

with w_n served as the held constant weight for each value of n .

4.4 Participation

During nearly two months of the competition, 90 individual participants registered for the challenge. In which eight teams have signed the terms and conditions upon the use of the provided VieCap4H dataset. Out of all the participated teams, seven teams are selected based on their results submitted on the private test leaderboard. One team has eventually withdrawn their results from the challenge, therefore there are six teams competing for the top three prizes. All the teams in the final round submitted their technical reports explaining in detail their method and findings from the challenge. Table 4 summarizes the competition participation.




Figure 5: BLEU scores of submissions in top 8 on public and private test splits.

	Public Test	Private Test
# valid submissions	830	97
Best result	0.306	0.329

Table 5: Statistics of the results.

4.5 Outcomes

There were 927 valid submissions recorded on our challenge submission site. The best results measured by averaged BLEU scores (BLEU-0, BLEU-1, BLEU-2 and BLEU-3) on the public test and private test splits are 0.306 and 0.329 (See Figure 5). We provide statistics of the results in Table 5.

5 Conclusion

In this paper, we introduced VieCap4H, a novel image captioning dataset in Vietnamese, which provides over 10,000 image-captions pairs. The images in the VieCap4H dataset are crawled from healthcare related news while the captions are done by human annotators. The dataset serves as a reliable benchmark to advancing research on automatic caption generation in Vietnamese, specifically in the healthcare context. We also reported details of the VieCap4H challenge which was held as part of the eighth annual workshop on Vietnamese Language and Speech Processing (VLSP 2021). The challenge attracted over 90 individuals within only two months, revealing an enormous interest of the local research community in the task.

Acknowledgements

This work is partially supported by Vingroup Innovation Foundation (VINIF) in project code VINIF.2020.DA14. We would like to thank the students from the VNU University of Science who participated to the process of data annotation.

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.
- Jyoti Aneja, Aditya Deshpande, and Alexander G Schwing. 2018. Convolutional image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5561–5570.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Uniter: Learning universal image-text representations. *ECCV20*.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*.
- Yang Feng, Lin Ma, Wei Liu, and Jiebo Luo. 2019. Unsupervised image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4125–4134.
- Albert Gordo, Jon Almazan, Jerome Revaud, and Diane Larlus. 2017. End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision*, 124(2):237–254.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. 2019. Image captioning: Transforming objects into words. *Advances in neural information processing systems*.

- MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. 2019. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6):1–36.
- 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*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270.
- Van Thanh Le. 2021. Pretrained gpt-2 on vietnamese news. <https://huggingface.co/imthanhlv/gpt2news>.
- 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.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. *International Conference on Computer Vision (ICCV)*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*.
- Dat Quoc Nguyen and Anh Tuan Nguyen. 2020. Phobert: Pre-trained language models for vietnamese. pages 1037–1042.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. In *NAACL 2018 - Conference of the North American Chapter of the Association for Computational Linguistics*.
- 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.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7008–7024.
- Mingxing Tan and Quoc Le. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2016. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):652–663.
- Xuan-Son Vu, Thanh-Son Nguyen, Duc-Trong Le, and Lili Jiang. 2020. [Multimodal review generation with privacy and fairness awareness](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 414–425, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Xuan-Son Vu, Thanh Vu, Son N Tran, and Lili Jiang. 2019. Etnlp: A visual-aided systematic approach to select pre-trained embeddings for a downstream task. In *Proceedings of the International Conference Recent Advances in Natural Language Processing (RANLP)*.
- Jiahong Wu, He Zheng, Bo Zhao, Yixin Li, Baoming Yan, Rui Liang, Wenjia Wang, Shipei Zhou, Guosen Lin, Yanwei Fu, et al. 2017. Ai challenger: A large-scale dataset for going deeper in image understanding. *arXiv preprint arXiv:1711.06475*.
- Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. Scene graph generation by iterative message passing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5419.
- Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. 2017. Boosting image captioning with attributes. In *Proceedings of the IEEE international conference on computer vision*, pages 4894–4902.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4651–4659.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.

- Haoyu Zhang, Jianjun Xu, and Ji Wang. 2019. Pretraining-based natural language generation for text summarization. *Proceedings of the 23rd Conference on Computational Natural Language Learning*, pages 789–797.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.