

Implicit Kernel Attention

Kyungwoo Song¹, Yohan Jung², Dongjun Kim², Il-Chul Moon^{*2}

¹ Department of AI, University of Seoul

² Department of ISysE, Korea Advanced Institute of Science and Technology (KAIST)
kyungwoo.song@uos.ac.kr, {becre1776,dongjoun57,icmoon}@kaist.ac.kr

Abstract

Attention computes the dependency between representations, and it encourages the model to focus on the important selective features. Attention-based models, such as Transformer and graph attention network (GAT), are widely utilized for sequential data and graph-structured data. This paper suggests a new interpretation and generalized structure of the attention in Transformer and GAT. For the attention in Transformer and GAT, we derive that the attention is a product of two parts: 1) the RBF kernel to measure the similarity of two instances and 2) the exponential of L2 norm to compute the importance of individual instances. From this decomposition, we generalize the attention in three ways. First, we propose implicit kernel attention with an implicit kernel function instead of manual kernel selection. Second, we generalize L2 norm as the L_p norm. Third, we extend our attention to structured multi-head attention. Our generalized attention shows better performance on classification, translation, and regression tasks.

Introduction

Attention (Bahdanau, Cho, and Bengio 2014) is widely utilized to improve model performance as well as to explain the model mechanisms. The attention in Transformer computes the dot-product between query and key, which are the linear projections of hidden features. The attention in GAT also utilizes the dot-product between a weight vector and a concatenation of hidden representations. It is well known that the dot-product of two vectors is a product of two distinct terms: 1) the cosine of the angle between two vectors, which computes the similarity; and 2) the individual norm of two vectors, which measures the magnitude of each vector. This opens a question on how to analyze the attention with the explicitly separated terms of the similarity and the magnitude, and how to generalize the separation in attentions.

We provide a new interpretation of attention, and this interpretation leads to generalized attention by formalizing the attention weight into a multiplication of two terms: similarity and magnitude. We derive the explicit separation, so the derivation reveals that the attention in Transformer and GAT is a product of 1) the Radial Basis Function (RBF) kernel

between instances and 2) the exponential of L^2 norm for each instance.

This paper proposes an implicit kernel function that generalizes the RBF kernel function, embedded in the attention. Given that attention uses a kernel implicitly, we note that the kernel function measures the similarity under the inductive bias embedded in the kernel. Traditionally, the modelers manually selected kernels with the expert knowledge for a specific domain or a dataset. We construct an implicit kernel function instead of an explicit kernel to find a data-adaptive kernel, which provides better task performances under the data context. We formulate the implicit kernel function by constructing the spectral density depending on the dataset because a kernel construction can be interpreted as the spectral density estimation (Reed and Simon 1975; Yaglom 1987). This paper further develops the formulation by introducing multi-head implicit attention with a structured spectral density estimation.

The magnitude term is the second component in the decomposition of attention, and the magnitude separately measures the importance of each instance pair by an exponential of their L^2 norms. As L^2 norm in the attention is rigid without considering the dataset property, we extend L^2 as the L^p norm with a hyper-parameter p to control the scale of magnitude terms and attention weight sparsity. The p controls the growth rate of the magnitude terms, which are the scale of individual importance. By treating p as a hyper-parameter, we can impose an inductive bias to focus on the relative importance between inputs (similarity) on translation tasks; or the individual importance on each input (magnitude) for the classification task. Besides, we find that the decrement of p is eventually related to the sparsity of the attention weights, empirically and theoretically.

In summary, we derive the decomposition of attention as a product of similarity and magnitude. Under the decomposition, we generalize the attention with an implicit kernel function and a L^p norm. Further, we extend our works to structured multi-head attention with the implicit kernel functions. Our generalized implicit attention is exchangeable with the attention in Transformer and GAT without increasing the algorithmic complexity in order.

*Corresponding author.

Preliminary

Transformer The attention layers in Transformer compute the importance of each feature to consider the semantics of the given sequence. Transformer calculates the importance weight by the dot-product of $q_i = h_i W_Q$ and $k_j = h_j W_K$ with the hidden feature h , the query q_i , the key k_j , and linear projection matrix W_Q, W_K (Vaswani et al. 2017). The attention weight α_{ij} in Eq. 1 utilizes the softmax function with a scaling $1/\sqrt{d_k}$ where d_k is the dimension of k_j

$$\alpha_{ij} = \frac{\exp(q_i^T k_j / \sqrt{d_k})}{\sum_l \exp(q_i^T k_l / \sqrt{d_k})} \quad (1)$$

Graph Attention Network GAT (Veličković et al. 2018) adopts the attention in Eq. 2 to aggregate the relevant neighborhood’s features. GAT transforms each hidden feature, h_i and h_j , with a linear projection matrix, W . GAT aggregates features by concatenating the transformed hidden features. After the concatenation, GAT computes the dot-product between a weight vector, a , and aggregated features. GAT adopts LeakyReLU (Maas, Hannun, and Ng 2013), and it should be noted that LeakyReLU has a different slope parameter, c , for the positive region and the negative region. GAT utilizes the softmax function for the neighbor of a node i , N_i .

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(a^T [Wh_i || Wh_j]))}{\sum_{l \in N_i} \exp(\text{LeakyReLU}(a^T [Wh_i || Wh_l]))} \quad (2)$$

Multi-Head Attention Both Transformer and GAT extend the attention to Multi-Head Attention (MHA) to capture the diverse features by introducing the different linear projection matrix. The original MHA assumes the independence between multi-heads, and this paper provides a dependence structure by using the spectral densities of multi-heads.

Kernel

Spectral Density Due to the dual property given by Bochner theorem (Reed and Simon 1975), a stationary kernel, $k(x, x')$, is selected uniquely by a corresponding spectral density, $p(w)$, and vice versa, as in Eq. 3. The hypothesis class of the spectral density function determines the adaptiveness of a kernel, and the explicit specification on the density function class could restrict the adaptiveness. Hence, we provide an implicit density function class through an implicit generative model. For instance, the generator in Generative Adversarial Network (GAN) (Goodfellow et al. 2014) is able to represent spectral density, flexibly (Li et al. 2019). Generally, Eq. 3 with a spectral density, $p(w)$, is intractable, so we approximate Eq. 3 by the Monte Carlo (MC) integration with R sampled spectral points, w_r , as shown in Eq. 4 (Ton et al. 2018).

$$k(x, x') = \int_{\mathcal{R}^d} \exp(iw^T (x - x')) p(w) dw \quad (3)$$

$$\approx \frac{1}{R} \sum_{r=1}^R \exp(iw_r^T (x - x')) \quad (4)$$

Kernel in Attention Tsai et al. (Tsai et al. 2019) utilize a kernel to formulate the attention weights. They replace the $\exp(q_i^T k_j / \sqrt{d_k})$ in scaled dot-product attention with a manually selected kernel, such as the RBF kernel, without an explicit decomposition. Because they do not derive the decomposition explicitly, they only utilize the RBF kernel for a similarity term without a magnitude term. We name it RBF-only for the comparison.

The concurrent works (Katharopoulos et al. 2020; Choromanski et al. 2020b,a) focus on efficiency and propose linear time complexity attention based on the relationship between kernel and attention. This paper focuses on the new interpretation and generalization of attention to improve the capacity of attention in Transformer and GAT. We emphasize the improved capacity of attention with linear time complexity can be reached by combining our approach with the concurrent works.

Copula

The original MHA in Transformer and GAT lacks an explicit model to handle the dependency between multiple attentions. Therefore, we formulate a copula-augmented spectral density estimation to consider the dependency between the spectral densities in each head. The copula is a cumulative distribution function (CDF) defined on the unit cube with uniform marginals (Nelsen 2007). Formally, a copula, C , is defined as $C(u_1, \dots, u_d) = P(U_1 \leq u_1, \dots, U_d \leq u_d)$, where the marginal distribution on U_i is uniformly defined on $[0, 1]$. By Sklar’s theorem (Sklar 1959), we can represent a joint cumulative distribution of x_1, \dots, x_d as a marginal distribution of each random variables and their copula, $F(x_1, \dots, x_d) = C[F_1(x_1), \dots, F_d(x_d)]$. We impose the dependencies between attentions in MHA by estimating each head’s spectral density jointly with a copula augmented inference.

Methodology

Decomposition of Transformer and GAT

This section derives the decomposition of attention in Transformer and GAT as the product of two distinct terms: 1) similarity term and 2) magnitude term. The similarity measures the relative importance, and the magnitude computes the individual importance. We derive the decomposition for the scaled dot-product attention in Proposition 1.

Proposition 1. Let α_{ij} be $\frac{1}{Z_1(\alpha)} \exp\left(\frac{q_i^T k_j}{\sqrt{d_k}}\right)$ where $Z_1(\alpha)$ is a normalizing constant. Then α_{ij} has the form:

$$\frac{1}{Z_1(\alpha)} \times \underbrace{\exp\left(\frac{-\|q_i - k_j\|_2^2}{2\sqrt{d_k}}\right)}_{\text{similarity}} \times \underbrace{\exp\left(\frac{\|q_i\|_{p=2}^2 + \|k_j\|_{p=2}^2}{2\sqrt{d_k}}\right)}_{\text{magnitude}}$$

The scaled dot-product attention has the similarity term, $f_{RBF} = \exp\left(\frac{-\|q_i - k_j\|_2^2}{2\sqrt{d_k}}\right)$, which is the RBF kernel with fixed length-scale hyper-parameters, $\sqrt{d_k}$. Additionally, the scaled dot-product attention yields the magnitude term, $\exp\left(\frac{\|q_i\|_{p=2}^2 + \|k_j\|_{p=2}^2}{2\sqrt{d_k}}\right)$, and it measures the individual importance of each instance as $\|q_i\|_{p=2}^2$ and $\|k_j\|_{p=2}^2$ with L^2

norm. A large similarity term and a large magnitude term induce the large attention weight for the given pair of query and key.

On GAT, we formulate the decomposition for the attention in Eq. 2. The decomposition in Proposition 2 is slightly different from the decomposition of Transformer. For an attention weight between node i and node j , GAT computes the two similarity terms: 1) the similarity between q_i and ca , 2) the similarity between ca and k_j . From the decomposition, we can interpret the weight vector, a , as the medium that connects q_i and k_j .

Proposition 2. Define $\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(a^T [Wh_i || Wh_j]))}{Z_2(\alpha)}$, $q_i = [Wh_i || \vec{0}]$, and $k_j = [\vec{0} || Wh_j]$ where $Z_2(\alpha)$ is a normalizing constant, and c be the slope parameter in Leaky-ReLU. Then α_{ij} has the form:

$$\frac{1}{Z_2(\alpha)} \times \underbrace{\exp\left(\frac{-\|q_i - ca\|_2^2}{2}\right) \times \exp\left(\frac{-\|ca - k_j\|_2^2}{2}\right)}_{\text{similarity}} \times \underbrace{\exp\left(\frac{2\|ca\|_{p=2}^2 + \|q_i\|_{p=2}^2 + \|k_j\|_{p=2}^2}{2}\right)}_{\text{magnitude}}$$

Variations of Implicit Kernel Attention

From the factorization, we propose a new attention method, which formulates an implicit kernel function, and which utilizes a L^p norm. Attention in Transformer and GAT uses a fixed single Gaussian spectral density, $p(w)$, and its corresponding RBF kernel in Figure 1a. In contrast, we propose implicit kernel attention (IKA) that estimates the spectral density implicitly to find an appropriate kernel depending on the dataset in Figure 1b. Second, we interpret p in L^p norm as a hyper-parameter, so we define IKA with norm (IKAN). Besides, we propose a deterministic and simple but effective model, IKAN-direct, to optimize spectral points w , directly in Figure 1c. Lastly, we propose multi-head IKAN (MIKAN), which adopts a copula-augmented inference to estimate a structured spectral density of MHA, jointly in Figure 1d. Figure 1e represents the structure overview of our models.

IKA: Implicit Kernel Attention

RBF measures the similarity by the Euclidean distance. Recent works (Li and Dunson 2019; Chen et al. 2018) conjecture that the Euclidean metric might not be proper in the hidden feature space. Furthermore, RBF includes the exponential term, and the range of exponential is a positive real-value without zero. Therefore, the sparsity cannot be achieved in the attention weight because of the positive nature.

By following (Reed and Simon 1975; Yaglom 1987), the implicit spectral density allows us to construct an implicit kernel, so this paper proposes learning an implicit spectral density from the dataset. To formulate an implicit spectral density, we construct a base distribution, $p(z)$, such as a

Gaussian distribution, and we sample z from $p(z)$. Then, we transform the sampled z with a flexible function such as a neural network to construct the spectral density, $p(w)$. It should be noted that the spectral density should be symmetric for the real-valued kernel, which we add constraints in Eq. 8.

Under the implicit spectral density, the log marginal likelihood, $\log p(y|h)$, is intractable where y is an output variable and h is a hidden feature. Therefore, we derive the evidence lower bound (ELBO) of the log marginal likelihood with an inference network, $q(z|h)$. We maximize the ELBO in Eq. 5 with a re-parameterization method to backpropagate the gradients (Kingma and Welling 2014).

$$\mathcal{L} = \mathbb{E}_{q(z|h)} [\log p(y, z|h)] - \mathbb{E}_{q(z|h)} [\log q(z|h)] \quad (5)$$

For MHA in Transformer or GAT, we independently sample z from $q(z|h)$ for each head by formulating $q(z|h) = \prod_{m=1}^M q(z^{(m)}|h^{(m)})$. We construct each inference network, $q(z^{(m)}|h^{(m)})$, with a neural network parameterized by ψ_1 as in Eq. 6. We formulate $z^{(m)}$ as a concatenation of $\tilde{z}^{(m)}$ and $-\tilde{z}^{(m)}$ to preserve the symmetric of a base distribution.

$$\mu^{(m)}, \log \sigma^{(m)} = NN_{\psi_1}(h^{(m)}) \quad (6)$$

$$\tilde{z}^{(m)} \sim \mathcal{N}(\mu^{(m)}, (\sigma^{(m)})^2 \mathbf{I}), \quad z^{(m)} = [\tilde{z}^{(m)} || -\tilde{z}^{(m)}] \quad (7)$$

A symmetric base distribution enables us to construct an unconstrained neural network NN_{ψ_2} to transform z into w . After sampling $z^{(m)}$, we transform it with NN_{ψ_2} to construct the spectral density in Eq. 8. To preserve the symmetricity of the spectral density, we use the absolute value of $|z^{(m)}|$ as an input, and we multiply the sign of $z^{(m)}$ (Li et al. 2019).

$$w^{(m)} = \text{sign}(z^{(m)}) \times NN_{\psi_2}(|z^{(m)}|) \quad (8)$$

Eq. 9 and Eq. 10 denote the formulation of un-normalized attention weights, $\tilde{\alpha}_{ij}^{(m)}$, with the implicit kernel function f for Transformer and GAT, respectively. From the un-normalized attention weights, $\tilde{\alpha}_{ij}^{(m)}$; we compute attention weights, $\alpha_{ij}^{(m)}$, with normalization in Eq. 11. The implicit kernel function, f , is a similarity term of attention in Proposition 1 and 2. The rest of the structure is the same as the original Transformer and GAT with a neural network NN_{ψ_3} in Eq. 11.

$$\tilde{\alpha}_{ij}^{(m)} = f(q_i^{(m)}, k_j^{(m)}, w^{(m)}) \times \exp\left(\frac{\|q_i^{(m)}\|_2^2 + \|k_j^{(m)}\|_2^2}{2\sqrt{d_k}}\right) \quad (9)$$

$$\tilde{\alpha}_{ij}^{(m)} = f(q_i^{(m)}, ca^{(m)}, w^{(m)}) \times f(ca^{(m)}, k_j^{(m)}, w^{(m)}) \times \exp\left(\frac{2\|ca^{(m)}\|_2^2 + \|q_i^{(m)}\|_2^2 + \|k_j^{(m)}\|_2^2}{2}\right) \quad (10)$$

$$\hat{y} = NN_{\psi_3}(\alpha^{(1:M)}, h^{(1:M)}), \quad \alpha_{ij}^{(m)} = \frac{\tilde{\alpha}_{ij}^{(m)}}{\sum_l \tilde{\alpha}_{il}^{(m)}} \quad (11)$$

IKAs: Implicit Kernel Attention with Stationary Kernel
This section explains the construction of kernel function f ,




Figure 1: Visualization of the attention and our models. To capture the spectral density flexibly, we first sample z from the base distribution $p(z)$ and transform it to w with a flexible function such as a neural network. The IKA_S , IKA_{NS} and IKAN in (b) estimate the base distribution and the spectral density in each head depending on the dataset, unlike the scaled dot-product attention. IKAN-direct in (c) optimizes spectral points w directly, and MIKAN in (d) estimates the base distribution $p(z^{(1:M)})$ jointly. We provide a structure overview of our model in (e).

and we omit the multi-head related notation, m , for simplicity. We approximate a continuous stationary kernel with the MC integration with R sampled spectral points, w_r , and its random Fourier feature map, ϕ_r , by taking the real part in Eq. 4 (Ton et al. 2018; Rahimi and Recht 2008; Jung, Song, and Park 2020).

$$f(q_i, k_j, w) = \frac{1}{R} \sum_{r=1}^R \phi_r(q_i)^T \phi_r(k_j) \quad (12)$$

$$\phi_r(q_i) = \begin{pmatrix} \cos(w_r^T q_i) \\ \sin(w_r^T q_i) \end{pmatrix}, \phi_r(k_j) = \begin{pmatrix} \cos(w_r^T k_j) \\ \sin(w_r^T k_j) \end{pmatrix} \quad (13)$$

We generalize the attention by replacing the RBF in the scaled dot-product attention with an implicit kernel function, f as shown in Eq. 9. We name this attention as IKA_S . Similarly, we can construct the implicit kernel, f , for GAT in Eq. 10 by following Eq. 12 and 13.

In spite of the above kernel construction, the output of the kernel function in Eq. 12 might have a negative value, so it contradicts with the non-negative property of attention weights. As a result, we square Eq. 12, f^2 , to ensure the non-negativity of IKA . Proposition 3 claims that the squared implicit kernel function still generalizes the RBF kernel that is utilized in the scaled dot-product attention.

Proposition 3. Let $k_{RBF}^{(l)}$ be the RBF kernel function with a lengthscale l , and f be its approximation with R random Fourier features by Eq. 12 and 13. If we sample spectral points, w , from a Gaussian distribution $\mathcal{N}(0, \frac{1}{2l^2}\mathbf{I})$, then $\lim_{R \rightarrow \infty} f^2 = k_{RBF}^{(l)}$.

IKA_{NS} : Implicit Kernel Attention with Non-Stationary Kernel Similar to the stationary kernel, we can approximate a continuous non-stationary kernel with the random Fourier feature map ϕ_r in Eq. 14-16 (Ton et al. 2018; Yaglom 1987). We formulate a new attention, IKA_{NS} , induced by a non-stationary kernel function in Eq. 14-16. Different from stationary kernel, we sample two types of R spectral points, $w_{1,r}$ and $w_{2,r}$, from two distinct implicit spectral density estimators. Similarly, we construct an implicit kernel, f , for

GAT in Eq. 10 by following Eq. 14-16.

$$f(q_i, k_j, w) = \frac{1}{4R} \sum_{r=1}^R \phi_r(q_i)^T \phi_r(k_j) \quad (14)$$

$$\phi_r(q_i) = \begin{pmatrix} \cos(w_{1,r}^T q_i) + \cos(w_{2,r}^T q_i) \\ \sin(w_{1,r}^T q_i) + \sin(w_{2,r}^T q_i) \end{pmatrix} \quad (15)$$

$$\phi_r(k_j) = \begin{pmatrix} \cos(w_{1,r}^T k_j) + \cos(w_{2,r}^T k_j) \\ \sin(w_{1,r}^T k_j) + \sin(w_{2,r}^T k_j) \end{pmatrix} \quad (16)$$

The error bound for a non-stationary kernel approximation is not well explored. To validate IKA_{NS} , we show that the sufficient enough samples from a spectral density guarantee the non-stationary kernel approximation in Proposition 4.

Proposition 4. Let $k(x, x')$ be a non-stationary kernel, $f(x, x')$ be an approximated kernel with R sampled spectral points, w_1 and w_2 , and \mathcal{M} be a compact subset of \mathcal{R}^d with a diameter D . Let $\sigma_1^2 = \mathbb{E}[w_1^T w_1]$, $\sigma_2^2 = \mathbb{E}[w_2^T w_2]$ and $\mathcal{E} = \sup_{x, x' \in \mathcal{M}} |f(x, x') - k(x, x')|$. Then,

- i) $P[\mathcal{E} \geq \epsilon] \leq 2^8 \times \left(\frac{D\sqrt{\sigma_1^2 + \sigma_2^2}}{\epsilon} \right)^2 \times \exp\left(\frac{-R\epsilon^2}{2(d+1)}\right)$
- ii) $\mathcal{E} \leq \epsilon$ with any constant probability when $R = \Omega\left(\frac{d}{\epsilon^2} \log \frac{D\sqrt{\sigma_1^2 + \sigma_2^2}}{\epsilon}\right)$

Similar to the IKA_S , we use f^2 instead of f to ensure the non-negativity of the attention weights.

IKAN : Implicit Kernel Attention with L^p

IKAN In addition to the implicit kernel in IKA_{NS} , we generalize the L^2 norm in the magnitude term in Proposition 1 and 2 to be the L^p , which determines the importance of individual representations. The optimal p might be different for each dataset and task, so we select p through experiments. There are tasks where the similarity is more important than the magnitude and vice versa. Treating p as a hyper-parameter that determines the scale of magnitude term imposes an inductive bias on the model. Furthermore, we analyze the relationship between the magnitude and the sparsity of attention weights. For the analysis, we define attention weights, $\alpha_{ij}(p)$, whose magnitude utilizes $\|x\|_p = (|x_1|^p + \dots + |x_d|^p)^{1/p}$ instead

of $\|x\|_2$. It should be noted that $\|x\|_p$ is an absolutely homogeneous function for $p > 0$. We show that $\alpha_{ij}(p)$ becomes sparse when p goes to zero in Proposition 5.

Proposition 5. Let $\mathbb{1}_{A_p^{(i)}}$ be the multi-dimensional indicator function with n -th component to be 1 if $n \in A_p^{(i)}$ and 0 otherwise, where $A_p^{(i)}$ is the set of index for maximal elements: $A_p^{(i)} = \{l | l = \operatorname{argmax}_j \alpha_{ij}(p)\}$. Then, $\lim_{p \rightarrow 0^+} \alpha_{ij}(p) = \mathbb{1}_{A_p^{(i)}} / |A_p^{(i)}|$

IKAN-direct We propose a simple and effective alternative that sets spectral points, w , as a learnable parameter, such as a weight in a neural network. IKAN-direct optimizes spectral points, w , directly without sampling. It corresponds to handling the spectral density as a mixture of delta distribution.

MIKAN: Multi-Head Attention with IKAN

The dependency modeling between heads in MHA is necessary to diversify the attention weights. Similar to the individual calculation of the attention in Transformer and GAT, the variations of IKA estimate the spectral density in each head independently. Therefore, we introduce $q(z|h) = c_q(F_1(z^{(1)}), \dots, F_M(z^{(M)})) \prod_{m=1}^M q(z^{(m)}|h^{(m)})$, where c_q is a copula density; F_m is a CDF of each $z^{(m)}$; M is the number of attention head. This alternative variational distribution of z , $q(z|h)$, develops IKA to be structured as Multi-head IKAN (MIKAN) by introducing the joint structure through c_q . MIKAN maximizes ELBO in Eq. 5 with a Monte Carlo estimation by sampling from $q(z|h)$ (Kingma and Welling 2014). MIKAN estimates the base distribution without the mean-field assumption by introducing a copula-augmented posterior. IKAN is a special case of MIKAN if we fix c_q as a uniform distribution (Tran, Blei, and Airoldi 2015). c_q can be any copula density, and we set c_q as a Gaussian copula with a covariance, Σ .

Model Complexity

Table 1 denotes the time complexity of computing the attention weights, α_{ij} , from a hidden representation, h , where T is the length of a sequence or the number of nodes.

If we consider a task with a long sequence with a high T , there is no asymptotic complexity increment in order, so all variations of IKA and baselines fall into the complexity of $O(T^2)$. Also, it should be noted that the complexity of $O(T^2R)$ is controllable because R is the number of samples determined by a modeler.

Results

Sentence Classification

We compare our models with the scaled dot-product attention in Transformer (Vaswani et al. 2017), and RBF-only (Tsai et al. 2019) on five popular datasets (Kim 2014; Wang et al. 2020). We perform ten-fold cross validations by following the experimental settings (Wang et al. 2020). RBF-only represents the scaled dot-product attention with a RBF kernel without a magnitude term. To validate the performance for

	Model	Complexity
Baselines	Transformer	$O(T^2d + Td^2)$
	RBF-only	$O(T^2d + Td^2)$
Ours	Expsin	$O(T^2d + Td^2)$
	Linear	$O(T^2d + Td^2)$
	IKAS	$O(T^2R + Td^2 + TdR)$
	IKANS	$O(T^2R + Td^2 + TdR)$
	IKAN	$O(T^2R + Td^2 + TdR)$
	IKAN-direct	$O(T^2R + Td^2 + TdR)$
	MIKAN	$O(T^2R + Td^2 + TdR + dM^3)$

Table 1: Time Complexity for computing attention weights where T, d, R, M denotes the length of sequence, hidden dimension, the number of sample size, and the number of attention heads respectively.

the different kernels, we provide additional models, Expsin and Linear. Expsin and Linear denote the scaled dot-product attention that uses periodic kernel and linear kernel, instead of RBF kernel, respectively. For the magnitude term, Expsin and Linear utilize L^2 norm that is the same as the scaled dot-product attention.

Table 2 indicates that the appropriate kernel is different for each dataset. The RBF kernel in Transformer performs better than other baselines on MR and SUBJ. On the other hand, Expsin shows better performance on CR and SST, and Linear is superior to RBF on TREC. However, IKAS and IKANS show consistently better performance than RBF-only, Expsin, and Linear. This consistent tendency demonstrates the importance of data-adaptive implicit kernel empirically, instead of a manual kernel selection. Compared to IKANS, IKAN shows better performances, and this supports the importance of adapting the magnitude with p as a hyper-parameter. IKAN-direct, which optimizes w directly without sampling, shows lower performance than IKAN, but it still performs better than the scaled dot-product attention in Transformer. MIKAN with copula augmentation performs best except the SUBJ dataset.

The second column of Table 2 represents the number of parameters for each model on the MR dataset. Expsin and Linear have the same parameters as the scaled dot-product in Transformer. For the implicit kernel-based model, we adopt an implicit spectral density estimator, and it requires additional parameters. Because we utilize the shared spectral density estimator for each layer, the implicit kernel-based models require a marginal increment in the number of parameters.

Figure 2a and 2b show attention weights for a sentence in MR. In Figure 2a, our implicit kernel-based attention captures the important words *kiss* and *waste*. Besides, MIKAN focuses on the additional important word, *just*. Figure 2b represents the attention weights of the first four heads among eight heads in Transformer. Each row and column of the square matrix corresponds to the individual word in the given sentence. All first four heads in IKAN focus on the same last word *waste*. The first, second, and fourth heads in IKAN also focus on the words *kiss* together. On the other hand, MIKAN that adopts structured MHA shows the relatively diverse attention weights across the multi-heads.

Model	# Param.	MR	CR	SST	SUBJ	TREC
Transformer*	-	73.1%	76.2%	76.1%	86.3%	83.4%
Transformer	27.9M	74.45±0.94%	77.98±3.16%	76.25±1.60%	90.32±0.68%	84.12±0.77%
RBF-only	27.9M	73.98±0.48%	78.25±3.02%	76.00±1.82%	89.82±1.16%	84.24±0.90%
Expsin	27.9M	73.81±0.98%	78.78±3.26%	76.99±1.21%	89.10±0.28%	84.76±0.86%
Linear	27.9M	73.70±0.49%	78.57±2.86%	76.67±1.12%	89.22±0.85%	84.76±1.41%
IKAS	28.1M	76.06±0.48%	78.83±3.26%	77.37±1.68%	90.66±0.87%	84.88±1.29%
IKANS	28.2M	76.42±0.82%	79.05±2.26%	77.49±1.84%	90.86±0.60%	85.32±0.94%
IKAN	28.2M	76.63±0.89%	79.21±3.05%	77.49±1.84%	91.48±0.61%	85.84±0.43%
IKAN-direct	28.3M	75.97±0.54%	79.26±3.05%	77.10±1.68%	90.98±0.53%	85.16±0.38%
MIKAN	28.2M	76.70±1.03%	80.63±3.24%	77.69±1.85%	90.86±0.61%	85.92±0.44%

Table 2: Accuracy for the text classification task. * denotes the reported performance in (Wang et al. 2020). # Param. denotes the number of parameters for the MR dataset.




Figure 2: (a): Attention weights for the given sentence on MR, whose task is the sentiment classification of a movie review. The sentiment of the given sentence "just a kiss is a just a waste" is negative. To identify the sentiment, the model should focus on the *just*, *kiss*, *just*, and *waste*, and MIKAN captures all important words. (b): Attention weights matrix of first four heads among eight heads for IKAN and MIKAN. MIKAN shows relatively diverse attention weights across the multi-head because of the structured model with the copula.

Translation

We compare our models and baselines on IWSLT14 De-En. We include MIKAN and IKAN-direct because MIKAN is a generalized model of IKA and IKAN. Table 3 shows the results of three runs for Transformer, IKAN-direct and

Model	BLEU
Beam Search Optimization*	26.36
Actor-Critic*	28.53
Neural PBMT + LM*	30.08
Minimum Risk Training*	32.84
Variational Attention*	33.69
Transformer	34.44 ± 0.07
IKAN-direct	34.59 ± 0.09
MIKAN	34.70 ± 0.09

Table 3: BLEU for IWSLT14 De-En. * denotes the reported performance in (Deng et al. 2018).

MIKAN. Our models perform better than other models, Beam Search Optimization (Wiseman and Rush 2016), Actor-Critic (Bahdanau et al. 2017), Neural PBMT + LM (Jang, Gu, and Poole 2017), Minimum Risk Training (Edunov et al. 2018), Variational Attention (Deng et al. 2018), and Transformer.

We analyze the changes in similarity and magnitude terms in Figure 3. Figure 3a and 3b show the change in the encoder self-attention layer (self.); and Figure 3c and 3d represent the changes in the decoder-encoder cross attention layer (cross.). We visualize the changes of maximum similarity and average magnitude from the first epoch (start) to the last epoch (end). Figure 3b and 3d show that Transformer heavily depends on the magnitude term, instead of similarity. We see the same phenomena in the cross attention layer, where we hypothesize that the similarity is important for the alignment given the pair-wise dependency. This phenomenon occurs at the start of training, and it gets worse during the training. On the other hand, IKAN and MIKAN represent the stable scale of similarity and magnitude because of the implicit kernel and the L^p norm.

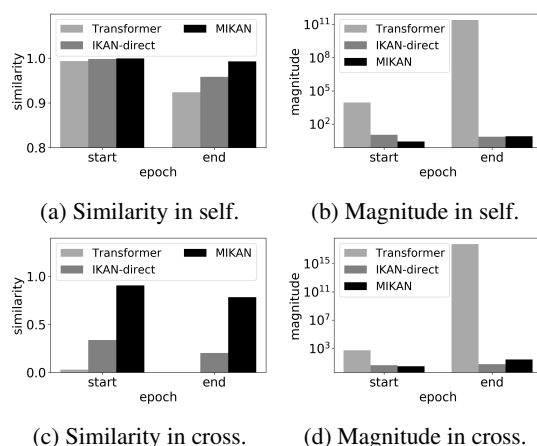


Figure 3: Similarity and magnitude of encoder self-attention layer (self.) and decoder-encoder attention layer (cross.).

Model	CO2	Passenger	Housing	Concrete	Parkinsons
GP (RBF)	0.060 ± 0.003	0.174 ± 0.052	0.234 ± 0.045	0.178 ± 0.012	0.075 ± 0.006
GP (SM)	0.057 ± 0.002	0.171 ± 0.048	0.191 ± 0.044	0.170 ± 0.021	0.072 ± 0.004
Transformer	0.057 ± 0.002	0.153 ± 0.043	0.126 ± 0.038	0.120 ± 0.017	0.040 ± 0.005
IKAN-direct	0.055 ± 0.002	0.110 ± 0.039	0.118 ± 0.038	0.111 ± 0.018	0.018 ± 0.004
MIKAN	0.055 ± 0.002	0.142 ± 0.044	0.108 ± 0.029	0.108 ± 0.015	0.037 ± 0.004

Table 4: RMSE on the UCI Regression dataset.




Figure 4: (a), (b): A histogram for attention weights in MIKAN for different p . (c) shows that RMSE of MIKAN varies depending on p .

Regression

We apply Transformer on the UCI regression dataset, and replace the scaled dot-product attention with our models. We perform 10-fold cross-validations, and Table 4 reports the performance following (Tompkins et al. 2019). Our models perform better than baselines. Afterward, we perform the qualitative analysis to clarify the relationship between the p and the sparsity of attention weights. Figure 4a and 4b visualize a histogram of MIKAN attention weights for different p s on the Housing dataset. As p goes to zero, most attention weights have either zero or one. Figure 4c shows the sensitive analysis with respect to p . The RMSE varies depending on p , but all results show better performance than the baselines.

Node Classification

We compare GAT and our generalized attention in Proposition 2. We perform ten-fold cross validations, and Table 5 represents the average accuracy for DeepWalk (Perozzi, Al-Rfou, and Skiena 2014), ICA (Geetor and Lu 2003), Chebyshev (Defferrard, Bresson, and Vandergheynst 2016), GCN (Kipf and Welling 2017), GAT (Veličković et al. 2018).

Model	Cora	Citeseer	Pubmed
DeepWalk*	67.2%	43.2%	65.3%
ICA*	75.1%	69.1%	73.9%
Chebyshev*	81.2%	69.8%	74.4%
GCN*	81.5%	70.3%	79.0%
GAT*	83.0±0.7%	72.5±0.7%	79.0±0.3%
IKAN-direct	83.3±0.7%	72.8±1.0%	79.2±0.4%
MIKAN	83.4±0.6%	72.9±0.8%	79.1±0.6%

Table 5: Accuracy for the node classification task. * denotes the reported performance in (Veličković et al. 2018).

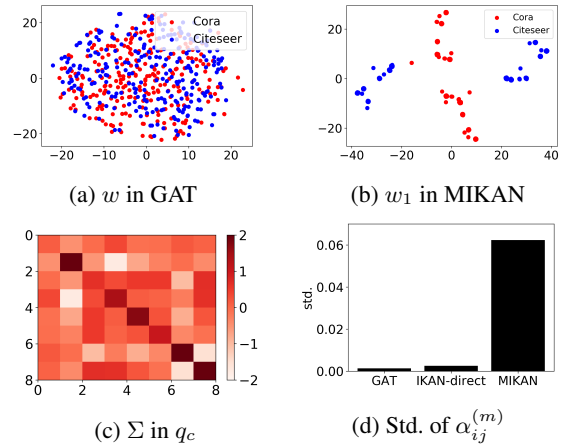


Figure 5: (a), (b): Sampled spectral points for GAT and MIKAN. (c): Learned full covariance Σ of q_c in MIKAN. (d): Standard deviation of attention weights across the heads.

Figure 5a and 5b show the spectral density in GAT and MIKAN by using t-SNE (Maaten and Hinton 2008), respectively. The attention in GAT depends on the RBF kernel, and its spectral density is fixed as a Gaussian distribution. However, MIKAN estimates the appropriate spectral density and its corresponding kernel depending on the dataset. Figure 5c shows the covariance, Σ , in q_c of MIKAN, and Figure 5d represents the standard deviation of attention weights across the heads. A large standard deviation denotes that each head has different attention weights. The results support that MIKAN has relatively diverse attention weights across the heads by imposing the dependency between heads with the copula.

Conclusion

This work provides the new interpretation of the attention as a product of similarity with the RBF kernel and the magnitude with the exponential of L^2 norm. We analyze the property of the kernel and the norm in the attention, theoretically and empirically. From the derivation and analysis, we generalize the attention with an implicit kernel function and a L^p norm. Furthermore, we propose the copula-augmented spectral density estimation for the dependence modeling in MHA to capture the diverse context. Our generalized attention can substitute the original attention in Transformer and GAT with the same algorithmic time complexity, and our experiments show better performance from our generalized attention.

Ethical Impact

We generalize the attention in Transformer and GAT, and we believe that our model is useful for capturing the underlying pattern or context of the given dataset. We expect that our proposed models contribute to both machine learning and social science. For the machine learning society, attention is widely used in many domains, such as natural language processing and vision. We propose a new direction to improve and interpret the attention in Transformer and GAT. For the social science aspect, we can apply our model to the recommendation, psychotherapy, and political ideal point estimation of legislators by capturing the context or patterns of user log history. We honor the AAAI Publications Ethics and Malpractice Statement, as well as the AAAI Code of Professional Conduct.

Acknowledgments

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2018R1C1B600865213)

References

- Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A. C.; and Bengio, Y. 2017. An Actor-Critic Algorithm for Sequence Prediction. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Chen, N.; Klushyn, A.; Kurle, R.; Jiang, X.; Bayer, J.; and Smagt, P. 2018. Metrics for deep generative models. In *International Conference on Artificial Intelligence and Statistics*, 1540–1550.
- Choromanski, K.; Likhoshesterov, V.; Dohan, D.; Song, X.; Davis, J.; Sarlos, T.; Belanger, D.; Colwell, L.; and Weller, A. 2020a. Masked Language Modeling for Proteins via Linearly Scalable Long-Context Transformers. *arXiv preprint arXiv:2006.03555*.
- Choromanski, K.; Likhoshesterov, V.; Dohan, D.; Song, X.; Gane, A.; Sarlos, T.; Hawkins, P.; Davis, J.; Mohiuddin, A.; Kaiser, L.; et al. 2020b. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, 3844–3852.
- Deng, Y.; Kim, Y.; Chiu, J.; Guo, D.; and Rush, A. 2018. Latent alignment and variational attention. In *Advances in Neural Information Processing Systems*, 9712–9724.
- Edunov, S.; Ott, M.; Auli, M.; Grangier, D.; and Ranzato, M. 2018. Classical Structured Prediction Losses for Sequence to Sequence Learning. In *Proceedings of NAACL-HLT*, 355–364.
- Geertor, L.; and Lu, Q. 2003. Link-based classification. In *Twelfth International Conference on Machine Learning (ICML 2003), Washington DC*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparametrization with Gumbel-Softmax. In *International Conference on Learning Representations (ICLR 2017)*. OpenReview.net.
- Jung, Y.; Song, K.; and Park, J. 2020. Approximate Inference for Spectral Mixture Kernel. *arXiv preprint arXiv:2006.07036*.
- Katharopoulos, A.; Vyas, A.; Pappas, N.; and Fleuret, F. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, 5156–5165. PMLR.
- Kim, Y. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1746–1751.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In Bengio, Y.; and LeCun, Y., eds., *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Li, C.-L.; Chang, W.-C.; Mroueh, Y.; Yang, Y.; and Póczos, B. 2019. Implicit Kernel Learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2007–2016.
- Li, D.; and Dunson, D. B. 2019. Geodesic distance estimation with spherelets. *arXiv preprint arXiv:1907.00296*.
- Maas, A. L.; Hannun, A. Y.; and Ng, A. Y. 2013. Rectifier nonlinearities improve neural network acoustic models. In *in ICML Workshop on Deep Learning for Audio, Speech and Language Processing*. Citeseer.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9(Nov): 2579–2605.
- Nelsen, R. B. 2007. *An introduction to copulas*. Springer Science & Business Media.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710.
- Rahimi, A.; and Recht, B. 2008. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, 1177–1184.

- Reed, M.; and Simon, B. 1975. *II: Fourier Analysis, Self-Adjointness*, volume 2. Elsevier.
- Sklar, M. 1959. Distribution functions in dimensions and their margins. *Publ. inst. statist. univ. Paris* 8: 229–231.
- Tompkins, A.; Senanayake, R.; Morere, P.; and Ramos, F. 2019. Black Box Quantiles for Kernel Learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 1427–1437.
- Ton, J.-F.; Flaxman, S.; Sejdinovic, D.; and Bhatt, S. 2018. Spatial mapping with Gaussian processes and nonstationary Fourier features. *Spatial statistics* 28: 59–78.
- Tran, D.; Blei, D.; and Airoldi, E. M. 2015. Copula variational inference. In *Advances in Neural Information Processing Systems*, 3564–3572.
- Tsai, Y.-H. H.; Bai, S.; Yamada, M.; Morency, L.-P.; and Salakhutdinov, R. 2019. Transformer Dissection: An Unified Understanding for Transformer’s Attention via the Lens of Kernel. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4335–4344.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Wang, B.; Zhao, D.; Lioma, C.; Li, Q.; Zhang, P.; and Simonsen, J. G. 2020. Encoding word order in complex embeddings. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Wiseman, S.; and Rush, A. M. 2016. Sequence-to-Sequence Learning as Beam-Search Optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1296–1306.
- Yaglom, A. M. 1987. Correlation Theory of Stationary and Related Random Functions. *Volume I: Basic Results*. 526.