

Can attention enable MLPs to catch up with CNNs?

Meng-Hao Guo¹, Zheng-Ning Liu¹, Tai-Jiang Mu¹, Dun Liang¹, Ralph R. Martin², and Shi-Min Hu¹ (✉)

© The Author(s) 2021.

In the first week of May 2021, researchers from four different institutions: Google, Tsinghua University, Oxford University, and Facebook shared their latest work [1–4] on [arXiv.org](https://arxiv.org) at almost the same time, each proposing new learning architectures, consisting mainly of linear layers, claiming them to be comparable or superior to convolutional-based models. This sparked immediate discussion and debate in both academic and industrial communities as to whether MLPs are sufficient, many thinking that learning architectures are returning to MLPs. Is this true?

In the following, we give a brief history of learning architectures, including multilayer perceptrons (MLPs), convolutional neural networks (CNNs), and transformers. We then examine what the four newly proposed architectures have in common. Finally, we give our views on challenges and directions for new learning architectures, hoping to inspire future research.

1 Learning architectures for visual tasks

Multilayer perceptrons (MLPs) [5] consist of an input layer and an output layer, possibly with multiple hidden layers in between. Layers are typically fully connected using linear transformations and activation functions. MLPs were the basis for neural networks before deep convolutional neural networks (DCNNs) took over, and greatly improved the power

of computers to handle problems of classification and regression. However, MLPs are computationally costly and prone to overfitting, due to their large numbers of parameters. MLPs are also poor at capturing local structures in the input, since the linear transformations between layers always take the output from the previous layer as a whole. However, we note that the capabilities of MLPs were not fully explored when they were proposed, both because of limited computer performance, and unavailability of massive data for training.

To learn local structures in the input while maintaining computational efficiency, convolutional neural networks (CNNs) were proposed. In 1998, LeCun et al. [6] presented LeNet, which greatly improved the accuracy of handwritten digit recognition using a five-layer convolutional neural network. Later, AlexNet [7] led to wide acceptance of CNNs in the research community: it was much larger than previous CNNs like LeNet, and beat all other competitors by a significant margin in the ImageNet Large Scale Visual Recognition Challenge in 2012[®]. Since then, many more models with ever deeper architectures have been developed, with many providing more accurate results than humans in various realms, resulting in profound paradigm changes in both scientific research, and engineering and commercial applications.

Notwithstanding advances in computing power and amounts of training data, the key success of CNNs lies in the inductive bias they introduce: they assume that information has spatial locality and can thus reduce the number of network parameters by making use of a sliding convolution with shared weights. However, the side-effect of this approach is that the receptive fields of CNNs are limited, making CNNs less able to learn

1 BNRist, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: M.-H. Guo, gmh20@mails.tsinghua.edu.cn; Z.-N. Liu, liu-zn17@mails.tsinghua.edu.cn; T.-J. Mu, taijiang@tsinghua.edu.cn; D. Liang, liangd16@mails.tsinghua.edu.cn; S.-M. Hu, shimin@tsinghua.edu.cn (✉).

2 Cardiff University, Cardiff CF243AA, UK. E-mail: ralph@cs.cf.ac.uk.

Manuscript received: 2021-05-18; accepted: 2021-05-23

[®] <https://image-net.org/challenges/LSVRC/2012/>

long-range dependencies. To enlarge the receptive field, a larger convolutional kernel is required, or other special strategies must be employed, such as dilated convolutions [8]. Simply composing a large kernel from several small kernels is not a suitable approach for enlarging the receptive field of CNNs [9].

Recently, the Transformer neural network architecture was proposed [10] for sequential data, with great success in natural language processing [11, 12], and more recently, in vision [13–17]. The attention mechanism is at the core of Transformer, which is readily capable of learning long-range dependencies between any two positions in the input data in the form of an attention map. However, this additional freedom and reduced inductive bias mean that effectively training Transformer-based architectures requires huge amounts of data. For best results, such models should be first pre-trained on a very large dataset, such as GPT-3 [18] or ViT [13].

2 Linear layer based architectures

2.1 Four recent architectures

To avoid the drawbacks of the aforementioned learning architectures, and, with the aim of achieving better results at lower computational cost, very recently, four architectures motivated by self-attention and transformers were proposed almost simultaneously [1–4]. Their common aim is to take full advantage of linear layers. We briefly summarize these architectures below; also see Fig. 1. All four employ transposition to model interactions at all scales. Residual connections and normalization are also utilized in a similar way to ensure stable training.

2.2 MLP-Mixer

MLP-Mixer [1] takes S non-overlapping image patches of resolution $P \times P$ as input. Each patch is first projected to a C -dimensional embedding via a shared-weight linear layer: this representation of the input image is thus a matrix, $\mathbf{X} \in \mathbb{R}^{S \times C}$.

Next, \mathbf{X} is fed into a sequence of identical *mixer layers*, each of which is composed of a *token-mixing* MLP block and a *channel-mixing* MLP block, mixing information from all patches, and from all channels, respectively. We may express the computation as

$$\mathbf{U} = \mathbf{X} + f_2(\sigma(f_1(\text{Norm}(\mathbf{X})^T)))^T \quad (1)$$

$$\mathbf{Y} = \mathbf{U} + f_4(\sigma(f_3(\text{Norm}(\mathbf{U}))) \quad (2)$$

where f_1, \dots, f_4 are linear layers, and σ denotes GELU (nonlinear) activation [19]. Layer normalization [20] is employed. $\mathbf{U} \in \mathbb{R}^{S \times C}$ is the intermediate matrix after per-channel feature aggregation: a shared-weight mapping $\mathbb{R}^S \mapsto \mathbb{R}^S$ of the column vectors in \mathbf{X} . Similarly, two linear transformations are performed per patch, giving the output \mathbf{Y} .

2.3 External attention

External attention [2] reveals the relation between self-attention and linear layers. It first simplifies self-attention as in Eq. (4), where $\mathbf{F} \in \mathbb{R}^{N \times d}$ is the input feature map.

$$\mathbf{A} = \text{softmax}(\mathbf{F}\mathbf{F}^T) \quad (3)$$

$$\mathbf{F}_{\text{out}} = \mathbf{A}\mathbf{F} \quad (4)$$

Then an external memory unit $\mathbf{M} \in \mathbb{R}^{S \times d}$ is introduced to replace \mathbf{F} -to- \mathbf{F} attention by \mathbf{M} -to- \mathbf{F} attention as below:

$$\mathbf{A} = (\alpha)_{i,j} = \text{Norm}(\mathbf{F}\mathbf{M}^T) \quad (5)$$

$$\mathbf{F}_{\text{out}} = \mathbf{A}\mathbf{M} \quad (6)$$

Finally, like self-attention, it uses two different memory units \mathbf{M}_k and \mathbf{M}_v as the key and the value to increase the capability of the network respectively. The overall computation of external attention is as below:

$$\mathbf{A} = \text{Norm}(\mathbf{F}\mathbf{M}_k^T) \quad (7)$$

$$\mathbf{F}_{\text{out}} = \mathbf{A}\mathbf{M}_v \quad (8)$$

Because $\mathbf{F}\mathbf{M}_k^T$ is a matrix multiplication, it is linear in \mathbf{F} , so Eq. (8) can be written as

$$\mathbf{F}_{\text{out}} = f_2(\text{Norm}(f_1(\mathbf{F}))) \quad (9)$$

The final output is then obtained by adding an identity mapping as below:

$$\mathbf{F}_{\text{out}} = \mathbf{F} + f_2(\text{Norm}(f_1(\mathbf{F}))) \quad (10)$$

Based on this external attention, Guo et al. [2] also designed a multi-head external attention and provided an all MLP architecture EAMLP.

2.4 Feed-forward-only model

The feed-forward-only model [3] replaces the attention layers in Transformer [10] by simple feed-forward layers on the token dimension. It firstly uses linear layers on the channel dimension and then adopts linear layers on the token dimension in a linear block. Given an input $\mathbf{X} \in \mathbb{R}^{N \times C}$, the computation can be expressed in detail as

$$\mathbf{U} = \mathbf{X} + f_2(\sigma(f_1(\text{LayerNorm}(\mathbf{X})))) \quad (11)$$

$$\mathbf{Y} = \mathbf{U} + f_4(\sigma(f_3(\text{LayerNorm}(\mathbf{U}^T))))^T \quad (12)$$

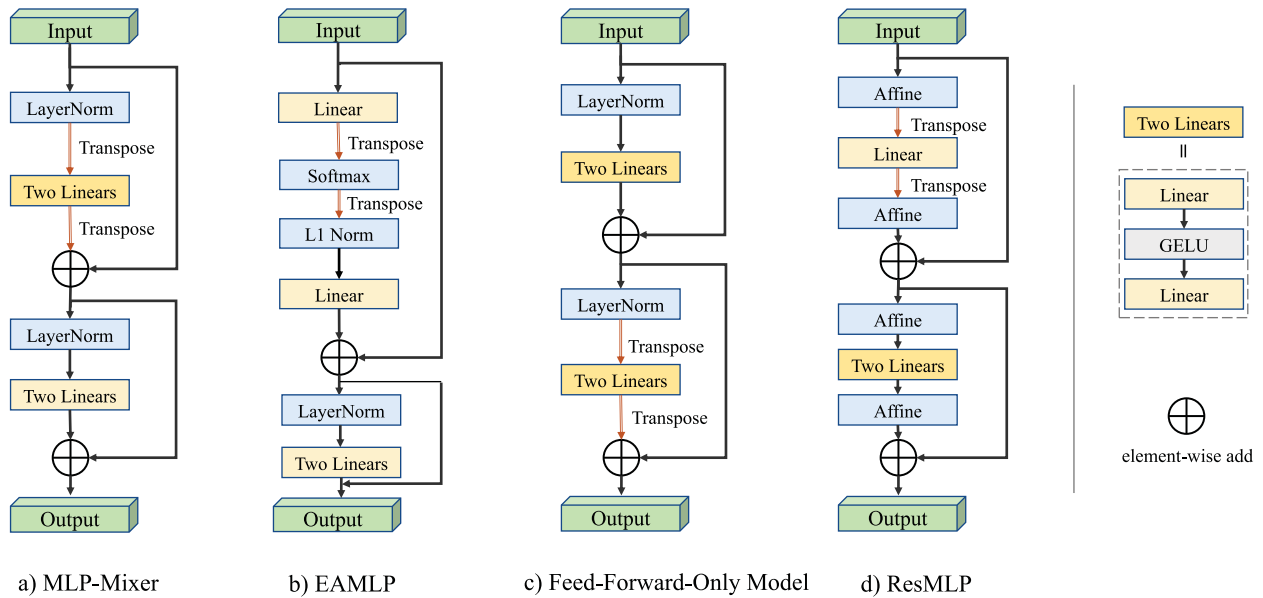


Fig. 1 Basic blocks of four recent architectures in which linear layers predominate.

2.5 ResMLP

ResMLP [4] also separately aggregates information in per-patch-style and per-channel-style, and can be formulated as follows:

$$U = X + \text{Norm}(f_1(\text{Norm}(X)^T)^T) \tag{13}$$

$$Y = U + \text{Norm}(f_3(\sigma(f_2(\text{Norm}(U)))))) \tag{14}$$

A major difference of ResMLP is that it uses an affine transformation in the role of a normalization layer. This affine transformation is parameterized by two learnable vectors to scale and shift the input component-wise:

$$\text{Aff}_{\alpha,\beta}(X) = \text{Diag}(\alpha)x + \beta \tag{15}$$

Note that no statistics of the input are used in the above, and thus it can be integrated in the linear layers during inferencing for further speed.

3 Common themes

We now examine the above approaches, to see what they have in common.

3.1 Long distance interactions

As in self-attention, interactions between different patches are taken into account by these four methods. MLP-Mixer, ResMLP, and the Feed-forward-only model use linear layers acting on the token dimension to allow different patches to communicate with each

other. External attention adopts softmax and L1 normalization to perform a similar role. Unlike CNNs, these models can consider long distance interactions between patches and automatically select suitable and irregular receptive fields.

3.2 Local semantic information

Unlike each word in natural language, single pixels have very little semantic information and their interactions with other pixels are not directly informative. It is thus important to extract meaningful information before using MLPs. MLP-Mixer, ResMLP, and the Feed-forward-only model divide the image into 16×16 local patches to obtain semantic information. External attention adopts a T2T module [16] or a CNN backbone to provide rich semantics before passing information to linear layers.

3.3 Residual connections

Residual connections [21] solve the problem of vanishing gradients and stabilize the training process, so they are commonly used in deep convolutional neural networks. They also benefit architectures based around linear layers and are adopted by all the above models.

3.4 Reduced inductive bias

Localised processing in CNNs results in inductive bias, which can decrease accuracy given sufficient

training data. The recently introduced architectures use linear layers on single tokens independently, or process all tokens equally, resulting in lower inductive bias than CNNs.

4 Challenges and future directions

These promising recently introduced architectures have simple network structure and fast inferencing throughput. However, on ImageNet, their results are currently 5%–10% less accurate than those provided by the best CNNs or Transformer networks. They also do not significantly outperform light-weight networks in the trade-off between accuracy and speed. Thus additional research is needed if the potential of such architectures is to be realised.

We suggest possible directions for future work below, and make other observations about these architectures:

- All linear layers process image patches in a direct or indirect manner to extract local features, thereby reducing computational cost. Dividing images into non-overlapping patches again introduces inductive bias. On one hand, CNNs capture local structure extremely well, but lack the ability to handle long range interactions. On the other hand, these four architectures provide a good way to process long range interactions. It seems natural to try to combine the advantages of both architectures.
- More reasonable global interaction and fewer inductive bias. Self-attention has been successfully used in many kinds of fields to capture long range dependency. However, whether there is a more effective global interaction method is still worth exploring.
- One main goal of these four methods is to avoid the use of the self-attention mechanism. The successful configurations used for this purpose in Transformer could be employed in these linear architectures. For example, Transformer can use multi-head attention, and a similar multi-head mechanism could be employed by these methods to improve model capability.
- Residual connections play a key role in all these methods, indicating that the network structure is crucial. Because these new architectures are simpler than CNNs, better backbones are needed.
- Due to the simplicity of these new architectures,

they can easily tackle irregular data structures used in various applications, including point clouds, graphs, etc. Furthermore, this flexibility promises the ability to make cross-modal models, with a unified network backbone for all modes of data.

- An additional benefit is that all computations are matrix multiplications, which can be highly optimized in deep learning frameworks and readily performed on hardware. This simplicity can promote deployment in industry and commerce, and also reduce energy consumption.

5 Conclusions

Overall, the new architectures motivated by self-attention separately apply linear layers in the element (token) dimension and channel dimension to learn long range interactions between any two positions in the feature matrix, while traditional MLPs mix these two dimensions together as a long vector, with too much freedom for effective learning. We conclude that the new architectures do not simply reuse traditional MLPs, but are a significant advance over them.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Project No. 61521002).

References

- [1] Tolstikhin, I.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; Zhai, X.; Unterthiner, T.; Yung, J.; Keysers, D.; Uszkoreit, J.; Lucic, M. et al. MLP-mixer: An all-MLP architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.
- [2] Guo, M. H.; Liu, Z. N.; Mu, T. J.; Hu, S. M. Beyond self-attention: External attention using two linear layers for visual tasks. *arXiv preprint arXiv:2105.02358*, 2021.
- [3] Melas-Kyriazi, L. Do you even need attention? A stack of feed-forward layers does surprisingly well on imageNet. *arXiv preprint arXiv:2105.02723*, 2021.
- [4] Touvron, H.; Bojanowski, P.; Caron, M.; Cord, M.; El-Nouby, A.; Grave, E.; Izacard, G.; Joulin, A.; Synnaeve, G.; Verbeek, J. et al. ResMLP: Feedforward networks for image classification with data-efficient training. *arXiv preprint arXiv:2105.03404*, 2021.
- [5] Rumelhart, D. E.; Hinton, G. E.; Williams, R. J. Learning internal representations by error propagation. In: *Parallel Distributed Processing: Explorations in the*

- Microstructure of Cognition, Volume 1, Foundations.* Rumelhart, D. E.; McClelland, J. L. Eds. MIT Press, 318–362, 1986.
- [6] LeCun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* Vol. 86, No. 11, 2278–2324, 1998.
- [7] Krizhevsky, A.; Sutskever, I.; Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Communications of the ACM* Vol. 60, No. 6, 84–90, 2017.
- [8] Yu, F.; Koltun, V. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2016.
- [9] Peng, C.; Zhang, X. Y.; Yu, G.; Luo, G. M.; Sun, J. Large kernel matters—Improve semantic segmentation by global convolutional network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1743–1751, 2017.
- [10] Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, 6000–6010, 2017.
- [11] Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving language understanding with unsupervised learning. Technical Report. OpenAI, 2018.
- [12] Devlin, J.; Chang, M.; Lee, K.; Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Vol. 1, 4171–4186, 2019.
- [13] Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S. et al. An image is worth 16x16 words: Transformers for image recognition at scale. In: Proceedings of the International Conference on Learning Representations, 2021.
- [14] Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; Zagoruyko, S. End-to-end object detection with transformers. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12346.* Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J. M. Eds. Springer Cham, 213–229, 2020.
- [15] Guo, M. H.; Cai, J. X.; Liu, Z. N.; Mu, T. J.; Martin, R. R.; Hu, S. M. PCT: Point cloud transformer. *Computational Visual Media* Vol. 7, No. 2, 187–199, 2021.
- [16] Yuan, L.; Chen, Y.; Wang, T.; Yu, W.; Shi, Y.; Jiang, Z.; Tay, F. E.; Feng, J.; Yan, S. Tokens-to-token ViT: Training vision transformers from scratch on imageNet. *arXiv preprint arXiv:2101.11986*, 2021.
- [17] Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; Jégou, H. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2021.
- [18] Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A. Shyam, P.; Sastry, G.; Askell, A. et al. Language models are few-shot learners. In: Proceedings of the 34th Conference on Neural Information Processing Systems, 2020.
- [19] Hendrycks, D.; Gimpel, K. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.
- [20] Ba, J. L.; Kiros, J. R.; Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [21] He, K. M.; Zhang, X. Y.; Ren, S. Q.; Sun, J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778, 2016.



Meng-Hao Guo is a Ph.D. candidate supervised by Prof. Shi-Min Hu in the Department of Computer Science and Technology at Tsinghua University, Beijing, China. His research interests include computer graphics, computer vision, and machine learning.



Zheng-Ning Liu received his bachelor degree in computer science from Tsinghua University in 2017. He is currently a Ph.D. candidate in computer science at Tsinghua University. His research interests include 3D computer vision, 3D reconstruction, and computer graphics.



Tai-Jiang Mu is currently an assistant researcher at Tsinghua University, where he received his B.S. and Ph.D. degrees in computer science in 2011 and 2016, respectively. His research interests include computer vision, robotics, and computer graphics.



Dun Liang is a Ph.D. candidate in computer science at Tsinghua University, where he received his B.S. degree in 2016. His research interests include computer graphics, visual media learning, and high-performance computing.



Ralph R. Martin received his Ph.D. degree from Cambridge University in 1983. He is an emeritus professor of Cardiff University with over 300 papers and 14 books in geometry processing, computer graphics, and computer vision. He is a sometime Fellow of the Learned Society of Wales, the Institute of Mathematics and its Applications, and the British Computer Society. He is currently an Associate Editor-in-Chief of *Computational Visual Media*.



Shi-Min Hu is currently a professor in computer science at Tsinghua University. He received his Ph.D. degree from Zhejiang University in 1996. His research interests include geometry processing, image & video processing, rendering, computer animation, and CAD. He has published more than 100 papers in journals and refereed conferences. He is Editor-in-Chief of *Computational Visual Media*, and on the editorial boards

of several journals, including *Computer Aided Design* and *Computer & Graphics*.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.