

PowerMLP: An Efficient Version of KAN

Ruichen Qiu^{1,2}, Yibo Miao^{2,3}, Shiwen Wang³, Yifan Zhu^{2,3}, Lijia Yu^{4,5}, Xiao-Shan Gao^{2,3*}

¹School of Advanced Interdisciplinary Sciences, UCAS, Beijing 100049, China

²Academy of Mathematics and Systems Science, CAS, Beijing 100190, China

³University of Chinese Academy of Sciences, Beijing 101408, China

⁴Institute of Software, CAS, Beijing 100190, China

⁵State Key Laboratory of Computer Science

qiuruichen20@mailsucas.ac.cn, miaoyibo@amss.ac.cn, xgao@mmrc.iss.ac.cn

Abstract

The Kolmogorov-Arnold Network (KAN) is a new network architecture known for its high accuracy in several tasks such as function fitting and PDE solving. The superior expressive capability of KAN arises from the Kolmogorov-Arnold representation theorem and learnable spline functions. However, the computation of spline functions involves multiple iterations, which renders KAN significantly slower than MLP, thereby increasing the cost associated with model training and deployment. The authors of KAN also noted that “*the biggest bottleneck of KANs lies in their slow training. KANs are usually 10x slower than MLPs, given the same number of parameters.*” To address this issue, we propose a novel MLP-type neural network PowerMLP that employs simpler non-iterative spline function representation, offering approximately the same training time as MLP while theoretically demonstrating stronger expressive power than KAN. Furthermore, we compare the FLOPs of KAN and PowerMLP, quantifying the faster computation speed of PowerMLP. Our comprehensive experiments demonstrate that PowerMLP generally achieves higher accuracy and a training speed about 40 times faster than KAN in various tasks.

1 Introduction

A long-standing problem of deep learning is the identification of more effective neural network architectures. The Kolmogorov-Arnold Network (KAN), introduced by Liu et al. (2024), presents a new architecture. Unlike traditional MLP that places activation functions on nodes, KAN employs learnable univariate spline functions as activation functions placed on edges. The expressive power of KAN is derived from the Kolmogorov-Arnold representation theorem and the property of spline functions. Due to its exceptional expressiveness, KAN achieves high accuracy and interpretability in multiple tasks such as function fitting and PDE solving. Significant performance improvements using KAN have been observed in time series prediction (Inzirillo and Genet 2024), graph data processing (Kiamari et al. 2024), and explainable natural language processing (Galitsky 2024).

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: PowerMLPs define a strictly larger function space than KANs over \mathbb{R}^n (Corollary 3), and define the same function space over $[-E, E]^n$ for any $E \in \mathbb{R}_+$ (Corollary 7), where n is the input dimension. $\mathcal{P}_{d,w,k,p}$ is the set of all PowerMLPs with depth d , width w , k -th power ReLU activation function, and p nonzero parameters. $\mathcal{K}_{d,w,k,G,p}$ is the set of all KANs with depth d , width w , using (k, G) -spline (see Eq. (3)), and p nonzero parameters.

Unfortunately, despite the impressive performance of KAN, it faces a critical drawback: slow inference and training speeds, which increase the cost associated with training and deploying the model. Liu et al. (2024) highlight in the KAN paper that “*the biggest bottleneck of KANs lies in their slow training. KANs are usually 10x slower than MLPs, given the same number of parameters.*” The inefficient training and inference latency of KAN plays a crucial role in ensuring a positive user experience and low computational resource requirements.

Upon examining the structure of KAN, we contribute the primary cause of the slow computation speed to the spline activation function. Specifically, a k -order spline function is a linear combination of k -order B-splines, and each k -order B-spline requires construction through $\mathcal{O}(k^2)$ iterations by the de Boor-Cox formula in KAN’s paper, resulting in slow computation. To address this issue, we need to identify a more efficient method to obtain B-splines, moving away from the recursive computation of the de Boor-Cox formula.

The k -order B-spline was initially defined as k -th divided difference of the truncated power function (Curry and Schoenberg 1947). Consequently, B-splines can also be expressed as linear combinations of powers of ReLUs (Greville 1969). Inspired by this, we introduce a novel MLP-type neural network by incorporating a basis function into the MLP with powers of ReLU as activation function, termed **PowerMLP**, as shown in Figure 2. From Figure 1, Pow-

erMLPs define a strictly larger function space than KANs over \mathbb{R}^n and define the same function space over $[-E, E]^n$, indicating that PowerMLP can serve as a viable substitute for KAN. Intuitively, the B-spline can be represented without iterative recursions in PowerMLP, leading to faster computation. In addition, the activation functions of PowerMLP are on the nodes instead of the edges, inhering the fast training advantages of MLP. We further choose Floating Point Operations (FLOPs) as the metric and demonstrate that the FLOPs of KAN are more than 10 times those of PowerMLP, theoretically explaining why PowerMLP is significantly faster than KAN.

We conducted extensive experiments to validate the advantages of PowerMLP in various experimental settings, including those of KAN and more complicated tasks such as image classification and language processing. The results show that PowerMLP trains about 40 times faster than KAN and achieves the best performance in most experiments.

2 Related Work

Network Architecture. Multilayer Perceptron (MLP) is a basic form of neural networks used primarily for supervised learning tasks (Haykin 1998). One of the activation functions commonly used with MLP is the rectified linear unit (ReLU) (Glorot, Bordes, and Bengio 2011), which is further extended to LeakyReLU (Maas et al. 2013), PReLU (He et al. 2015), RePU (Li, Tang, and Yu 2020a,b), and GeLU (Hendrycks and Gimpel 2023). LeCun et al. (1989) introduced CNNs specialized for images and Vaswani (2017) introduced Transformers, which have become a cornerstone for large language models. PowerMLP can be obtained from the RePU network (Li, Tang, and Yu 2020a,b) by adding a basis function to each layer.

Kolmogorov-Arnold Network. KAN (Liu et al. 2024) is a new network architecture using learnable activation functions on edges. In small-scale AI + Science tasks, KAN outperforms MLP in terms of both accuracy and interpretability. Additionally, KAN also achieves remarkable performance in time series prediction (Xu, Chen, and Wang 2024; Inzirillo and Genet 2024), graph-structured data processing (Kiamari et al. 2024; Carlo et al. 2024), and explainable natural language processing (Galitsky 2024). Meanwhile, notable approaches are taken to improve KAN’s interpretability and performance, including incorporating wavelet functions into KAN (Bozorgasl and Chen 2024), replacing the basis function with rational functions (Aghaei 2024), and faster implementation by approximation of radial basis functions (Li 2024). However, KAN and these variants have a computational speed significantly slower than MLP, since they have a structure similar to that of KAN.

PowerMLP vs. KAN. PowerMLP can be considered as an MLP-type representation of KAN. In Section 6 of (Liu et al. 2024), “KANs’ limitations and future directions” are discussed. Most of these limitations of KAN can be eliminated or mitigated by PowerMLP. (1) In algorithmic aspects, the authors noted that “KANs are usually 10x slower than MLPs.” PowerMLP, which is about 40 times faster than

KAN in our experiments, eliminates this limitation. (2) In mathematical aspects, the authors call for a more general approximation result beyond the depth-2 Kolmogorov-Arnold representations. This result is given in Corollary 8 using our proposed PowerMLP. (3) In the application aspects, the authors call for “integrating KANs into current architectures.” This can be done by replacing the MLP in these architectures with PowerMLP more naturally than KAN.

3 Preliminaries

In this section, we introduce the preliminary knowledge of Kolmogorov-Arnold Network (KAN) (Liu et al. 2024).

3.1 Spline Function

The following *de Boor-Cox formula* (Cox 1972; de Boor 1978) for B-spline is used to define KAN.

Definition 1 (B-spline). *Let $t := (t_j)$ be a nondecreasing sequence of real numbers, called knot sequence. The zero-order B-spline on (t_j, t_{j+1}) is defined as¹*

$$B_{j,0,t}(x) = \begin{cases} 1, & t_j \leq x < t_{j+1}, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Then the j -th k -order normalized B-spline for the knot sequence t is defined recursively as

$$B_{j,k,t}(x) = \frac{x - t_j}{t_{j+k} - t_j} B_{j,k-1,t}(x) + \frac{t_{j+k+1} - x}{t_{j+k+1} - t_{j+1}} B_{j+1,k-1,t}(x), \quad \text{for } k \geq 1. \quad (2)$$

To maintain consistency with KAN’s setting, let $t = (t_{-k}, \dots, t_{-1}, t_0, t_1, \dots, t_G, t_{G+1}, \dots, t_{G+k})$ be an *increasing* knot sequence, called a (k, G) -grid. Through a linear combination of B-splines on t , we provide the definition of spline functions as follows.

Definition 2 (Spline Function). *Let t be a (k, G) -grid. A k -order spline function for the knot sequence t is given by the following linear combination of B-splines*

$$\text{spline}_{k,G}(x) = \sum_{j=-k}^{G-1} c_j B_{j,k,t}(x), \quad (3)$$

where $c_j \in \mathbb{R}$ are the coefficients of the spline function. For simplicity, (3) is called a (k, G) -spline function.

3.2 Kolmogorov-Arnold Networks

A KAN network (Liu et al. 2024) is a composition of L layers: given an input vector $\mathbf{x} = (x_1, \dots, x_{n_0}) \in \mathbb{R}^{n_0}$, the output of KAN is

$$\text{KAN}(\mathbf{x}) = (\Phi_{L-1} \circ \dots \circ \Phi_1 \circ \Phi_0)(\mathbf{x}),$$

where Φ_ℓ is the function matrix corresponding to the ℓ -th layer. The dimension of the input vector of Φ_ℓ is denoted as n_ℓ , and Φ_ℓ is defined below:

$$\Phi_\ell(\cdot) = \begin{pmatrix} \phi_{\ell,1,1}(\cdot) & \phi_{\ell,1,2}(\cdot) & \cdots & \phi_{\ell,1,n_\ell}(\cdot) \\ \phi_{\ell,2,1}(\cdot) & \phi_{\ell,2,2}(\cdot) & \cdots & \phi_{\ell,2,n_\ell}(\cdot) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{\ell,n_{\ell+1},1}(\cdot) & \phi_{\ell,n_{\ell+1},2}(\cdot) & \cdots & \phi_{\ell,n_{\ell+1},n_\ell}(\cdot) \end{pmatrix},$$

¹Note that de Boor-Cox formula defines the function in Eq. (1) as order 1. To keep the same mark with KAN, we adjust it to 0.

where $\phi_{\ell,q,p}$ is a residual activation function:

$$\phi_{\ell,q,p}(x_p) = u_{\ell,q,p}b(x_p) + v_{\ell,q,p}\text{spline}_{k,G}(x_p). \quad (4)$$

$b(x)$ is a non-parameter basis function² similar to residual shortcut, which is included to improve training. $\text{spline}_{k,G}(x)$ is defined in Definition 2. The KAN defined above, called a KAN of k -order, has depth L , width $W = \max_{i=0}^{L-1} n_i$, and $O(W^2L(G+k))$ parameters.

4 PowerMLP

While using spline functions as activations enables KAN to achieve excellent performance, computing spline functions involves multiple iterations, leading to a high number of FLOPs (see Table 1). This results in KAN’s computation speed being slower than that of MLP, thereby increasing the cost associated with training and deploying the model.

To replace KAN with a network structure that offers similar expressiveness but faster computation, we present a simpler representation of spline functions, avoiding the recursive calculations of the de Boor-Cox formula. In Section 4.1, we introduce a novel network, named **PowerMLP**. Through theorems in Sections 4.2 and 4.3, we demonstrate that KAN and PowerMLP define the same function space over bounded intervals, indicating their interchangeability. Furthermore, in Section 4.4, we compare the FLOPs of KAN and PowerMLP, revealing that PowerMLP theoretically achieves faster speeds than KAN. Proofs are given in Appendix A.

4.1 PowerMLP

Referring back to the de Boor-Cox formula in Eq. (2), we attribute the slow computation speed of KAN to the $\mathcal{O}(k^2)$ iterations of calculation required for constructing k -order B-spline. According to the initial definition of k -order B-spline from the k -th divided difference of the truncated power function (Curry and Schoenberg 1947) and subsequent work by Greville (1969), we express B-spline through a non-iterative approach as a linear combination of powers of ReLUs. Thus, we introduce a novel MLP network that incorporates a basis function into the MLP with powers of ReLU as the activation, termed **PowerMLP**.

The k -th power of ReLU (Mhaskar 1993; Li, Tang, and Yu 2020a) is defined as:

$$\sigma_k(x) = (\text{ReLU}(x))^k = (\max(0, x))^k \quad k \in \mathbb{Z}_+. \quad (5)$$

The fully connected feedforward neural network with the k -th power of ReLU as the activation function is referred to as *ReLU- k MLP*. We integrate a basis function into ReLU- k MLP and define PowerMLP following the form of KAN.

Definition 3 (PowerMLP). *A PowerMLP is a neural network composition of L layers:*

$$\begin{aligned} \text{PowerMLP}(\mathbf{x}) &= (\Psi_{L-1} \circ \dots \circ \Psi_1 \circ \Psi_0)(\mathbf{x}), \quad \text{where} \\ \Psi_\ell(\mathbf{x}_\ell) &= \begin{cases} \alpha_\ell b(\mathbf{x}_\ell) + \sigma_k(\omega_\ell \mathbf{x}_\ell + \gamma_\ell), & \text{for } \ell < L-1, \\ \omega_{L-1} \mathbf{x}_{L-1} + \gamma_{L-1}, & \text{for } \ell = L-1. \end{cases} \end{aligned} \quad (6)$$

²Basis function is $b(x) = x/(1 + e^{-x})$ in KAN’s paper.




Figure 2: Structure of a 3-layer PowerMLP. The first two layers are calculated by: (1) affine transformation, (2) k -th power of ReLU activation, (3) addition with a basis function. The last layer contains only an affine transformation.

$\alpha_\ell \in \mathbb{R}^{m_{\ell+1} \times m_\ell}$, $\omega_\ell \in \mathbb{R}^{m_{\ell+1} \times m_\ell}$, $\gamma_\ell \in \mathbb{R}^{m_\ell \times 1}$ are trainable parameters. $b(\mathbf{x})$ is a basis function that performs the same operation as the basis function in KAN on each component of \mathbf{x} , and σ_k is the ReLU- k activation function. A PowerMLP using σ_k as the activation function is called a k -order PowerMLP. Width and depth of the PowerMLP are defined to be width = $\max_{l=0}^{L-1} \{m_l\}$, depth = L . Refer to Figure 2 for an illustration.

Through non-parameter activation function σ_k and non-iteration calculation, PowerMLP computes faster than KAN. Furthermore, they can represent each other within bounded intervals, a condition usually met in practical scenarios.

4.2 PowerMLP Can Represent KAN

In this section, we show that any KAN can be represented by a PowerMLP. We first give the following connection between B-spline and the k -th power of ReLU σ_k .

Lemma 1 (Represent the B-spline with k -th power of ReLU). *If $t_u \neq t_v (\forall u \neq v)$, then the k -order B-spline on the knot sequence $t = (t_j, \dots, t_{j+k+1})$ can be represented as a linear combination of σ_k functions:*

$$B_{j,k,t}(x) = \sum_{i=j}^{j+k+1} \frac{t_{j+k+1} - t_j}{\prod_{l=j}^{j+k+1, l \neq i} (t_l - t_i)} \sigma_k(x - t_i).$$

Since each spline activation function in KAN is a linear combination of B-splines, we can represent it as a linear combination of ReLU- k . By performing this operation on each component of the input vector \mathbf{x} and incorporating a basis function, we represent a KAN layer as a PowerMLP layer. Since KAN and PowerMLP are both compositions of layers, it means that a KAN can be represented as a PowerMLP. We have the following theorem.

Theorem 2 (KAN is a Subset of PowerMLP). *Fix the input dimension of networks as n . Let $\mathcal{K}_{d,w,k,G,p}$ be the set of all KAN networks with depth d , width w , p nonzero parameters, using (k, G) -spline; and $\mathcal{P}_{d,w,k,p}$ be the set of all PowerMLPs with depth d , width w , k -th power of ReLU and p nonzero parameters. Then it holds*

$$\mathcal{K}_{d,w,k,G,p} \subset \mathcal{P}_{d,w^2(G+k),k,p}. \quad (7)$$

By Theorem 2 and the fact that a spline function is zero outside of certain interval (see Eq. (1)) while PowerMLPs include all polynomials (Li, Tang, and Yu 2020a), we have

Corollary 3. *PowerMLPs define a strictly larger function space than KANs over \mathbb{R}^n .*

4.3 KAN Can Represent PowerMLP over Intervals

In this section, we prove the inclusion relationship in another direction. A PowerMLP layer

$$\mathbf{z} = \sigma_k(\omega\mathbf{x} + \gamma) + \alpha b(\mathbf{x}) \quad (8)$$

can be decomposed into 3 operations: (1) an affine transformation: $\mathbf{y} = \omega\mathbf{x} + \gamma$; (2) a ReLU- k activation: $\mathbf{u} = \sigma_k(\mathbf{y})$; (3) an addition with basis function: $\mathbf{z} = \mathbf{u} + \alpha b(\mathbf{x})$. By Lemmas 4 and 5, operations (1) and (2) can be represented by spline functions, while operation (3) can be easily achieved.

Lemma 4 (Affine Transformation). *Consider an affine transformation on \mathbb{R} : $\mathcal{A}(x) = \omega x + \gamma$. For any G , we can find a (k, G) -grid $t = (t_{-k}, \dots, t_{-1}, t_0, t_1, \dots, t_G, t_{G+1}, \dots, t_{G+k})$, and a k -order spline function*

$$\text{spline}_{k,G}(x) = \sum_{j=-k}^{G-1} c_j B_{j,k,t}(x),$$

where $c_j = (\sum_{i=j+1}^{j+k} t_i/k)\omega + \gamma$, $k > 0$, such that $\mathcal{A}(x) = \text{spline}_{k,G}(x)$ for $t_0 \leq x \leq t_G$.

Lemma 5 (ReLU- k Function). *We can find a $(k, 2)$ -grid $t = (t_{-k}, \dots, t_{-1}, t_0, 0, t_2, t_3, \dots, t_{k+2})$ and a k -order spline function defined on t*

$$\text{spline}_{k,2}(x) = \sum_{j=-k}^1 \left[\left(\prod_{l=j+1}^{j+k} \sigma_1(t_l) \right) B_{j,k,t}(x) \right],$$

such that $\sigma_k(x) = \text{spline}_{k,2}(x)$ for $t_0 \leq x \leq t_2$.

By the above lemmas, we can use two layers of spline functions to represent operations (1) and (2) successively, as shown in Figure 3. For the first layer, we set the coefficients before the basis function to be 0. By Lemma 4, we can represent affine transformation $\mathcal{A}(x) = \omega x + \gamma$ by activation functions $\phi_{1,q,p}(x_p)$ (see Eq. (4)). Then, for $1 \leq p \leq m$, we take the activation function as follows:

$$\phi_{1,q,p}(x_p) = \begin{cases} \omega_{q,p}x_p + \gamma_{q,p}, & \text{for } 1 \leq q \leq m, \\ \delta_{q-m,p}x_p, & \text{for } 1 \leq q-m \leq n, \end{cases}$$

where δ_{ij} equals to 1 if $i = j$ and 0 otherwise. Thus, we have $y_q = \sum_{p=1}^n \omega_{q,p}x_p + \gamma_{q,p}$ for $1 \leq q \leq m$ and $y_q = x_{q-m}$ for $m+1 \leq q \leq m+n$. In the second layer, we represent the addition of ReLU- k activation and the basis function. By Lemma 5, spline function can represent σ_k . So for $1 \leq r \leq m$, we set $\phi_{2,r,q}(y_q)$ as follows:

$$\phi_{2,r,q}(y_q) = \begin{cases} \delta_{r,q}\sigma_k(y_q), & \text{for } 1 \leq q \leq m, \\ \alpha_{r,q-m}b(y_q), & \text{for } 1 \leq q-m \leq n. \end{cases}$$

By direct computation, we show that the output of the two-layer KAN is the same as our PowerMLP layer in Eq. (8). Hence, we derive Theorem 6 and Corollary 7. Then by (Li, Tang, and Yu 2020a, Theorem 3.3), we obtain a general approximation result for KAN in Corollary 8.




Figure 3: Represent a PowerMLP layer with a 2-layer KAN. δ_{ij} equals to 1 if $i = j$ and 0 otherwise. The first layer represents the affine transformation $y_q = \sum_{p=1}^n \omega_{q,p}x_p + \gamma_{q,p}$ for $1 \leq q \leq m$ and keeps $y_q = x_{q-m}$ for $m+1 \leq q \leq m+n$. The second layer represents the ReLU- k activation and adds the basis function: $z_r = \sigma_k(y_r) + \sum_{q=m+1}^{m+n} \alpha_{r,q-m}b(y_q)$.

Theorem 6 (PowerMLP is a subset of KAN over interval). *Use notations in Theorem 2. For any $E \in \mathbb{R}_+$, it holds*

$$\mathcal{P}_{d,w,k,p} \subset \mathcal{K}_{2d,2w,k,2,\mathcal{O}(kp)} \text{ over } [-E, E]^n. \quad (9)$$

Corollary 7. *PowerMLPs and KANs define the same function space over $[-E, E]^n$ for any $E \in \mathbb{R}_+$.*

Corollary 8. *Let f be a continuous, first-order differentiable function on $[-1, 1]^n$, which satisfies $\sum_{i=1}^n \int_{[-1,1]^n} (\partial_{x_i} f(x))^2 dx \leq 1$. Then, for any $\epsilon \in (0, 1)$, a 2-order KAN K requires at most $\mathcal{O}(n \log_2 \frac{1}{\epsilon})$ layers and $\mathcal{O}(\epsilon^{-n})$ nonzero parameters to ensure $\|K - f\|_{L_2} \leq \epsilon$.*

4.4 FLOPs: Comparison of Computing Cost

In this section, we show that PowerMLP exhibits significantly faster training and inference speeds compared to KAN in terms of FLOPs metric. FLOPs, an acronym for *Floating Point Operations*, is a metric used to quantify the computational complexity of neural networks, especially in frameworks like PyTorch (Molchanov et al. 2017). It measures the number of floating-point operations required to perform one forward pass through the network.

Following Yu, Yu, and Wang (2024), we consider FLOPs for any arithmetic operations like $+$, $-$, \times , \div to be 1, and for Boolean operations to be 0. Meanwhile, any operation of comparing two numbers is set to be 0 FLOPs, which means that the FLOPs of ReLU function are 0. We denote FLOPs of basis function as λ . Then we can calculate FLOPs of one layer of MLP (with ReLU), KAN and PowerMLP below:

$$\mathcal{F}_{\text{MLP}} = 2d_{in}d_{out},$$

$$\mathcal{F}_{\text{KAN}} = d_{in}d_{out}(9kG + 13.5k^2 + 2G - 2.5k + 3) + \lambda d_{in},$$

$$\mathcal{F}_{\text{PowerMLP}} = 4d_{in}d_{out} + (k-1)d_{out} + \lambda d_{in}$$

where d_{in} and d_{out} denote the input and output dimensions of the layer. The KAN layer uses (k, G) -grid, and the PowerMLP layer is k -order. Refer to Appendix A for details.

Then we compare FLOPs of three network layers with same number of parameters, given by the ratio of FLOPs \mathcal{F} to numbers of parameters:

$$r_{\text{MLP}} = \frac{2d_{in}d_{out}}{d_{in}d_{out} + d_{out}},$$

$$r_{\text{KAN}} = \frac{d_{in}d_{out}(9kG + 13.5k^2 + 2G - 2.5k + 3) + \lambda d_{in}}{d_{in}d_{out}(k + G + 2)},$$

$$r_{\text{PowerMLP}} = \frac{4d_{in}d_{out} + (k - 1)d_{out} + \lambda d_{in}}{2d_{in}d_{out} + d_{out}}.$$

Assuming that d_{in} and d_{out} increase at the same rate, r_{MLP} and r_{PowerMLP} tend towards values less than 2, while r_{KAN} approaches to numbers larger than 20 for $k \geq 3, G \geq 3$ ($k = 3, G = 3$ are the smallest values used in KAN’s paper). **Thus, PowerMLP shares a close computing speed with MLP, and is over 10 times faster than KAN under the FLOPs metric.** We give an example of FLOPs of KAN, MLP and PowerMLP with almost the same number of parameters in Table 1. Training speed comparisons in experimental settings are in Section 5.3.

Network	Shape	#Params	FLOPs
KAN ($G = 3, k = 3$)	[2, 1, 1]	24	564
MLP (ReLU)	[2, 6, 1]	25	36
PowerMLP ($k = 3$)	[2, 4, 1]	25	40

Table 1: Comparison of KAN, MLP and PowerMLP. With almost the same parameters, MLP and PowerMLP have much fewer FLOPs than KAN.

5 Experiments

In Section 4, PowerMLPs are shown to define the same function space as KANs over bounded intervals and achieve faster computation. In this section, we employ several experiments to validate these theoretical findings and demonstrate the advantages of PowerMLP.

Four experiments are conducted. (1) We consider AI for science tasks in the KAN paper (Liu et al. 2024) under the same settings, showing that PowerMLP performs better. (2) More complex tasks such as machine learning, natural language processing, and image classification are considered. We show that PowerMLP outperforms KAN in all tasks. (3) We compare training and convergence time of KAN and PowerMLP, validating that PowerMLP can be much faster. (4) We conduct an ablation experiment to show that both the basis function and ReLU- k activation are needed for the performance of PowerMLP.

All KANs in the paper are the latest version (0.2.5) up to 2024-8-14. More details are given in Appendix B.³

5.1 AI for Science Tasks

Function Fitting PowerMLP is tested on a regression task for 16 special functions in KAN’s experiments (Liu et al. 2024). For KAN, MLP, and PowerMLP, we choose two sizes for the networks: **(1) Small size:** KAN with $k = 3, G = 3$, shape [2, 1, 1], and 24 parameters; MLP with shape [2, 6, 1]

³Code is available at <https://github.com/Iri-sated/PowerMLP>.

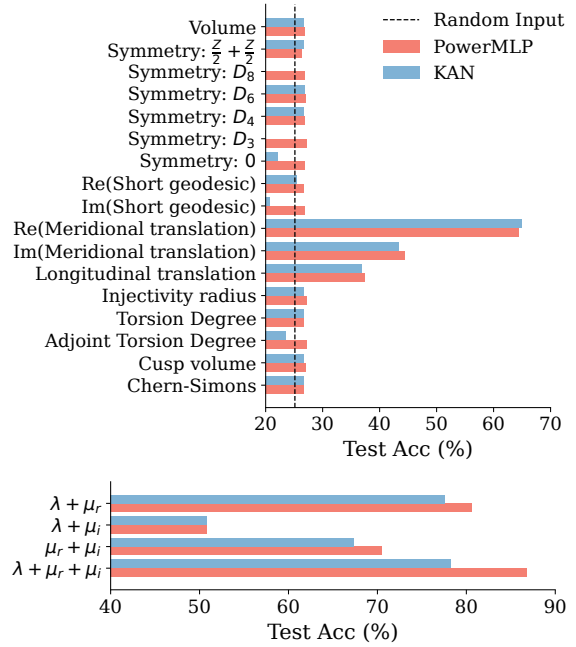


Figure 4: In the upper figure, PowerMLP can correctly find that 3 of 17 geometric invariants have influence on the output. Additionally, PowerMLP outperforms KAN in 15 of 17 input cases while KAN fails to converge with Symmetry D_3 or D_8 as input. In the bottom figure, trained on part or all of the 3 influencing geometric invariants, PowerMLP achieves much higher test accuracy than KAN in 3 cases.

and 25 parameters; 3-order PowerMLP with shape [2, 4, 1] and 25 parameters; **(2) Large size:** KAN with $k = 3, G = 100$, shape [2, 2, 1, 1] and 735 parameters; MLP with shape [2, 32, 18, 1] and 709 parameters; 3-order PowerMLP with shape [2, 32, 8, 1] and 689 parameters.

Results are given in Table 2. We see that PowerMLP achieves the best results on 11/10 (small size/large size) out of 16 cases. This is attributed to that PowerMLP has ReLU- k for stronger power of expression than MLP and can better converge than KAN with a simpler structure of non-parameter activations.

Knot Theory Davies et al. (2021) used MLP to discover highly non-trivial relations among 17 *geometric invariants* and *signatures* in knot theory. With 17 geometric invariants as inputs and the corresponding signature as output, they trained MLPs and found a strong connection between signatures and 3 of all 17 geometric invariants: the *longitudinal translation* λ , *real and image part of meridional translation* μ_r, μ_i . Based on this observation, they proved a theorem and explained the connection theoretically, which is an interesting example of AI for math tasks (Davies et al. 2022).

Liu et al. (2024) reproduced the same experiment with KANs. Following their settings, we conduct the experiment with our PowerMLP. In Figure 4, we show the test accuracy of using a single geometric invariant as input each time to predict the signature. Most geometric invariants are close

Function Name	Small Size: ~ 25 parameters			Large Size: ~ 700 parameters		
	KAN	MLP	PowerMLP	KAN	MLP	PowerMLP
JE	4.63×10^{-3}	4.98×10^{-3}	5.79×10^{-4}	1.04×10^{-4}	5.88×10^{-4}	7.23×10^{-5}
IE1	1.34×10^{-2}	5.79×10^{-3}	3.43×10^{-3}	4.52×10^{-5}	5.57×10^{-4}	3.37×10^{-5}
IE2	1.16×10^{-2}	4.71×10^{-3}	1.73×10^{-3}	1.18×10^{-3}	5.98×10^{-4}	3.20×10^{-5}
B1	7.71×10^{-1}	3.94×10^{-2}	3.93×10^{-2}	1.70×10^{-2}	5.47×10^{-3}	2.77×10^{-3}
B2	7.94×10^{-2}	6.02×10^{-2}	7.75×10^{-2}	1.77×10^{-3}	4.62×10^{-3}	2.61×10^{-3}
MB1	2.29×10^0	3.76×10^{-2}	3.48×10^{-2}	1.70×10^{-2}	5.11×10^{-3}	3.75×10^{-3}
MB2	7.97×10^{-1}	1.04×10^{-2}	7.47×10^{-3}	9.44×10^{-5}	1.01×10^{-3}	7.42×10^{-5}
AL ($m = 0$)	1.09×10^{-1}	8.14×10^{-2}	7.21×10^{-2}	1.88×10^{-3}	7.45×10^{-3}	6.72×10^{-3}
AL ($m = 1$)	1.25×10^{-1}	7.49×10^{-2}	6.94×10^{-2}	1.40×10^{-2}	1.20×10^{-2}	1.02×10^{-2}
AL ($m = 2$)	2.30×10^{-1}	1.10×10^{-1}	9.95×10^{-2}	1.67×10^{-3}	8.83×10^{-3}	7.29×10^{-3}
SH ($m = 0, n = 1$)	4.05×10^{-5}	2.02×10^{-3}	2.59×10^{-5}	8.39×10^{-6}	1.45×10^{-4}	7.77×10^{-6}
SH ($m = 1, n = 1$)	1.92×10^{-2}	5.57×10^{-3}	1.20×10^{-2}	7.03×10^{-5}	4.16×10^{-4}	2.60×10^{-4}
SH ($m = 0, n = 2$)	5.94×10^{-5}	4.46×10^{-3}	4.79×10^{-4}	1.02×10^{-5}	3.17×10^{-4}	2.55×10^{-5}
SH ($m = 1, n = 2$)	3.25×10^{-2}	1.81×10^{-2}	2.37×10^{-3}	2.27×10^{-3}	8.90×10^{-4}	1.45×10^{-4}
SH ($m = 2, n = 2$)	1.17×10^{-2}	1.59×10^{-2}	2.16×10^{-2}	1.09×10^{-2}	8.32×10^{-2}	3.66×10^{-5}

Table 2: Fitting Special Functions. JE: Jacobian elliptic functions, IE1/IE2: Incomplete elliptic integral of the first/second kind, B1/B2: Bessel function of the first/second kind, MB1/MB2: Modified Bessel function of the first/second kind, AL: Associated Legendre function, SH: spherical harmonics. All the values are test RMSE loss, the less the better. The best results are marked as boldface. PowerMLP achieves the best results on 11/10 (small/large size) out of 16 cases.

to random inputs, but the longitudinal translation, real and image part of the meridional translation show relationships with the signature. For better comparison, we also show results of a KAN with the same depth and almost same number of parameters. KAN performs worse than PowerMLP in 15 of 17 geometric invariants. In particular, KAN fails to converge with Symmetry D_3 or D_8 as input.

Furthermore, in Figure 4, we show test accuracy of PowerMLPs trained on all or part of the three relevant geometric invariants λ , μ_i and μ_r . With a test accuracy of 86.74%, PowerMLP successfully validates the connection among signature and λ , μ_i , μ_r . Furthermore, with a test accuracy of 80.62%, PowerMLP can also find the connection among λ , μ_r and the signature discovered by KAN. More importantly, compared to the close test accuracy of KAN and PowerMLP in single geometric invariant input, PowerMLP achieves much higher test accuracy than KAN in 3 out of 4 cases of combination input. This indicates that PowerMLP can better utilize the correlation between inputs.

5.2 More Complex Tasks

We perform three more complex tasks. We use three networks with approximately the same number of parameters.

Machine Learning For basic machine learning tasks, we conduct two experiments on Titanic and Income (Becker and Kohavi 1996), which are classification tasks of small input dimension. In Figure 5, we show the test accuracy of three networks, and PowerMLP outperforms MLP and KAN.

Natural Language Processing We conduct two experiments on SMS Spam Collection (Spam) (Gómez Hidalgo et al. 2006) and AG_NEWS (Zhang, Zhao, and LeCun 2015) dataset, which are text classification tasks. We use TF-IDF transformation (Ramos et al. 2003) to convert text into vectors, and *networks need to deal with high-dimensional*

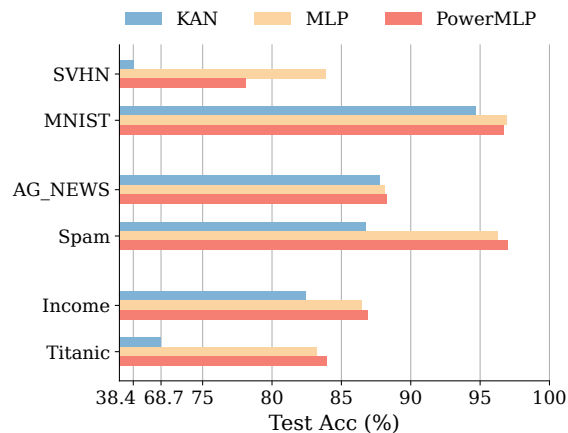


Figure 5: Test accuracy on multiple classification tasks.

sparse inputs. In Figure 5, we show the test accuracy of three networks, and PowerMLP outperforms MLP and KAN.

Image Classification For image classification tasks, we conduct two experiments on MNIST (LeCun et al. 1998) and SVHN (Netzer et al. 2011) datasets. We convert SVHN data to greyscale images, and flatten image tensor into 1-dimension vectors. In this task, *networks need to deal with high-dimensional inputs with strong connection among different input dimensions*. In Figure 5, we show the test accuracy of three networks, and PowerMLP outperforms KAN.

In summary, with almost the same number of parameters and the same depth, PowerMLP achieves better accuracy than KAN in all tasks. PowerMLP outperforms MLP in machine learning (Income, Titanic) and language processing (AG_NEWS, Spam), while performing worse in image classification (SVHN, MNIST).

Tasks	#Params	Time(s)		
		KAN	MLP	PowerMLP
JE	~25	17.73	0.583	0.738
IE1	~25	25.32	0.430	0.674
IE2	~25	31.91	0.584	0.709
B1	~25	29.68	0.422	0.646
B2	~25	35.20	0.559	0.696
Titanic	~100	35.07	0.529	0.591
Spam	~800	62.06	0.568	0.655
SVHN	$\sim 1.3 \times 10^5$	89.82	1.53	2.39

Table 3: Training times on 8 tasks. Training times of PowerMLP are about 40 times smaller than KAN on average.




Figure 6: Time of convergence. MLP and PowerMLP converge much faster than KAN.

5.3 Training Time

In Table 3, eight tasks are considered. The first five tasks are function regression tasks in Section 5.1, and the last three tasks are machine learning, language processing, and image classification tasks in Section 5.2. For better comparison, the experiments are on a single NVIDIA GeForce RTX 4090 GPU, repeated each task 10 times to take an average, and networks in each task are trained with the same hyperparameters. From Table 3, among all tasks and in different numbers of parameters, the training time of PowerMLP is close to MLP, which is about 40 times less than KAN. This is consistent with our theoretical analysis in Section 4.4.

To be more comprehensive, in Figure 6, we give the variation curves of the test RMSE loss relative to the training time in one training progress. Two curves with similar colors represent two training processes of the same network. We can see that MLP and PowerMLP also converge much faster than KAN in terms of training time.

5.4 Ablation Study

We show that both the basis function and the ReLU- k are useful in PowerMLP. In Figure 7, we compare three networks, PowerMLP, PowerMLP without basis function, and PowerMLP without ReLU- k activation (ReLU instead) on a function regression task for $f(x, y) = x \exp(-y)$.

The upper graph shows the RMSE loss on the test set with variant network depth for fixed width of 4. Although PowerMLPs without basis function perform well in shallow networks, they fail to converge for large depth. The bottom graph shows the influence of the number of parameters,

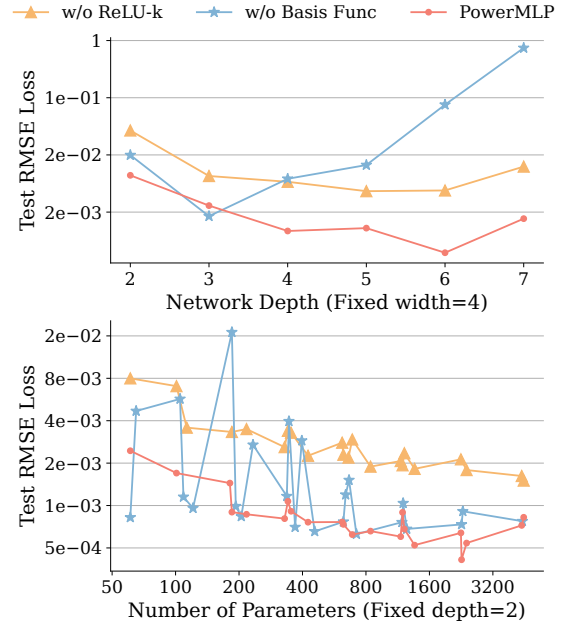


Figure 7: Ablation study. Basis function enhances training stability, while ReLU- k improves expressive ability.

with a fixed network depth of 2. PowerMLPs without basis function achieve nearly the same test RMSE loss with PowerMLP for some networks, while performing much worse in other cases, indicating that basis function enhances training stability. In all situations, PowerMLPs perform better than PowerMLPs without ReLU- k , consistent with ReLU- k 's better expressive ability.

6 Conclusion

In this paper, we introduce a novel neural network architecture, PowerMLP, which employs the powers of ReLU as activation and expresses linear combination as inner product to represent spline function. PowerMLP can be viewed as a more efficient and stronger version of KAN. In terms of expressiveness, PowerMLPs define a larger or equal function space than KANs. In terms of efficiency, PowerMLPs can be trained with over 10 times fewer FLOPs than KANs. We conducted comprehensive experiments to demonstrate that PowerMLPs achieve higher accuracy in most cases and can be trained about 40 times faster compared to KANs.

Limitations and Future Works. Despite its powerful expressive abilities, a limitation of our proposed PowerMLP is that it fails to achieve exceptional performance in complicated computer vision tasks and long-text processing. This stems from PowerMLP's relatively simple architecture, which lacks specialized structures such as convolutional layers or attention mechanisms. However, given that PowerMLP shares a similar underlying architecture with traditional MLP, it is feasible to substitute the MLP parts in CNNs or Transformers with PowerMLP naturally and expect that more complicated problems can be solved better by using PowerMLP as a basic building block.

Acknowledgements

This work is partially supported by NKRDPC grant 2018YFA0704705, NSFC grant 12288201, and the Beijing Natural Science Foundation grant QY23156.

References

- Aghaei, A. A. 2024. rKAN: Rational Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2406.14495*.
- Becker, B.; and Kohavi, R. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- Bozorgasl, Z.; and Chen, H. 2024. Wav-KAN: Wavelet Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2405.12832*.
- Carlo, G. D.; Mastropietro, A.; Anagnostopoulos, A.; and Aris. 2024. Kolmogorov-Arnold Graph Neural Networks. *arXiv preprint arXiv:2406.18354*.
- Cox, M. G. 1972. The Numerical Evaluation of B-Splines. *IMA Journal of Applied Mathematics*, 10(2): 134–149.
- Curry, H. B.; and Schoenberg, I. J. 1947. On spline distributions and their limits—the polya distribution functions. *Bulletin of the American Mathematical Society*, 53(11): 1114–1114.
- Davies, A.; Juhász, A.; Lackenby, M.; and Tomasev, N. 2022. The signature and cusp geometry of hyperbolic knots. *arXiv preprint arXiv:2111.15323*.
- Davies, A.; Veličković, P.; Buesing, L.; Blackwell, S.; Zheng, D.; Tomašev, N.; Tanburn, R.; Battaglia, P.; Blundell, C.; Juhász, A.; et al. 2021. Advancing mathematics by guiding human intuition with AI. *Nature*, 600(7887): 70–74.
- de Boor, C. 1978. *A Practical Guide to Splines*, volume 27. New York: Springer.
- Galitsky, B. A. 2024. Kolmogorov-Arnold network for word-level explainable meaning representation. *Preprints*.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep Sparse Rectifier Neural Networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, 315–323.
- Gómez Hidalgo, J. M.; Bringas, G. C.; Sánz, E. P.; and García, F. C. 2006. Content based SMS spam filtering. In *Proceedings of the 2006 ACM symposium on Document engineering*, 107–114.
- Greville, T. 1969. *Theory and Applications of Spline Functions*. Army. Mathematics Research Center, Madison, Wis. Publication. Academic Press. ISBN 9780123029508.
- Haykin, S. 1998. *Neural Networks: A Comprehensive Foundation*. USA: Prentice Hall PTR, 2nd edition. ISBN 0132733501.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034.
- Hendrycks, D.; and Gimpel, K. 2023. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415*.
- Inzirillo, H.; and Genet, R. 2024. SigKAN: Signature-Weighted Kolmogorov-Arnold Networks for Time Series. *arXiv preprint arXiv:2406.17890*.
- Kiamari, M.; Kiamari, M.; Kiamari, M.; and Krishnamachari, B. 2024. GKAN: Graph Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2406.06470*.
- LeCun, Y.; Boser, B.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W.; and Jackel, L. D. 1989. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4): 541–551.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Li, B.; Tang, S.; and Yu, H. 2020a. Better Approximations of High Dimensional Smooth Functions by Deep Neural Networks with Rectified Power Units. *Communications in Computational Physics*, 27(2): 379–411.
- Li, B.; Tang, S.; and Yu, H. 2020b. PowerNet: Efficient Representations of Polynomials and Smooth Functions by Deep Neural Networks with Rectified Power Units. *J. Math. Study*, 53(2): 159–191.
- Li, Z. 2024. Kolmogorov-Arnold Networks are Radial Basis Function Networks. *arXiv preprint arXiv:2405.06721*.
- Liu, Z.; Wang, Y.; Vaidya, S.; Ruehle, F.; Halverson, J.; Soljačić, M.; Hou, T. Y.; and Tegmark, M. 2024. KAN: Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2404.19756v4*.
- Maas, A. L.; Hannun, A. Y.; Ng, A. Y.; et al. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, 3.
- Mhaskar, H. N. 1993. Approximation properties of a multilayered feedforward artificial neural network. *Advances in Computational Mathematics*, 1: 61–80.
- Molchanov, P.; Tyree, S.; Karras, T.; Aila, T.; and Kautz, J. 2017. Pruning Convolutional Neural Networks for Resource Efficient Inference. *arXiv preprint arXiv:1611.06440*.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*.
- Ramos, J.; et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, 29–48.
- Vaswani, A. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Xu, K.; Chen, L.; and Wang, S. 2024. Kolmogorov-Arnold Networks for Time Series: Bridging Predictive Power and Interpretability. *arXiv preprint arXiv:2406.02496*.
- Yu, R.; Yu, W.; and Wang, X. 2024. KAN or MLP: A Fairer Comparison. *arXiv preprint arXiv:2407.16674*.
- Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.