

Miti-DETR: Object Detection based on Transformers with Mitigatory Self-Attention Convergence

Wenchi Ma¹, Tianxiao Zhang², Guanghui Wang^{3*}

¹ Machine Learning Group, Rekor Systems, Inc., Columbia, MD, USA, 21046

² Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA, 66045

³ Department of Computer Science, Ryerson University, Toronto, ON Canada, M5B 2K3

Abstract

Object Detection with Transformers (DETR) and related works reach or even surpass the highly-optimized Faster-RCNN baseline with self-attention network architectures. Inspired by the evidence that pure self-attention possesses a strong inductive bias that leads to the transformer losing the expressive power with respect to network depth, we propose a transformer architecture with a mitigatory self-attention mechanism by applying possible direct mapping connections in the transformer architecture to mitigate the rank collapse so as to counteract feature expression loss and enhance the model performance. We apply this proposal in object detection tasks and develop a model named Miti-DETR. Miti-DETR reserves the inputs of each single attention layer to the outputs of that layer so that the “non-attention” information has participated in any attention propagation. The formed residual self-attention network addresses two critical issues: (1) stop the self-attention networks from degenerating to rank-1 to the maximized degree; and (2) further diversify the path distribution of parameter update so that easier attention learning is expected. Miti-DETR significantly enhances the average detection precision and convergence speed towards existing DETR-based models on the challenging COCO object detection dataset. Moreover, the proposed transformer with the residual self-attention network can be easily generalized or plugged in other related task models without specific customization.

Introduction

The attention mechanism has been effectively used in transformer networks (Vaswani et al. 2017), not only in the application of long-range sequential knowledge, such as natural language processing (Devlin et al. 2018), speech recognition (Luo et al. 2021), but also in computer vision tasks (Gajurel, Zhong, and Wang 2021; Carion et al. 2020; Dai et al. 2021; Zhu et al. 2020; Sajid et al. 2021b), where DETR has achieved competitive performance as an end-to-end object

detector (Carion et al. 2020). Attention mechanism, transformer networks and DETR, thus, have become the research focuses, where the inner workings of transformers and attention, the training and optimization challenge of DETR, etc., have been regarded shedding light for future works.

The attention-mechanism based transformer networks realize the most generalized deep learning model in terms of computer vision and image tasks. In transformer networks, one pixel in an image cares about all the other pixels in that image so that any single region obtains and integrates relevance with all other regions, which is in comparison with CNN that any single pixel cares about its immediate neighborhood and then what the neighborhood as a whole cares about is its immediate neighborhood. This could be a good explanation of why DETR can be on par with state-of-the-art classifiers in terms of classification accuracy. It also explains the strong inductive bias of self-attention. While the research in (Dong, Cordonnier, and Loukas 2021) demonstrates that pure self-attention networks (SANs) would lead to the loss of expressive power doubly exponentially with respect to network depth, and the output converges with a cubic rate to a rank one matrix that has identical rows.

Inspired by the analysis that skips connections play a key role in mitigating rank collapse in transformers, we propose the Miti-DETR detector model where residual self-attention network architecture is introduced. Specifically, the inputs of a multi-head self-attention layer are short connected to its outputs, as shown in Figure 1. This connection skips the Multi-layer perceptions (MLP), which is usually considered rendering the model less sensitive to its input perturbations (Cranko et al. 2018). Thus, the relative “non-attention” features are integrated to avoid the outputs dramatically converging to rank one matrix. In summary, this work makes two critical contributions to the current DETR models:

- The training and optimization challenges of DETR could be better solved based on the network model itself, which is data-independent and could be easily applied in the related models. It is verified that the proposed model could significantly speed up the training convergence so as to avoid the extremely long time training schedule.

*This work was partly supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grant no. RGPIN-2021-04244, and the United States Department of Agriculture (USDA) under Grant no. 2019-67021-28996. Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

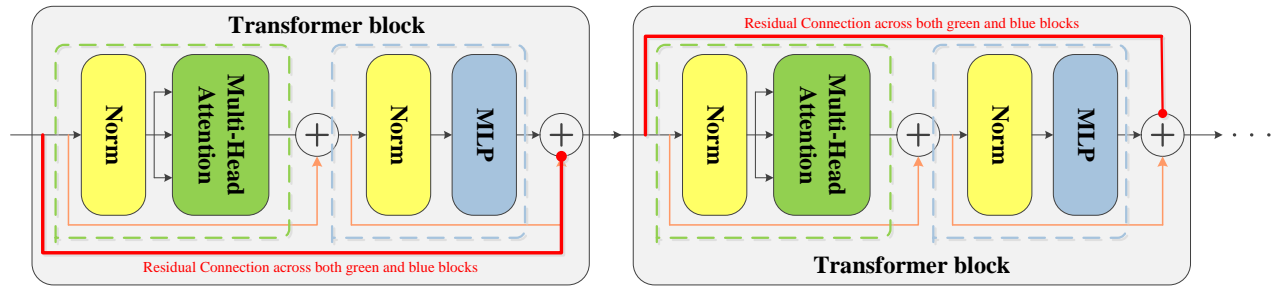


Figure 1: Residual Attention Network in Transformer

- The proposed model can further enhance the object detection performance of DETR models. From the perspective of models themselves, we address the limitations of DETR models where the transformer networks tend to lose efficient feature expression power by proposing the residual self-attention network. Thus, the proposed Miti-DETR could better stop the outputs from degeneration and achieve nearly 3% higher performance than the original DETR models.

We evaluate Miti-DETR on one of the most popular and challenging object detection datasets, COCO, and compare its performance with the traditional DETR-related models. Our experiments show that our model is capable of further enhancing the DETR models' performance. Specifically, Miti-DETR brings the superiority of DETR detecting large objects into full play, enabled by the protection of global expression power. Code is available on <https://github.com/wenchima/Miti-DETR>.

Related Work

This work is built on prior researches of attention mechanism (Vaswani et al. 2017), feature mapping and propagation (He et al. 2016) and object detection with transformer (Carion et al. 2020).

Attention Mechanism

Attention-based architectures have become ubiquitous in machine learning, which brings about better learning for long-sequence and large-range knowledge (Bahdanau, Cho, and Bengio 2014) (Ramachandran et al. 2019). They have permeated machine learning applications across data domains, such as natural language processing (Devlin et al. 2018), speech recognition (Luo et al. 2021), and computer vision (Bello et al. 2019) (Carion et al. 2020) (Sajid et al. 2021a). Attention mechanisms are neural network layers that aggregate information from the entire input sequence (Bahdanau, Cho, and Bengio 2014). They allow modeling of dependencies without regard to their distance in the input or output sequences (Bahdanau, Cho, and Bengio 2014) (Kim et al. 2017) and the early such attention mechanisms models mostly are applied in conjunction with the recurrent network (Parikh et al. 2016).

Self-attention is an attention mechanism that relates different positions in a single sequence so as to com-

pute a sequence representation (Cheng, Dong, and Lapata 2016) (Lin et al. 2017). End-to-end memory networks are based on a recurrent attention mechanism rather than sequence aligned recurrence, which shows advantages on simple-language question answering and language modeling tasks (Sukhbaatar et al. 2015). Transformer, currently, is the first transduction model which entirely relies on self-attention to compute representations of its input and output without using sequence aligned RNNs or convolution (Vaswani et al. 2017). They introduce self-attention layers to Non-Local Neural Networks (Wang et al. 2018). One advantage of attention-based models is the global computations and superior memory, making them more suitable compared with RNNs on long sequences. Transformers now have shown its replacing role then RNNs in many problems in natural language processing, speech processing and computer vision (Parmar et al. 2018) (Synnaeve et al. 2019).

Recently, researchers find that pure self-attention networks (SANs), for example, transformers with skip connections and multi-layer perceptrons (MLPs) disabled, lose expressive power doubly exponentially with respect to network depth. They prove that the output converges with a cubic rate to a rank one matrix with identical rows (Dong, Cordonnier, and Loukas 2021). Their analysis verifies that skip connections are the key in mitigating rank collapse, and MLPs can slow down the convergence by increasing their Lipschitz constant. This research inspires our deep thoughts towards the current object detection models with transformer. We try to excavate the inner specialty of the transformer so as to solve the existing problem, especially in object detection applications.

Object Detection with Transformer

Previously, the mainstream object detection models make predictions relative to some initial guesses (Ma et al. 2021). Two-stage detectors (Ren et al. 2015; Zhang et al. 2020) predict bounding boxes relative to proposals, and single-stage methods make predictions based on anchors (Ma et al. 2020; Ma, Li, and Wang 2020) or other features (Li et al. 2021; Zhang, Ma, and Wang 2021). The performance of these models heavily depends on the set of initial guesses. Anchor-free detection technologies assign positive and negative samples to feature maps by a grid of object centers (Law and Deng 2018).

DETR is recently proposed that successfully apply transformer in object detection that is conceptually simpler without handcrafted process by direct set prediction (Carion et al. 2020). DETR utilizes a simple architecture, by combining convolutional neural networks (CNNs) and Transformer encoder-decoders. Deformable DETR is proposed to improve the problems of slow-convergence and limited feature spatial resolution by making its attention modules only attend to a small set of key points around a reference (Zhu et al. 2020). UP-DETR is inspired by the pre-training transformers in natural language processing and proposes the random query patch detection to unsupervisedly pre-train the transformer of DETR (Dai et al. 2021), which boosts the performance significantly. While these two models both solve the problem from the data end, one by ImageNet pre-training, the other through multi-scale feature representation. Compared with the previous works, Miti-DETR tries to solve the existing problems from the inner property of transformer so that it is data-independent and more practical, which is meaningful as an improvement towards DETR.

Miti-DETR

The proposed Miti-DETR model maintains the simple architecture as the original DETR, as depicted in Figure 2. It contains three main components: a CNN working as the feature extractor; the transformer encoder and decoder with the proposed residual self-attention network; and the final feed-forward networks (FFN) working for the object detection prediction. We adopt the effective bipartite matching loss for direct prediction (Carion et al. 2020). In this section, we mainly discuss why we need to introduce the proposed residual connection, and how to build up the corresponding transformer architecture, and finally, we show the proof of why this new architecture could bring about mitigatory self-attention convergence.

Attention Network Loses Rank

The attention mechanism has become ubiquitous by its outstanding performance of learning long-range knowledge both in timing and spatial sequence (Bahdanau, Cho, and Bengio 2014) (Vaswani et al. 2017). However, it has been certified that the pure-attention networks (SANs), by disabling the skip connections and multi-layer perceptions (MLPs), tends to losing expressive power with the network getting deep, leading to the output converging to a rank one matrix with cubic rate (Dong, Cordonnier, and Loukas 2021). This can be expressed as below (Dong, Cordonnier, and Loukas 2021).

$$\|res(SAN(\mathbf{X}))\|_{1,\infty} \leq \left(\frac{4\alpha}{\sqrt{d_{qk}}} \right)^{\frac{3L-1}{2}} \|res(\mathbf{X})\|_{1,\infty}^{3L} \quad (1)$$

which agrees to a double exponential rate of convergence. The residual item in the above expression works as,

$$res(\mathbf{X}) = \mathbf{X} - \mathbf{1}\mathbf{x}^T, \text{ where } \mathbf{x} = \operatorname{argmin}_{\mathbf{x}} \|\mathbf{X} - \mathbf{1}\mathbf{x}^T\| \quad (2)$$

Here the input \mathbf{X} is a $n \times d_{in}$ matrix that includes n tokens. \mathbf{W}_{QK}^l and \mathbf{W}_V^l are the corresponding value weight matrices. It is noted that the bound in Equation 1 ensures $\|res(SAN(\mathbf{X}))\|_{1,\infty}$ converge for all the small residual's inputs whenever $4\alpha \leq \sqrt{d_{qk}}$. The detailed proofs can be found in (Dong, Cordonnier, and Loukas 2021).

Skip Connection and MLP Counteract Degeneration

There is a natural but pertinent question: If the (pure) attention network degenerates to a rank one matrix with the increase of depth, why do attention-based transformer networks work in applications? It is verified that the presence of skip connections is crucial that prevents the SAN from degenerating to rank one and the Multi-layer perceptrons (MLPs) help control the convergence rate (Dong, Cordonnier, and Loukas 2021). This argument was originally proved by the discussion of the bounds for the residual.

We denote the output expression of an MLP with depth L and width H as

$$\mathbf{X}^{l+1} = f_l \left(\sum_{h \in [H]} \mathbf{P}_h \mathbf{X}^l \mathbf{W}_h \right), \quad (3)$$

where \mathbf{P}_h is the $n \times n$ row-stochastic matrix. \mathbf{W}_h is the weights matrix.

$\lambda_{l,1,\infty}$ here is used to represent the Lipschitz constant of f_l concerning $l_{1,\infty}$ norm. The upper bound for the residual is derived in the following (Dong, Cordonnier, and Loukas 2021):

$$\|res(\mathbf{X}^L)\|_{1,\infty} \leq \left(\frac{4\alpha H \lambda}{\sqrt{d_{qk}}} \right)^{\frac{3L-1}{2}} \|res(\mathbf{X})\|_{1,\infty}^{3L} \quad (4)$$

which agrees with a double exponential rate of convergence.

Thus, the convergence rate can be adjusted by the MLPs' Lipschitz constants $\lambda_{f,1,\infty}$, which certifies that more powerful MLPs bring about slower convergence. This shows the hard struggle between self-attention layers and the MLPs whose nonlinearity contributes to increasing the rank (Dong, Cordonnier, and Loukas 2021).

Residual Self-attention Network

By observing the current transformer network, we find that the skip connections are across "independent modules". Here the "independent modules" refer to the multi-head self-attention network and the feed-forward fully connection network in transformer encoder and decoder layer architectures. Based on the analysis towards the functions of SANs, skip-connection and MLPs, we can define the features' attention levels and provide the sequence. Considering one single transformer layer, we use \mathbf{X}^l to denote its inputs, $SAN(\mathbf{X}^l)$ to denote the outputs of SAN network, and \mathbf{X}^{l+1} as the outputs after MLPs. If g is defined as the mapping function of features' attention level, we can derive the following conclusion.

$$g(SAN(\mathbf{X}^l)) \geq g(\mathbf{X}^{l+1}) \geq g(\mathbf{X}^l) \quad (5)$$




Figure 2: Architecture Streamline of Miti-DETR

Combined with the advantage brought by skip connections, we believe that the skip connections that dramatically diversify the path distribution are the structural factor that explains the degeneration counteraction. This structural factor brings in the diversity of attention levels. Essentially, we believe that it is the diversity of features' attention levels that mitigate the strong inductive bias towards "token uniformity" of the self-attention networks. Inspired by this idea, we propose the residual self-attention network, in which we reserve the inputs of each single attention layer to the outputs of this layer so that the "non-attention" information could participate in the feature attention propagation. Thus, the diversity of feature attention levels is maximized. The corresponding architecture is depicted in Figure 1. In the sequence, we provide proof for the convergence property of the new self-attention network with the proposed residual connection.

The output of the l_{th} layer attention network, in this case, can be expressed as

$$\mathbf{X}^{l+1} = f_l \left(\sum_{h \in [H]} \mathbf{P}_h \mathbf{X}^l \mathbf{W}_h \right) + \mathbf{X}^l, \quad (6)$$

then the l_{th} layer output not considering the residual connection is,

$$\mathbf{X}^{l+1} - \mathbf{X}^l = f_l \left(\sum_{h \in [H]} \mathbf{P}_h \mathbf{X}^l \mathbf{W}_h \right). \quad (7)$$

We follow the definition of the residual in Equation 2,

$$res(\mathbf{X}^{l+1} - \mathbf{X}^l) = (\mathbf{X}^{l+1} - \mathbf{X}^l) - \mathbf{1}\mathbf{x}^T, \quad (8)$$

where

$$\mathbf{x} = \underset{\mathbf{x}}{\operatorname{argmin}} \|(\mathbf{X}^{l+1} - \mathbf{X}^l) - \mathbf{1}\mathbf{x}^T\|, \quad (9)$$

while the proposed residual connection skips both the multi-head attention layer and the MLPs, the actual output residual should be

$$res(\mathbf{X}^{l+1}) = \mathbf{X}^{l+1} - \mathbf{1}\mathbf{x}^T, \quad (10)$$

where

$$\mathbf{x} = \underset{\mathbf{x}}{\operatorname{argmin}} \|(\mathbf{X}^{l+1} - \mathbf{X}^l) - \mathbf{1}\mathbf{x}^T\|. \quad (11)$$

Thus, $res(\mathbf{X}^{l+1})$ can also be expressed in the following equation

$$res(\mathbf{X}^{l+1}) = \mathbf{X}^l + \epsilon, \quad (12)$$

where

$$\epsilon < \mathbf{X}^l \text{ and } \epsilon = \underset{\epsilon}{\operatorname{argmin}} \|\mathbf{X}^{l+1} - \mathbf{X}^l - \epsilon\|. \quad (13)$$

With the introduced residual connection, $res(\mathbf{X}^{l+1})$ anchors at the inputs \mathbf{X}^l of the current layer. Its fluctuation depends on the perturbation ϵ of the inputs and the original convergence performance of the attention network. We compare the above three equations, $res(\mathbf{X}^{l+1} - \mathbf{X}^l)$ still satisfies the Equation 4 by a convergence upper bound, and

$$\|res(\mathbf{X}^{l+1})\| > \|res(\mathbf{X}^{l+1} - \mathbf{X}^l)\|. \quad (14)$$

Thus, the proposed residual self-attention network brings about an anchor for the convergence of each attention layer, which propagates through all attention layers in the transformer architecture and helps render the model keep sensitive to the input perturbation, which further slows down the convergence and counteracts the rank collapse. This research conclusion is verified by concrete learning performance in the experiment section.

Transformer with Mitigatory Convergence in Object Detection

Based on the above analysis, the proposed residual attention architecture enables the transformer to have mitigatory convergence performance theoretically. We follow the same architecture streamline as the DETR model but apply our proposed transformer network as the corresponding encoder-decoder transformer. Specifically, Miti-DETR consists of a convolutional CNN backbone, self-encoder and decoder attention network (Carion et al. 2020) and prediction feed-forward network (FFNs).

In this work, we design the residual self-attention network in every transformer encoder and decoder layer, as the illustration shows in the transformer block in Figure 2. The residual connection bridges the inputs and outputs of a single transformer layer in a short connection way. This path bypasses the composite module of multi-head self-attention network and the feed-forward fully connection network, forming the new "non-attention" feature propagation. Then we concatenate the original inputs of the current layer and the outputs from the MLPs and normalization layer at the top of the transformer layer. Thus, this proposed new transformer network works as an independent module that can be




Figure 3: Comparison of the learning process distributions on COCO: (left to right) Learning Loss, Test Error, and Average Recall.

implemented in any deep learning framework that provides a common CNN backbone and a transformer architecture implementation.

Experiments

We show that Miti-DETR achieves competitive results compared to the original DETR and UP-DETR in quantitative evaluation on COCO (Lin et al. 2014). Then we provide a detailed analysis towards the training and learning progress, with insights and qualitative results. Then, we provide the detection accuracy results of Miti-DETR on COCO and shows its leading performance at multiple measuring criterion. To show the advantage of speeding up convergence and effective optimization, we compare Miti-DETR and UP-DETR (Dai et al. 2021) and present the detection results and analysis as well. The corresponding experimental settings and implementation details come below.

Implementation Details

We train the related models in the work with AdamW (Loshchilov and Hutter 2017). The transformer weights are initialized with Xavier init (Glorot and Bengio 2010) and the backbone is the ImageNet-pretrained ResNet model (He et al. 2016). In this work, we report the results with the backbone of ResNet-50, which is a relatively basal option. All the other hyperparameters in this experiment strictly follow the setting of DETR (Carion et al. 2020).

We use the training schedule of 300 epochs with a learning rate drop of 10 after 200 epochs. All the training images are passed over the model for a single epoch. We train the related models on 4 P100 GPUs, which means each GPU processes two images at the same time in this setting. We apply the pretrained backbone network of DETR R50 provided by DETR official code.¹

Dataset

We conduct the experiments on the COCO2017 detection dataset (Lin et al. 2014), which contains 118K training images and 5K validation images. On average, there are 7 instances, up to 63 instances in a single image in the train-

ing set, including small and large objects in the same images (Carion et al. 2020). We report Average Precision (AP) as bounding box AP, under the integral metric with multiple thresholds. For the comparison with state-of-the-art models, we report the validation AP from the highest epoch performance.

Training and Learning

Typically, transformers are trained with Adam or Adagrad optimizers with very long training schedules and dropout, which is true for DETR models as well. Despite this disadvantage, both UP-DETR and Miti-DETR are designed to make a transformer-based detector with faster convergence. The difference is that UP-DETR focuses on solving the problem by pre-training the transformer network, while Miti-DETR works on handling this problem by attention mechanism and transformer network itself. We report the training loss curves and test error curves of all the experimental models during the learning process.

Detection Results

Setup. The models in the experiment settings are fine-tuned on COCO train2017 (approximate of 118k images) and evaluated on val2017. A comprehensive comparison, including AP, AP₅₀, AP₇₅, AP_S, AP_M and AP_L, is reported. Moreover, we also show the curves of training loss, test error, and average recall in the learning process and make a comparison among the related models. In order to measure the general performance of the related models considering the trade-off of model size, efficiency and performance, we present the quantitative results of the corresponding properties, including average evaluation time (averaged on 10 times' evaluation) and the number of model parameters.

Results. In Figure 3, Miti-DETR outperforms DETR for the entire learning process in terms of the convergence speed, showing a clear advantage, especially after the 150 epoch schedule. The statistic of test error in the middle figure in Figure 3 is the remaining value after deducting Average Precision (AP) from 1 at each epoch. Moreover, DETR shows an unstable learning state, even divergence, between epoch

¹<https://github.com/facebookresearch/detr>




Figure 4: Visualization Detection Results Comparison. Images in the first row and the second row are the results from DETR and Miti-DETR respectively.

Model	Backbone	Epoch	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
DETR	R50	300	37.6	57.8	39.3	18.0	40.6	55.7
UP-DETR	R50	300	33.1	50.9	34.2	14.6	35.0	50.1
Miti-DETR	R50	300	40.5	60.4	42.7	19.7	43.9	59.3

Table 1: Detection accuracy on COCO.

Model	Backbone	Avg Eval Time (s)	#Params	Accuracy (AP)
DETR	R50	484	41302368	37.7
UP-DETR	R50	496	41302880	33.1
Miti-DETR	R50	486	41302368	40.5

Table 2: General performance comparison in terms of efficiency and accuracy

150 and epoch 200, the end of the first learning rate schedule. While both UP-DETR and Miti-DETR appear to have very stable convergence procedures during the entire process and all of the loss, test error and recall curves show a consistent trend. This is in accord with the theory that the original transformer tends to rank collapse. It is noted that even if both UP-DETR and Miti-DETR seem to be helpful for stabilizing the convergence of DETR, Miti-DETR has a sharply faster convergence performance.

We can also see from the results that although DETR returns to the normal track of fast convergence after the 200 epoch schedule, the unstable state could be the reason that lowers down the final converging performance of DETR. After the learning rate is reduced at epoch 200, the Miti-DETR averagely keeps the 0.9 test error lower than that of DETR. A similar situation applies to the average recall curves. This experiment suggests the proposed residual self-attention network could effectively improve the convergence of DETR, where the unstable learning procedure caused by the rank-1 trend is avoided by the residual connection across composite network modules in transformer, basically in an end-to-end method without further pre-processing.

Table 1 shows the AP statistic results of DETR, UP-DETR and the proposed Miti-DETR. As shown in the table, Miti-DETR generally leads DETR by 3% in terms of AP. For AP₇₅, the highest threshold, Miti-DETR even surpasses DETR by nearly 4%. This indicates that Miti-DETR generally prominently improves the detection quality of DETR. More detected bounding boxes have higher IoU with the ground truth. Considering the size property of objects, Miti-DETR yields significant advantages over DETR, ranging from small, medium, and large thresholds. Although Miti-DETR has no more than 2% higher than DETR in terms of AP_S, this is still a big enhancement considering small objects are hard to handle for current object detectors, and it’s an existing problem of DETR as well. As for UP-DETR, it could not obtain better results under our experimental settings, where the limited number of epoch and GPU bring about insufficient training compared with the experiments in the UP-DETR paper (Dai et al. 2021). From this perspective, the proposed Miti-DETR is more robust in different experimental settings. Combined with Figure 3, it is worth noting that the stable learning process effectively contributes to the

final detection accuracy of DETR and the Miti-DETR shows a promising research direction for future research based on the transformer network itself.

Figure 4 presents the visualization detection results in the comparison between the DETR and Miti-DETR. It can be seen that Miti-DETR has less false detection than DETR, such as the tie in the first set of images and the coachman in the second group of images. On the other hand, Miti-DETR seems to have less leak detection, such as the boats and the aircraft in the last two groups of images, respectively.

We also provide the superior performance comparison of Miti-DETR in terms of the trade-off of the model size, running speed and detection performance in Table 2. The statistic of evaluation time is based on the entire COCO evaluation dataset, calculating the processing time on the 5K images. Miti-DETR keeps the same model size with the same number of parameters and nearly the same inference time. While for UP-DETR, it increases the model size and lowers down the running speed. Thus, we can conclude that Miti-DETR shows a better practical significance by solving the problem based on the model itself.

Conclusion

We have presented Miti-DETR, a DETR object detector based on the transformer with mitigatory self-attention convergence. The model outperforms DETR by a large advantage. The proposed core technique, the residual self-attention network, is verified capable of preventing the attention network from losing rank based on the transformer network itself. Miti-DETR is straightforward to implement and has a flexible structure where the residual self-attention network could be extensible to other attention-mechanism models or tasks. In addition, it achieves significantly better performance on small objects than DETR, thanks to the effective processing of global information and the stable convergence procedure. This work could inspire further research towards the working principles of attention mechanism and transformer network, especially for the effective and efficient processing towards the global attention. Accordingly, current models could be more productive with the training process and a more comprehensive method of combining global information and local features based on the transformer network itself is expected.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bello, I.; Zoph, B.; Vaswani, A.; Shlens, J.; and Le, Q. V. 2019. Attention augmented convolutional networks. In *Proceedings of CVPR*, 3286–3295.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229. Springer.
- Cheng, J.; Dong, L.; and Lapata, M. 2016. Long short-term memory-networks for machine reading. *arXiv preprint arXiv:1601.06733*.
- Cranko, Z.; Kornblith, S.; Shi, Z.; and Nock, R. 2018. Lipschitz networks and distributional robustness. *arXiv preprint arXiv:1809.01129*.
- Dai, Z.; Cai, B.; Lin, Y.; and Chen, J. 2021. Up-detr: Unsupervised pre-training for object detection with transformers. In *Proceedings of the CVPR*, 1601–1610.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dong, Y.; Cordonnier, J.-B.; and Loukas, A. 2021. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. *arXiv preprint arXiv:2103.03404*.
- Gajurel, K.; Zhong, C.; and Wang, G. 2021. A Fine-Grained Visual Attention Approach for Fingerspelling Recognition in the Wild. *arXiv preprint arXiv:2105.07625*.
- Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of AISTATS*, 249–256.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Kim, Y.; Denton, C.; Hoang, L.; and Rush, A. M. 2017. Structured attention networks. *arXiv preprint arXiv:1702.00887*.
- Law, H.; and Deng, J. 2018. Cornernet: Detecting objects as paired keypoints. In *Proceedings of ECCV*, 734–750.
- Li, K.; Wang, N. Y.; Yang, Y.; and Wang, G. 2021. SGNet: A Super-class Guided Network for Image Classification and Object Detection. *arXiv preprint arXiv:2104.12898*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 740–755. Springer.
- Lin, Z.; Feng, M.; Santos, C. N. d.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Luo, H.; Zhang, S.; Lei, M.; and Xie, L. 2021. Simplified self-attention for transformer-based end-to-end speech recognition. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, 75–81. IEEE.
- Ma, W.; Li, K.; and Wang, G. 2020. Location-aware box reasoning for anchor-based single-shot object detection. *IEEE Access*, 8: 129300–129309.
- Ma, W.; Tu, X.; Luo, B.; and Wang, G. 2021. Semantic Clustering based Deduction Learning for Image Recognition and Classification. *Pattern Recognition*.
- Ma, W.; Wu, Y.; Cen, F.; and Wang, G. 2020. Mdfn: Multi-scale deep feature learning network for object detection. *Pattern Recognition*, 100: 107149.
- Parikh, A. P.; Täckström, O.; Das, D.; and Uszkoreit, J. 2016. A decomposable attention model for natural language inference. *arXiv preprint arXiv:1606.01933*.
- Parmar, N.; Vaswani, A.; Uszkoreit, J.; Kaiser, L.; Shazeer, N.; Ku, A.; and Tran, D. 2018. Image transformer. In *International Conference on Machine Learning*, 4055–4064. PMLR.
- Ramachandran, P.; Parmar, N.; Vaswani, A.; Bello, I.; Levskaya, A.; and Shlens, J. 2019. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Sajid, U.; Chen, X.; Sajid, H.; Kim, T.; and Wang, G. 2021a. Audio-Visual Transformer Based Crowd Counting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2249–2259.
- Sajid, U.; Chow, M.; Zhang, J.; Kim, T.; and Wang, G. 2021b. Parallel Scale-wise Attention Network for Effective Scene Text Recognition. *arXiv preprint arXiv:2104.12076*.
- Sukhbaatar, S.; Szlam, A.; Weston, J.; and Fergus, R. 2015. End-to-end memory networks. *arXiv preprint arXiv:1503.08895*.
- Synnaeve, G.; Xu, Q.; Kahn, J.; Likhomanenko, T.; Grave, E.; Pratap, V.; Sriram, A.; Liptchinsky, V.; and Collobert, R. 2019. End-to-end asr: from supervised to semi-supervised learning with modern architectures. *arXiv preprint arXiv:1911.08460*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7794–7803.
- Zhang, T.; Ma, W.; and Wang, G. 2021. Six-channel Image Representation for Cross-domain Object Detection. *arXiv preprint arXiv:2101.00561*.
- Zhang, T.; Zhang, X.; Yang, Y.; Wang, Z.; and Wang, G. 2020. Efficient Golf Ball Detection and Tracking Based on Convolutional Neural Networks and Kalman Filter. *arXiv preprint arXiv:2012.09393*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.