

Efficient Visual Representation Learning with Heat Conduction Equation

Zhemín Zhang¹, Xun Gong^{1,2*}

¹School of Computing and Artificial Intelligence, Southwest Jiaotong University

²Engineering Research Center of Sustainable Urban Intelligent Transportation, Ministry of Education
zheminzhang@my.swjtu.edu.cn, xgong@swjtu.edu.cn

Abstract

Foundation models, such as CNNs and ViTs, have powered the development of image representation learning. However, general guidance to model architecture design is still missing. Inspired by the connection between image representation learning and heat conduction, we model images by the heat conduction equation, where the essential idea is to conceptualize image features as temperatures and model their information interaction as the diffusion of thermal energy. Based on this idea, we find that many modern model architectures, such as residual structures, SE block, and feed-forward networks, can be interpreted from the perspective of the heat conduction equation. Therefore, we leverage the heat equation to design new and more interpretable models. As an example, we propose the Heat Conduction Layer and the Refinement Approximation Layer inspired by solving the heat conduction equation using Finite Difference Method and Fourier series, respectively. The main goal of this paper is to integrate the overall architectural design of neural networks into the theoretical framework of heat conduction. Nevertheless, our Heat Conduction Network (HcNet) still shows competitive performance, e.g., HcNet-T achieves 83.0% top-1 accuracy on ImageNet-1K while only requiring 28M parameters and 4.1G MACs. The code is publicly available at: <https://github.com/ZhemínZhang1/HcNet>.

1 Introduction

The dominance of convolutional neural networks (CNNs) in computer vision has been long-standing. Transformers [Vaswani *et al.*, 2017], originated from natural language processing (NLP), have recently demonstrated state-of-the-art performance in visual learning. Despite tremendous successes, the full self-attention has computational and memory requirement that is quadratic in the sequence length, thereby limiting its applicability to high-resolution vision tasks. Moreover, modeling 2-D images as sequences impairs

*Corresponding Author

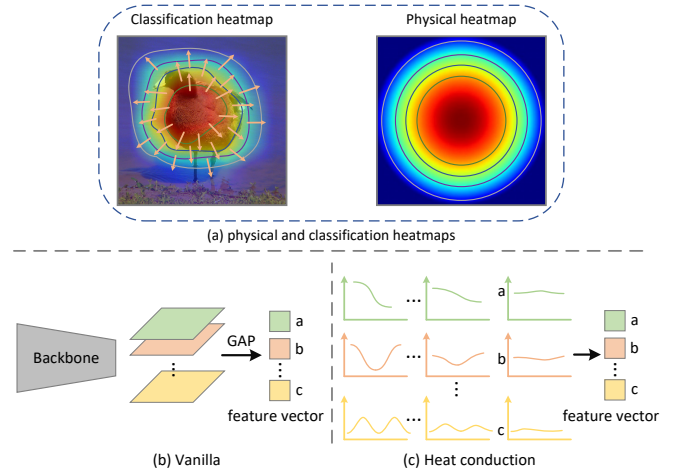


Figure 1: The connection between vanilla image representation learning and heat conduction. (a) Comparison of physical and classification heatmaps. (b) Vanilla feature extraction process. GAP denotes global average pooling and a, b, and c are scalars. (c) Thermal diffusion process. The temperature distribution gradually approaches a steady state over time. This process is similar to information compression and abstraction in vanilla feature extraction. Note that, for the sake of clarity, we use 1D temperature distributions as examples, and 2D temperature distributions are similar.

interpretability. Recently, using classical physical models such as state space [Gu and Dao, 2024; Liu *et al.*, 2024], gravitational field [Xu *et al.*, 2022], and heat conduction equation [Chen *et al.*, 2022] to inspire architecture design has become an effective way to explore new model architectures.

Our motivation is primarily driven by two key observations. First, the physical heatmap computed from heat equation and the classification heatmap learned from supervised learning are very similar, as shown in Figure 1(a). This may be due to the propagation of thermal energy through the collision of neighboring particles during heat conduction, resulting in similar temperatures in local space, which is similar to the property that neighboring regions tend to contain related information in images. Secondly, the vanilla feature extraction process maps the image to a low-dimensional feature space, which can be regarded as the compression and abstraction of the image. In this process, the model retains the most

important and discriminative global information in the image, while ignoring redundant details, as shown in Figure 1(b). This mechanism is similar to heat conduction, where high-frequency information gradually decays, and low-frequency global information becomes dominant over time, as shown in Figure 1(c). Based on the above inspiration, connecting heat conduction and visual semantic propagation, we propose HcNet. This work shows that many modern model components, such as residual structures [He *et al.*, 2016], SE block [Hu *et al.*, 2018] and feed-forward network (FFN) [Vaswani *et al.*, 2017], can be interpreted in terms of the heat conduction equation. From the perspective of the heat equation, the success of deep neural networks is mainly due to their ability to efficiently approximate heat conduction processes with complex initial conditions.

By establishing a connection between model architecture and the heat conduction equation, we can take advantage of the rich knowledge in the heat conduction equation to guide us in designing novel and more interpretable models. As an example, we propose the Heat Conduction Layer and the Refinement Approximation Layer inspired by solving the heat conduction equation using Finite Difference Method and Fourier series, respectively. Both layers can be explained mathematically using the concept of discrete heat conduction equation. Building upon the aforementioned components, we design a general vision backbone with a hierarchical architecture, named as Heat Conduction Network (HcNet). Compared to vision backbones with various architectures, HcNet consistently achieves competitive performance on image classification, object detection, and semantic segmentation across model scales. Specifically, HcNet-Base achieves 84.3% top-1 accuracy on ImageNet-1K, surpassing Swin [Liu *et al.*, 2021] by 0.8%. In addition, HcNet has lower computational cost compared to ViT-based models.

Our contributions can be summarized as follows:

- We propose the Heat Conduction Layer, a network layer inspired by the finite difference solution of the heat conduction equation, which follows the heat conduction principle to transfer information with low computational complexity and high interpretability.
- We propose the Refinement Approximation Layer, a network layer inspired by the Fourier series solution of the heat conduction equation, which improves the fitting ability of the model by linearly combining the features of different channels. The Refinement Approximation Layer explains the role of the FFN layer in terms of Fourier series approximation.
- Unlike previous work that applied the heat conduction equation to a specific module of the model, HcNet attempts to guide the design of all model components through the heat conduction equation. This approach helps to introduce rich knowledge in partial differential equations into neural networks.

2 Related Work

Convolution Neural Networks (CNNs). In the early eras, ConvNet [LeCun *et al.*, 1998] serves as the de-facto standard

network design for computer vision. The distinctive characteristics of CNNs are encapsulated in the convolution kernels, which enjoy high computational efficiency given specifically designed GPUs. Therefore, many convolutional neural architectures [Radosavovic *et al.*, 2020; Tan and Le, 2019; Radosavovic *et al.*, 2020; Tan and Le, 2021] have been proposed as the vision backbone for various visual applications. To address the born limitations of local receptive fields, recent works [Chen *et al.*, 2024; Ding *et al.*, 2024] have been focused on developing large convolution kernels to obtain larger receptive fields.

Vision Transformers (ViTs). Transformers were proposed by Vaswani *et al.* [Vaswani *et al.*, 2017] for machine translation, and have since become the state-of-the-art method in many NLP tasks. The pioneering work, Vision Transformer [Dosovitskiy *et al.*, 2021] demonstrates that pure Transformer-based architectures can also achieve very competitive results in computer vision. One challenge for Vision Transformer-based models is data efficiency. Although ViT can perform better than convolutional networks with hundreds of millions of images for pre-training, such a data requirement is difficult to meet in many cases. To improve data efficiency, many recent works [Liu *et al.*, 2021; Tian *et al.*, 2023; Zhu *et al.*, 2023; Wang *et al.*, 2024; Zhang *et al.*, 2024] have focused on introducing the locality and hierarchical structure into ViT. On the other hand, many works have blended attention model with convolution, and accordingly propose a series of hybrid Vision Transformers [Hatamizadeh *et al.*, 2023; Yu *et al.*, 2023].

Physics Inspired Models (PIMs). The principles of physics bring a novel perspective to foundation model research, taking advantage of the rich knowledge in physics to guide us in designing new and potentially more effective models. The diffusion model [Xu *et al.*, 2022], motivated by the gravitational field, interprets the data points as electrical charges on a hyperplane and devises a backward ODE to generate the image. Mamba [Gu and Dao, 2024] models sequences based on state space, and Vision Mamba [Zhu *et al.*, 2024; Liu *et al.*, 2024] introduces state space into image modeling inspired by mamba. QB-Heat [Chen *et al.*, 2022] extends the heat equation into high dimensional feature space to predict surrounding image patches for masked image modeling self-supervised task.

Unlike previous work that applied the heat conduction equation to a specific module of the model, HcNet attempts to guide the design of all model components through the heat conduction equation.

3 Methodology

3.1 Preliminaries: Heat Conduction Equation

Compared with the simple homogeneous diffusion process, the anisotropic heat equation can control the diffusion intensity in different directions, thereby enabling a more precise simulation of the complex heat conduction process. This capability establishes a foundation for the model to effectively capture complex semantic information within images. Considering that images are typically anisotropic, we employ the

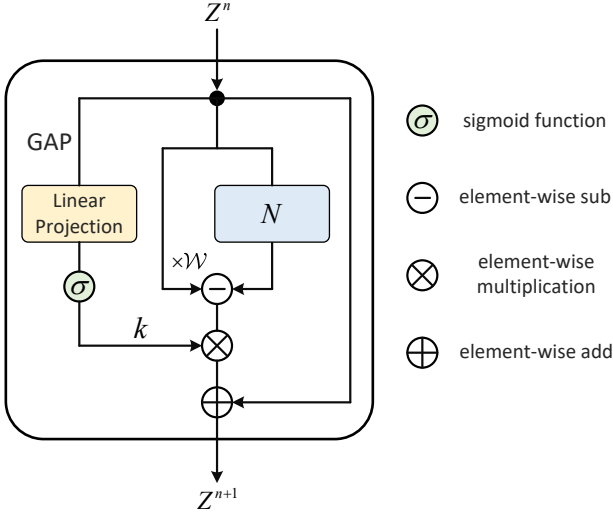


Figure 2: Heat Conduction Layer. N denotes calculating the weighted sum of neighboring points.

anisotropic heat equation [Widder, 1976] for modeling:

$$\frac{\partial u}{\partial t} = \alpha_x \frac{\partial^2 u}{\partial x^2} + \alpha_y \frac{\partial^2 u}{\partial y^2} \quad (1)$$

where $u(x, y, t)$ is the temperature of point (x, y) at time t within a 2-D region, α_x and α_y denote the thermal diffusivity [Bird, 2002] of the material in the x - and y -directions, respectively.

Discretization. We use the Finite Difference Method (FDM) to solve the Heat Conduction Equation approximately. Divide the spatial domain into a grid of points. Let the grid points be (x_i, y_j) , where $i = 1, 2, \dots, N_x$ and $j = 1, 2, \dots, N_y$. The time domain is also discretized into time steps t^n , where $n = 0, 1, 2, \dots, N_t$. Substitute the finite difference approximations into the heat conduction equation:

$$\begin{aligned} \frac{u_{i,j}^{n+1} - u_{i,j}^n}{\Delta t} = & \frac{\alpha_{x,1}(u_{i+1,j}^n - u_{i,j}^n) + \alpha_{x,2}(u_{i-1,j}^n - u_{i,j}^n)}{\Delta x^2} \\ & + \frac{\alpha_{y,1}(u_{i,j+1}^n - u_{i,j}^n) + \alpha_{y,2}(u_{i,j-1}^n - u_{i,j}^n)}{\Delta y^2} \end{aligned} \quad (2)$$

Rearrange to solve for $u_{i,j}^{n+1}$:

$$\begin{aligned} u_{i,j}^{n+1} = & u_{i,j}^n + \Delta t \left(\frac{\alpha_{x,1}(u_{i+1,j}^n - u_{i,j}^n) + \alpha_{x,2}(u_{i-1,j}^n - u_{i,j}^n)}{\Delta x^2} \right. \\ & \left. + \frac{\alpha_{y,1}(u_{i,j+1}^n - u_{i,j}^n) + \alpha_{y,2}(u_{i,j-1}^n - u_{i,j}^n)}{\Delta y^2} \right) \end{aligned} \quad (3)$$

3.2 Heat Conduction Layer

Suppose $Z \in \mathbb{R}^{C \times H \times W}$ is the image feature map in networks, C is the number of channels, H is the height of the feature, and W is the width of the feature. We draw an analogy between feature information propagation in feature maps and temperature conduction, treating each channel as a separate 2-D heat conduction process, denoted as $z \in \mathbb{R}^{H \times W}$.

The heat conduction layer embodies this concept of information conduction through temperature. Based on this idea, we first replace temperature u in the original heat equation with feature z , and further set the thermal diffusivity α as a learnable parameter, so that the model can adaptively adjust the diffusion behavior according to the local structure of the feature map. Since the model works with discrete image patch features, it naturally leads to a grid of points where $\Delta x = \Delta y = \Delta d$. Substituting into Eq.3 yields:

$$\begin{aligned} z_{i,j}^{n+1} = & z_{i,j}^n + \frac{\Delta t}{\Delta d^2} \left(w_1 z_{i+1,j}^n + w_2 z_{i-1,j}^n + w_3 z_{i,j+1}^n + w_4 z_{i,j-1}^n \right. \\ & \left. - (w_1 + w_2 + w_3 + w_4) z_{i,j}^n \right) \end{aligned} \quad (4)$$

where $w_1 = \alpha_{x,1}$, $w_2 = \alpha_{x,2}$, $w_3 = \alpha_{y,1}$, $w_4 = \alpha_{y,2}$, w_i is a learnable parameter. Eq.4 is further simplified as follows:

$$z_{i,j}^{n+1} = z_{i,j}^n + k (N(z_{i,j}^n) - \mathcal{W} z_{i,j}^n), \quad k = \frac{\Delta t}{\Delta d^2} \quad (5)$$

where $N(z_{i,j}^n)$ denotes calculating the weighted sum of the neighboring points of $z_{i,j}^n$, $\mathcal{W} = w_1 + w_2 + w_3 + w_4$. Based on Eq.5, we propose the Heat Conduction Layer, which considers the layer dimension of the model as the time step, i.e., the feature passes through a Heat Conduction Layer and t^n becomes t^{n+1} . Meanwhile, inspired by Mamba [Gu and Dao, 2024], we make k input-dependent to increase the context-aware ability of the model. Heat Conduction Layer:

$$\begin{aligned} k^n &= \sigma(\text{Linear}(\text{GAP}(Z^n))) \\ h^n &= N(z^n) - \mathcal{W} z^n \\ z^{n+1} &= z^n + k^n h^n \end{aligned} \quad (6)$$

where $Z^n \in \mathbb{R}^{C \times H \times W}$, $\text{GAP}: \mathbb{R}^{C \times H \times W} \mapsto \mathbb{R}^C$ denotes global average pooling, Linear denotes the mapping functions like fully connected layer, and k is a scalar. σ is a sigmoid function, it ensures $k > 0$ and limits the size of k , thereby improving the stability of the FDM solution. From the principles of heat conduction, we naturally derive the residual structure [He *et al.*, 2016] and the SE block [Hu *et al.*, 2018], used in modern model architectures. The overall structure of the Heat Conduction Layer is shown in Figure 2.

Eq.6 illustrates the calculation process for a single channel in the feature map Z^{n+1} , and the other channels are calculated in the same way.

3.3 Preliminaries: Fourier Series

Principle of Superposition. Linear partial differential equations follow the superposition principle. For the heat conduction equation, this principle states that if $u_1(x, y, t)$, $u_2(x, y, t)$, \dots , $u_n(x, y, t)$ are both solutions to the heat equation, any linear combination of these solutions is also a solution:

$$u(x, y, t) = \sum_{i=1}^n a_i u_i(x, y, t) \quad (7)$$

where u_1, u_2, \dots, u_n satisfy Eq.1, $u(x, y, t)$ is also a solution for any constants a_1, a_2, \dots, a_n . The principle of superposition allows us to construct complex solutions to the heat equation from simpler ones.




Figure 3: Refinement Approximation Layer. We take a feature map with two channels as an example; feature maps with more channels are processed in a similar manner.

Fourier Series. The Fourier series is a powerful tool for solving the heat equation. The Fourier series solution of the heat equation in a two-dimensional rectangular space is as follows:

$$u(x, y, t) = \sum_{m=1}^M \sum_{n=1}^N B_{mn} \sin\left(\frac{m\pi x}{L}\right) \sin\left(\frac{n\pi y}{L}\right) e^{-\frac{k(m^2+n^2)\pi^2 t}{L^2}} \quad (8)$$

where m and n denote constant frequencies, B_{mn} is the Fourier coefficient, and L is the length of the sides of the rectangular region. Further simplification of Eq.8:

$$u(x, y, t) = \sum_{m=1}^M \beta_m g_m(x, y, t) \quad (9)$$

where $g_m(x, y, t) = \sin\left(\frac{m\pi x}{L}\right) \sin\left(\frac{n\pi y}{L}\right) e^{-\frac{k(m^2+n^2)\pi^2 t}{L^2}}$, $\beta_m = B_{mn}$, β_m is a constant. From Eq.9, we can conclude that the Fourier series fits the complex heat conduction process by linear combination of basis functions $\{g_m(x, y, t)\}_{m=1}^M$ at different frequencies, and becomes a more precise approximation as more terms are included. That is, as M increases, it can represent more intricate patterns and capture much finer details.

3.4 Refinement Approximation Layer

In *Heat Conduction Layer* subsection, we have illustrated the heat conduction of features on a single channel. To effectively extract visual features, it is also necessary to integrate

information from multiple channels. We consider each channel as a separate 2-D heat conduction process. Therefore, each channel can be approximated using the Fourier series. According to Eq.9:

$$z_c(x, y, t^n) = \sum_{h=1}^H \beta_h g_h(x, y, t^n) \quad (10)$$

where $z_c(x, y, t^n)$ denotes the featured heat distribution of the c -th channel of the t^n time-step. The accuracy of the approximation can be improved by increasing the number of terms in the Fourier series:

$$\begin{aligned} \hat{z}_c(x, y, t^n) &= \sum_{h=1}^H \beta_h g_h(x, y, t^n) + \sum_{m=H}^M \beta_m g_m(x, y, t^n) \\ &= z_c(x, y, t^n) + \sum_{m=H}^M \beta_m g_m(x, y, t^n) \end{aligned} \quad (11)$$

Based on Eq.11, we propose the Refinement Approximation Layer. Mathematically, considering the input as $Z^n = \{z_c(x, y, t^n)\}_{c=1}^C$. Firstly, based on Eq.7, we use $\{z_c(x, y, t^n)\}_{c=1}^C$ to construct more new solutions as candidate terms $\{z_m(x, y, t^n)\}_{m=1}^{M-H}$. It is worth noting that the simple linear combination cannot generate new frequencies, so we perform a nonlinear transformation of $\{z_m(x, y, t^n)\}_{m=1}^{M-H}$ to introduce new frequencies to obtain new basis functions $\{g_m\}_{m=1}^{M-H}$. This can be done by a nonlinear activation function, taking the GELU [Hendrycks and Gimpel, 2023] as an example:

$$\text{GELU}(z) \approx 0.5z \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (z + 0.044715z^3) \right) \right) \quad (12)$$

where higher-order terms (z^2, z^3) will bring new frequencies, take z^2 as an example:

$$\begin{aligned} z(x, y, t^n)^2 &= B_{mn}^2 \sin^2\left(\frac{m\pi x}{L}\right) \sin^2\left(\frac{n\pi y}{L}\right) e^{-2\frac{k(m^2+n^2)\pi^2 t^n}{L^2}} \\ &= \frac{B_{mn}^2}{4} e^{-2\frac{k(m^2+n^2)\pi^2 t^n}{L^2}} \left(1 - \cos\left(\frac{2m\pi x}{L}\right) - \cos\left(\frac{2n\pi y}{L}\right) \right. \\ &\quad \left. + \frac{1}{2} \cos\left(\frac{2m\pi x}{L} + \frac{2n\pi y}{L}\right) + \frac{1}{2} \cos\left(\frac{2m\pi x}{L} - \frac{2n\pi y}{L}\right) \right) \end{aligned} \quad (13)$$

New frequencies, such as $2m$ and $2n$, are generated in Eq.13. Although nonlinear transformations can generate new higher-order harmonics, they also introduce undesirable components (e.g., DC component) and make it challenging to precisely control the amplitude of specific Fourier series terms. To address this issue, we introduce convolution operations as filters, enabling the model to adaptively select Fourier series terms with specific frequencies and amplitudes. Specifically, we input the results of the nonlinear transformation into a depth-wise separable convolution to obtain the final basis functions $\{g_m\}_{m=1}^{M-H}$. Extensive research has demonstrated that convolutional kernels can be regarded as filters [Mairal et

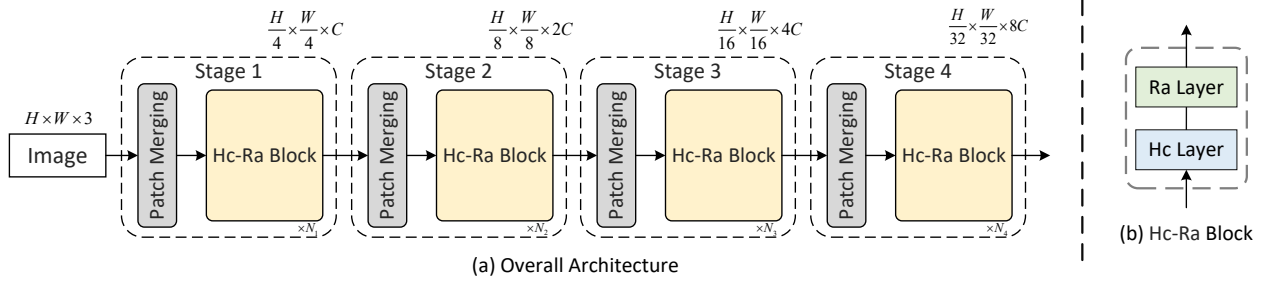


Figure 4: (a) **The overall framework of HcNet.** HcNet adopts hierarchical architecture with 4 stages, where stage [1, 2, 3, 4] have $[N_1, N_2, N_3, N_4]$ blocks, respectively. (b) **The architecture of Hc-Ra block.** Compared with transformer block, it does not draw the residual structure, as it is inherently incorporated within the heat conduction layer or the refinement approximation layer.

al., 2014; Zoumpourlis *et al.*, 2017]. Then, based on Eq.11, each channel uses a separate set of β_m to linearly combine g_m to obtain the new term $\sum_{m=H}^M \beta_m g_m(x, y, t^n)$. Finally, add the new term to the input $z_c(x, y, t^n)$ to get $\hat{z}_c(x, y, t^n)$, where the number of terms increases from H to M . Refinement Approximation Layer:

$$\begin{aligned} z_m^n &= \text{Linear}(Z^n) \\ g_m^n &= \tau(z_m^n) \\ z_M^n &= \text{Linear}(g_m^n) \\ \hat{z}^n &= z^n + z_M^n \end{aligned} \quad (14)$$

where Z^n denotes the feature at the n -th time-step. $\tau(\cdot)$ denotes the nonlinear transformation, which consists of a nonlinear activation function and a depth-wise convolutional layer. z_M^n denotes the new term $\sum_{m=H}^M \beta_m g_m(x, y, t_n)$, and the overall flow is shown in Figure 3. The Refinement Approximation Layer increases the number of Fourier series terms in each channel, so as the depth of the model increases, the model’s ability to capture complex patterns and dynamics becomes stronger.

From the perspective of the heat equation, the role of the nonlinear activation function is to generate Fourier series terms with new frequencies, while the role of the MLP layer is to linearly combine the basis functions to approximate any function.

Version	Blocks				Dims			
	S_1	S_2	S_3	S_4	D_1	D_2	D_3	D_4
HcNet-T	5	5	14	5	64	128	320	512
HcNet-S	6	6	18	6	64	192	384	768
HcNet-B	10	10	28	10	96	192	384	768

Table 1: Detailed settings of HcNet series.

3.5 HcNet Model

Equipped with the above Heat Conduction layer and Refinement Approximation layer, the Hc-Ra block is shown in Figure 4 (b). We take advantage of the Hc-Ra block by adopting a hierarchical structure similar to hierarchical transformer

General Arch.	Model	Params	MACs	Top-1 acc.
CNNs	RegNetY-4G [Radosavovic <i>et al.</i> , 2020]	21M	4.0G	80.0
	RegNetY-16G [Radosavovic <i>et al.</i> , 2020]	84M	16.0G	82.9
	ConvNeXt-T [Liu <i>et al.</i> , 2022]	29M	4.5G	82.1
	ConvNeXt-S [Liu <i>et al.</i> , 2022]	50M	8.7G	83.1
	ConvNeXt-B [Liu <i>et al.</i> , 2022]	89M	15.4G	83.8
ViTs	ViT-B [Dosovitskiy <i>et al.</i> , 2020]	86M	55.4G	77.9
	Swin-T [Liu <i>et al.</i> , 2021]	28M	4.6G	81.3
	Swin-S [Liu <i>et al.</i> , 2021]	50M	8.7G	83.0
	Swin-B [Liu <i>et al.</i> , 2021]	88M	15.4G	83.5
	BiFormer-T [Zhu <i>et al.</i> , 2023]	13M	2.2G	81.4
	BiFormer-S [Zhu <i>et al.</i> , 2023]	26M	4.5G	83.8
	BiFormer-B [Zhu <i>et al.</i> , 2023]	57M	9.8G	84.3
	CrossFormer-S [Wang <i>et al.</i> , 2024]	31M	4.9G	82.5
	CrossFormer-B [Wang <i>et al.</i> , 2024]	52M	9.2G	83.4
	CrossFormer-L [Wang <i>et al.</i> , 2024]	92M	16.1G	84.0
	QFormer-T [Zhang <i>et al.</i> , 2024]	29M	4.6G	82.5
PIMs	QFormer-S [Zhang <i>et al.</i> , 2024]	51M	8.9G	84.0
	QFormer-B [Zhang <i>et al.</i> , 2024]	90M	15.7G	84.1
	Vim-S [Zhu <i>et al.</i> , 2024]	26M	5.3G	81.4
	Vim-B [Zhu <i>et al.</i> , 2024]	98M	19.0G	83.2
	VMamba-T [Liu <i>et al.</i> , 2024]	31M	4.9G	82.6
	VMamba-S [Liu <i>et al.</i> , 2024]	50M	8.7G	83.6
	VMamba-B [Liu <i>et al.</i> , 2024]	89M	15.4G	83.9
	HcNet-T	28M	4.1G	83.0
	HcNet-S	51M	8.6G	83.9
	HcNet-B	85M	14.2G	84.3

Table 2: Comparison of different types of models on ImageNet-1K. **PIMs** denote physics inspired models.

variants [Liu *et al.*, 2021; Zhu *et al.*, 2023]. Figure 4 (a) shows the overall framework of HcNet. Specifically, HcNet has 4 stages, each stage contains a patch merging layer and several Hc-Ra blocks. As the network gets deeper, the input features are spatially downsampled by a certain ratio through the patch merging layer. The channel dimension is expanded twice to produce a hierarchical image representation. Specifically, the spatial downsampling ratio is set to 4 in the first stage and 2 in the last three stages. The outputs of the patch merging layer are fed into the subsequent Hc-Ra block, and the number of tokens is kept constant.

Based on this architecture, we build three different variants of HcNet, including HcNet-T, HcNet-S, and HcNet-B, as shown in Table 1.

Backbone	Params	MACs	Mask R-CNN 1× schedule						Mask R-CNN 3× schedule + MS					
			AP ^{box}	AP ^{box} ₅₀	AP ^{box} ₇₅	AP ^{mask}	AP ^{mask} ₅₀	AP ^{mask} ₇₅	AP ^{box}	AP ^{box} ₅₀	AP ^{box} ₇₅	AP ^{mask}	AP ^{mask} ₅₀	AP ^{mask} ₇₅
ResNet-50 [He <i>et al.</i> , 2016]	44M	260G	38.0	58.6	41.4	34.4	55.1	36.7	41.0	61.7	44.9	37.1	58.4	40.1
Swin-T [Liu <i>et al.</i> , 2021]	48M	264G	43.7	66.6	47.6	39.8	63.3	42.7	46.0	68.2	50.2	41.6	65.1	44.8
Focal-T [Yang <i>et al.</i> , 2021]	49M	291G	44.8	67.7	49.2	41.0	64.7	44.2	47.1	69.4	51.9	42.7	66.5	45.9
CrossFormer-S [Wang <i>et al.</i> , 2024]	50M	301G	45.4	68.0	49.7	41.4	64.8	44.6	—	—	—	—	—	—
QFormer-T [Zhang <i>et al.</i> , 2024]	49M	—	45.9	68.5	50.3	41.5	65.2	44.6	47.5	69.6	52.1	42.7	66.4	46.1
VMamba-T [Liu <i>et al.</i> , 2024]	42M	262G	46.5	68.5	50.7	42.1	65.5	45.3	48.5	69.9	52.9	43.2	66.8	46.3
HcNet-T	35M	242G	46.0	68.3	50.2	41.9	65.4	44.7	47.6	69.3	52.3	42.8	66.5	45.8
ResNet-101 [He <i>et al.</i> , 2016]	63M	336G	40.4	61.1	44.2	36.4	57.7	38.8	42.8	63.2	47.1	38.5	60.1	41.3
Swin-S [Liu <i>et al.</i> , 2021]	69M	354G	44.8	66.6	48.9	40.9	63.2	44.2	48.2	69.8	52.8	43.2	67.0	46.1
CSWin-S [Dong <i>et al.</i> , 2022]	54M	342G	47.9	70.1	52.6	43.2	67.1	46.2	50.0	71.3	54.7	44.5	68.4	47.7
CrossFormer-B [Wang <i>et al.</i> , 2024]	71M	408G	47.2	69.9	51.8	42.7	66.6	46.2	—	—	—	—	—	—
QFormer-S [Zhang <i>et al.</i> , 2024]	70M	—	—	—	—	—	—	—	49.5	71.2	54.2	44.2	68.3	47.6
VMamba-S [Liu <i>et al.</i> , 2024]	64M	357G	48.2	69.7	52.5	43.0	66.6	46.4	49.7	70.4	54.2	44.0	67.6	47.3
HcNet-S	50M	311G	47.6	69.3	52.2	43.1	66.5	46.3	49.8	71.1	54.6	44.7	68.1	47.7
RegNeXt-101-64 [Xie <i>et al.</i> , 2017]	101M	493G	42.8	63.8	47.3	38.4	60.6	41.3	44.4	64.9	48.8	39.7	61.9	42.6
Swin-B [Liu <i>et al.</i> , 2021]	107M	496G	46.9	—	—	42.3	—	—	48.5	69.8	53.2	43.4	66.8	46.9
CSWin-B [Dong <i>et al.</i> , 2022]	97M	526G	48.7	70.4	53.9	43.9	67.8	47.3	50.8	72.1	55.8	44.9	69.1	48.3
VMamba-B [Liu <i>et al.</i> , 2024]	108M	485G	49.2	70.9	53.9	43.9	67.7	47.6	—	—	—	—	—	—
HcNet-B	90M	473G	48.8	70.2	53.7	44.3	67.9	47.5	50.4	71.9	55.6	45.0	69.3	48.5

Table 3: Object detection and instance segmentation performance on the COCO val2017 with the Mask R-CNN framework. The models have been pre-trained on ImageNet-1K. The resolution used to calculate MACs is 800×1280.

4 Experiments

To demonstrate the effectiveness of the HcNet, we conduct experiments on ImageNet-1K [Deng *et al.*, 2009]. To further evaluate the generalization and robustness of our backbone, we also conduct experiments on ADE20K [Zhou *et al.*, 2017] for semantic segmentation, and COCO [Lin *et al.*, 2014] for object detection. Finally, we perform comprehensive ablation studies to analyze each component of the HcNet.

4.1 Classification on the ImageNet-1K

Implementation details. This setting mostly follows Swin [Liu *et al.*, 2021]. We use the PyTorch toolbox [Paszke *et al.*, 2019] to implement all our experiments. We employ an AdamW [Kingma and Ba, 2014] optimizer for 300 epochs using a cosine decay learning rate scheduler and 20 epochs of linear warm-up. A batch size of 256, an initial learning rate of 0.001, and a weight decay of 0.05 are used. ViT-B/16 uses an image size 384×384 and others use 224×224. We include most of the augmentation and regularization strategies of Swin transformer [Liu *et al.*, 2021] in training.

Results. Table 2 compares the performance of the proposed HcNet with the state-of-the-art CNN, Vision Transformer and PIM backbones on ImageNet-1K. Surprisingly, despite the simple heat conduction layer, HcNets can still achieve highly competitive performance compared with CNNs, ViTs and PIMs. For example, HcNet-T reaches the top-1 accuracy of 83.0 while only requiring 28M parameters and 4.1G MACs. Compared with ViT-B and Swin-T, HcNet-T outperforms them by 5.1 and 1.7, respectively. With similar parameters and MACs, HcNet performs better than other PIMs. Experimental results show that physics inspired models have the potential to achieve reasonable performance, and their interpretability is unmatched by other models.

4.2 Object Detection and Instance Segmentation

Implementation details. We use the Mask R-CNN [He *et al.*, 2017] framework to evaluate the performance of the proposed HcNet backbone on the COCO benchmark for object detection and instance segmentation. We pretrain the backbones on the ImageNet-1K dataset and apply the finetuning strategy used in Swin Transformer [Liu *et al.*, 2021] on the COCO training set.

Results. Table 3 reports the results of the Mask R-CNN framework with “1×” (12 training epoch) and “3× +MS” (36 training epoch with multi-scale training) schedule. With a 12-epoch fine-tuning schedule, HcNet-T/S/B models achieve object detection mAPs of 46.0%/47.6%/48.8%, outperforming Swin-T/S/B by 2.3%/2.8%/1.9% bAP, respectively. HcNets have similar performance in instance segmentation. Overall, for COCO object detection and instance segmentation, HcNets achieve competitive performance.

4.3 Semantic Segmentation

Implementation details. We further investigate the capability of HcNet for Semantic Segmentation on the ADE20K [Zhou *et al.*, 2017] dataset. Here we employ the widely-used UperNet [Xiao *et al.*, 2018] as the basic framework and followed Swin’s [Liu *et al.*, 2021] experimental settings. In Table 4, we report both the single-scale (SS) and multi-scale (MS) mIoU for better comparison. The default input resolution is 512 × 512.

Results. As shown in Table 4, HcNet can attain competitive performance on semantic segmentation. This further indicates the great potential of physics inspired models.

4.4 Visualization Analysis

To comprehensively evaluate the quality of the learned visual representations, we employ the GradCAM technique [Sel-

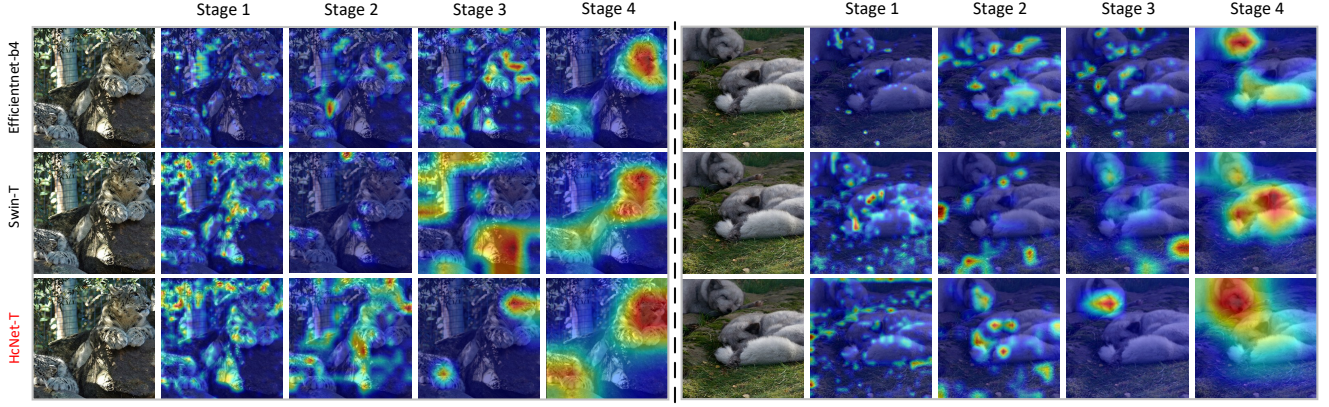


Figure 5: Grad-CAM visualization of the models trained on ImageNet-1K. The visualized images are randomly selected from the test dataset.

Backbone	Params	MACs	SS mIoU	MS mIoU
ResNet-50 [He <i>et al.</i> , 2016]	67M	953G	42.1	42.8
ConvNeXt-T [Liu <i>et al.</i> , 2022]	60M	939G	46.0	46.7
Swin-T [Liu <i>et al.</i> , 2021]	60M	945G	44.4	45.8
CSWin-T [Dong <i>et al.</i> , 2022]	60M	959G	49.3	—
CrossFormer-S [Wang <i>et al.</i> , 2024]	62M	979G	47.6	48.4
QFormer-T [Zhang <i>et al.</i> , 2024]	61M	—	46.9	48.1
Vim-S [Zhu <i>et al.</i> , 2024]	46M	—	44.9	—
VMamba-T [Liu <i>et al.</i> , 2024]	62M	948G	48.3	48.6
HcNet-T	50M	851G	48.1	48.3
ResNet-101 [He <i>et al.</i> , 2016]	85M	1030G	42.9	44.0
ConvNeXt-B [Liu <i>et al.</i> , 2022]	122M	1170G	49.1	49.9
Swin-B [Liu <i>et al.</i> , 2021]	121M	1188G	48.1	49.7
CSWin-B [Dong <i>et al.</i> , 2022]	109M	1222G	51.1	—
CrossFormer-L [Wang <i>et al.</i> , 2024]	125M	1258G	50.4	51.4
QFormer-B [Zhang <i>et al.</i> , 2024]	123M	—	49.5	50.6
VMamba-B [Liu <i>et al.</i> , 2024]	122M	1170G	51.0	51.6
HcNet-B	96M	1095G	51.0	51.5

Table 4: Comparison of the segmentation performance of different backbones on the ADE20K. All backbones are pretrained on ImageNet-1K.

varaju *et al.*, 2017] to generate activation maps for visual representations at various stages of efficientnet-b4 [Tan and Le, 2019], Swin-T, and HcNet-T. These activation maps provide insight into the significance of individual pixels in depicting class discrimination for each input image. As shown in Figure 5, the activation maps generated by HcNet are significantly different from the other two types of models, and identifying more semantically meaningful regions in deeper layers.

4.5 Ablation Studies

The experiments of ablation studies are conducted on ImageNet-1K. We discuss the ablation below according to the following aspects.

The way to get k . $k = \frac{\Delta t}{\Delta d^2}$ is the ratio of the time step to the spatial step. To demonstrate the effectiveness of input-dependent, we conduct the following experiments on ImageNet-1K: (1) fix k , (2) treat k as a learnable parameter for each layer, and (3) k in each layer depends on the in-

Settings	HcNet-T	HcNet-B
Fixed $k = 0.1$	81.3	82.7
Fixed $k = 1.0$	81.2	82.8
Fixed $k = 3.0$	81.0	82.5
k as a learnable parameter	82.2	83.3
k is input-dependent	83.0	84.3

 Table 5: Compare different ways of getting k . The table shows the Top-1 acc.

	HcNet-T	HcNet-B
without filter	82.1	83.5
with filter	83.0	84.3

Table 6: Impact of filters. The table shows the Top-1 acc.

put. As shown in Table 5, setting k to be input-dependent improves the performance significantly.

Filters. The generated new frequency terms are filtered using depth-wise convolution within the Refinement Approximation Layer. In this subsection, we explore the impact of depth-wise convolution on the model’s performance. The model performance while keeping the same number of parameters is shown in Table 6, and the results prove the necessity of filtering the Fourier series terms that introduce new frequencies after nonlinear transformation.

5 Conclusion

In this paper, we draw a relatively comprehensive connection between modern model architectures and the heat conduction equation. Such connection enables us to design new and more interpretable neural network architectures. Benefiting from this connection, we propose the Heat Conduction Layer and the Refinement Approximation Layer inspired by solving the heat conduction equation using Finite Difference Method and Fourier series, respectively. Extensive experiments demonstrate that HcNet achieves competitive performance across various visual tasks, offering new insights for the development of physics inspired model architecture design.

Acknowledgments

This work is partially supported by National Natural Science Foundation of China (62376231), Tangshan Basic Research Science and Technology Program (23130230E), Tangshan Lunan District Science and Technology Program, Fundamental Research Funds for the Central Universities (2682025ZTPY052).

References

- [Bird, 2002] R Byron Bird. Transport phenomena. *Appl. Mech. Rev.*, 55(1):R1–R4, 2002.
- [Chen *et al.*, 2022] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Lu Yuan, Zicheng Liu, and Youzuo Lin. Self-supervised learning based on heat equation, 2022.
- [Chen *et al.*, 2024] Honghao Chen, Xiangxiang Chu, Yongjian Ren, Xin Zhao, and Kaiqi Huang. Pelk: Parameter-efficient large kernel convnets with peripheral convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5557–5567, June 2024.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [Ding *et al.*, 2024] Xiaohan Ding, Yiyuan Zhang, Yixiao Ge, Sijie Zhao, Lin Song, Xiangyu Yue, and Ying Shan. Unireplknet: A universal perception large-kernel convnet for audio video point cloud time-series and image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5513–5524, June 2024.
- [Dong *et al.*, 2022] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12124–12134, June 2022.
- [Dosovitskiy *et al.*, 2020] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR*, abs/2010.11929, 2020.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [Gu and Dao, 2024] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2024.
- [Hatamizadeh *et al.*, 2023] Ali Hatamizadeh, Greg Heinrich, Hongxu Yin, Andrew Tao, Jose M. Alvarez, Jan Kautz, and Pavlo Molchanov. Fastervit: Fast vision transformers with hierarchical attention, 2023.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [He *et al.*, 2017] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [Hendrycks and Gimpel, 2023] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023.
- [Hu *et al.*, 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing.
- [Liu *et al.*, 2021] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, October 2021.
- [Liu *et al.*, 2022] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986, June 2022.
- [Liu *et al.*, 2024] Yue Liu, Yunjie Tian, Yuzhong Zhao, Hongtian Yu, Lingxi Xie, Yaowei Wang, Qixiang Ye, and Yunfan Liu. Vmamba: Visual state space model, 2024.
- [Mairal *et al.*, 2014] Julien Mairal, Piotr Koniusz, Zaid Harchaoui, and Cordelia Schmid. Convolutional kernel networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca

- Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- [Radosavovic et al., 2020] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollar. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [Selvaraju et al., 2017] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [Tan and Le, 2019] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019.
- [Tan and Le, 2021] Mingxing Tan and Quoc Le. Efficientnetv2: Smaller models and faster training. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10096–10106. PMLR, 18–24 Jul 2021.
- [Tian et al., 2023] Yunjie Tian, Lingxi Xie, Zhaozhi Wang, Longhui Wei, Xiaopeng Zhang, Jianbin Jiao, Yaowei Wang, Qi Tian, and Qixiang Ye. Integrally pre-trained transformer pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18610–18620, June 2023.
- [Vaswani et al., 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [Wang et al., 2024] Wenxiao Wang, Wei Chen, Qibo Qiu, Long Chen, Boxi Wu, Binbin Lin, Xiaofei He, and Wei Liu. Crossformer++: A versatile vision transformer hinging on cross-scale attention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3123–3136, 2024.
- [Widder, 1976] David Vernon Widder. *The heat equation*, volume 67. Academic Press, 1976.
- [Xiao et al., 2018] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [Xie et al., 2017] Saining Xie, Ross Girshick, Piotr Dollar, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [Xu et al., 2022] Yilun Xu, Ziming Liu, Max Tegmark, and Tommi Jaakkola. Poisson flow generative models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 16782–16795. Curran Associates, Inc., 2022.
- [Yang et al., 2021] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. *CoRR*, abs/2107.00641, 2021.
- [Yu et al., 2023] Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. Metaformer baselines for vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [Zhang et al., 2024] Qiming Zhang, Jing Zhang, Yufei Xu, and Dacheng Tao. Vision transformer with quadrangle attention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3608–3624, 2024.
- [Zhou et al., 2017] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [Zhu et al., 2023] Lei Zhu, Xinjiang Wang, Zhaghan Ke, Wayne Zhang, and Rynson W.H. Lau. Biformer: Vision transformer with bi-level routing attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10323–10333, June 2023.
- [Zhu et al., 2024] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 62429–62442. PMLR, 21–27 Jul 2024.
- [Zoumpourlis et al., 2017] Georgios Zoumpourlis, Alexandros Dumanoglou, Nicholas Vretos, and Petros Daras. Non-linear convolution filters for cnn-based learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.