Signal Processing for Implicit Neural Representations

Dejia Xu[*]
dejia@utexas.edu

Peihao Wang[*]
peihaowang@utexas.edu

Yifan Jiang
yifanjiang97@utexas.edu

Zhiwen Fan
zhiwenfan@utexas.edu

Zhangyang Wang
atlaswang@utexas.edu

The University of Texas at Austin
https://vita-group.github.io/INSP/

Abstract

Implicit Neural Representations (INRs) encoding continuous multi-media data via multi-layer perceptrons has shown undebatable promise in various computer vision tasks. Despite many successful applications, editing and processing an INR remains intractable as signals are represented by latent parameters of a neural network. Existing works manipulate such continuous representations via processing on their discretized instance, which breaks down the compactness and continuous nature of INR. In this work, we present a pilot study on the question: *how to directly modify an INR without explicit decoding?* We answer this question by proposing an implicit neural signal processing network, dubbed **INSP-Net**, via differential operators on INR. Our key insight is that spatial gradients of neural networks can be computed analytically and are invariant to translation, while mathematically we show that any continuous convolution filter can be uniformly approximated by a linear combination of high-order differential operators. With these two knobs, INSP-Net instantiates the signal processing operator as a weighted composition of computational graphs corresponding to the high-order derivatives of INRs, where the weighting parameters can be data-driven learned. Based on our proposed INSP-Net, we further build the first Convolutional Neural Network (CNN) that implicitly runs on INRs, named **INSP-ConvNet**. Our experiments validate the expressiveness of INSP-Net and INSP-ConvNet in fitting low-level image and geometry processing kernels (e.g. blurring, deblurring, denoising, inpainting, and smoothening) as well as for high-level tasks on implicit fields such as image classification.

1 Introduction

The idea that our visual world can be represented continuously has attracted increasing popularity in the field of implicit neural representations (INR). Also known as coordinate-based neural representations, INRs learn to encode a coordinate-to-value mapping for continuous multi-media data. Instead of storing the discrete signal values in a grid of pixels or voxels, INRs represent discrete data as samples of a continuous manifold. Using multi-layer perceptrons, INRs bring practical benefits to various computer vision applications, such as image and video compression [1, 2, 3], 3D shape representation [4, 5, 6, 7, 8, 9, 10, 11], inverse problems [12, 2, 13, 14], and generative models [15, 16, 17, 18, 19, 20, 21, 22].

Despite their recent success, INRs are not yet amenable to flexible editing and processing as the standard images could do. The encoded coordinate-to-value mapping is too complex to comprehend and the parameters stored in multi-layer perceptrons (MLPs) remains less explored. One direction of existing approaches enables editing on INRs by training them with conditional input. For example,

[*]Equal Contribution.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

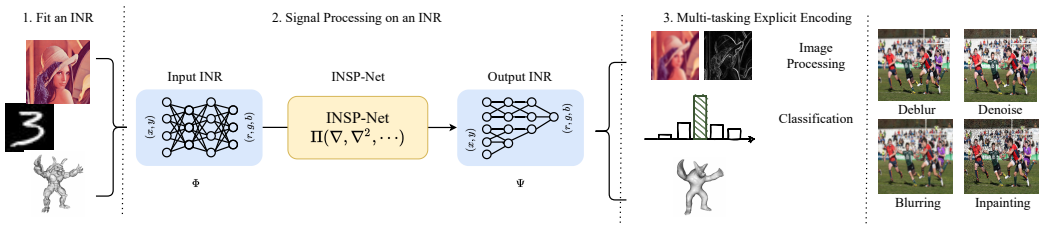

Figure 1: An illustration of implicit neural signal processing. Given an INR representing digital signals, our INSP-Net is capable of direct signal processing without needing to explicitly decode it. Our model first constructs derivative computation graphs of the original INR and then generates a linear combination of them into a new INR. It can be later decoded into discretized forms such as image pixels. The framework is capable of fitting low-level image processing kernels as well as performing high-level processing such as image classification.

[23, 24, 25, 20, 21, 26] utilize conditional codes to indicate different characteristics of the scene including shape and color. Another main direction benefits from existing image editing techniques and operates on discretized instances of continuous INRs such as pixels or voxels. However, such solutions break down the continuous characteristic of INR due to the prerequisite of decoding and discretizing before editing and processing.

In this paper, we conduct the first pilot study on the question: *how to generally modify an INR without explicit decoding?* The major challenge is that one cannot directly interpret what the parameters in an INR stand for, not to mention editing them correctly. Our key motivation is that spatial gradients can be served as a favorable tool to tackle this problem as they can be computed analytically, and possess desirable invariant properties. Theoretically, we prove that any continuous convolution filter can be uniformly approximated by a linear combination of high-order differential operators. Based on the above two rationales, we propose an Implicit Neural Signal Processing Network, dubbed **INSP-Net**, which processes INR utilizing high-order differential operators. The proposed INSP-Net is composed of an inception fusion block connecting computational graphs corresponding to derivatives of INRs. The weights in the branchy part are loaded from the INR being processed, while the weights in the fusion block are parameters of the operator, which can be either hand-crafted or learned by the data-driven algorithm. Even though we are not able to perform surgery on neural network parameters, we can implicitly process them by retrofitting their architecture and reorganizing the spatial gradients.

We further extend our framework to build the first Convolutional Neural Network (CNN) operating directly on INRs, dubbed INSP-ConvNet. Each layer of INSP-ConvNet is constructed by linearly combining the derivative computational graphs of the former layers. Nonlinear activation and normalization are naturally supported as they are element-wise functions. Data augmentation can be also implemented by augmenting the input coordinates of INRs. Under this pipeline (shown in Fig. 1), we demonstrate the expressiveness of our INSP-Net framework in fitting low-level image processing kernels including edge detection, blurring, deblurring, denoising, and image inpainting. We also successfully apply our INSP-ConvNet to high-level tasks on implicit fields such as classification.

Our main contributions can be summarized as follows:

- We propose a novel signal processing framework, dubbed **INSP-Net**, that operates on INRs analytically and continuously by closed-form high-order differential operators[2]. Repeatedly cascading the computational paradigm of INSP-Net, we also build a convolutional network, called INSP-ConvNet, which directly runs on implicit fields for high-level tasks.
- We illustrate the advantage of adopting differential operators by revealing their inherent group invariance. Furthermore, we rigorously prove that the convolution operator in the continuous regime can be uniformly approximated by a linear combination of the gradients.
- Extensive experiments demonstrate the effectiveness of our approach in both low-level processing (e.g. edge detection, blurring, deblurring, denoising, image inpainting, and smoothening) and high-level processing such as image classification.

[2]By saying "closed-form", we mean the computation follows from an analytical mathematical expression.

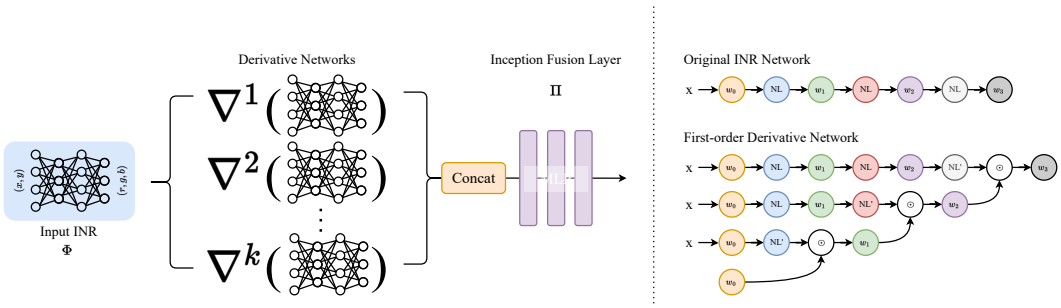

Figure 2: The left image provides an overview of our INSP-Net framework. Each layer combines the high-order derivative computational graphs of the original INR network. The right image illustrates the weight sharing scheme in calculating the derivative sub-networks.

2 Preliminaries: Implicit Neural Representation

Implicit Neural Representation (INR) parameterizes continuous multi-media signals or vector fields with neural networks. Formally, we consider an INR as a continuous function $\Phi : \mathbb{R}^m \to \mathbb{R}$ that maps low-dimension spatial/temporal coordinates to the value space[3]. For example, to represent 2D image signals, the domain of Φ is (x, y) spatial coordinates, and the range of Φ are the pixel intensities. The typical use of INR is to solve a feasibility problem where Φ is sought to satisfy a set of N constraints $\{\mathcal{C}(\Phi, a_j|\Omega_j)\}_{j=1}^N$, where $\mathcal{C}$ is a functional that relates function Φ to some observable quantities a_j evaluating over a measurable domain $\Omega_j \subseteq \mathbb{R}^m$. This problem can be cast into an optimization problem that minimizes deviations from each of the constraints:

$$\Phi^* = \arg\min_{\Phi} \sum_{j=1}^{N} \|\mathcal{C}(\Phi, a_j|\Omega_j)\|_2. \tag{1}$$

For instance, we can let $\mathcal{C} = \Phi(\boldsymbol{x}_j) - a_j$ with $\Omega_j = \{\boldsymbol{x}_j\}$, then our objective boils down to a point-to-point supervision which memorizes a signal into Φ [27]. When functional $\mathcal{C}$ is a combination of differential operators taking values in a point set, i.e., $\mathcal{C}(a(\boldsymbol{x}), \Phi(\boldsymbol{x}), \nabla\Phi(\boldsymbol{x}), \cdots), \forall \boldsymbol{x} \in \Omega_j$, Eq. 1 is objective to solving a bunch of differential equations [28, 7, 29]. Note that in this paper, without particular specification, the gradients are all computed with respect to the input coordinate $\boldsymbol{x}$. $\mathcal{C}$ can also form an integral equation system over some intervals Ω_j [12]. In practice of computer vision, we reconstruct a signal by capturing sparse observations $\mathcal{D} = \{(\Omega_j, a_j)\}_{j=1}^N$ from unknown function Φ, and dynamically sampling a mini-batch from $\mathcal{D}$ to minimize Eq. 1 to obtain a feasible Φ.

A handy parameterization of function Φ is a fully-connected neural network, which enables solving Eq. 1 via gradient descent through a differentiable $\mathcal{C}$. Common INR networks consist of pure Multi-Layer Perceptrons (MLP) with periodic activation functions. Fourier Feature Mapping (FFM) [27] places a sinusoidal transformation before the MLP, while Sinusoidal Representation Network (SIREN) [28] replaces every piece-wise linear activation with a sinusoidal function. Below we give a unified formulation of INR networks:

$$\Phi(\boldsymbol{x}) = \boldsymbol{W}_n(\phi_{n-1} \circ \phi_{n-2} \circ \cdots \circ \phi_1)(\boldsymbol{x}), \quad \phi_i(\boldsymbol{x}) = \sigma_i(\boldsymbol{W}_i\boldsymbol{x} + \boldsymbol{b}_i), \tag{2}$$

where $\boldsymbol{W}_i \in \mathbb{R}^{d_{i-1} \times d_i}$, $\boldsymbol{b}_i \in \mathbb{R}^{d_i}$ are the weight matrix and bias of the i-th layer, respectively, n is the number of layers, and $\sigma_i(\cdot)$ is an element-wise nonlinear activation function. For FFM architecture, $\sigma_i = \sin(\cdot)$ when $i = 1$ denotes the positional encoding layer [12, 30] and otherwise $\sigma_i = \mathrm{ReLU}(\cdot)$. For SIREN, $\sigma_i = \sin(\cdot)$ for every layer $i = 1, \cdots, n-1$.

3 Implicit Representation Processing via Differential Operators

Digital Signal Processing (DSP) techniques have been widely applied in computer vision tasks, such as image restoration [31], signal enhancement [32] and geometric processing [33]. Even modern deep

[3]Without loss of generality, here we simplify Φ to be a scalar field, i.e., the range of Φ is one-dimensional.

learning models are consisting of the most basic signal processing operators. Suppose we already acquire an Implicit Neural Representation (INR) $\Phi : \mathbb{R}^m \to \mathbb{R}$, now we are interested in whether we can run a signal processing program on the implicitly represented signals. One straightforward solution is to rasterize the implicit field with a 2D/3D lattice and run a typical kernel on the pixel/voxel grids. However, this decoding strategy produces a finite resolution and discretizes signals, which is memory inefficient and unfriendly to modeling fine details. In this section, we introduce a computation paradigm that can process an INR analytically with spatial/temporal derivatives. We show that our proposed method serves as a universal operator that can represent any continuous convolutional kernels.

3.1 Computational Paradigm

It has not escaped our notice that spatial/temporal gradients on INRs $\nabla^k\Phi$ can be computed analytically due to the differentiable characteristics of neural networks. Inspired by this, we propose an Implicit Neural Signal Processing (INSP) framework that composes a class of closed-form operators for INRs using functional combinations of high-order derivatives.

We denote our proposed signal processing operator by $\mathcal{A}$ built upon high-order derivatives. Given an acquired INR Φ, we denote the resultant INR processed by operator $\mathcal{A}$ as $\Psi = \mathcal{A}\Phi : \mathbb{R}^m \to \mathbb{R}$. To evaluate point $\boldsymbol{x} \in \mathbb{R}^m$ of processed INR, we propose the following computational paradigm:

$$\Psi(\boldsymbol{x}) := \mathcal{A}\Phi(\boldsymbol{x}) = \Pi\left(\Phi(\boldsymbol{x}), \nabla\Phi(\boldsymbol{x}), \nabla^2\Phi(\boldsymbol{x}), \cdots, \nabla^k\Phi(\boldsymbol{x}), \cdots\right), \tag{3}$$

where $\Pi : \mathbb{R}^M \to \mathbb{R}$ can be arbitrary continuous functions, which can be either handcrafted or learned from data. To learn an operator $\mathcal{A}$ from data, we represent Π by Multi-Layer Perceptrons (MLP) with parameters $\boldsymbol{\theta}$. Here we can slightly abuse the notation of ∇^k to be a flattened vector of high-order derivatives *without multiplicity* since differential operators defined over continuous functions form a commutative ring. The input dimension of Π depends on the highest order of used derivatives. Suppose we compute derivatives up to K-th order, then $M = \sum_{k=0}^{K} \binom{k+m-1}{k} = (K+1)\binom{K+m}{K+1}/m$, where $\binom{k+m-1}{k}$ is the number of distinctive k-th order differential operators [4]. Intuitively, directional derivatives encode (local) neighboring information, which can have similar effects of a convolution. As we will show in Sec. 3.2, Π can construct both shift-invariant and rotation-invariant operators, which introduces favorable inductive bias to images and 3D geometry processing. More importantly, we rigorously prove that Eq. 3 is also a universal approximator of arbitrary convolutional operators.

We note that $\Psi(\boldsymbol{x})$ as a whole can also be regarded as a neural network. Recall the architecture of $\Phi(\boldsymbol{x})$ in Eq. 2, its k-th order derivative is another computational graph parameterized by $\boldsymbol{W}_i$ and $\boldsymbol{b}_i$ that maps $\boldsymbol{x}$ to $\nabla^k\Phi(\boldsymbol{x})$. For example, the first-order gradient will have the following form:

$$\nabla\Phi(\boldsymbol{x}) = \hat{\phi}_{n-1} \circ (\phi_{n-2} \circ \cdots \circ \phi_1)(\boldsymbol{x}) \odot \cdots \odot \hat{\phi}_2 \circ \phi_1(\boldsymbol{x}) \odot \boldsymbol{W}_1, \tag{4}$$

where $\hat{\phi}_i(\boldsymbol{y}) = \boldsymbol{W}_i^\top \sigma'_{i-1}(\boldsymbol{W}_{i-1}\boldsymbol{y} + \boldsymbol{b}_{k-1})$, and $\sigma'_i(\cdot)$ is the first-order derivative of $\sigma_i(\cdot)$. Since $\hat{\phi}_i$ shares the weights with ϕ_i, $\nabla\Phi$ is represented by a closed-form computational network re-using the weights from Φ, which we refer to as the first-order *derivative network*. The higher-order derivatives should induce the derivative network of similar forms. Therefore, the processed INR Ψ will have an Inception-like architecture, namely, a multi-branch structure connecting the original INR network and weight-sharing derivative subnetworks followed by a fusion layer Π. We call the entire model ($\Psi = \mathcal{A}\Phi$ or Eq. 3) an *Implicit Neural Signal Processing Network* or an *INSP-Net*. Note that the only parameters of INSP-Net $\boldsymbol{\theta}$ are located at the last fusion layer, and can be trained in an end-to-end manner.

We illustrate an INSP-Net in Fig. 2 where the color indicates the weight-sharing scheme. A similar weight-sharing scheme is also adopted in AutoInt [34]. In practice, we employ auto-differentiation in PyTorch [35] to automatically create such derivatives networks and reassemble them parallelly to constitute the architecture of an INSP-Net. When inputting an INR, we load the weights of the INR to our model following the weight-sharing scheme, and then we obtain an INSP-Net, which implicitly and continuously represents the processed INR $\Psi(\boldsymbol{x})$. To effectively express high-order derivatives, we choose SIREN as the base model [28].

3.2 Theoretical Analysis

In this section, we provide a theoretical justification for the design of our INSP-Net. We will focus on discussing the latent invariance property and the expressive power of INSP-Net.

[4] This is equal to the number of monic monomials over $\mathbb{R}^m$ with degree k.

Translation and Rotation Invariance. Group invariance has been shown to be a favorable inductive bias for image [36], video [37], and geometry processing [38]. It has also been well-known that group invariance is an intrinsic property of Partial Differential Equations (PDEs) [39, 40]. Since our INSP-Net is built using differential operators, we are motivated to reveal its hidden invariance property to demonstrate its advantage in processing visual signals.

In this section, we only consider two transformation groups: translation group $\mathbb{T}(m)$ and the special orthogonal group $\mathbb{SO}(m)$ (a.k.a. rotation group). Elements $T_{\boldsymbol{v}} \in \mathbb{T}(m)$ in translation group shift the function Φ by some offset $\boldsymbol{v} \in \mathbb{R}^m$. The shifted function can be denoted as $\Phi \circ T_{\boldsymbol{v}}(\boldsymbol{x}) = \Phi(\boldsymbol{x} + \boldsymbol{v})$. Similarly, elements in rotation group perform a coordinate transformation on function Φ by a rotation matrix $\boldsymbol{R} \in \mathbb{SO}(m)$. The transformed function can be written as $\Phi \circ \boldsymbol{R}(\boldsymbol{x}) = \Phi(\boldsymbol{R}\boldsymbol{x})$. Group invariance means deforming the input space of a function first and then processing it via an operator is equivalent to directly applying the transformation to the processed function. For a more rigorous argument, $\mathcal{A}$ is said to be translation-invariant if $\forall T_{\boldsymbol{v}} \in \mathbb{T}(m)$, $\Psi(\boldsymbol{x} + \boldsymbol{v}) = \mathcal{A}[\Phi \circ T_{\boldsymbol{v}}](\boldsymbol{x})$. Likewise, $\mathcal{A}$ is rotation-invariant if $\forall \boldsymbol{R} \in \mathbb{SO}(m)$ we have $\Psi(\boldsymbol{R}\boldsymbol{x}) = \mathcal{A}[\Phi \circ \boldsymbol{R}](\boldsymbol{x})$. Below we provide Theorem 1 to characterize the invariance property for our model.

Theorem 1. *Given function $\Pi : \mathbb{R}^M \to \mathbb{R}$, the composed operator $\mathcal{A}$ (Eq. 3) can satisfy:*

1. *shift invariance for every Π.*
2. *rotation invariance if Π has the form: $\Pi(\boldsymbol{y}) = f(\|\boldsymbol{y}\|_2)$ for some $f : \mathbb{R} \to \mathbb{R}$.*

We prove Theorem 1 in Appendix A. Our Theorem 1 implies that operator $\mathcal{A}$ is inherently shift-invariant. This is due to the shift-invariant intrinsics of differential operators as we show in the proof. Rotation invariance is not guaranteed in general. However, if one carefully designs Π, it can also be achieved via our framework. Moreover, we also suggest a feasible solution to constructing a rotation-invariant operator $\mathcal{A}$ in Theorem 1. In our construction, Π first isotropically pools over the squares of all directional derivatives, and then maps the summarized information through another scalar function f. We refer interested readers to [39] for more group invariance in differential forms.

Universal Approximation. Convolution, formally known as the linear shift-invariant operator, has served as one of the most prevalent signal processing tools in the vision domain. Given two (real-valued) signals f and g, we denote their convolution as $g \star f = f \star g$. In this section, we examine the expressiveness of our INSP-Net (Eq. 3) by showing it can represent any convolutional filter. To draw this conclusion, we first present an informal version of our main results as follows:

Theorem 2. *(Informal statement) For every real-valued function $g : \mathbb{R}^m \to \mathbb{R}$, there exists a polynomial $p(x_1, \cdots, x_m)$ with real coefficients such that $p(\nabla) f$ can uniformly approximate $g \star f$ by arbitrary precision for all real-valued signals f.*

The formal statement and proof can be found in Appendix B. Theorem 2 involves the notion of polynomials in partial differential operators (see more details in Appendix B). $p(\nabla)f$ in turn can be written as a *linear* combination of high-order derivatives of f (a special case of Eq. 3 when Π is linear). The key step to prove Theorem 2 is applying Stone-Weierstrass approximation theorem on the Fourier domain. However, we note that functions obtained by the Fourier transform are generally complex functions. The prominence of our proof is that we can constrain the range of the polynomial coefficients into the real domain, which makes it implementable via a common deep learning infrastructure. The implication of Theorem 2 is that the mapping between convolution and derivative is as simple as a linear transformation. Recent works [41, 42, 43] show the converse argument that derivatives can be approximated via a linear combination of *discrete* convolution. Theorem 2 establishes the equivalence between differential operator and convolution in the *continuous* regime. In our proof, k-th order derivatives correspond to k-th order monomial in the spectral domain. Fitting convolution using derivatives amounts to approximating spectrum via polynomials. This implies higher degree of polynomial induces closer approximation. Since $p(\nabla)$ is not difficult to be approximated by a neural network $\Pi_{\boldsymbol{\theta}}$, we can easily derive the next result Corollary 3.

Corollary 3. *For every real-valued function g, there exists a neural network $\Pi_{\boldsymbol{\theta}}$ such that $\Psi = \mathcal{A}\Phi$ (Eq. 3) can uniformly approximate $g \star \Phi$ by arbitrary precision for every real-valued signals Φ.*

As we discussed in Theorem 1, $\mathcal{A}$ are constantly shift-invariant. This means when approximating a convolutional kernel, the trajectory of $\mathcal{A}$ is restricted into the shift-invariant space. Moreover, we emphasize that INSP-Net is far more expressive than convolutional kernels since $\Pi_{\boldsymbol{\theta}}$ can also fit any nonlinear continuous functions due to the universal approximation theorem [44, 45, 46].

3.3 Building CNNs for Implicit Neural Representations

Convolutional Neural Networks (CNN) are capable of extracting informative semantics by only piling up basic signal processing operators. This motivates us to build CNNs based on INSP-Net that can directly run on INRs for high-level downstream tasks. In fact, to simulate *exact* convolution, our Theorem 2 suggests simplify $\Pi_{\boldsymbol{\theta}}$ to a linear mapping. Then our former computational paradigm Eq. 3 is changed to:

$$\Psi(\boldsymbol{x}) := p(\nabla)\Phi(\boldsymbol{x}) = \theta_0\Phi(\boldsymbol{x}) + \boldsymbol{\theta}_1^\top \nabla\Phi(\boldsymbol{x}) + \boldsymbol{\theta}_2^\top \nabla^2\Phi(\boldsymbol{x}) + \cdots + \boldsymbol{\theta}_k^\top \nabla^k\Phi(\boldsymbol{x}) + \cdots, \tag{5}$$

where $\boldsymbol{\theta}_k \in \mathbb{R}^{\binom{k+m-1}{k}}$ are parameters of the operator $p(\nabla)$. We name this special case of Eq. 3 as INSP-Conv. One plausible implementation of INSP-Conv is to employ a one-layer MLP to represent $\Pi_{\boldsymbol{\theta}}$. When $\mathcal{A} = p(\nabla)$, INSP-Conv preserves both linearity and shift-invariance when evolving during the training. We propose to repeatedly apply INSP-Conv with non-linearity to INRs that mimics a CNN-like architecture. We name this class of CNNs composed by multi-layer INSP-Conv (Eq. 5) as INSP-ConvNet. Previous works [47, 48] extracting semantic features from INR either lack local information by point-wisely mapping INR's intermediate representation to a semantic space or explicitly rasterize INR into regular grids. To the best of our knowledge, it is the first time that one can run a CNN directly on an implicit representation thanks to closed-formness of INSP-Net. The overall architecture of INSP-ConvNet can be formulated as:

$$\mathrm{ConvNet}[\Phi](\boldsymbol{x}) = \mathcal{A}^{(L)} \cdot \sigma \circ \mathcal{A}^{(L-1)} \cdot \sigma \circ \cdots \circ \mathcal{A}^{(2)} \cdot \sigma \circ \mathcal{A}^{(1)} \cdot \Phi(\boldsymbol{x}), \tag{6}$$

where σ is an element-wise non-linear activation, L is the number of INSP-Net layers, and Φ is an input INR. We use the symbol $\cdot$ to denote operator functioning, and $\circ$ to denote function composition. Due to page limit, we defer detailed introduction to INSP-ConvNet to Appendix C.

4 Related Work

4.1 Implicit Neural Representation

Implicit Neural Representation (INR) represents signals by continuous functions parameterized by multi-layer perceptrons (MLPs) [28, 27], which is different from traditional discrete representations (e.g., pixel, mesh). Compared with other representations, the continuous implicit representations are capable of representing signals at infinite resolution and have become prevailing to be applied upon image fitting [28], image compression [1, 49] and video compressing [3]. In addition, INR has been applied to more efficient and effective shape representation [4, 5, 6, 7, 8, 9, 10, 11], texture mapping [50, 51], inverse problems [12, 2, 13, 14] and generative models [15, 16, 17, 18, 19, 20, 21, 22]. There are also efforts speeding up the fitting of INRs [52] and improving the representation efficiency [53]. Nowadays, editing and manipulating multi-media objects gains increasing interest and demand [54]. Thus, signal processing on implicit neural representation is essentially an important task worth investigating.

4.2 Editable Implicit Fields

Editing implicit fields has recently attracted much research interest. Several methods have been proposed to allow editing the reconstructed 3D scenes by rearranging the objects or manipulating the shape and appearance. One line of work alters the structure and color of objects by conditioning latent codes for different characteristics of the scene [25, 20, 21, 26]. Another direction involves discretizing the continuous fields. By converting the implicit fields into pixels or voxels, traditional image and voxel editing techniques [55, 56] can be applied effortlessly. These approaches, however, are not capable of directly performing signal processing on continuous INRs. Functa [54] can use a latent code to control implicit funtions. NID [57] represents neural fields as a linear combination of implicit functional basis, which enables editing by change of sparse coefficients. However, such editing scheme suffers from limited flexibility. Recently, NFGP [58] proposes to use neural fields for geometry processing by exploring various geometric regularization. INS [59] distills stylized features into INRs via the neural style transfer framework [60]. Our INSP-Net that makes smart use of closed-form differential operators does not require neither additional per-scene fine-tuning nor discretization to grids.

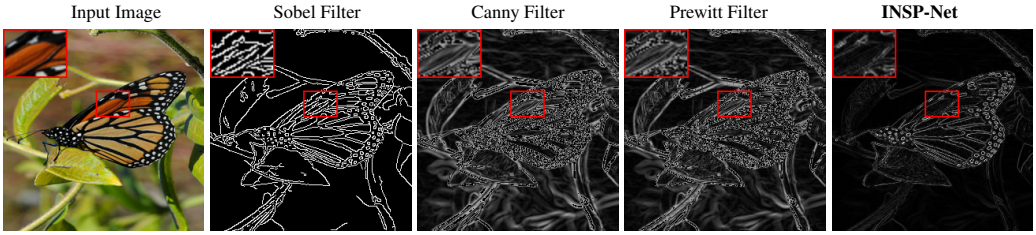

Figure 3: Edge detection. We fit the natural images with SIREN and use our INSP-Net to process implicitly into a new INR that can be decoded into edge maps.

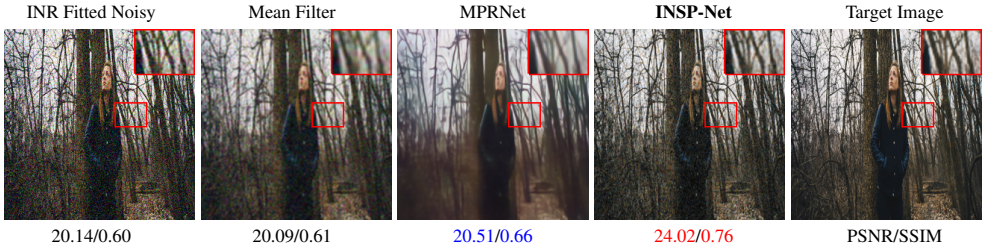

Figure 4: Image denoising. We fit the noisy images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into natural clear images.

4.3 PDE based Image Processing

Partial differential equations (PDEs) have been successfully applied to many tasks in image processing and computer vision, such as image enhancement [61, 62, 63], segmentation [64, 40], image registration [65], saliency detection [66] and optical flow computation [67]. Early traditional PDEs are written directly based on mathematical and physical understanding of the PDEs (e.g., anisotropic diffusion [61], shock filter [62] and curve evolution based equations [68, 69, 70]). Variational design methods [63, 71, 70] start from an energy function describing the desired properties of output images and compute the Euler-Lagrange equation to derive the evolution equations. Learning-based attempts [40, 66] build PDEs from image pairs based on the assumption (without proof) that PDEs could be written as linear combinations of fundamental differential invariants. Although it might be feasible to let INRs solve this bunch of signal processing PDEs, one needs to per-case re-fit an INR with an additional temporal axis, which is presumably sampling inefficient. The multi-layer structure appearing in INSP-Net can be viewed as an unfolding network [72, 73] of the Euler method to solve time-variant PDEs [74]. We elaborate this connection in Appendix D.

5 Experiments

In this section, we evaluate the proposed **INSP** framework on several challenging tasks, using different combinations of Π. First, we build low-level image processing filters using either hand-crafted or learnable Π. Then, we construct convolutional neural networks with our INSP-ConvNet framework and validate its performance on image classification. More results and implementation details are provided in the Appendix E F.

5.1 Low-Level Vision for Implicit Neural Images

For low-level image processing, we operate on natural images from Set5 dataset [75], Set14 dataset [76], and DIV-2k dataset [77]. Originally designed for super-resolution, the images are diverse in style and content. Note that the unprocessed images presented in figures are the images decoded from unprocessed INRs.

Since our method operates directly on INRs, we firstly fit the images with INRs and then feed the INRs into our framework. The final output is another INR which can be decoded into desired images. The training set of our method consists of 90 examples of INRs, where each INR is built on SIREN [28] architectures.

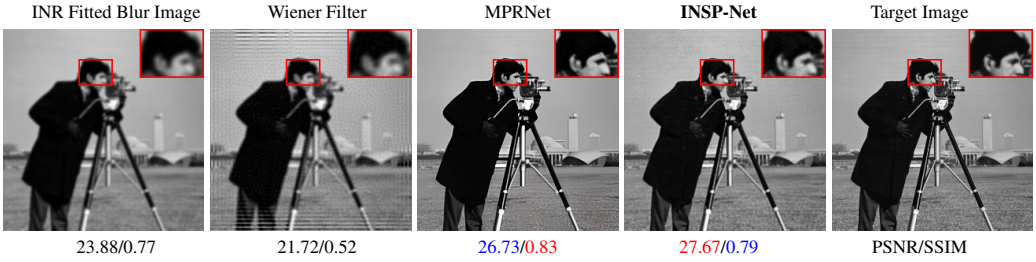

Figure 5: Image deblurring. We fit the blurred images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into clear natural images.

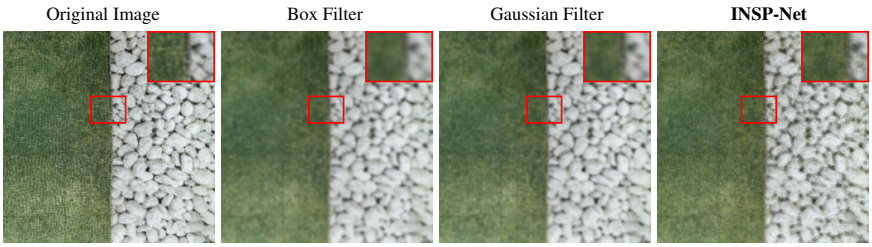

Figure 6: Image blurring. We fit the natural images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into blurred images.

Edge Detection Since the edges correspond to gradients in the images, using gradients of INRs to obtain edges is straightforward. θ_1 is set to 1 while other coefficients are set to 0. We provide visual comparisons against Sobel filter [78], Canny detector [79] and Prewitt operator [80] in Fig. 3.

Image Denoising For classical image denoising filters, we compare against the median filter and mean filter. We use MPRNet [81] as a learning-based baseline. The input noisy images are synthesized using additive gaussian noise. Visual results are provided in Fig. 4.

Image Blurring Image blurring is a low-pass filtering operation. We provide a visual comparison against classical filters including 3×3 box filter and 3×3 gaussian filter. The target images used for training our INSP-Net are the results of the Gaussian filter. Visual results are provided in Fig. 6.

Image Deblurring We compare the proposed method with both traditional algorithms (e.g., wiener filter [82]) and learning-based algorithms(e.g., MPRNet [81]). We synthesize blurry images using Gaussian filters. As shown in Fig. 5, Wiener Filter produce severe artifacts and MPRNet successfully reconstructs clear textures. INSP-Net is capable of generating competitive results against MPRNet and outperforms the Wiener Filter.

Image Inpainting We conduct two kinds of experiments in image inpainting, to inpaint 30% random masked pixels or to remove the texts ("INSP-Net"). Comparison methods include mean filter, median filter, and LaMa [83]. LaMa is a learning-based method using Fourier convolution for inpainting. As shown in Fig. 7, mean filter and median filter partially restore the masked pixels, but severely hurt the visual quality of the rest parts. Also, they can not handle the text region. LaMa successfully removes the text and inpaint the masked pixels. Our proposed method largely outperforms the filter-based algorithms and performs as well as the LaMa.

5.2 Geometry Processing on Signed Distance Function

We demonstrate that the proposed **INSP** framework is not only capable of processing images, but also capable of processing geometry. Signed Distance Function (SDF) [25] is adopted to represent geometries in this section. We first fit an SDF from a point cloud following the training loss proposed in [28, 7]. Then we train our INSP-Net to simulate a Gaussian-like filter similar to image blurring. Afterwards, we apply the trained INSP-Net to process the specified INR. When visualization, we use marching cube algorithm to extract meshes from SDF. We choose Thai Statue, Armadillo, and Dragon from Stanford 3D Scanning Repository [84, 85, 86, 87] to demonstrate our results. Fig. 8

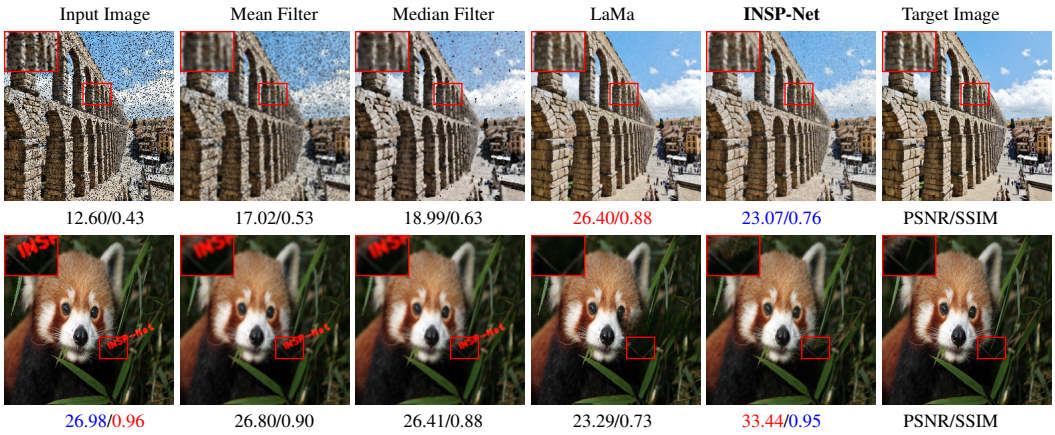

Figure 7: Image inpainting. We fit the input images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into natural images. Note that LaMa requires explicit masks to select the regions for inpainting and the masks are roughly provided.

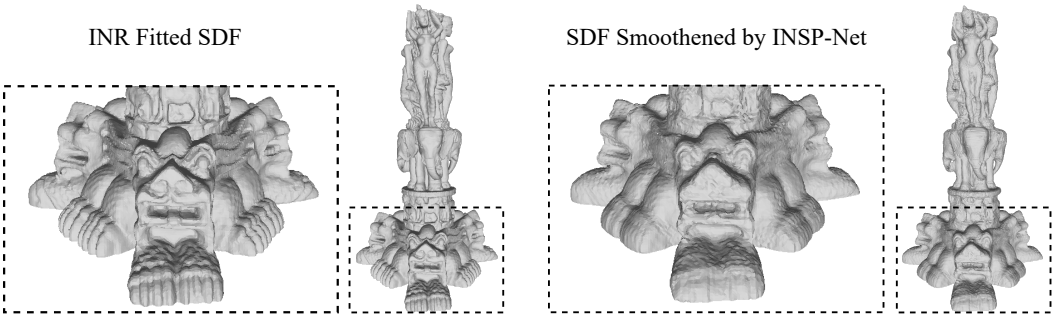

Figure 8: Left: unprocessed geometry decoded from an unprocessed INR. Right: smoothened geometry decoded from the output INR of our INSP-Net. Best view in a zoomable electronic copy.

exhibits our results on Thai Statue. Our method is able to smoothen the surface of the geometry and erase high-frequency details acting as if a low-pass filter. We defer more results to Fig. 14.

5.3 Classification on Implicit Neural Representations

We demonstrate that the proposed **INSP** framework is not only capable to express low-level image processing filters, but also supports high-level tasks such as image classification. To achieve this goal, we construct a 2-layer INSP-ConvNet. The INSP-ConvNet consists of 2 INSP-Net layers. Each of them decomposes the INR via the differential operator and combines them with learnable Π. We build another 2-layer depthwise ConvNets running on pixels as the baseline for a fair comparison, since it has comparable expressiveness to our INSP-ConvNet in theory. We also build a PCA + SVM method and an MLP classifier that directly classify INRs according to (vectorized) weight matrices.

We evaluate the proposed INSP-ConvNet on MNIST (28×28 resolution) and CIFAR-10 (32×32 resolution) datasets, respectively. For each dataset, we will firstly fit each image into an implicit representation using SIREN [28]. Both experiments take 1000 epochs to optimize with AdamW optimizer [88] and a learning rate of 10^{-4}. Results are shown in Tab. 1.

Accuracy	Depthwise CNN	PCA + SVM	MLP classifier	INSP-ConvNet
MNIST	87.6%	11.3%	9.8%	88.1%
CIFAR-10	59.5%	9.4%	10.1%	62.5%

Table 1: Quantitative Results of Image Classification. All methods except "Depthwise CNN" operate on the parameter of INR directly, while "Depthwise CNN" operates on images decoded from INR.

We categorize Depthwise CNN as explicit method, which requires to extract the image grids from INRs before classification. PCA + SVM and MLP classifier working on the network parameter space can be regarded as two straightforward implicit baselines. We find that traditional classifiers can hardly classify INR on weight space due to its high-dimensional unstructured data distribution. Our method, however, can effectively leverage the information implicitly encoded in INRs by exploiting their derivatives. As a consequence, INSP-ConvNet can achieve classification accuracy on-par with CNN-based explicit method, which validates the learning representation power of INSP-ConvNet.

6 Conclusion

Contribution. We present INSP-Net framework, an implicit neural signal processing network that is capable of directly modifying an INR without explicit decoding. By incorporating differential operators on INR, we can instantiate the INR signal operator as a composition of computational graphs approximating any continuous convolution filter. Furthermore, we make the first effort to build a convolutional neural network that implicitly runs on INRs. While all other methods run on discrete grids, our experiment demonstrates our INSP-Net can achieve competitive results with entirely implicit operations.

Limitations. (Theory) Our theory only guarantees the expressiveness of convolution by allowing infinite sequence approximation. Construction of more expressive operators and more effective parameterization for convolution remain widely open questions. (Practice) INSP-Net requires the computation of high-order derivatives which is neither memory efficient nor numerically stable. This hinders the scalability of our INSP-ConvNet that requires recursive computation of derivatives. Addressing how to reconstruct INRs in a scalable manner is beyond the scope of this paper. All INRs used in our experiments are fitted by per-scene optimization.

Acknowledgement

Z. Wang is in part supported by an NSF Scale-MoDL grant (award number: 2133861).

References

[1] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021.

[2] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34, 2021.

[3] Yunfan Zhang, Ties van Rozendaal, Johann Brehmer, Markus Nagel, and Taco Cohen. Implicit neural video compression. *arXiv preprint arXiv:2112.11312*, 2021.

[4] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7154–7164, 2019.

[5] Matan Atzmon, Niv Haim, Lior Yariv, Ofer Israelov, Haggai Maron, and Yaron Lipman. Controlling neural level sets. *Advances in Neural Information Processing Systems*, 32, 2019.

[6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019.

[7] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.

[8] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.

[9] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5379–5389, 2019.

[10] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.

[11] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision*, pages 523–540. Springer, 2020.

[12] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020.

[13] Vincent Sitzmann, Semon Rezchikov, William T Freeman, Joshua B Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *arXiv preprint arXiv:2106.02634*, 2021.

[14] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul Srinivasan, and Jonathan T Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. *arXiv preprint arXiv:2111.13679*, 2021.

[15] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021.

[16] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W Taylor, and Joshua M Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14304–14313, 2021.

[17] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*, 2021.

[18] Zekun Hao, Arun Mallya, Serge Belongie, and Ming-Yu Liu. Gancraft: Unsupervised 3d neural rendering of minecraft worlds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14072–14082, 2021.

[19] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6351–6361, 2021.

[20] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021.

[21] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020.

[22] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. Cips-3d: A 3d-aware generator of gans based on conditionally-independent pixel synthesis. *arXiv preprint arXiv:2110.09788*, 2021.

[23] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5773–5783, 2021.

[24] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. *arXiv preprint arXiv:2112.05139*, 2021.

[25] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.

[26] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5932–5941, 2019.

[27] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *arXiv preprint arXiv:2006.10739*, 2020.

[28] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020.

[29] Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

[30] Ellen D Zhong, Tristan Bepler, Bonnie Berger, and Joseph H Davis. Cryodrgn: reconstruction of heterogeneous cryo-em structures using neural networks. *Nature Methods*, 18(2):176–185, 2021.

[31] Dejia Xu, Yihao Chu, and Qingyan Sun. Moiré pattern removal via attentive fractal network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 472–473, 2020.

[32] Yaqian Xu, Wenqing Zheng, Jingchen Qi, and Qi Li. Blind image blur assessment based on markov-constrained fcm and blur entropy. In *2019 IEEE international conference on image processing (icip)*, pages 4519–4523. IEEE, 2019.

[33] Peng-Shuai Wang, Xiao-Ming Fu, Yang Liu, Xin Tong, Shi-Lin Liu, and Baining Guo. Rolling guidance normal filter for geometric processing. *ACM Transactions on Graphics (TOG)*, 34(6):1–9, 2015.

[34] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14556–14565, 2021.

[35] 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, 2019.

[36] Jan J Koenderink and Andrea J van Doorn. Image processing done right. In *European Conference on Computer Vision*, pages 158–172. Springer, 2002.

[37] Hae-Kwang Kim and Jong-Deuk Kim. Region-based shape descriptor invariant to rotation, scale and translation. *Signal Processing: Image Communication*, 16(1-2):87–93, 2000.

[38] Niloy J Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. Symmetry in 3d geometry: Extraction and applications. In *Computer Graphics Forum*, volume 32, pages 1–23. Wiley Online Library, 2013.

[39] Peter J Olver. *Applications of Lie groups to differential equations*, volume 107. Springer Science & Business Media, 2000.

[40] Risheng Liu, Zhouchen Lin, Wei Zhang, and Zhixun Su. Learning pdes for image restoration via optimal control. In *European Conference on Computer Vision*, pages 115–128. Springer, 2010.

[41] Bin Dong, Qingtang Jiang, and Zuowei Shen. Image restoration: Wavelet frame shrinkage, nonlinear evolution pdes, and beyond. *Multiscale Modeling & Simulation*, 15(1):606–660, 2017.

[42] Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. Pde-net: Learning pdes from data. In *International Conference on Machine Learning*, pages 3208–3216. PMLR, 2018.

[43] Zichao Long, Yiping Lu, and Bin Dong. Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399:108925, 2019.

[44] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 1989.

[45] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 1991.

[46] Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.

[47] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021.

[48] Suhani Vora, Noha Radwan, Klaus Greff, Henning Meyer, Kyle Genova, Mehdi SM Sajjadi, Etienne Pot, Andrea Tagliasacchi, and Daniel Duckworth. Nesf: Neural semantic fields for generalizable semantic segmentation of 3d scenes. *arXiv preprint arXiv:2111.13260*, 2021.

[49] Brandon Yushan Feng and Amitabh Varshney. Signet: Efficient neural representation for light fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14224–14233, 2021.

[50] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4531–4540, 2019.

[51] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019.

[52] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P Srinivasan, Jonathan T Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2846–2855, 2021.

[53] Jaeho Lee, Jihoon Tack, Namhoon Lee, and Jinwoo Shin. Meta-learning sparse implicit neural representations. *Advances in Neural Information Processing Systems*, 34:11769–11780, 2021.

[54] Emilien Dupont, Hyunjik Kim, SM Eslami, Danilo Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you should treat it like one. *arXiv preprint arXiv:2201.12204*, 2022.

[55] James W. Hennessey, Wilmot Li, Bryan Russell, Eli Shechtman, and Niloy J. Mitra. Transferring image-based edits for multi-channel compositing. *ACM Transactions on Graphics*, 36(6), 2017.

[56] Jerry Liu, Fisher Yu, and Thomas Funkhouser. Interactive 3d modeling with a generative adversarial network. In *2017 International Conference on 3D Vision (3DV)*, pages 126–134. IEEE, 2017.

[57] Peihao Wang, Zhiwen Fan, Tianlong Chen, and Zhangyang Wang. Neural implicit dictionary learning via mixture-of-expert training. In *International Conference on Machine Learning*, pages 22613–22624. PMLR, 2022.

[58] Guandao Yang, Serge Belongie, Bharath Hariharan, and Vladlen Koltun. Geometry processing with neural fields. *Advances in Neural Information Processing Systems*, 34, 2021.

[59] Zhiwen Fan, Yifan Jiang, Peihao Wang, Xinyu Gong, Dejia Xu, and Zhangyang Wang. Unified implicit neural stylization. *arXiv preprint arXiv:2204.01943*, 2022.

[60] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.

[61] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on pattern analysis and machine intelligence*, 12(7):629–639, 1990.

[62] Stanley Osher and Leonid I Rudin. Feature-oriented image enhancement using shock filters. *SIAM Journal on numerical analysis*, 27(4):919–940, 1990.

[63] Xue-Cheng Tai, Stanley Osher, and Randi Holm. Image inpainting using a tv-stokes equation. In *Image Processing based on partial differential equations*, pages 3–22. Springer, 2007.

[64] Zhouchen Lin, Wei Zhang, and Xiaoou Tang. Designing partial differential equations for image processing by combining differental invariants, 2009.

[65] Lars Hömke, Claudia Frohn-Schauf, Stefan Henn, and Kristian Witsch. Total variation based image registration. In *Image processing based on partial differential equations*, pages 343–361. Springer, 2007.

[66] Risheng Liu, Junjie Cao, Zhouchen Lin, and Shiguang Shan. Adaptive partial differential equation learning for visual saliency detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3866–3873, 2014.

[67] Adam Rabcewicz. Clg method for optical flow estimation based on gradient constancy assumption. In *Image Processing Based on Partial Differential Equations*, pages 57–66. Springer, 2007.

[68] Guillermo Sapiro. *Geometric partial differential equations and image analysis*. Cambridge university press, 2006.

[69] Frédéric Cao. *Geometric curve evolution and image processing*. Springer Science & Business Media, 2003.

[70] Bart M Haar Romeny. *Geometry-driven diffusion in computer vision*, volume 1. Springer Science & Business Media, 2013.

[71] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.

[72] Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th international conference on international conference on machine learning*, pages 399–406, 2010.

[73] Jialin Liu and Xiaohan Chen. Alista: Analytic weights are as good as learned weights in lista. In *International Conference on Learning Representations (ICLR)*, 2019.

[74] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

[75] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie line Alberi Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *Proceedings of the British Machine Vision Conference*, pages 135.1–135.10. BMVA Press, 2012.

[76] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730. Springer, 2010.

[77] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

[78] I Sobel. An isotropic $3\times$ 3 image gradient operator, presentation at stanford ai project (1968), 2014.

[79] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 1986.

[80] Judith MS Prewitt et al. Object enhancement and extraction. *Picture processing and Psychopictorics*, 10(1):15–19, 1970.

[81] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021.

[82] Saeed V Vaseghi. *Advanced digital signal processing and noise reduction*. John Wiley & Sons, 2008.

[83] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint arXiv:2109.07161*, 2021.

[84] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.

[85] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. *ACM Transactions on Graphics (TOG)*, 22(3):749–758, 2003.

[86] Venkat Krishnamurthy and Marc Levoy. Fitting smooth surfaces to dense polygon meshes. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 313–324, 1996.

[87] Greg Turk and Marc Levoy. Zippered polygon meshes from range images. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 311–318, 1994.

[88] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[89] Kathrin Schacke. On the kronecker product. *Master's thesis, University of Waterloo*, 2004.

[90] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.

[91] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxim: Multi-axis mlp for image processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5769–5780, 2022.

A Proof of Theorem 1

To begin with, we give the formal definitions of translation and rotation group, along with the notion of shift invariance and rotation invariance.

Definition A.1. *(Translation Group & Shift Invariance) Translation group $\mathbb{T}(m)$ is a transformation group isomorphic to m-dimension Euclidean space, where each group element $T_{\boldsymbol{v}}$ transforms a vector $\boldsymbol{x} \in \mathbb{R}^m$ by $T_{\boldsymbol{v}}(\boldsymbol{x}) = \boldsymbol{x} + \boldsymbol{v}$. An operator $\mathcal{A}$ is said to be shift-invariant if $\mathcal{A}(\Phi \circ T_{\boldsymbol{v}})(\boldsymbol{x}) = \mathcal{A}\Phi(T_{\boldsymbol{v}}(\boldsymbol{x})) = \mathcal{A}\Phi(\boldsymbol{x} + \boldsymbol{v})$.*

Definition A.2. *(Rotation Group & Rotation Invariance) Rotation group $\mathbb{SO}(m)$ is a transformation group also known as the special orthogonal group, where each group element $\boldsymbol{R} \in \mathbb{R}^{m\times m}$ satisfying $\boldsymbol{R}^\top \boldsymbol{R} = \boldsymbol{I}$ transforms a vector $\boldsymbol{x} \in \mathbb{R}^m$ by $\boldsymbol{R}(\boldsymbol{x})$. An operator $\mathcal{A}$ is said to be rotation-invariant if $\mathcal{A}(\Phi \circ \boldsymbol{R})(\boldsymbol{x}) = \mathcal{A}\Phi(\boldsymbol{R}\boldsymbol{x})$.*

Proof. (Shift Invariance) To show the shift invariance of our model Eq. 3, it is equivalent to show any differential operators are shift-invariant. For the first-order derivatives (gradients), we consider arbitrary shift operator $T_{\boldsymbol{v}} \in \mathbb{T}$, by chain rule we will have:

$$\nabla[\Phi \circ T_{\boldsymbol{v}}](\boldsymbol{x}) = \left[\frac{d(\boldsymbol{x} + \boldsymbol{v})}{d\boldsymbol{x}}\right]^\top \nabla\Phi(\boldsymbol{x} + \boldsymbol{v}) = \nabla\Phi(\boldsymbol{x} + \boldsymbol{v}), \tag{7}$$

where the Jacobian matrix of $T_{\boldsymbol{v}}(\boldsymbol{x})$ is an identity matrix. Eq. 7 implies that gradient operator is shift-invariant. By induction, any high-order differential operators must also be shift-invariant:

$$\nabla^k[\Phi \circ T_{\boldsymbol{v}}](\boldsymbol{x}) = \nabla^k\Phi(\boldsymbol{x} + \boldsymbol{v}), \tag{8}$$

Therefore, we can conclude $\Pi(\Phi, \nabla\Phi, \nabla^2\Phi, \cdots)$ is shift-invariant for any Π combining derivatives in any form.

(Rotation Invariance) By Lemma A.1, given arbitrary function $\Phi : \mathbb{R}^m \to \mathbb{R}$, and for every rotation matrix $\boldsymbol{R} \in \mathbb{SO}(m)$, we can compute the k-th derivatives as:

$$\operatorname{vec}\left(\nabla^k[\Phi \circ \boldsymbol{R}](\boldsymbol{x})\right) = \boldsymbol{R}^{\top\otimes k} \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right). \tag{9}$$

Then adopting properties of Kronecker product [89], the norm of $\nabla^k[\Phi \circ \boldsymbol{R}](\boldsymbol{x})$ can be written as:

$$\|\nabla^k[\Phi \circ \boldsymbol{R}](\boldsymbol{x})\|_F^2 = \operatorname{Tr}\left[\operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right)^\top \boldsymbol{R}^{\otimes k}\boldsymbol{R}^{\top\otimes k} \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right)\right] \tag{10}$$

$$= \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right)^\top \left(\boldsymbol{R}^{\otimes k-1} \otimes \boldsymbol{R}\right)\left(\boldsymbol{R}^{\top\otimes k-1} \otimes \boldsymbol{R}^\top\right) \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right) \tag{11}$$

$$= \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right)^\top \left(\left(\boldsymbol{R}^{\otimes k-1}\boldsymbol{R}^{\top\otimes k-1}\right) \otimes \boldsymbol{I}\right) \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right) \tag{12}$$

$$= \cdots = \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right)^\top \boldsymbol{I}^{\otimes k} \operatorname{vec}\left(\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\right) = \|\nabla\Phi(\boldsymbol{R}\boldsymbol{x})\|_F^2 \tag{13}$$

where Eq. 11 is due to the fact $(\boldsymbol{A} \otimes \boldsymbol{B})^\top = \boldsymbol{A}^\top \otimes \boldsymbol{B}^\top$, Eq. 12 is because of $(\boldsymbol{A} \otimes \boldsymbol{B})(\boldsymbol{C} \otimes \boldsymbol{D}) = \boldsymbol{A}\boldsymbol{C} \otimes \boldsymbol{B}\boldsymbol{D}$, and Eq. 13 is yielded by applying the orthogonality of $\boldsymbol{R}$ and repeating step Eq. 12 to Eq. 13. Therefore, for every integer $k > 0$, operator $\|\nabla^k\Phi(\boldsymbol{x})\|_2^2$ is rotation-invariant. Hence, $\Pi = f\left(\left\|\left[\Phi(\boldsymbol{x}) \quad \nabla\Phi(\boldsymbol{x}) \quad \nabla^2\Phi(\boldsymbol{x}) \quad \cdots\right]\right\|_F\right) = f\left(\sqrt{\sum_{k=0}\|\nabla^k\Phi(\boldsymbol{x})\|_2^2}\right)$ is also rotation-invariant. □

Below we supplement the Lemma A.1 used to prove Theorem 1.

Lemma A.1. *Suppose given function $f : \mathbb{R}^m \to \mathbb{R}$ and arbitrary linear transformation $\boldsymbol{A} \in \mathbb{R}^{m\times m}$, then $\operatorname{vec}\left(\nabla^k[f \circ \boldsymbol{A}](\boldsymbol{x})\right) = \boldsymbol{A}^{\top\otimes k} \operatorname{vec}\left(\nabla^k f(\boldsymbol{A}\boldsymbol{x})\right)$ for $\forall k \geq 0$.*

Proof. $\operatorname{vec}\left(\nabla^k[f \circ \boldsymbol{A}](\boldsymbol{x})\right) = \boldsymbol{A}^{\top\otimes k} \operatorname{vec}\left(\nabla^k f(\boldsymbol{A}\boldsymbol{x})\right)$ trivially holds for $k = 0, 1$. Then we prove Lemma A.1 by induction. Suppose the $(j-1)$-th case satisfies the equality: $\operatorname{vec}\left(\nabla^{j-1}[f \circ \boldsymbol{A}](\boldsymbol{x})\right) = \boldsymbol{A}^{\top\otimes j-1} \operatorname{vec}\left(\nabla^{j-1} f(\boldsymbol{A}\boldsymbol{x})\right)$, then consider the j-th case:

$$\nabla^j[f \circ \boldsymbol{A}](\boldsymbol{x}) = \nabla \operatorname{vec}\left(\nabla^{j-1}[f \circ \boldsymbol{A}](\boldsymbol{x})\right) = \nabla\boldsymbol{A}^{\top\otimes j-1} \operatorname{vec}\left(\nabla^{j-1} f(\boldsymbol{A}\boldsymbol{x})\right) = \boldsymbol{A}^\top\nabla^j f(\boldsymbol{A}\boldsymbol{x})\boldsymbol{A}^{\otimes j-1},$$

where the first equality is done by reshaping the m^j tensor to be an $m \times m^{j-1}$ Jacobian matrix, the second equality is due to the induction hypothesis, and the third equality is an adoption of chain rule. Due to the fact $\operatorname{vec}(\boldsymbol{A}\boldsymbol{B}\boldsymbol{C}) = (\boldsymbol{C}^\top \otimes \boldsymbol{A})\operatorname{vec}(\boldsymbol{B})$, we have $\operatorname{vec}\left(\nabla^j[f \circ \boldsymbol{A}](\boldsymbol{x})\right) = \boldsymbol{A}^{\top\otimes j} \operatorname{vec}\left(\nabla^j f(\boldsymbol{A}\boldsymbol{x})\right)$. Then by induction, we can conclude the proof. □

B Proof of Theorem 2

For a sake of clarity, we first introduce few notations in algebra and real analysis. We use $C^k(\mathcal{X}, \mathbb{R})$ to denote the k-th differentiable functions defined over domain $\mathcal{X}$, $W^{k,p}(\mathcal{X}, \mathbb{R})$ to denote the k-th differentiable and L^p integrable Sobolev space over domain $\mathcal{X}$. We use notation $\mathcal{A}[f]$ to denote the image of function (say f) under the transformation of an operator (say $\mathcal{A}$). We use symbol $\circ$ to denote function composition (e.g., $f \circ g(\boldsymbol{x}) = f(g(\boldsymbol{x}))$). We use dot-product $\cdot$ between two functions (say f and g) to represent element-wise multiplication of function values (say $f \cdot g(\boldsymbol{x}) = f(\boldsymbol{x}) \cdot g(\boldsymbol{x})$). Besides, we list the following definitions and assumptions:

Definition B.1. *(Polynomial) We use $\mathbb{R}[x_1, \cdots, x_m]$ to represent the multivariate polynomials in terms of $x_1, \cdots, x_m$ with real coefficients. We write a (monic) multivariate monomial $m(x_1, \cdots, x_m) = x_1^{n_1} x_2^{n_2} \cdots x_m^{n_m}$ as $m(\boldsymbol{x}) = \boldsymbol{x}^{\boldsymbol{n}}$ where $\boldsymbol{n} = [n_1 \quad \cdots \quad n_m] \in \mathbb{N}^m$. Then we denote a polynomial as $p(\boldsymbol{x}) = a_1 \boldsymbol{x}^{\boldsymbol{n}_1} + \cdots + a_d \boldsymbol{x}^{\boldsymbol{n}_d} \in \mathbb{R}[x_1, \cdots, x_m]$ where $\boldsymbol{x}^{\boldsymbol{n}_i}$ denotes the i-th multivariate monomial and $a_i \in \mathbb{R}$ is the corresponding coefficient.*

Definition B.2. *(Differential Operator) Suppose a compact set $\mathcal{X} \subseteq \mathbb{R}^m$. we denote $\mathcal{D}^{\boldsymbol{n}} : C^\infty(\mathcal{X}, \mathbb{R}) \to C^\infty(\mathcal{X}, \mathbb{R})$ as the high-order differential operator associated with indices $\boldsymbol{n} \in \mathbb{N}^m$:*

$$\mathcal{D}^{\boldsymbol{n}}[f] = \frac{\partial^{\|\boldsymbol{n}\|_1}}{\partial x_1^{n_1} \cdots \partial x_m^{n_m}} f. \tag{14}$$

Definition B.3. *We define polynomial in gradient operator as: $p(\nabla) = p\left(\frac{\partial}{\partial x_1}, \cdots, \frac{\partial}{\partial x_m}\right) = a_1 \mathcal{D}^{\boldsymbol{n}_1} + \cdots + a_d \mathcal{D}^{\boldsymbol{n}_d} \in \mathbb{R}[x_1, \cdots, x_m]$ where $\mathcal{D}^{\boldsymbol{n}_i}$ denotes the $\boldsymbol{n}_i$-th order partial derivative (Definition B.2) and $a_i \in \mathbb{R}$ is the corresponding coefficient.*

Remark B.1. *The mapping between $p(\boldsymbol{x})$ and $p(\nabla)$ is a ring homomorphism from polynomial ring $\mathbb{R}[x_1, \cdots, x_m]$ to the ring of endomorphism defined over $C^\infty(\mathcal{X}, \mathbb{R})$.*

Definition B.4. *(Fourier Transform) Given real-valued function $f : \mathbb{R}^m \to \mathbb{C}$ that satisfies Dirichlet condition[5], then Fourier transform $\mathcal{F}$ is defined as:*

$$\mathcal{F}[f](\boldsymbol{w}) = \int_{\mathbb{R}^m} f(\boldsymbol{x}) \exp(-2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{x}. \tag{15}$$

Inverse Fourier transform $\mathcal{F}^{-1}$ exists and has the form of:

$$f(\boldsymbol{x}) = \int_{\mathbb{R}^m} \mathcal{F}[f](\boldsymbol{w}) \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{w}. \tag{16}$$

Definition B.5. *(Convolution) Given two real-valued functions $f : \mathbb{R}^m \to \mathbb{R}$ and $g : \mathbb{R}^m \to \mathbb{R}$, convolution between f and g is defined as:*

$$(f \star g)(\boldsymbol{x}) = \int_{\mathbb{R}^m} f(\boldsymbol{x} - \boldsymbol{\xi}) g(\boldsymbol{\xi}) d\boldsymbol{\xi}. \tag{17}$$

Then we denote $f \star g = \mathcal{T}_g[f]$, where $\mathcal{T}_g$ represents a convolutional operator associated with the function g.

We make the following mild assumptions on the signals and convolutional operators, which are widely satisfied by the common signals and systems.

Assumption B.6. *(Band-limited Signal Space) Define the signal space $\mathcal{S}$ as a Sobolev space $W^{\infty,1}(\mathbb{R}^m, \mathbb{R})$ of real-valued functions such that for $\forall f \in \mathcal{S}$:*

- *(I) $f \in C^\infty(\mathbb{R}^m, \mathbb{R})$ is continuous and smooth over $\mathbb{R}^m$.*
- *(II) f satisfies the Dirichlet condition.*
- *(III) f has a limited width of spectrum: there exists a compact subset $\mathcal{W} \subset \mathbb{R}^m$ such that $|\mathcal{F}[f](\boldsymbol{w})| = 0$ if $\boldsymbol{w} \notin \mathcal{W}$, and $\int_{\mathcal{W}} |\mathcal{F}[f](\boldsymbol{w})| d\boldsymbol{w} < \infty$.*

[5] Dirichlet condition guarantees Fourier transform exists: (1) The function is L_1 integrable over the entire domain. (2) The function has at most a countably infinite number of infinte minima or maxma or discontinuities over the entire domain.

Assumption B.7. *(Convolution Space) Define a convolutional operator space $\mathcal{T}$ such that $\forall \mathcal{T}_g \in \mathcal{T}$:*

(IV) $g : \mathbb{R}^m \to \mathbb{R}$ *is real-valued function.*

(V) $\mathcal{F}[g] \in C(\mathbb{R}^m, \mathbb{R})$ *has a continuous spectrum.*

Before we prove Theorem 2, we enumerate the following results as our key mathematical tools:

First of all, we note the following well-known result without a proof.

Lemma B.1. *(Convolution Theorem) For every $\mathcal{T}_g \in \mathcal{T}$, it always holds that $\mathcal{F} \circ \mathcal{T}_g[f](\boldsymbol{w}) = \mathcal{F}[f](\boldsymbol{w}) \cdot \mathcal{F}[g](\boldsymbol{w})$.*

Next, we present Stone-Weierstrass Theorem as our Lemma B.2 as below.

Lemma B.2. *(Stone-Weierstrass Theorem) Suppose $\mathcal{X}$ is a compact metric space. If $\mathcal{A} \subset C(\mathcal{X}, \mathbb{R})$ is a unital sub-algebra which separates points in $\mathcal{X}$. Then $\mathcal{A}$ is dense in $C(\mathcal{X}, \mathbb{R})$.*

A straightforward corollary of Lemma B.2 is the following Lemma B.3.

Lemma B.3. *Let $\mathcal{X} \subset \mathbb{R}^m$ be a compact subset of $\mathbb{R}^m$. For every $\epsilon > 0$, there exists a polynomial $p(\boldsymbol{x}) \in \mathbb{R}[x_1, \cdots, x_m]$ such that $\sup_{\boldsymbol{x} \in \mathcal{X}} |f(\boldsymbol{x}) - p(\boldsymbol{x})| < \epsilon$.*

Proof. Proved by checking polynomials $\mathbb{R}[x_1, \cdots, x_m]$ form a unital sub-algebra separating points in $\mathcal{X}$, and equipping $C(\mathcal{X}, \mathbb{R})$ with the distance metric $d(f, h) = \sup_{\boldsymbol{x} \in \mathcal{X}} |f(\boldsymbol{x}) - g(\boldsymbol{x})|$. □

We also provide the following Lemma B.4 to reveal the spectrum-domain symmetry for real-valued signals.

Lemma B.4. *Suppose f is a continuous real-valued function satisfying Dirichlet condition. Then $\mathcal{F}[f](\boldsymbol{w}) = \mathcal{F}[f](-\boldsymbol{w})^*$, i.e., the spectrum of real-valued function is conjugate symmetric.*

Proof. By the definition of Fourier transform (Definition B.4):

$$\mathcal{F}[f](-\boldsymbol{w}) = \int_{\mathbb{R}^m} f(\boldsymbol{x}) \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{x} = \int_{\mathbb{R}^m} f(\boldsymbol{x})^* \exp(-2\pi i \boldsymbol{w}^\top \boldsymbol{x})^* d\boldsymbol{x} \tag{18}$$

$$= \left[\int_{\mathbb{R}^m} f(\boldsymbol{x}) \exp(-2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{x} \right]^* = \mathcal{F}[f](\boldsymbol{w})^*, \tag{19}$$

where Eq. 18 holds because f is a real-valued function. □

We present Lemma B.5 as below, which reflects the effect of differential operators on the spectral domain.

Lemma B.5. *Suppose $f \in C^\infty(\mathbb{R}^m, \mathbb{R})$ is a smooth real-valued function satisfying Dirichlet condition. Then $\mathcal{F} \circ \mathcal{D}^{\boldsymbol{n}}[f](\boldsymbol{w}) = (2\pi i)^{\|\boldsymbol{n}\|_1} \boldsymbol{w}^{\boldsymbol{n}} \cdot \mathcal{F}[f](\boldsymbol{w})$ for every $\boldsymbol{n} \in \mathbb{N}^m$.*

Proof. We first show the case of first-order partial derivative. Suppose $h \in C(\mathbb{R}^m, \mathbb{R})$ is L^1 integrable (then Fourier transform exists).

$$\frac{\partial}{\partial x_i} h(\boldsymbol{x}) = \frac{\partial}{\partial x_i} \int_{\mathbb{R}^m} \mathcal{F}[h](\boldsymbol{w}) \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{w} \tag{20}$$

$$= \int_{\mathbb{R}^m} \mathcal{F}[h](\boldsymbol{w}) \frac{\partial}{\partial x_i} \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{w} \tag{21}$$

$$= 2\pi i \int_{\mathbb{R}^m} w_i \mathcal{F}[h](\boldsymbol{w}) \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) d\boldsymbol{w}. \tag{22}$$

Then we apply the Fourier transform to Eq. 20, we can obtain:

$$\mathcal{F} \circ \frac{\partial}{\partial x_i}[h](\boldsymbol{w}) = 2\pi i w_i \mathcal{F}[h](\boldsymbol{w}). \tag{23}$$

Note that $f \in W^{\infty,1}(\mathbb{R}^m, \mathbb{R})$ ensures all its partial derivatives are differentiable and absolutely integrable. We can recursively apply $\frac{\partial}{\partial x_i}$ to f for n_i times for each $i \in [m]$, and use Eq. 23 above to conclude the proof. □

Below is the formal statement of our Theorem 2 and its detailed proof.

Theorem B.6. *For every $\mathcal{T}_g \in \mathcal{T}$ and arbitrarily small $\epsilon > 0$, there exists a polynomial $p(\boldsymbol{x}) \in \mathbb{R}[x_1, \cdots, x_m]$ such that $\sup_{\boldsymbol{x}\in\mathbb{R}^m} |\mathcal{T}_g[f](\boldsymbol{x}) - p(\nabla)[f](\boldsymbol{x})| < \epsilon$ for all $f \in \mathcal{S}$.*

Proof. For every $f \in \mathcal{S}$ and $\mathcal{T}_g \in \mathcal{T}$, by Lemma B.1, one can rewrite:

$$\mathcal{F} \circ \mathcal{T}_g[f](\boldsymbol{w}) = \mathcal{F}[f](\boldsymbol{w}) \cdot \mathcal{F}[g](\boldsymbol{w}) := \hat{f}(\boldsymbol{w})\hat{g}(\boldsymbol{w}), \tag{24}$$

where we use $\hat{f} : \mathbb{R}^m \to \mathbb{C}$ and $\hat{g} : \mathbb{R}^m \to \mathbb{C}$ to denote the Fourier transform of f and g, respectively. We can construct an invertible mapping ϕ by letting:

$$\phi[\hat{f}](\boldsymbol{w}) = \Re\{\hat{f}(\boldsymbol{w})\} - \Im\{\hat{f}(\boldsymbol{w})\}, \tag{25}$$

$$\phi^{-1}[\tilde{f}](\boldsymbol{w}) = \frac{\tilde{f}(\boldsymbol{w}) + \tilde{f}(-\boldsymbol{w})}{2} - i\frac{\tilde{f}(\boldsymbol{w}) - \tilde{f}(-\boldsymbol{w})}{2}, \tag{26}$$

which is also known as the Hartley transform. By Lemma B.4 (with Assumption (I) (IV)), $\tilde{f} := \phi[\hat{f}]$ and $\tilde{g} := \phi[\hat{g}]$ are both real-valued functions.

Since $\hat{f}$ is only supported in $\mathcal{W}$ (by Assumption (III)), we only consider $\tilde{g}$ within the compact subset $\mathcal{W}$. By Lemma B.3 (with Assumption (V)), there exists a polynomial $\tilde{p}(\boldsymbol{w}) \in \mathbb{R}[w_1, \cdots, w_m] = \tilde{a}_0 + \tilde{a}_1\boldsymbol{w}_{\boldsymbol{n}_1} + \cdots + \tilde{a}_d\boldsymbol{w}^{\boldsymbol{n}_d}$ such that $\sup_{\boldsymbol{w}\in\mathcal{W}} |\tilde{g}(\boldsymbol{w}) - \tilde{p}(\boldsymbol{w})| < \epsilon/2C$ for every $\epsilon > 0$, where d is the number of monomials in $\tilde{p}$, $\tilde{a}_0, \tilde{a}_1, \cdots, \tilde{a}_d \in \mathbb{R}$ are corresponding coefficients, and $C > 0$ is some constant.

Applying ϕ^{-1} to $\tilde{p}$, we will obtain a new (complex-valued) polynomial $\hat{p} := \phi^{-1}[\tilde{p}] \in \mathbb{C}[w_1, \cdots, w_m]$ such that:

$$\Re\{\hat{p}(\boldsymbol{w})\} = \frac{\tilde{p}(\boldsymbol{w}) + \tilde{p}(-\boldsymbol{w})}{2}, \qquad \Im\{\hat{p}(\boldsymbol{w})\} = \frac{\tilde{p}(\boldsymbol{w}) - \tilde{p}(-\boldsymbol{w})}{2}. \tag{27}$$

We observe that the coefficients of $\hat{p}$ satisfy: $\hat{a}_k = \tilde{a}_k$ if $\|\boldsymbol{n}_k\|_1$ is even and $\hat{a}_k = i\tilde{a}_k$ if $\|\boldsymbol{n}_k\|_1$ is odd. Then we bound the difference between $\hat{g}$ and $\hat{p}$ for every $\boldsymbol{w} \in \mathcal{W}$:

$$|\hat{g}(\boldsymbol{w}) - \hat{p}(\boldsymbol{w})| = \left| \left(\frac{\tilde{f}(\boldsymbol{w}) + \tilde{f}(-\boldsymbol{w})}{2} - \frac{\tilde{p}(\boldsymbol{w}) + \tilde{p}(-\boldsymbol{w})}{2} \right) \right. \tag{28}$$

$$\left. - i \left(\frac{\tilde{f}(\boldsymbol{w}) - \tilde{f}(-\boldsymbol{w})}{2} - \frac{\tilde{p}(\boldsymbol{w}) - \tilde{p}(-\boldsymbol{w})}{2} \right) \right| \tag{29}$$

$$\leq \frac{1}{2} \left(\left| \tilde{f}(\boldsymbol{w}) - \tilde{p}(\boldsymbol{w}) \right| + \left| \tilde{f}(-\boldsymbol{w}) - \tilde{p}(-\boldsymbol{w}) \right| \right) \tag{30}$$

$$+ \frac{1}{2} \left(\left| \tilde{p}(\boldsymbol{w}) - \tilde{f}(\boldsymbol{w}) \right| + \left| \tilde{p}(-\boldsymbol{w}) - \tilde{f}(-\boldsymbol{w}) \right| \right) \tag{31}$$

$$\leq \frac{\epsilon}{C}. \tag{32}$$

In the meanwhile, by Lemma B.5 (with Assumption (I), (II)), $\mathcal{F} \circ \mathcal{D}^{\boldsymbol{n}}[f](w) = (2\pi i)^{\|\boldsymbol{n}\|_1}\boldsymbol{w}^{\boldsymbol{n}} \cdot \mathcal{F}[f](\boldsymbol{w})$ for every $\boldsymbol{n} \in \mathbb{N}^m$. Define a sequence $q_{\boldsymbol{n}}(\boldsymbol{w}) = (2\pi i)^{\|\boldsymbol{n}\|_1}\boldsymbol{w}^{\boldsymbol{n}}$, then partial derivatives of f in terms of $\boldsymbol{n} \in \mathbb{N}^m$ can be written as:

$$\mathcal{F} \circ \mathcal{D}^{\boldsymbol{n}}[f](\boldsymbol{w}) = q_{\boldsymbol{n}}(\boldsymbol{w}) \cdot \mathcal{F}[f](\boldsymbol{w}). \tag{33}$$

Next we decompose polynomial $\hat{p}$ in terms of $q_{\boldsymbol{n}}$. Let $a_k = \hat{a}_k/(2\pi i)^{\|\boldsymbol{n}_k\|_1}$, then $\hat{p}(\boldsymbol{w}) = a_0 + a_1 q_{\boldsymbol{n}_1}(\boldsymbol{w}) + \cdots + a_d q_{\boldsymbol{n}_d}(\boldsymbol{w})$. We note that $\{a_k, \forall k \in [d]\}$ must be real numbers since $\hat{a}_k$ is real/imaginary when $\|\boldsymbol{n}_k\|_1$ is even/odd, which coincides with $(2\pi i)^{\|\boldsymbol{n}_k\|_1}$.

By linearity of inverse Fourier transform and Eq. 33, element-wisely multiplying $\sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k}$ to $\hat{f}$ will lead to a transform on the spatial domain:

$$\mathcal{F}^{-1}\left[\sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k} \cdot \mathcal{F}[f] \right] = \sum_{k=0}^{d} a_k \mathcal{D}^{\boldsymbol{n}_k} f := p(\nabla)[f], \tag{34}$$

Algorithm 1 Forward pass of INSP-ConvNet

1: **Input**: An INR network $\Phi(\boldsymbol{x}) : \mathbb{R}^m \to \mathbb{R}$, convolutional operator weights $\boldsymbol{\theta}^{(l)} \in \mathbb{R}^M$ and an input coordinate $\boldsymbol{x}$.
2: **Output**: Value at $\boldsymbol{x}$ of INR $\mathrm{ConvNet}[\Phi]$ processed by INSP-ConvNet.
3: $y^{(0)} \leftarrow \Phi(\boldsymbol{x})$
4: **for** $l = 1, \cdots, L$ **do**
5: $\quad \hat{y}^{(l)} \leftarrow \left[y^{(l-1)} \quad \frac{\partial y^{(l-1)}}{\partial \boldsymbol{x}}^\top \quad \frac{\partial^2 y^{(l-1)}}{\partial \boldsymbol{x}^2}^\top \quad \cdots \quad \frac{\partial^K y^{(l-1)}}{\partial \boldsymbol{x}^K}^\top \right] \boldsymbol{\theta}^{(l)}$ ▷ Convolutional layer
6: $\quad y^{(l)} \leftarrow \mathrm{ReLU}(\mathrm{InstanceNorm1D}(\hat{y}^{(l)}))$ ▷ Non-linearity and normalization
7: **end for**
8: **return** $y^{(L)}$.

where we define polynomial $p(\nabla) := a_0 + a_1 \mathcal{D}^{\boldsymbol{n}_1} + \cdots + a_d \mathcal{D}^{\boldsymbol{n}_d} \in \mathbb{R}\left[\frac{\partial}{\partial x_1}, \cdots, \frac{\partial}{\partial x_m}\right]$ over the ring of partial differential operators (Definition B.3). Now we bound the difference between $\mathcal{T}_g[f]$ and $p(\nabla)[f]$ for every $f \in \mathcal{S}$ and $\boldsymbol{x} \in \mathbb{R}^m$:

$$|\mathcal{T}_g[f](\boldsymbol{x}) - p(\nabla)[f](\boldsymbol{x})| = \left| \int_{\mathcal{W}} \exp(2\pi i \boldsymbol{w}^\top \boldsymbol{x}) \hat{f}(\boldsymbol{w}) \left(\hat{g}(\boldsymbol{w}) - \sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k}(\boldsymbol{w}) \right) d\boldsymbol{w} \right| \tag{35}$$

$$\leq \int_{\mathcal{W}} \left| \hat{f}(\boldsymbol{w}) \left(\hat{g}(\boldsymbol{w}) - \sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k}(\boldsymbol{w}) \right) \right| d\boldsymbol{w} \tag{36}$$

$$\leq \left(\sup_{\boldsymbol{w} \in \mathcal{W}} \left| \hat{g}(\boldsymbol{w}) - \sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k}(\boldsymbol{w}) \right| \right) \left(\int_{\mathcal{W}} \left| \hat{f}(\boldsymbol{w}) \right| d\boldsymbol{w} \right) \tag{37}$$

$$\leq \epsilon, \tag{38}$$

where Eq. 37 follows from Hölder's inequality, and Eq. 38 is obtained by substituting the upper bound of difference $|\hat{g}(\boldsymbol{w}) - \sum_{k=0}^{d} a_k q_{\boldsymbol{n}_k}(\boldsymbol{w})|$ and letting C equal to the L^1 norm of $\hat{f}(\boldsymbol{w})$ (by Assumption (III)). □

C Implementation Details of INSP-ConvNet

We have formulated exact convolution form and INSP-Conv in Sec. 3.3. We provide a pseudocode to illustrate the forward pass of INSP-ConvNet in Algorithm 1. Below we elaborate each main component:

Convolutional Layer. Each $\mathcal{A}^{(l)}$ represents an implicit convolution layer. We follow the closed-form solution in Eq. 5 to parameterize $\mathcal{A}^{(l)}$ with $\boldsymbol{\theta}^{(l)}$. We point out that $\mathrm{ConvNet}[\Phi]$ also corresponds to a computational graph, which can continuously map coordinates to the output features. To construct this computational graph, we recursively call for gradient networks of the previous layer until the first layer. For example, $\mathcal{A}^{(l)}$ will request the gradient network of $\mathcal{A}^{(l-1)} \cdot \sigma \circ \cdots \circ \mathcal{A}^{(1)} \cdot \Phi$, and then $\mathcal{A}^{(l-1)}$ will request the gradient network of the rest part. This procedure will proceed until the first layer, which directly returns the derivative network of Φ. Kernels in CNNs typically perform multi-channel convolution. However, this is not memory friendly to gradient computing in our framework. To this end, we run channel-wise convolution first and then employ a linear layer to mix channels [90].

Nonlinear Activation and Normalization. Nonlinear activation and normalization are naturally element-wise functions. They are point-wisely applied to the output of an INSP-Net and participate the computational graph construction process. This corresponds to the line 6 of Algorithm 1.

Training Recipe. Given a dataset $\mathcal{D} = \{(\Phi_i, y_i)\}$ with a set of pre-trained INRs Φ_i and their corresponding labels y_i, our goal is to learn a $\mathrm{ConvNet}[\cdot]$ that can process each example. In contrast to standard ConvNets that are designed for grid-based images, the computational graph of INSP-ConvNet contains parameters of both the input INR Φ_i and learnable kernels $\mathcal{A}^{(l)}$. During the

training stage, we randomly sample a mini-batch (Φ_i, y_i) from $\mathcal{D}$ to optimize INSP-ConvNet. The corresponding loss will be evaluated according to the network output, and then back-propagate the calculated gradients to the learnable parameters in $\mathcal{A}^{(l)}$, using the stochastic gradient descent optimization. Along the whole process, the parameters of Φ_i is fixed and only the parameters in $\mathcal{A}^{(l)}$ is optimized. Standard data augmentations are included by default, including rotation, zoom in/out, etc. In practice, we implement these augmentations by using affine transformation on the coordinates of INRs.

D Connection with PDE based Signal Processing

Partial Differential Equation (PDE) has been successfully applied to image processing domain as we discussed in Sec. 4.3. In this section, we focus on their connection with our INSP-Net. We summarize the methods of this line of works [61, 40, 66] in the following formulation:

$$\frac{\partial \Psi(\boldsymbol{x}, t)}{\partial t} = M_t \left[\Psi(\boldsymbol{x}, t), \nabla_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t), \nabla^2_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t), \cdots \right], \tag{39}$$

where $M_t(\cdot)$ is a time-variant function that remaps the direct output and high-order derivatives of function Ψ. For heat diffusion, M_t boils down to be an stationary isotropic combination of second-order derivatives. In [61], M_t is chosen to be a gradient magnitude aware diffusion operator running on divergence operators. [40, 66] degenerate M_t to a time-dependent linear mapping of pre-defined invariants of the maximal order two. We note that Eq. 39 can be naturally solved with INRs, as INRs are amenable to solving complicated differential equation shown by [28]. One straightforward solution is to parameterize M_t by another time-dependent coordinate network [27] and enforce the boundary condition $\Psi(\boldsymbol{x}, 0) = \Phi(\boldsymbol{x})$ and minimize the difference between the two hands of the Eq. 39. However, foreseeable problem falls in sampling inefficiency over the time axis. Suppose we discretize the time axis into small intervals $0 = t_0 < t_1 < \cdots < t_N$, then Eq. 39 has a closed-form solution given M_t by Euler method:

$$\Psi(\boldsymbol{x}, t_{n+1}) = \int_{t_n}^{t_{n+1}} M_t \left[\Psi(\boldsymbol{x}, t), \nabla_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t), \nabla^2_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t), \cdots \right] dt + \Psi(\boldsymbol{x}, t_n) \tag{40}$$

$$\approx M_{t_n} \left[\Psi(\boldsymbol{x}, t_n), \nabla_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t_n), \nabla^2_{\boldsymbol{x}} \Psi(\boldsymbol{x}, t_n), \cdots \right] (t_{n+1} - t_n) + \Psi(\boldsymbol{x}, t_n). \tag{41}$$

One can see Eq. 41 can be regarded as a special case of our model Eq. 3, where we absorb M_{t_n}, time interval $t_{n+1} - t_n$, and the residual term $\Psi(\boldsymbol{x}, t_n)$ into one Π. Considering our multi-layer model INSP-ConvNet (see Sec. 3.3), we can analogize t_n to the layer number, and then solving Eq. 39 at time t_N is approximately equal to forward passing an N-layer INSP-ConvNet.

E More Experiment Details

We implement our INSP framework using PyTorch. The gradients are obtained directly using the autograd package from PyTorch. All learnable parameters are trained with AdamW optimizer and a learning rate of 1e-4. For low-level image processing kernels, images are obtained from Set5 dataset [75], Set14 dataset [76], and DIV-2k dataset [77]. These datasets are original collected for super-resolution task, so the images are diverse in style and content. In our experiments, we construct SIREN [28] on each image. For efficiency, we resized the images to 256×256. We use 90 images to construct the INRs used for training, and use the other images for evaluation.

For image classification, we construct a 2-layer INSP-ConvNet framework. Each INSP layer constructs the derivative computational graphs of the former layers and combines them with learnable Π. The INSP-layer is capable of approximating a convolution filter. For a fair comparison, we build another 2-layer depthwise convolutional network running on image pixels as the baseline. Both our INSP-ConvNet and the ConvNet running on pixels are trained with the same hyper-parameters. Both experiments take 1000 training epochs, with a learning rate of $1e-4$ using AdamW optimizer.

F More Experimental Results

Additional Visualization. In this section, we provide more experimental results. Fig. 9 provides comparisons on edge detection task. Fig. 10 shows image denoising results. Fig. 11 demonstrates

	PSNR	SSIM	LPIPS
Input (decoded from INR)	20.51	0.47	0.40
MPRNet [81]	23.95	0.72	0.36
MAXIM [91]	24.64	0.74	0.33
Mean Filter	22.57	0.60	0.43
INSP-Net	23.86	0.65	0.38

Table 2: Quantitative result of image denoising on 100 testing images from DIV-2k dataset [77], where the synthetic noise is rgb gaussian noise. The noise is similar to the ones seen during the training of MPRNet and MAXIM, so they obtain better performance with the help of a much wider training set.

image deblurring results. Fig. 12 shows image blurring results. Fig. 13 shows image inpainting results. Fig. 14 presents additional results on geometry smoothening.

Additional Quantitative Results. We also provide quantitative comparisons on the test set in Tab. 2. The test set consists of 100 INRs fitted from 100 images in DIV-2k dataset [77]. In Tab. 2, their performance is better when the synthetic noise becomes three-channel Gaussian noise. The synthetic noise is similar to those seen during the training process of MPRNet [81] and MAXIM [91], so they benefit from their much wider training set.

Audio Signal Processing. We additionally validate the ability of our INSP framework by processing audio signals. We add synthetic Gaussian noise onto the audio and use it to fit a SIREN. The noisy audio decoded from the INR is shown in Fig. 15(b). Then we use our INSP-Net to implicitly process it to a new INR that can be further decoded into denoised audio. It's decoded result is shown in Fig. 15(c). We also provide visualization of the denoising effect in Fig. 15(f).

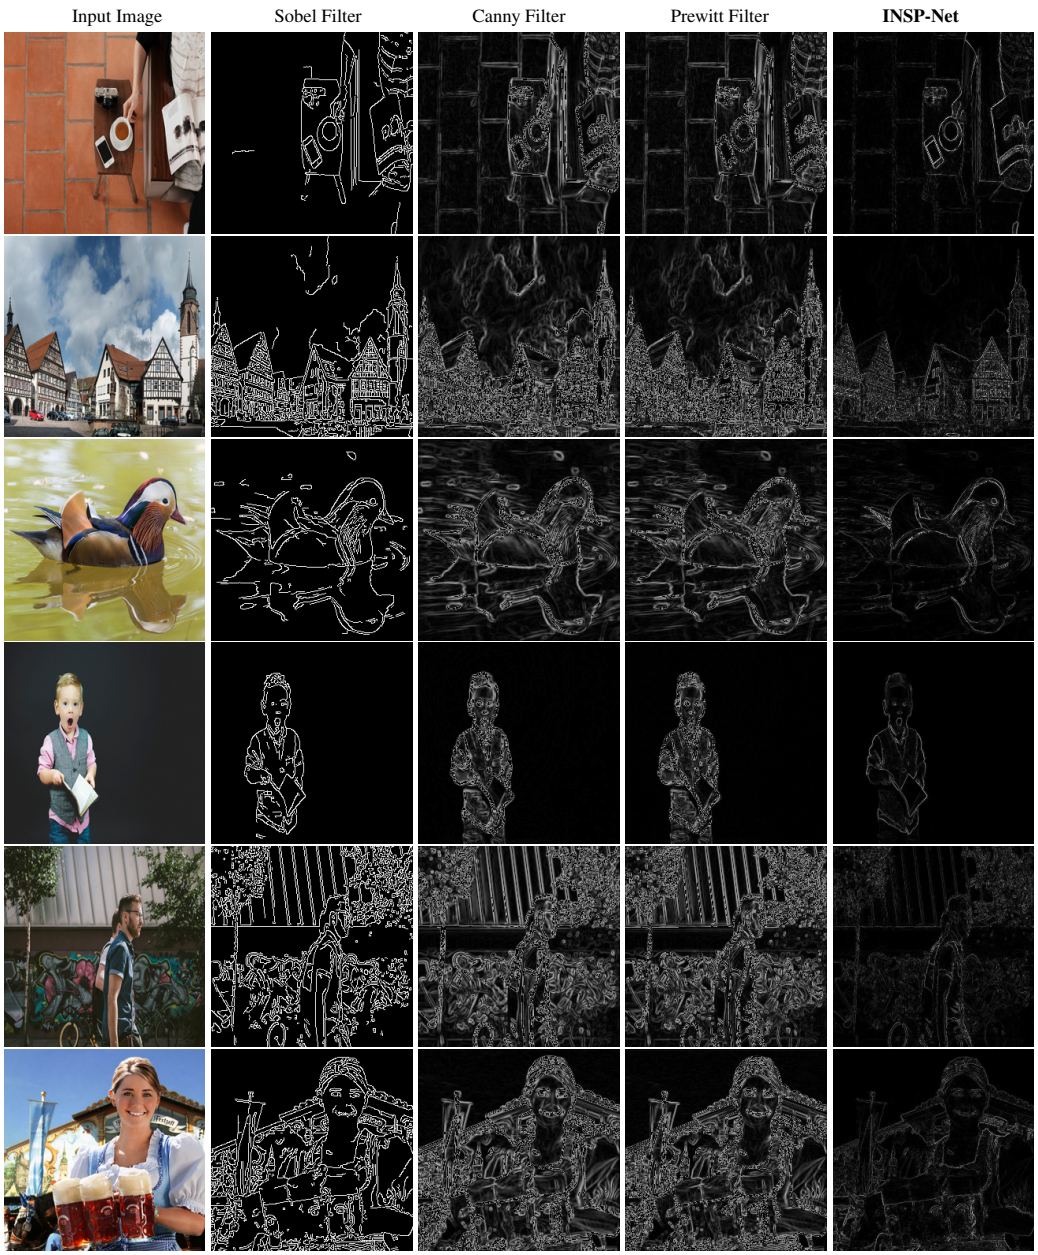

Figure 9: Edge detection. We fit the natural images with SIREN and use our INSP-Net to process implicitly into a new INR that can be decoded into edge maps.

Figure 10: Image denoising. We fit the noisy images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into natural clear images.

Figure 11: Image deblurring. We fit the blurred images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into clear natural images.

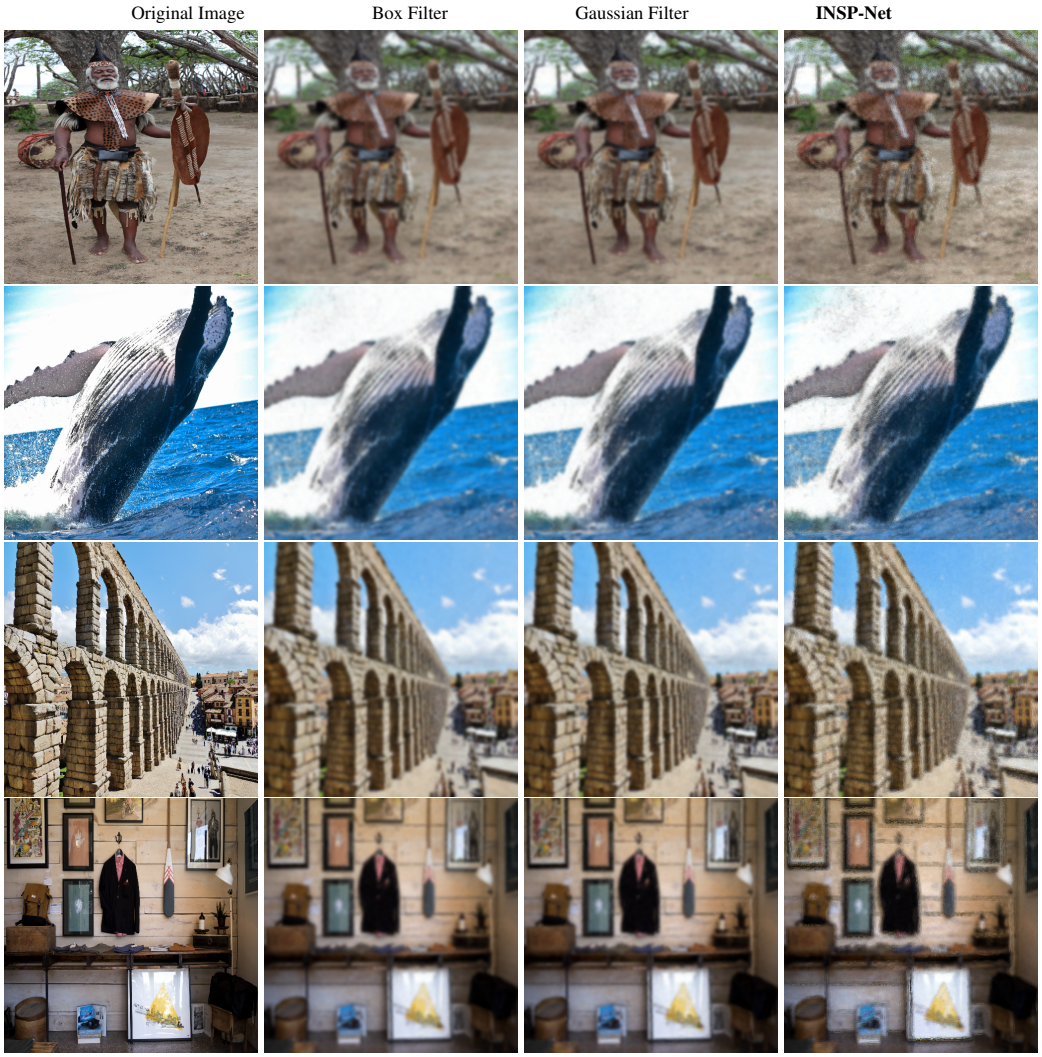

Figure 12: Image blurring. We fit the natural images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into blurred images.

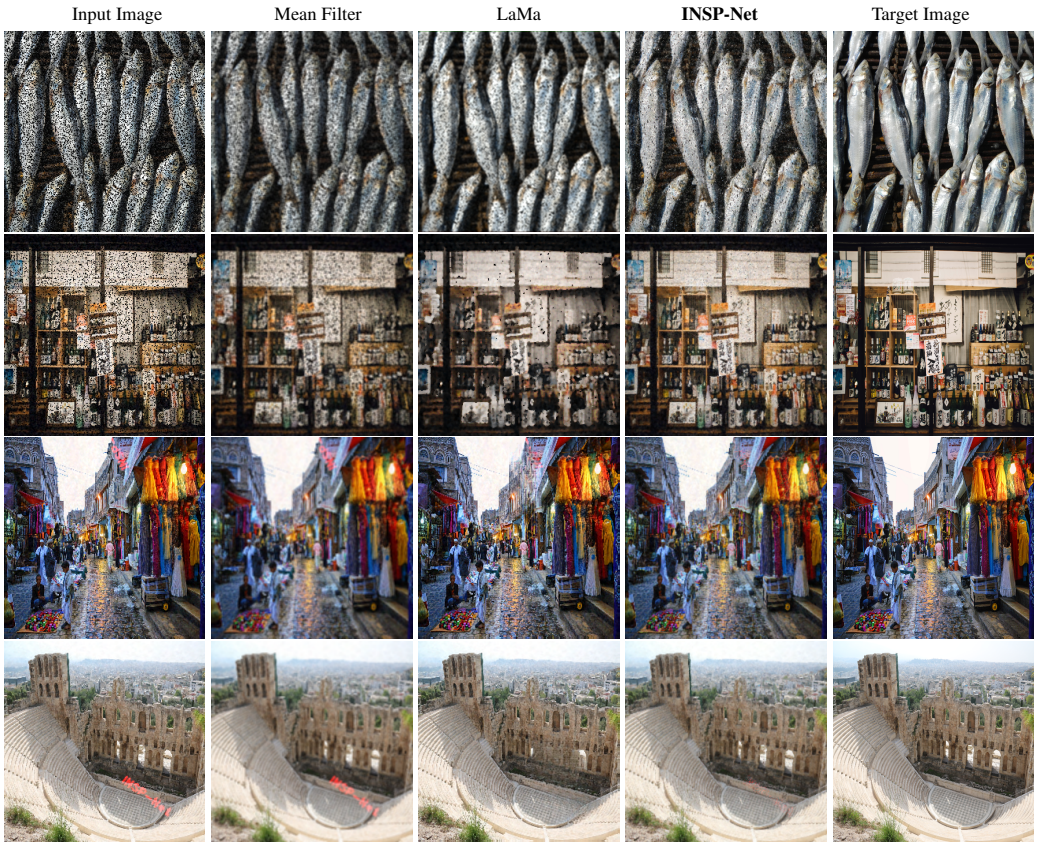

Figure 13: Image inpainting. We fit the input images with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into natural images. Note that LaMa requires explicit masks to select the regions for inpainting and the masks are roughly provided. The first two rows contain input images with random pixels erased. The last two rows contain input images with text contamination.

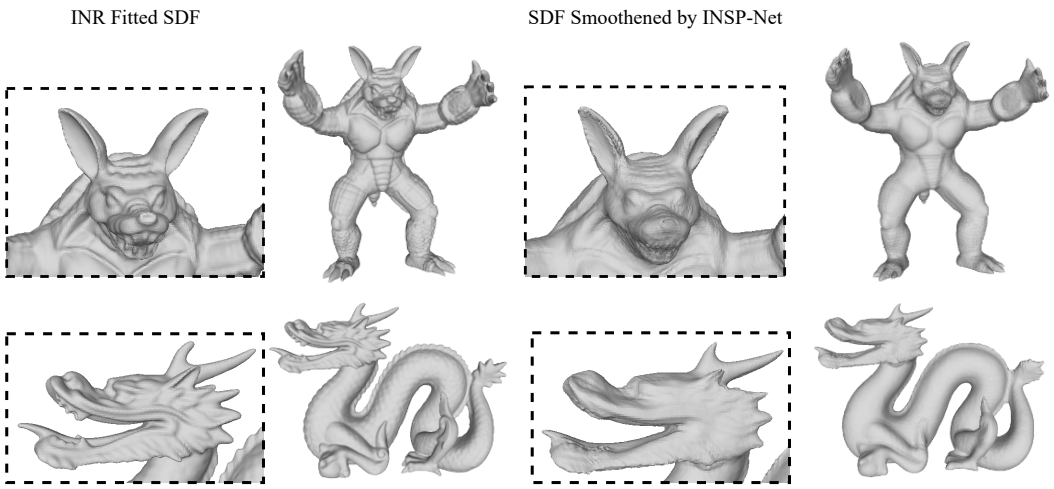

Figure 14: Additional results on geometry smoothening via INSP-Net. Left: unprocessed geometry decoded from an unprocessed INR. Right: smoothened geometry decoded from the output INR of our INSP-Net. Best view in a zoomable electronic copy.

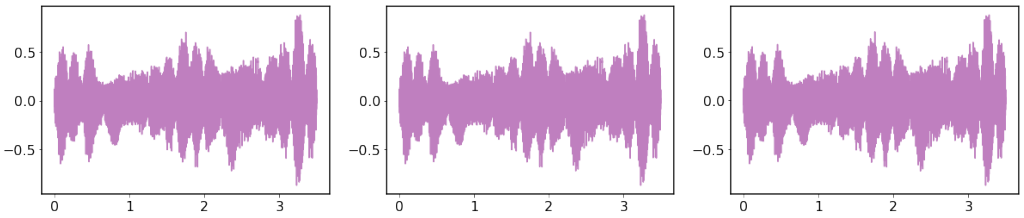

(a) Ground truth clear audio. (b) Input noisy audio (decoded from INR). (c) Output audio of INSP-Net.

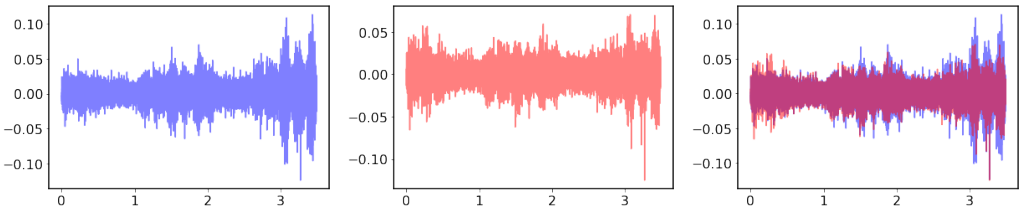

(d) Difference between input and ground truth. (e) Difference between output and ground truth. (f) Contrast between input (blue) and output (red) differences.

Figure 15: Audio denoising. We fit the noisy audio with SIREN and train our INSP-Net to process implicitly into a new INR that can be decoded into denoised audio.