

# BUILDING BLOCKS FOR A COMPLEX-VALUED TRANSFORMER ARCHITECTURE

Florian Eilers\*, Xiaoyi Jiang†

Faculty of Mathematics and Computer Science, University of Münster, Münster, Germany

## ABSTRACT

Most deep learning pipelines are built on real-valued operations to deal with real-valued inputs such as images, speech or music signals. However, a lot of applications naturally make use of complex-valued signals or images, such as MRI or remote sensing. Additionally the Fourier transform of signals is complex-valued and has numerous applications. We aim to make deep learning directly applicable to these complex-valued signals without using projections into  $\mathbb{R}^2$ . Thus we add to the recent developments of complex-valued neural networks by presenting building blocks to transfer the transformer architecture to the complex domain. We present multiple versions of a complex-valued Scaled Dot-Product Attention mechanism as well as a complex-valued layer normalization. We test on a classification and a sequence generation task on the MusicNet dataset and show improved robustness to overfitting while maintaining on-par performance when compared to the real-valued transformer architecture.

**Index Terms**— Deep learning techniques, Complex-valued neural networks, Transformer architecture

## 1. INTRODUCTION

In recent years, many applications have benefited from the fast development and high quality results of deep learning methods. Most of these methods focus on real-valued pipelines for applications with real-valued signals, such as natural images or encodings of natural language processing. There is however a great amount of applications that naturally deal with complex-valued signals, such as MRI images [1, 2] or remote sensing [3] and the Fourier transform of real-valued signals [4, 5] or images [6, 7] and it has been shown that fully complex-valued architectures often (but not always [8]) deliver superior performance when dealing with complex-valued signals. The complex numbers come with an intrinsic algebraic structure that can not be captured by the simple isomorphism of  $\mathbb{C} \sim \mathbb{R}^2$ , especially because there is no natural way to define multiplication in  $\mathbb{R}^2$ , which, however, is an important part of many deep learning building blocks. [9] has provided a lot of those building blocks, such as complex-valued convolution, batch normalization and initialization. These building blocks are of great help for a large amount of current architectures, especially in image and signal processing. In many fields, architectures building on the idea of attention mechanisms have successfully been applied. Especially the immense success of the transformer architecture [10] has shown that attention based architectures can be superior and have since become standard in many applications. We seek to provide a solid generalization of the building blocks of the transformer architecture in the complex domain and show experimental evidence that it improves robustness

to overfitting while maintaining on-par performance when compared to the real-valued transformer architecture.

Our key contributions are: **1)** Newly developed building blocks consisting of: a. derivation of a complex-valued attention mechanism, generalizing the Scaled Dot-Product attention [10]; b. introduction of complex-valued layer normalization. **2)** Adaptation of building blocks from existing complex-valued neural networks for the transformer architecture. **3)** Demonstration of improved robustness to overfitting while maintaining on-par results compared to the real-valued model.

The combination of the first two contributions provide the foundation for a mathematically rigorous complex-valued transformer architecture. The source code for the full architecture and all experiments is available as a Pytorch module <sup>1</sup>.

## 2. RELATED WORK

Complex-valued neural networks have been researched for a long time [11, 12]. An early standard book and foundation for much research to come is the work by Hirose [13]. Recently, an increasing number of works in complex-valued neural networks have been published [14], driven by the interest in applications, which naturally deal with complex-valued signals: remote sensing [15, 16], MRI processing [17, 1] and frequency analysis through Fourier transform.

[9] provides building blocks for complex-valued neural networks. They present complex versions of linear layers, convolutional layers, batch normalization, initialization and different activation functions. They also comment on complex differentiability, referring to earlier works [18]. Complex-valued building blocks have been used to develop a multitude of architectures, such as complex-valued generative adversarial networks [2, 3], complex-valued convolutional recurrent networks [19] and a complex-valued U-net [4]. There has also been recent interest in optimizing computability on GPUs for complex-valued neural networks [20].

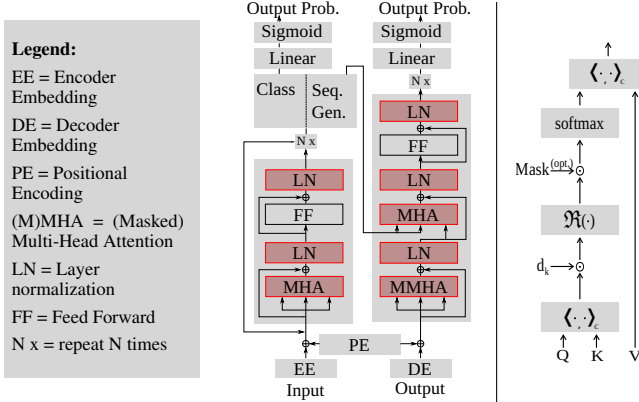
The transformer architecture [10] was a great success in natural language processing and has since become dominant in the field [21, 22]. It has spread into vision [23, 24], music [25] and more applications [26]. Additionally, there has been many works to improve and rework the architecture [27].

To the best of our knowledge, there are only two works concerning the design of a complex-valued transformer architecture. [5] proposes a complex-valued transformer, motivated by the multiplicative structure of the Dot-Product attention. They separate the product  $Q(K^T)$  into eight real-valued products and then apply real-valued attention to all summands separately. While being a well motivated choice, they use real-valued encoding matrices, making the network not fully complex-valued. Using complex-valued encoding matrices would lead to a total of 64 summands and an unreasonable computational blowup. While testing their framework against competitive

\*Member of the CiM IMPRS Graduate School Münster

†This work was partly supported by the Deutsche Forschungsgemeinschaft (DFG) CRC 1450 431460824 and European Unions Horizon 2020 under the Marie Skłodowska-Curie grant agreement No 778602 Ultracette.

<sup>1</sup><https://zivgitlab.uni-muenster.de/ag-pria/cv-transformer>



**Fig. 1:** Left: The general transformer architecture as introduced in [10]. Building blocks, whose complex-valued versions are derived in this paper are highlighted in red. Right: The complex-valued Scaled Dot-Product Attention as suggested in subsection 4.1.

models, such as real-valued transformers and LSTMs, they do not test against other definitions of complex-valued attention modules. Additionally, in their experimental part, they do not use an independent test set but rather just evaluate on the validation set. We answer the questions this work left open by incorporating their idea into our experiments and using it as one of many valid definitions for a complex-valued transformer. [28] proposes a complex-valued meta-learning framework for signal recognition. As a byproduct, they define a complex-valued attention. However, they do not evaluate different options and they do not utilize the Dot-Product in the complex domain. Additionally, they propose to use the complex variance for normalization instead of the more flexible covariance matrix. Their definition is, as one of many, incorporated in our framework.

There are some more works on different kinds of complex-valued attention modules [15, 16, 29]. These use different kinds of convolutional architectures but are not complex-valued versions of the Scaled Dot-Product attention [10].

### 3. TRANSFORMER

This section serves as a brief description to the transformer architecture as introduced in [10].

#### 3.1. Architecture

The architecture consists of an encoder and decoder module. The encoder module alone can be used for classification tasks while the full architecture can be used for sequence generation where the core idea is to input the original input into the encoder and the earlier outputs of the sequence generation into the decoder. Both modules start with an embedding and a positional encoding of their respective inputs. Afterwards the modules consist of Multi-Head Attention mechanism with residual connections followed by a layer normalization and a feed forward module - a small MLP with two linear layers and an activation. Details can be seen in Figure 1.

#### 3.2. Real-valued Attention

The Scaled Dot-Product Attention is the core of the Transformer architecture. The input consists of three matrices, called query  $Q$ , key

$K$  and value  $V$ . First the Dot-Product of the key and query is calculated, then scaled by the square-root of their equal dimensions  $\sqrt{d_k}$ . The output is normalized by the softmax function and afterwards multiplied with the values. Defining the softmax for a vector  $X$  of length  $n$  as

$$\text{softmax}(X) = \sigma(X) = \frac{\exp(X)}{\sum_{i=1}^n \exp(X_i)} \quad (1)$$

we can formulate the Scaled Dot-Product Attention as

$$\text{Att}(Q, K, V) = \sigma \left( \frac{Q(K)^T}{\sqrt{d_k}} \right) V \quad (2)$$

This core concept of Scaled Dot-Product Attention is extended to the more general concept of Multi-Head Attention (MHA) by applying learnable linear projections  $W^Q, W^K, W^V$  to the inputs and project it back with another learnable linear projection  $W^O$ :

$$\begin{aligned} \text{MHA}(Q, K, V) &= \text{Concat}(h_1, \dots, h_k) W_i^O, \\ \text{where } h_i &= \text{Att}(Q W_i^Q, K W_i^K, V W_i^V) \end{aligned} \quad (3)$$

For the training process, it is necessary to mask future events in the decoder input. This is obtained by addition of  $-\infty$  to the respective tokens before the softmax, which results in an attention score of 0.

## 4. COMPLEX-VALUED BUILDING BLOCKS

The general purpose of this section is the introduction of existing building blocks and the development of new ones where needed (subsections 4.1 and 4.2) for a mathematically rigorous extension of the transformer architecture to the complex domain. [5] has already introduced complex-valued fully connected feed forward layers. Since the position within a sequence is real-valued, we adopt the sine and cosine positional encoding as originally used in [10] but other positional encodings would be possible [30, 31].

### 4.1. Complex-valued Attention

When generalizing the Scaled Dot-Product Attention to the complex domain a problem arises: the max operation does not work in  $\mathbb{C}$  and the softmax does not either. However, the core idea behind the operation  $\sigma(XY^T)$  for  $X, Y \in \mathbb{R}^n$  is to define a similarity between  $X$  and  $Y$  which is then scaled to  $(0, 1)$  by the softmax non-linearity  $\sigma$ . The operation without softmax-rescaling can be described as the Dot-Product in  $\mathbb{R}^n$ . Using this concept in  $\mathbb{C}^n$  leads to:


$$\langle X, Y \rangle = \sum_{i=1}^n X_i \bar{Y}_i = \sum_{i=1}^n |X| |Y| \exp(i(\phi_{X_i} - \phi_{Y_i})) \quad (4)$$

When neglecting the magnitudes of  $X$  and  $Y$ , we get

$$\exp(i(\phi_{X_i} - \phi_{Y_i})) = \cos(\phi_{X_i} - \phi_{Y_i}) + i \sin(\phi_{X_i} - \phi_{Y_i}) \quad (5)$$

The real part of this term,  $\cos \phi_{X_i} - \phi_{Y_i}$ , maximizes at 1 for  $\phi_{X_i} - \phi_{Y_i} = 0$ , which is equivalent to  $X_i = Y_i$ . It strictly decreases, when  $|\phi_{X_i} - \phi_{Y_i}|$  growth, up to a minimum of  $-1$  at  $|\phi_{X_i} - \phi_{Y_i}| = \pi$ , which is equivalent to  $X_i = -Y_i$ . Thus we have that  $\mathcal{R}(\langle X, Y \rangle)$  measures the similarity of  $X$  and  $Y$  for every component and adds these up. Additionally, we get two desired properties for a similarity measure: Symmetry and rotational invariance. Symmetry holds because the conjugate symmetry of the Dot-Product does not change its real part:

$$\mathcal{R}\langle Q, K \rangle = \mathcal{R}\overline{\langle K, Q \rangle} = \mathcal{R}\langle K, Q \rangle \quad (6)$$



**Fig. 2:** Illustration of Dot-Product vs  $Q(K^T)$  in  $\mathbb{C}$ . 1st and 2nd column show behavior for the extreme cases of  $Q = K$  and  $Q = -K$ . 3rd and 4th column show rotational invariance of the Dot-Product vs behavior on rotations of  $Q(K^T)$ .

By rotational invariance we mean: If  $Q$  and  $K$  are (elementwise) both rotated by a fixed angle  $\alpha$ ,  $\langle Q, K \rangle$  does not change (Figure 2).

$$\exp(i((\phi_{X_i} + \alpha) - (\phi_{Y_i} + \alpha))) = \exp(i(\phi_{X_i} - \phi_{Y_i})) \quad (7)$$

Note that the equality  $\langle X, Y \rangle = Q(K^T)$  holds in the real domain, but not in the complex domain. While symmetry still holds when using  $Q(K)^T$ , the rotational invariance does not (Figure 2).

When taking the magnitude into account, this similarity is scaled by the factors  $|X|$  and  $|Y|$ , meaning that high values are obtained for vectors of high magnitude pointing in the same direction. We can now formulate the complex-valued Dot-Product Attention as:

$$\mathbb{C}Att(A, B) = \sigma \left( \frac{\mathcal{R}\langle Q, K \rangle}{\sqrt{d_k}} \right) V \quad (8)$$

This pipeline is presented on the right in Figure 1.

The motivation of Scaled Dot-Product attention leads to Equation 8. However, other possibilities to generalize the real-valued Scaled Dot-Product attention are using the absolute value with and without keeping the phase information and using both the real and the imaginary part. We define the following possibilities and test all these in section 5.  $|z|_{\mathbb{C}}$  denotes the absolute value of a complex number  $z$  and  $\text{sgn}(z)$  its sign (e.g.  $\frac{z}{|z|}$  if  $z \neq 0$  and 1 o/w):

$$AAtt(A, B) = \sigma \left( \frac{|\langle Q, K \rangle|_{\mathbb{C}}}{\sqrt{d_k}} \right) V \quad (9)$$

$$PAtt(A, B) = \sigma \left( \frac{|\langle Q, K \rangle|_{\mathbb{C}}}{\sqrt{d_k}} \right) \text{sgn}(\langle Q, K \rangle) V \quad (10)$$

$$\mathcal{R}IAtt(A, B) = \left( \sigma \left( \frac{\mathcal{R}\langle Q, K \rangle}{\sqrt{d_k}} \right) + i \sigma \left( \frac{\mathcal{I}\langle Q, K \rangle}{\sqrt{d_k}} \right) \right) V \quad (11)$$

Additionally, it is possible to replace the dot product  $\langle Q, K \rangle$  in every version with  $Q(K)^T$ . Using  $Q(K)^T$ ,  $AAtt$  and  $\mathbb{C}Att$  have been used before [28], we test these variants in section 5. Note  $K \mapsto \bar{K}$  is not linear in  $\mathbb{C}$  and thus cannot be learned directly by  $W^K$ .

Using any of these formulations of the complex-valued Scaled Dot-Product Attention the adoption of Multi-Head Attention as described in subsection 3.2 is straightforward. We can replace the learnable linear projections  $W^Q$ ,  $W^K$ ,  $W^V$  and  $W^O$  by complex-valued linear projections [9] and can then use the formulations as described in [10] and Equation 3. The necessary masking of future results in the training process as described in subsection 3.2 works in this framework by applying the mask after the respective mappings from the complex to the real domain (such as  $\mathcal{R}, \mathcal{I}, |\cdot|$ ).

We also compare to the approach of [5] which relies on splitting the product  $Q(K)^T$  into (real-valued) summands and applying real-valued attention per summand.

## 4.2. Complex-valued Layer Normalization

Normalization layers play a big role in the success of most neural network architectures. A complex-valued version has been proposed for batch normalization [9], however layer normalization is preferable for methods like LSTM or RNNs [32] as well as the transformer architecture [10]. It is insufficient to normalize the real and imaginary part of the complex-valued layer independently, since this may lead to very elliptic shapes of the output distribution [9]. This brings the need of a complex-valued version of layer normalization. The necessary building blocks are the complex-valued expected value and the covariance matrix. For some complex vector  $z \in \mathbb{C}^n$  these are defined as

$$\mathbb{E}(z) = \frac{1}{n} \sum_{i=1}^n z_i \quad (12)$$

$$\text{Cov}_{\mathbb{C}}(z) = \begin{pmatrix} \text{Var}(\mathcal{R}(z)) & \text{Cov}(\mathcal{R}(z), \mathcal{I}(z)) \\ \text{Cov}(\mathcal{R}(z), \mathcal{I}(z)) & \text{Var}(\mathcal{I}(z)) \end{pmatrix} \quad (13)$$

where  $\text{Var}$  and  $\text{Cov}$  denote the (real-valued) Variance and Covariance, respectively.

Let  $X$  be the output of a layer, the normalized output is then:

$$\begin{pmatrix} \mathcal{R}(\text{CLN}(X)) \\ \mathcal{I}(\text{CLN}(X)) \end{pmatrix} = \text{Cov}_{\mathbb{C}}^{-\frac{1}{2}}(X) \begin{pmatrix} \mathcal{R}(X - \mathbb{E}(X)) \\ \mathcal{I}(X - \mathbb{E}(X)) \end{pmatrix} \quad (14)$$

This compact form can easily be calculated with fast closed form solutions for the inverse and the square root of  $2 \times 2$  matrices. It is possible to manipulate the output distribution with learnable parameters. It can be shifted with a learnable parameter  $\beta \in \mathbb{C}$  and scaled with a learnable covariance matrix, a positive definite  $2 \times 2$  matrix. To ensure the positive definiteness of the resulting matrix, we utilize:

$$\begin{pmatrix} a & b \\ b & c \end{pmatrix} \text{ positive definite} \Leftrightarrow a > 0, c > 0, b^2 < ac \quad (15)$$

Thus, we can scale and shift the output  $\hat{X}$  of the layer normalization with 5 degrees of freedom by learning a covariance matrix  $\zeta$  and a shifting parameter  $\beta \in \mathbb{C}$  and get:

$$\begin{pmatrix} \mathcal{R}(\hat{X}) \\ \mathcal{I}(\hat{X}) \end{pmatrix} = \zeta^{\frac{1}{2}} \text{Cov}_{\mathbb{C}}^{-\frac{1}{2}}(X) \begin{pmatrix} \mathcal{R}(X - \mathbb{E}(X)) \\ \mathcal{I}(X - \mathbb{E}(X)) \end{pmatrix} + \beta \quad (16)$$

The output distribution then has covariance  $\zeta$  and expected value  $\beta$ .

## 5. EXPERIMENTAL RESULTS

Overall we perform two experiments: Automatic music transcription performed by the transformer encoder and a sequence generation task performed by the full transformer architecture. We compare the introduced methods for a complex-valued attention module as described in Equations 8-11 using the proposed Dot-Product as well as the version using  $Q(K^T)$ . Additionally, we compare to the approach of [5] and to the real-valued transformer as a baseline. For the latter, the real and imaginary part of the real-valued input was stacked alternating resulting in an input dimension of twice the original dimension. Both tasks are trained and evaluated on the MusicNet dataset [33]. The dataset consists of 330 pieces of music divided into 39438 samples consisting of 64 time steps, which are interpreted as one input token. These samples are split into 35111 training, 2030 validation and 3897 test samples, where the pieces of music between the splits do not overlap. We perform Fourier transform on the data

Architecture	Classification	Seq. generation
$\mathbb{C}$ -Transformer (ours)	14m	27m
Yang et al [5]	12m	20m
$\mathbb{R}$ -Transformer [10]	18m	33m

**Table 1:** Number of real-valued trainable parameters. For complex-valued parameters the real and imaginary parts count separately.

Attention	Classification		Sequence Generation	
	Dot-Prod	$Q(K^T)$	Dot-Prod	$Q(K^T)$
$\mathbb{C}$ Att	0.7164	0.7142	0.3272	0.3283
APAtt	0.6965	0.6926	0.2240	0.3231
AAtt	0.7117	0.7099	0.3172	0.3271
$\mathcal{RI}$ Att	0.7070	0.7059	0.3201	0.3236
Yang et al [5]	x	0.7088	x	0.3072
Real [10]	0.7109	x	0.0737	x

**Table 2:** Average precision results on test set for both tasks. Dot-Prod refers to the use of Dot-Product as described in subsection 4.1.

as preprocessing, as well as resampling as done in [5] with a method introduced by [34].

For both experiments the important hyperparameters are: Batch-size 35, 100 epochs, dropout 0.1, learning rate  $10^{-4}$ , embedding dimension 320, 6 layers with 8 attention heads and a hidden dimension in the feed forward module of 2048. As the encoder embedding we use a four layer complex-valued CNN followed by a fully connected layer. For the decoder embedding we used a fully connected embedding, since the input here are labels rather than a continuous signal. To test the impact of the convolutions on the encoder, we perform a small ablation study on just the proposed method (Equation 8) by removing the CNN in the encoder embedding.

### 5.1. Automatic music transcription task


For the multiclass classification problem, we classify into 128 classes, as offered in the dataset. The results in Table 2 show competitive behavior of most methods after 100 epochs. The best result by a slight margin is obtained by the proposed method as introduced in Equation 8. For most attention variants, it shows that the inner product version performs better, than using the product  $Q(K^T)$ . Additionally, all complex-valued architectures show improved robustness to overfitting (Figure 3), with no or minor decreases after longer training time, while the real transformer shows massive overfitting starting after 10 epochs.

### 5.2. Sequence generation task

For the sequence generation task, we split each sample into 43 input time steps and 21 time steps to be generated. The output to be generated are the notes of the missing 21 time steps of the samples in an iterative way, where in each iteration the input of the decoder is the output of the earlier iterations. The results show that the real transformer is not able to learn the sequence generation properly. The low training loss implies that the inability to learn is due to heavy overfitting. The robustness to overfitting already shown in subsection 5.1 seems to solve this problem, where all complex-valued methods learn reasonably well. The best performance is again obtained by the proposed method Equation 8. While some of the other methods introduced in this paper perform similarly well, the method of [5] performs noticeably worse.

	Classification	Seq. generation
$\mathbb{C}$ -Attention w/o conv.	0.5240	0.1652
$\mathbb{C}$ -Attention with conv.	0.7164	0.3272

**Table 3:** Average precision results on test set for both task as a small ablation study on the impact of the convolutional encoder.



**Fig. 3:** Training and validation Results. Average Precision Score (AVS) and Loss are displayed. Left: Classification, right: Sequence generation. "DP" denotes Dot-Product, "QK" denotes  $Q(K^T)$ .

### 5.3. Discussion

Table 3 shows that the complex-valued transformer architecture without convolutions is able to learn a meaningful solution, but the CNN in the encoder is necessary for state-of-the-art results.

Overall, the real transformer struggles with overfitting on both tasks. The complex-valued transformers improve in this regard while maintaining on-par (subsection 5.1) or superior (subsection 5.2) performance. This result is in line with earlier results, showing superior robustness to overfitting for, e.g., CNNs [35] or RCNNs [36]. We are the first ones to show this for the transformer architecture.

## 6. CONCLUSION

We presented building blocks for a complex-valued transformer architecture. That includes newly developed formulations of complex-valued attention mechanisms as well as a complex-valued layer normalization. We have shown that it improves robustness to overfitting on a classification and a sequence generation task, while maintaining competitive performance compared to the real-valued algorithm when applied to complex-valued signals. This opens up the opportunity to incorporate the transformer architecture into a more broad class of applications that naturally make use of complex-valued images. Additionally, the complex-valued Fourier transform of signals can now directly be used in the transformer architecture without using the isomorphism  $\mathbb{C} \rightarrow \mathbb{R}^2$  that results in a loss in robustness against overfitting. This work also serves as a base for the further development of complex-valued versions of extensions to the transformer [27] and other architectures that use the attention mechanism.

## 7. REFERENCES

- [1] E. Cole, J. Cheng, J. Pauly, and S. Vasanawala, "Analysis of deep complex-valued convolutional neural networks for MRI reconstruction and phase-focused applications," *Magnetic Resonance in Medicine*, vol. 86, no. 2, pp. 1093–1109, 2021.
- [2] B. Vasudeva, P. Deora, S. Bhattacharya, and P. M. Pradhan, "Compressed sensing mri reconstruction with Co-VeGAN: Complex-valued generative adversarial network," in *WACV*, 2022, pp. 1779–1788.
- [3] X. Li, Q. Sun, L. Li, X. Liu, H. Liu, L. Jiao, and F. Liu, "SSCV-GANs: Semi-supervised complex-valued GANs for PolSAR image classification," *IEEE Access*, vol. 8, pp. 146560–146576, 2020.
- [4] H.-S. Choi, J.-H. Kim, J. Huh, A. Kim, J.-W. Ha, and K. Lee, "Phase-aware speech enhancement with deep complex U-Net," in *ICLR*, 2018.
- [5] M. Yang, M. Q. Ma, D. Li, Y.-H. H. Tsai, and R. Salakhutdinov, "Complex transformer: A framework for modeling complex-valued sequence," in *ICASSP*, 2020, pp. 4232–4236.
- [6] Y. Yang and S. Soatto, "FDA: Fourier domain adaptation for semantic segmentation," in *CVPR*, 2020, pp. 4084–4094.
- [7] Q. Xu, R. Zhang, Y. Zhang, Y. Wang, and Q. Tian, "A Fourier-based framework for domain generalization," in *CVPR*, 2021, pp. 14383–14392.
- [8] D. Yin, C. Luo, Z. Xiong, and W. Zeng, "PHASEN: A phase-and-harmonics-aware speech enhancement network," in *AAAI*, 2020.
- [9] C. Trabelsi, O. Bilaniuk, Y. Zhang, D. Serdyuk, S. Subramanian, J. Felipe Santos, S. Mehri, N. Rostamzadeh, Y. Bengio, and C. J. Pal, "Deep complex networks," in *ICLR*, 2018.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
- [11] M. Kataoka, M. Kinouchi, and M. Hagiwara, "Music information retrieval system using complex-valued recurrent neural networks," in *IEEE SMC*, 1998, pp. 4290–4295.
- [12] Y. Kuroe, N. Hashimoto, and T. Mori, "On energy function for complex-valued neural networks and its applications," in *ICONIP*, 2002.
- [13] A. Hirose, *Complex-Valued Neural Networks: Theories and Applications*, World Scientific, 2003.
- [14] J. Bassey, L. Qian, and X. Li, "A survey of complex-valued neural networks," *arXiv:2101.12249*, 2021.
- [15] Y.-P. Zhang, Q. Zhang, Le Kang, Y. Luo, and L. Zhang, "End-to-end recognition of similar space cone-cylinder targets based on complex-valued coordinate attention networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2021.
- [16] S. Ren and F. Zhou, "Polsar image classification with complex-valued residual attention enhanced U-Net," in *IGARSS*, 2021, pp. 3045–3048.
- [17] P. Virtue, X. Yu Stella, and M. Lustig, "Better than real: Complex-valued neural nets for MRI fingerprinting," in *ICIP*, 2017, pp. 3953–3957.
- [18] A. Hirose and S. Yoshida, "Generalization characteristics of complex-valued feedforward neural networks in relation to signal coherence," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 4, pp. 541–551, 2012.
- [19] Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "DCCRN: Deep complex convolution recurrent network for phase-aware speech enhancement," in *INTERSPEECH*, 2020.
- [20] H. Zhang et al., "An optical neural chip for implementing complex-valued neural network," *Nature Communications*, vol. 12, pp. 457, 2021.
- [21] D. W. Otter, J. R. Medina, and J. K. Kalita, "A survey of the usages of deep learning for natural language processing," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 2, pp. 604–624, 2020.
- [22] D. Hu, "An introductory survey on attention mechanisms in NLP problems," in *IntelliSys*, 2019.
- [23] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.
- [24] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, "Transformers in vision: A survey," *ACM Computing Surveys*, vol. 54, pp. 200:1–200:41, 2021.
- [25] C.-Z. A. Huang, A. Vaswani, J. Uszkoreit, I. Simon, C. Hawthorne, N. Shazeer, A. M. Dai, M. D. Hoffman, M. Dinulescu, and D. Eck, "Music transformer: Generating music with long-term structure," in *ICLR*, 2019.
- [26] T. Lin, Y. Wang, X. Liu, and X. Qiu, "A survey of transformers," *AI Open*, vol. 3, pp. 111–132, 2022.
- [27] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler, "Efficient transformers: A survey," *ACM Computing Surveys*, vol. 55, no. 6, pp. 109:1–109:28, 2023.
- [28] Y. Dong, Y. Peng, M. Yang, S. Lu, and Q. Shi, "Signal transformer: Complex-valued attention and meta-learning for signal recognition," *arXiv:2106.04392*, 2021.
- [29] H.-W. Cho, S. Choi, Y.-R. Cho, and J. Kim, "Complex-valued channel attention and application in ego-velocity estimation with automotive radar," *IEEE Access*, vol. 9, pp. 17717–17727, 2021.
- [30] B. Wang, D. Zhao, C. Lioma, Q. Li, P. Zhang, and J. G. Simonsen, "Encoding word order in complex embeddings," in *ICLR*, 2020.
- [31] X. Chu, Z. Tian, B. Zhang, X. Wang, X. Wei, H. Xia, and C. Shen, "Conditional positional encodings for vision transformers," *arXiv:2102.10882*, 2021.
- [32] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv:1607.06450*, 2016.
- [33] J. Thickstun, Z. Harchaoui, and S. M. Kakade, "Learning features of music from scratch," in *ICLR*, 2017.
- [34] J. O. Smith, "Digital audio resampling home page," <https://ccrma.stanford.edu/~jos/resample/>, 2020.
- [35] N. Guberman, "On complex valued convolutional neural networks," *arXiv:1602.09046*, 2016.
- [36] F. Zhao, G. Ma, W. Xie, and H. Liu, "Semi-supervised recurrent complex-valued convolution neural network for polsar image classification," in *IGARSS*, 2019.