

Intrinsic Neural Fields: Learning Functions on Manifolds

Lukas Koestler^{*,1}, Daniel Grittner^{*,1}, Michael Moeller², Daniel Cremers¹,
and Zorah Lähler²

¹ Technische Universität München

² Universität Siegen

Abstract. Neural fields have gained significant attention in the computer vision community due to their excellent performance in novel view synthesis, geometry reconstruction, and generative modeling. Some of their advantages are a sound theoretic foundation and an easy implementation in current deep learning frameworks. While neural fields have been applied to signals on manifolds, *e.g.*, for texture reconstruction, their representation has been limited to extrinsically embedding the shape into Euclidean space. The extrinsic embedding ignores known intrinsic manifold properties and is inflexible wrt. transfer of the learned function. To overcome these limitations, this work introduces intrinsic neural fields, a novel and versatile representation for neural fields on manifolds. Intrinsic neural fields combine the advantages of neural fields with the spectral properties of the Laplace-Beltrami operator. We show theoretically that intrinsic neural fields inherit many desirable properties of the extrinsic neural field framework but exhibit additional intrinsic qualities, like isometry invariance. In experiments, we show intrinsic neural fields can reconstruct high-fidelity textures from images with state-of-the-art quality and are robust to the discretization of the underlying manifold. We demonstrate the versatility of intrinsic neural fields by tackling various applications: texture transfer between deformed shapes & different shapes, texture reconstruction from real-world images with view dependence, and discretization-agnostic learning on meshes and point clouds.

1 Introduction

Neural fields have grown incredibly popular for novel view synthesis since the breakthrough work by Mildenhall et al. [29]. They showed that neural radiance fields together with differentiable volume rendering can be used to reconstruct scenes and often yield photorealistic renderings from novel viewpoints. This inspired work in related fields, *e.g.*, human shape modeling [34], shape and texture generation from text [28], and texture representation on shapes [32,2], where neural fields are able to generate a wide variety of functions with high fidelity.

Authors with * contributed equally.

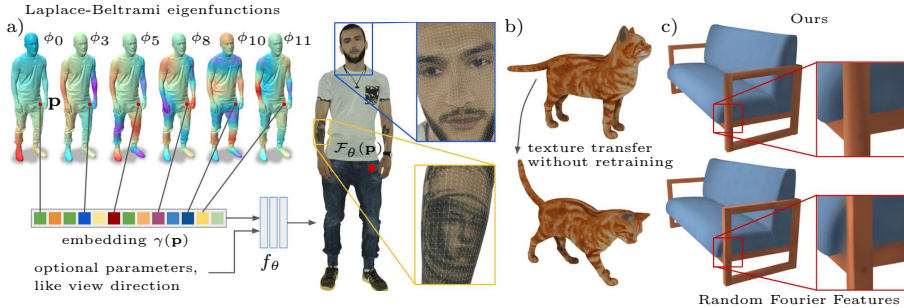


Fig. 1: (a) Overview of our method. We use the eigenfunctions ϕ_i of the Laplace-Beltrami operator (LBO) at each point as a point embedding $\gamma(\mathbf{p})$. This overcomes the spectral bias of the multilayer perceptron (MLP) f_θ , and hence the combined *intrinsic neural field* \mathcal{F}_θ can represent a high-frequency function on the surface. Notice that \mathbf{p} can be inside a triangle, and the function is clearly more detailed than the discretization (*insets*). (b) An intrinsic neural texture field trained on one shape (*top*) can be transferred to a new shape (*bottom*) without retraining. (c) Due to our intrinsic approach (LBO eigenfunctions) local geometry is maintained in close but separate parts, whereas an extrinsic approach (Random Fourier Features [51]) shows bleeding artifacts when trained with sparse supervision.

These methods use neural fields as functions from a point in Euclidean space to the quantity of interest. While this is valid for many applications, for others, the output actually lives on a general manifold. For example, texture mappings define a high-frequency color function on the surface of a 3D object. TextureFields [32] and Text2Mesh [28] solve this discrepancy by defining a mapping of each surface point to its Euclidean embedding and then learning the neural field there. Both show that they can achieve detail preservation above the discretization level, but the detour to Euclidean space has drawbacks. The Euclidean and geodesic distance between points can differ significantly. This is important on intricate shapes with fine geometric details that overlap because the local geometry prior is lost. Further, extrinsic representations cannot be used in the presence of surface deformations without retraining or applying heuristics.

Similar challenges have been solved in geometry processing by using purely intrinsic representations, most famously properties derived from the Laplace-Beltrami operator (LBO). Some of the main advantages of the LBO are its invariance under rigid and isometric deformations and reparametrization. We follow this direction by defining intrinsic neural fields on manifolds independent of the extrinsic Euclidean embedding and thus inherit the favorable properties of intrinsic representations. This is enabled by the fact that random Fourier features [51], an embedding technique that enabled the recent success of Euclidean neural fields, have an intrinsic analog based on the LBO. The result is a fully differentiable method that can learn high-frequency information on any 3D geometry representation that admits the computation of the LBO. A schematic overview of our method can be found in Figure 1. Our main theoretical and experimental **contributions** are:

- We introduce **intrinsic neural fields**, a novel and versatile representation for neural fields on manifolds. Intrinsic neural fields combine the advantages of neural fields with the spectral properties of the Laplace-Beltrami operator.
- We extend the **neural tangent kernel analysis** of [51] to the manifold setting. This yields a proof characterizing the stationarity of the kernel induced by intrinsic neural fields and insight into their spectral properties.
- We show that intrinsic neural fields can **reconstruct high-fidelity textures** from images with state-of-the-art quality.
- We demonstrate the versatility of intrinsic neural fields by tackling **various applications**: texture transfer between isometric and non-isometric shapes, texture reconstruction from real-world images with view dependence, and discretization-agnostic learning on meshes and point clouds.

The source code can be found at github.com/tum-vision/intrinsic-neural-fields.

This work studies how a *neural field* can be defined on a manifold. Current approaches use the *extrinsic Euclidean embedding* to define the neural field on a manifold in the extrinsic embedding space – we describe this approach in [Sec. 3.1](#). In contrast, our approach uses the well-known Laplace-Beltrami Operator (LBO), which we briefly introduce in [Sec. 3.2](#). The final definition of *intrinsic neural fields* is given in [Sec. 4](#). The experimental results are presented in [Sec. 5](#).

2 Related Work

This work investigates neural fields for learning on manifolds, and we will only consider directly related work in this section. We point interested readers to the following overview articles: neural fields in visual computing [59], advances in neural rendering [52], and an introduction into spectral shape processing [25].

Neural Fields. While representing 3D objects and scenes with coordinate-based neural networks, or neural fields, has already been studied more than two decades ago [16,37,36], the topic has gained increased interest following the breakthrough work by Mildenhall et al. [29]. They show that a Neural Radiance Field (NeRF) often yields photorealistic renderings from novel viewpoints. One key technique underlying this success is positional encoding, which transforms the three-dimensional input coordinates into a higher dimensional space using sines and cosines with varying frequencies. This encoding overcomes the low-frequency bias of neural networks [38,3] and thus enables high-fidelity reconstructions. The aforementioned phenomenon is analyzed using the neural tangent kernel [20] by Tancik et al. [51], and our analysis extends theirs from Euclidean space to manifolds. Simultaneously to Tancik et al., Sitzmann et al. [48] use periodic activation functions for neural scene representation, which is similar to the above-mentioned positional encoding [4]. Additionally, many other works [61,41,40,27,22,19,64,55,26] offer insights into neural fields and their embedding functions. Most notably, [14] introduces spectral features for transformers on graphs. However, none of these works considers neural fields on manifolds.

Neural Fields on Manifolds. Prior works [32,11,2,58,34,30,60,28,18] use neural fields to represent a wide variety of quantities on manifolds. Oechsle et al. [32] use the extrinsic embedding of the manifold to learn textures as multilayer perceptrons. Their Texture Fields serve as an important baseline for this work. NeuTex by Xiang et al. [58] combines neural, volumetric scene representations with a 2D texture network to facilitate interpretable and editable texture learning. To enable this disentanglement, their method uses mapping networks from the 3D space of the object to the 2D space of the texture and back. We compare with an adapted version of their method that utilizes the known geometry of the object. Baatz et al. [2] introduce NeRF-Tex, a combination of neural radiance fields (NeRFs) and classical texture maps. Their method uses multiple small-scale NeRFs to cover the surface of a shape and represent mesoscale structures, such as fur, fabric, and grass. Because their method focuses on mesoscale and artistic editing, we believe that extending the current work to their setting is an interesting direction for future research. Further, neural fields are used to represent quantities other than texture on manifolds. Palafox et al. [34] define a neural deformation field that maps points on a canonical shape to their location on the deformed shape. This model is applied to generate neural parametric models which can be used similarly to traditional models like SMPL [24]. Yifan et al. [60] decompose a neural signed distance function (SDF) into a coarse SDF and a high-frequency implicit displacement field. Morreale et al. [30] define neural surface maps, which can be used to define surface-to-surface correspondences among other applications. Text2Mesh [28] uses a coarse mesh and a textual description to generate a detailed mesh and associated texture as neural fields.

Intrinsic Geometry Processing. Intrinsic properties are a popular tool in geometry processing, especially in the analysis of deformable objects. The most basic intrinsic features are Gauss curvature and intrinsic point descriptors based on the Laplace-Beltrami operator (LBO). They have been heavily used since the introduction of the global point signature [42] and refined since then [50,1]. Intrinsic properties are not derived from a manifold’s embedding into its embedding space but instead arise from the pairwise geodesic distance on the surface. These are directly related to natural kernel functions on manifolds, *e.g.*, shown by the efficient approximation of the geodesic distance from the heat kernel [12]. Kernel functions as a measure of similarity are popular in geometry processing. They have been used in various applications, *e.g.*, in shape matching [54,9,23], parallel transport [46], and robustness wrt. discretization [53,44]. Manifold kernels naturally consider the local and global geometry [5], and our approach follows this direction by showing a natural extension of neural fields on manifolds.

3 Background

Differential geometry offers two viewpoints onto manifolds: intrinsic and extrinsic. The extrinsic viewpoint studies the manifold \mathcal{M} through its *Euclidean embedding* where each point $\mathbf{p} \in \mathcal{M}$ is associated with its corresponding point

in Euclidean space. In contrast, the intrinsic viewpoint considers only properties of points independent of the extrinsic embedding, such as, the geodesic distance between a point pair. Both can have advantages depending on the method and application. An intrinsic viewpoint is by design invariant against certain deformations in the Euclidean embedding, like rigid transformations but also pose variations that are hard to characterize in the extrinsic view.

3.1 Neural Fields for Euclidean Space

A Euclidean neural field $\mathcal{F}_\theta^E : \mathbb{R}^m \rightarrow \mathbb{R}^o$ is a neural network that maps points in Euclidean space to vectors and is parametrized by weights $\theta \in \mathbb{R}^p$. The network is commonly chosen to be a multilayer perceptron (MLP). Let $\mathcal{M} \subset \mathbb{R}^m$ be a manifold with a Euclidean embedding into \mathbb{R}^m . Naturally, the restriction of \mathcal{F}_θ^E to \mathcal{M} leads to a neural field on a manifold: $\mathcal{F}_\theta : \mathcal{M} \rightarrow \mathbb{R}^o$, $\mathcal{F}_\theta(x) = \mathcal{F}_\theta^E(x)$.

Natural signals, such as images and scenes, are usually quite complex and contain high-frequency variations. Due to spectral bias, standard neural fields fail to learn high-frequency functions from low dimensional data [51,48] and generate blurry reconstructions. With the help of the neural tangent kernel, it was proven that the composition $\mathcal{F}_\theta^E \circ \gamma$ of a higher dimensional Euclidean neural field and a random Fourier feature (RFF) encoding γ helps to overcome the spectral bias and, consequently, enables the neural field to better represent high-frequency signals. The RFF encoding $\gamma : \mathbb{R}^m \rightarrow \mathbb{R}^d$ with $d \gg m$ is defined as

$$\gamma(\mathbf{x}) = [a_1 \cos(\mathbf{b}_1^\top \mathbf{x}), a_1 \sin(\mathbf{b}_1^\top \mathbf{x}), \dots, a_{d/2} \cos(\mathbf{b}_{d/2}^\top \mathbf{x}), a_{d/2} \sin(\mathbf{b}_{d/2}^\top \mathbf{x})], \quad (1)$$

where the coefficients $\mathbf{b}_i \in \mathbb{R}^m$ are randomly drawn from the multivariate normal distribution $\mathcal{N}(\mathbf{0}, (2\pi\sigma)^2 \mathbf{I})$. The a_i are often set to one and $\sigma > 0$ is a hyperparameter that offers a trade-off between reconstruction fidelity and overfitting.

3.2 The Laplace-Beltrami Operator

In the following, we briefly introduce the Laplace-Beltrami operator (LBO) and refer the interested reader to [42] for more details. The LBO $\Delta_{\mathcal{M}}$ is the generalization of the Euclidean Laplace operator on general closed compact manifolds. Its eigenfunctions $\phi_i : \mathcal{M} \rightarrow \mathbb{R}$ and eigenvalues $\lambda_i \in \mathbb{R}$ are the non-trivial solutions of the equation $\Delta_{\mathcal{M}} \phi_i = \lambda_i \phi_i$. The eigenvalues are non-negative and induce a natural ordering which we will use for the rest of the paper. The eigenfunctions are orthonormal to each other, build an optimal basis for the space of square-integrable functions [35], and are frequency ordered allowing a low-pass filtering by projecting onto the first k eigenfunctions. Hence, a function $f : \mathcal{M} \rightarrow \mathbb{R} \in L^2(\mathcal{M})$ can be expanded in this basis:

$$f = \sum_{i=0}^{\infty} c_i \phi_i = \sum_{i=0}^{\infty} \langle f, \phi_i \rangle \phi_i \approx \sum_{i=0}^k \langle f, \phi_i \rangle \phi_i, \quad (2)$$

where the quality of the last \approx depends on the amount of high-frequency information in f . The projection onto the LBO basis is similar to the Fourier

transform, allowing the same operations, and thus we use the LBO basis as the replacement for Fourier features. In fact, if $[0, 1]^2$ is considered as a manifold, its LBO eigenfunctions with different boundary conditions are exactly combinations of sines and cosines. Furthermore, the eigenfunctions of the LBO are identical up to sign ambiguity for isometric shapes since the LBO is entirely intrinsic.

4 Intrinsic Neural Fields

We introduce *Intrinsic Neural Fields* based on the eigenfunctions of the Laplace-Beltrami operator (LBO) which can represent detailed surface information, like texture, directly on the manifold. In the presence of prior geometric information, it is more efficient than using an extrinsic embedding which is often mainly empty. Further, this representation is naturally invariant to translation, rotation, different surface discretization, as well as isometric deformations.

Definition 1 (Intrinsic Neural Field). *Let $\mathcal{M} \subset \mathbb{R}^m$ be a closed, compact manifold and ϕ_1, \dots, ϕ_d be the first d Laplace-Beltrami eigenfunctions of \mathcal{M} . We define an **intrinsic neural field** $\mathcal{F}_\theta : \mathcal{M} \rightarrow \mathbb{R}^o$ as*

$$\mathcal{F}_\theta(\mathbf{p}) = (f_\theta \circ \gamma)(\mathbf{p}) = f_\theta(a_1\phi_1(\mathbf{p}), \dots, a_d\phi_d(\mathbf{p})). \quad (3)$$

where $\gamma : \mathcal{M} \rightarrow \mathbb{R}^d, \gamma(\mathbf{p}) = (a_1\phi_1(\mathbf{p}), \dots, a_d\phi_d(\mathbf{p}))$, with $a_i \geq 0$ and $\lambda_i = \lambda_j \Rightarrow a_i = a_j$, is our embedding function and $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^o$ represents a neural network with weights $\theta \in \mathbb{R}^p$.

Within this work, we will use $a_i = 1$, which has proven sufficient in praxis, and multilayer perceptrons (MLPs) for f_θ , as this architectural choice is common for Euclidean neural fields [59]. A detailed description of the architecture can be found in the supplementary material. It is possible to choose different embedding functions γ but we choose the LBO eigenfunctions as they have nice theoretical properties (see Section 4.1) and are directly related to Fourier features.

In Fig. 2, we apply intrinsic neural fields to the task of signal reconstruction on a 1D manifold to give an intuition about how it works and what its advantages are. The results show that the neural tangent kernel (NTK) for intrinsic neural fields exhibits favorable properties, which we prove in Sec. 4.1. We show that we can represent high-frequency signals on manifold surfaces that go far beyond the discretization level. In Sec. 5, we apply the proposed intrinsic neural fields to a variety of tasks including texture reconstruction from images, texture transfer between shapes without retraining, and view-dependent appearance modeling.

4.1 Theory

In this section, we prove that the embedding function γ proposed in Definition 1 generalizes the stationarity result of [51] to certain manifolds. Stationarity is a desirable property if the kernel is used for interpolation, for example, in novel view synthesis [51, App. C]. Fourier features induce a stationary (shift-invariant)

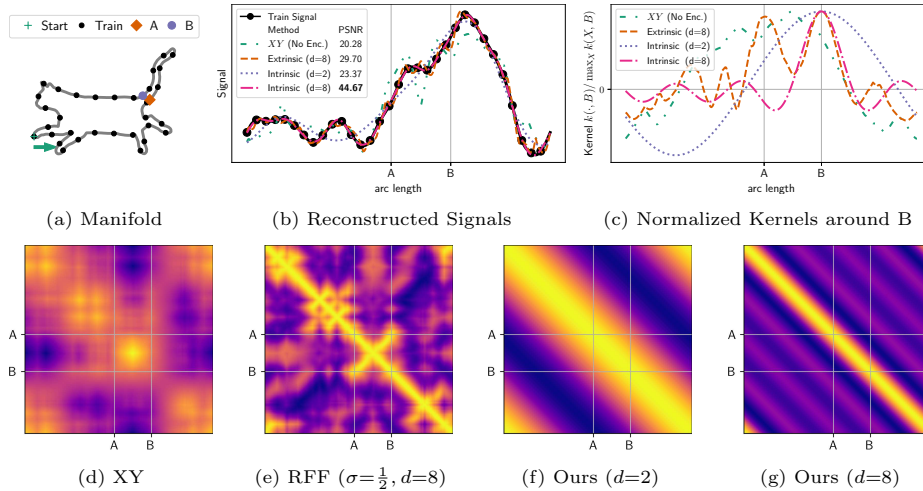


Fig. 2: Signal reconstruction. (2b) The target is sampled at 32 points and MLPs with three layers, 1024 channels, and different embeddings are trained using L^2 loss. The intrinsic neural field with $d=8$ eigenfunctions performs best. Using only two eigenfunctions with the extrinsic embedding and random Fourier features (RFF) [51] can capture the high-frequency details, but introduces artifacts when the Euclidean distance is not a good approximation of the geodesic distance, for example, at points A & B. (2d-2g) The second row of subfigures shows the pairwise neural tangent kernel (NTK) [20,31] between all points on the manifold. (2d) The NTK using the extrinsic Euclidean embedding is not maximal along the diagonal. (2e) For the NTK with RFF embedding the maximum is at the diagonal because each point’s influence is maximal onto itself. However, it has many spurious correlations between points that are close in Euclidean space but not along the manifold, for example, around B. (2f,2g) The NTK with our intrinsic embedding is localized correctly and is stationary (c.f. Thm. 1), which makes it most suitable for interpolation.



Fig. 3: Neural tangent kernels (NTKs) [20,31] with different embedding functions. The source S_1 lies directly inside the ear of the cat. (3a) The NTK using the extrinsic Euclidean embedding is not maximal at the source. (3b) The NTK using random Fourier features (RFF) [51] is localized correctly, but shows wrong behavior on the cat’s body. (3c) The NTK with our intrinsic embedding is localized correctly and adapts to the local and global geometry. (3d,3e) Additionally, the NTK with our intrinsic embedding is nearly shift-invariant, if the local geometry is approximately Euclidean: When the source is shifted from S_2 to S_3 the kernel is approximately shifted as well.

neural tangent kernel (NTK). Namely, the composed NTK for two points in Euclidean space $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$ is given by $k_{\text{NTK}}(\mathbf{x}, \mathbf{y}) = (\text{h}_{\text{NTK}} \circ \text{h}_\gamma)(\mathbf{x} - \mathbf{y})$ where $\text{h}_{\text{NTK}} : \mathbb{R} \rightarrow \mathbb{R}$ is a scalar function related to the NTK of the MLP and $\text{h}_\gamma : \mathbb{R}^m \rightarrow \mathbb{R}$ is a scalar function related to the Fourier feature embedding [51, Eqn. 7,8]. Extending this result to $\mathbf{p}, \mathbf{q} \in \mathcal{M}$ on a manifold is challenging because the point difference $\mathbf{p} - \mathbf{q}$ and, therefore, stationarity is not defined intrinsically.

Stationarity on Manifolds. While one could use the Euclidean embedding of the manifold to define the difference $\mathbf{p} - \mathbf{q}$, this would ignore the local connectivity and can change under extrinsic deformations. Instead, we use an equivalent definition from Bochner’s theorem which states that in Euclidean space any continuous, *stationary* kernel is the Fourier transform of a non-negative measure [39, Thm. 1]. This definition can be directly used on manifolds, and we define a kernel $k : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ to be **stationary** if it can be written as

$$k(\mathbf{p}, \mathbf{q}) = \sum_i \hat{k}(\lambda_i) \phi_i(\mathbf{p}) \phi_i(\mathbf{q}), \quad \hat{k}(\lambda_i) \geq 0 \quad \forall i, \quad (4)$$

where the function $\hat{k} : \mathbb{R} \rightarrow \mathbb{R}_0^+$ is akin to the Fourier transform. This implies that $\hat{k}(\lambda_i)$ and $\hat{k}(\lambda_j)$ for identical eigenvalues $\lambda_i = \lambda_j$ must be identical.

First, we want to point out that for inputs with $\|\mathbf{x}\| = \|\mathbf{y}\| = r$ the result of $k_{\text{NTK}}(\mathbf{x}, \mathbf{y}) = \text{h}_{\text{NTK}}(\langle \mathbf{x}, \mathbf{y} \rangle)$ shown by [20] for $r = 1$ and used in [51] still holds. We include this slight extension as Lemma 1 in the suppl. This is a prerequisite for the following theorem which requires the same setting as used in [20].

Theorem 1. *Let \mathcal{M} be \mathbb{S}^n or a closed 1-manifold. Let $(\lambda_i, \phi_i)_{i=1, \dots, d}$ be the positive, non-decreasing eigenvalues with associated eigenfunctions of the Laplace-Beltrami operator on \mathcal{M} . Let $a_i \geq 0$ be coefficients s.t. $\lambda_i = \lambda_j \Rightarrow a_i = a_j$, which define the embedding function $\gamma : \mathcal{M} \rightarrow \mathbb{R}^d$ with $\gamma(\mathbf{p}) = (a_1 \phi_1(\mathbf{p}), \dots, a_d \phi_d(\mathbf{p}))$. Then, the composed neural tangent kernel $k_{\text{NTK}} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ of an MLP with the embedding γ is stationary as defined in Eq. 4.*

Proof. Let $\mathcal{M} = \mathbb{S}^n$ and let \mathbf{H}_l^n be the space of degree l spherical harmonics on \mathbb{S}^n . Let $Y_{lm} \in \mathbf{H}_l^n$ be the m -th real spherical harmonic of degree l with $m = 1, \dots, \dim \mathbf{H}_l^n$. Notice that the spherical harmonics are the eigenfunctions of the LBO. We will use j to linearly index the spherical harmonics and $l(j)$ for the degree. Spherical harmonics of the same degree have the same eigenvalues, thus we use $c_{l(j)} = a_j = a_i$ for $\lambda_i = \lambda_j$ to denote the equal coefficients for same degree harmonics. First, the norm of the embedding function is constant:

$$\|\gamma(\mathbf{q})\|^2 = \sum_j c_{l(j)}^2 \phi_j^2(\mathbf{q}) = \sum_l c_l^2 \sum_{m=1}^{\dim \mathbf{H}_l^n} Y_{lm}^2(\mathbf{q}) \stackrel{(a)}{=} \sum_l c_l^2 Z_l(\mathbf{q}, \mathbf{q}) \stackrel{(b)}{=} \text{const.} \quad (5)$$

Here, $Z_l(\mathbf{q}, \mathbf{q})$ is the degree l zonal harmonic and (a,b) are properties of zonal harmonics [13, Lem. 1.2.3, Lem. 1.2.7]. Due to Eq. 5 and Lemma 1 (supp.)

$k_{\text{NTK}}(\gamma(\mathbf{p}), \gamma(\mathbf{q})) = h_{\text{NTK}}(\langle \gamma(\mathbf{p}), \gamma(\mathbf{q}) \rangle) \forall \mathbf{p}, \mathbf{q} \in \mathcal{M}$ holds. It follows:

$$\langle \gamma(\mathbf{p}), \gamma(\mathbf{q}) \rangle = \sum_j c_{l(j)}^2 \phi_j(\mathbf{p}) \phi_j(\mathbf{q}) = \sum_l c_l^2 \sum_{m=1}^{\dim \mathbf{H}_l^n} Y_{lm}(\mathbf{p}) Y_{lm}(\mathbf{q}) \quad (6)$$

$$\stackrel{(c)}{=} \sum_l c_l^2 Z_{\mathbf{p}}^l(\mathbf{q}) \stackrel{(d)}{=} \sum_l c_l^2 (1 + l/\alpha) C_l^\alpha(\langle \mathbf{p}, \mathbf{q} \rangle), \quad (7)$$

where $C_l^\alpha : [-1, 1] \rightarrow \mathbb{R}$ are the Gegenbauer polynomials which are orthogonal on $[-1, 1]$ for the weighting function $w_\alpha(z) = (1 - z^2)^{\alpha-1/2}$ with $\alpha = (n - 1)/2$ [13, B.2]. Equality (c) holds again due to [13, Lem. 1.2.3]. Equality (d) holds due to a property of Gegenbauer polynomials [13, Thm. 1.2.6], here $\langle \mathbf{p}, \mathbf{q} \rangle$ denotes the extrinsic Euclidean inner product. For the composed NTK we obtain

$$k_{\text{NTK}}(\gamma(\mathbf{p}), \gamma(\mathbf{q})) = h_{\text{NTK}} \left(\sum_l c_l^2 (1 + l/\alpha) C_l^\alpha(\langle \mathbf{p}, \mathbf{q} \rangle) \right). \quad (8)$$

We see that $k_{\text{NTK}}(\gamma(\mathbf{p}), \gamma(\mathbf{q}))$ is a function depending only on $\langle \mathbf{p}, \mathbf{q} \rangle$. Because the Gegenbauer polynomials are orthogonal on $[-1, 1]$, this function can be expanded with coefficients $\hat{c}_l \in \mathbb{R}$, which yields

$$k_{\text{NTK}}(\gamma(\mathbf{p}), \gamma(\mathbf{q})) = \sum_l \hat{c}_l (1 + l/\alpha) C_l^\alpha(\langle \mathbf{p}, \mathbf{q} \rangle) = \sum_l \hat{c}_l Z^l(\mathbf{p}, \mathbf{q}) \quad (9)$$

$$= \sum_l \hat{c}_l \sum_{m=1}^{\dim \mathbf{H}_l^n} Y_{lm}(\mathbf{p}) Y_{lm}(\mathbf{q}) = \sum_j \hat{c}_{l(j)} \phi_j(\mathbf{p}) \phi_j(\mathbf{q}). \quad (10)$$

The coefficients $\hat{c}_{l(j)}$ are non-negative as a consequence of the positive definiteness of the NTK [20, Prop. 2] and a classic result by Schoenberg [13, Thm. 14.3.3]. This shows that $k_{\text{NTK}}(\gamma(\mathbf{p}), \gamma(\mathbf{q}))$ is stationary as defined in Equation 4. \square

The adapted proof for 1-manifolds can be found in the supplementary. A qualitative example of the stationary kernels can be seen in Fig. 2. The theorem does not hold for general manifolds, however, our experiments with different γ (Tab. 1) indicate the eigenfunctions are still a superior choice for complex manifolds. We leave the theoretical explanation for this behaviour to future work.

5 Experiments

We refer to the supplementary material for all experimental details and hyperparameter settings as well as further results. To facilitate fair comparisons, all methods use the same hyperparameters like learning rate, optimizer, number of training epochs, and MLP architecture except when noted otherwise. For baselines using random Fourier features (RFF), we follow [51] and tune the standard deviation σ (c.f. Eq. 1) of the random frequency matrix to obtain optimal results.

5.1 Texture Reconstruction from Images

To investigate the representation power of the proposed intrinsic neural fields, we consider the task of texture reconstruction from posed images as proposed by Oechsle et al. [32] in Tab. 1 and Fig. 4. The input to our algorithms is a set of five 512×512 images with their camera poses and the triangle mesh of the shape. After fitting the intrinsic neural field to the data, we render images from 200 novel viewpoints and compare them to ground-truth images for evaluation.

For each pixel, we perform ray mesh intersection between the ray through the pixel and the mesh. The eigenfunctions of the Laplace-Beltrami operator are defined only on vertices of the mesh [45]. Within triangles, we use barycentric interpolation. We employ the mean L^1 loss across a batch of rays and the RGB color channels. The eigenfunction computation and ray-mesh intersection are performed once at the start of training. Hence, our training speed is similar to the baseline method that uses random Fourier features. Training takes approx. one hour on an Nvidia Titan X with 12 GB memory.

Comparison with State of the Art Methods. We compare against Texture Fields [32] enhanced with random Fourier features (RFF) [51]. Additionally, we compare against NeuTex [58], which uses a network to map a shape to the sphere and represents the texture on this sphere. We adapt NeuTex s.t. it takes advantage of the given geometry, see supplementary. Tab. 1 and Fig. 4 show that intrinsic neural fields can reconstruct texture with state-of-the-art quality. This is also true if the number of training epochs is decreased from 1000 to 200.

Ablation Study. We investigate the effect of different hyperparameters on the quality of the intrinsic neural texture field. The results in Tab. 2 show that the number of eigenfunctions is more important than the size of the MLP, which is promising for real-time applications. A model using only 64 eigenfunctions and $17k$ parameters³ still achieves a PSNR of 29.20 for the cat showing that intrinsic neural fields can be a promising approach for compressing manifold data. Additionally, we test the importance of the choice of Laplace-Beltrami eigenfunctions as γ for the results. Tab. 1 shows that popular point descriptors [50,43], that achieve great results in the difficult task of shape matching, perform worse within our framework. This indicates that, although we were not able to prove it, an extension of Thm. 1 likely holds on 2D manifolds.

5.2 Discretization-agnostic Intrinsic Neural Fields

For real-world applications, it is desirable that intrinsic neural fields can be trained for different discretizations of the same manifold. First, the training process of the intrinsic neural field should be robust to the sampling in the discretization. Second, it would be beneficial if an intrinsic neural field trained on one discretization could be transferred to another, which we show in Sec. 5.3.

³ For reference: a 80×80 3-channel color texture image has over $17k$ pixel values.

Table 1: Texture reconstruction from images. Our intrinsic neural fields show state-of-the-art performance (*first block*), even with fewer training epochs (*Ep. ↓*, *second block*). For a fair comparison, we improve Texture Fields by employing the same MLP architecture as our model and by using random Fourier features (*TF+RFF*). NeuTex has more parameters than our model but we increase the embedding size (*Em. ↑*) to match. We adapt NeuTex to take advantage of the given geometry (see supplementary). The methods are evaluated on novel views using PSNR, DSSIM [56], and LPIPS [63]. DSSIM and LPIPS are scaled by 100. The intrinsic representation shows better results than the extrinsic representation (*TF+RFF*) as well as when mapping to a textured sphere (*NeuTex*). The last block shows the performance of our method when using point descriptors HKS [50] and SHOT [43] instead of the proposed γ input.

		Em.	Ep.	cat			human		
				PSNR ↑	DSSIM ↓	LPIPS ↓	PSNR ↑	DSSIM ↓	LPIPS ↓
NeuTex [58]		63	1000	31.60	0.242	0.504	29.49	0.329	0.715
NeuTex Em. ↑		1023	1000	31.96	0.212	0.266	29.22	0.306	0.669
TF+RFF ($\sigma=4$) [32,51]		1023	1000	33.86	0.125	0.444	32.04	0.130	0.420
TF+RFF ($\sigma=16$)		1023	1000	34.19	0.105	0.167	31.53	0.193	0.414
TF+RFF ($\sigma=8$)		1023	1000	34.39	0.097	0.205	32.26	0.129	0.336
Intrinsic (Ours)		1023	1000	34.82	0.095	0.153	32.48	0.121	0.306
NeuTex Ep. ↓		1023	200	30.96	0.290	0.355	28.02	0.418	0.900
TF+RFF ($\sigma=8$) Ep. ↓		1023	200	34.07	0.116	0.346	31.85	0.142	0.427
Intrinsic (Ours) Ep. ↓		1023	200	34.79	0.100	0.196	32.37	0.126	0.346
Ours (HKS)		352	1000	23.40	1.219	2.877	22.26	0.904	2.347
Ours (SHOT)		352	1000	26.44	0.780	1.232	28.04	0.421	0.965
Ours (Efcts.)		352	1000	34.19	0.119	0.345	31.63	0.150	0.489

To quantify the discretization dependence of intrinsic neural fields, we follow the procedure proposed by Sharp et al. [44, Sec. 5.4] and rediscretize the meshes used in Sec. 5.1. The qualitative results in Fig. 5 and the quantitative results in Tab. 3 show that intrinsic neural fields work across various discretizations. Furthermore, Fig. 6 shows that transferring pre-trained intrinsic neural fields across discretizations is possible with minimal loss in visual quality.

5.3 Intrinsic Neural Field Transfer

One advantage of the Laplace-Beltrami operator is its invariance under isometries which allows transferring a pre-trained intrinsic neural field from one manifold to another. However, this theoretic invariance does not hold completely in practice, for example, due to discretization artifacts [21]. Hence, we employ functional maps [33] computed with Smooth Shells [15] to correct the transfer of eigenfunctions from source to target shape, see Fig. 6 for results. Specifically, transfer is possible between different discretizations, deformations [49] of the same shape, and even shapes from different categories. It is, of course, possible to generate similar results with extrinsic fields by calculating a point-to-point correspondence and mapping the coordinate values. However, functional maps are naturally low-dimensional, continuous, and differentiable. This makes them a beneficial choice in many applications, especially related to learning.

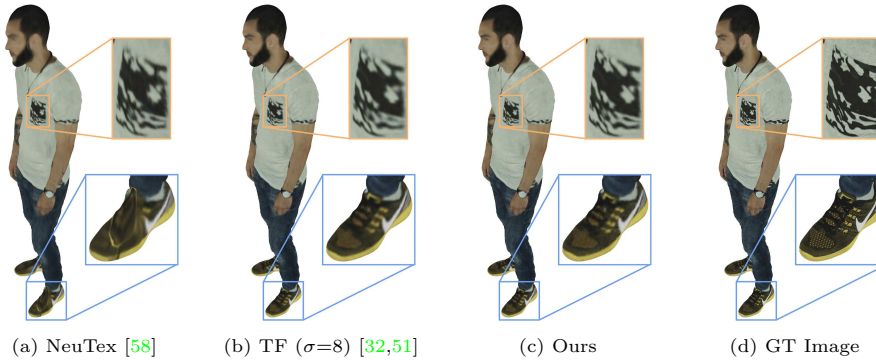


Fig. 4: Texture reconstruction from images. (4a) NeuTex uses a network to map from the shape to the sphere and represents the texture on the sphere, which yields distortions around the shoe. (4b) Texture Fields (TF) [32] with random Fourier Features (RFF) [51] learns the texture well and only around the breast pocket our method shows slightly better results. (4c) Intrinsic neural fields can reconstruct texture from images with state-of-the-art quality, which we show quantitatively in Tab. 1.




Fig. 5: Discretization-agnostic intrinsic neural fields. Our method produces identical results for a variety of triangular meshings and even point cloud data. For the point cloud, we use local triangulations [45, Sec. 5.7] for ray-mesh intersection. Pre-trained intrinsic neural fields can be transferred across discretizations as shown in Fig. 6.




Fig. 6: Intrinsic neural field transfer. (6a) The pre-trained intrinsic neural texture field from the source mesh is transferred to the target shapes using functional maps [33,15]. (6b,6c) The transfer across discretization (c.f. Fig. 5) and deformation gives nearly perfect visual quality. (6d,6e) As a proof of concept, we show artistic transfer to a different cat shape and a dog shape from the TOSCA dataset [8]. Both transfers work well but the transfer to the dog shows small visual artifacts in the snout area due to locally different geometry. Overall, the experiment shows the advantage of the intrinsic formulation which naturally incorporates field transfer through functional maps.

Table 2: Ablation study based on the texture reconstruction experiment (c.f. Sec. 5.1). The number of eigenfunctions is more important than the size of the MLP which is promising for real-time applications. A model using only 64 eigenfunctions and only 17k parameters still achieves a PSNR of 29.20 for the cat, which shows that intrinsic neural fields can be a promising approach for compressing manifold data.

	#Params	# ϕ	cat			human		
			PSNR \uparrow	DSSIM \downarrow	LPIPS \downarrow	PSNR \uparrow	DSSIM \downarrow	LPIPS \downarrow
Full model	329k	1023	34.82	0.095	0.153	32.48	0.121	0.306
Smaller MLP	140k	1023	34.57	0.108	0.205	32.20	0.134	0.379
Fewer eigenfunctions	83k	64	31.18	0.284	0.927	28.95	0.312	1.090
Smaller MLP & fewer efs	17k	64	29.20	0.473	1.428	26.72	0.493	1.766
Just 4 eigenfunctions	68k	4	22.84	1.367	3.299	20.60	1.033	2.756

Table 3: Discretization-agnostic intrinsic neural fields. We employ the procedure proposed by Sharp et al. [44, Sec. 5.4] to generate different discretizations of the original meshes (*orig*): uniform isotropic remeshing (*iso*), densification around random vertices (*dense*), refinement and subsequent quadric error simplification [17] (*qes*), and point clouds sampled from the surfaces with more points than vertices (*cloud* \uparrow) and with fewer points (*cloud* \downarrow). The discretizations are then used for texture reconstruction as in Sec. 5.1. For the point clouds, we use local triangulations [45, Sec. 5.7] for ray-mesh intersection. This table and the qualitative results in Fig. 5 show that intrinsic neural fields can be trained for a wide variety of discretizations. Furthermore, pre-trained intrinsic neural fields can be transferred across discretizations as shown in Fig. 6.

Method	cat						human					
	orig	iso	dense	qes	cloud \uparrow	cloud \downarrow	orig	iso	dense	qes	cloud \uparrow	cloud \downarrow
PSNR \uparrow	34.82	34.85	34.74	35.07	34.91	33.17	32.48	32.63	32.57	32.49	32.45	31.99
DSSIM \downarrow	0.095	0.093	0.096	0.096	0.096	0.130	0.121	0.117	0.120	0.121	0.123	0.135
LPIPS \downarrow	0.153	0.152	0.159	0.147	0.152	0.220	0.306	0.300	0.301	0.297	0.307	0.323

5.4 Real-world Data & View Dependence

We validate the effectiveness of intrinsic neural fields in a real-world setting on the BigBIRD dataset [47]. The dataset provides posed images and reconstructed meshes, and we apply a similar pipeline as in Sec. 5.1. However, the objects here are not perfectly Lambertian, and thus, view dependence must be considered. For this, we use viewing direction as an additional input to the network, as done in [29]. At first glance, using the viewing direction in its extrinsic representation opposes our intrinsic definition of neural fields. However, view dependence arises from extrinsic effects, such as lighting, which cannot be represented purely intrinsically. Fig. 7 shows that we can reconstruct high-quality textures from real-world data with imprecise calibration and meshes, quantitative results are shown in the supplementary. Decomposing the scene into intrinsic properties of the object, like the BRDF, and the influence of the environment, like light sources, is an interesting future application for our method, similar to what has been done in the context of neural radiance fields [10,62,6,57,7].



Fig. 7: Texture reconstruction from real-world data. (7b,7e) Intrinsic neural fields can reconstruct high quality textures from the real-world BigBIRD dataset [47] with imprecise calibration and imprecise meshes. (7a,7d) The baseline texture mapped meshes provided in the dataset show notable seams due to the non-Lambertian material, which are not present in our reconstruction that utilizes view dependence as proposed by [29].

6 Conclusion

Discussion. The proposed intrinsic formulation of neural fields outperforms the extrinsic formulation in the presented experiments. However, if the data is very densely sampled, and the kernel is thus locally limited, the extrinsic method can overcome many of its weaknesses. In practice, dense sampling often leads to increased runtime of further processing steps, and thus we consider our intrinsic approach still superior. Further, we provided the proof for a stationary NTK on n -spheres. Our experiments and intuition imply that even for general manifolds, it is advantageous how the NTK takes local geometry into account. The details leave an interesting direction for further theoretical analysis.

Conclusion. We present intrinsic neural fields, an elegant and direct generalization of neural fields for manifolds. Intrinsic neural fields can represent high-frequency functions on manifold surfaces independent of discretization by making use of the Laplace-Beltrami eigenfunctions. We introduce a new definition for stationary kernels on manifolds, and our theoretic analysis shows that the derived neural tangent kernel is stationary under specific conditions. We conduct experiments to investigate the capabilities of our framework on the application of texture reconstruction from a limited number of views. Furthermore, the learned functions can be transferred to new examples using functional maps without any retraining, and view-dependent changes can be incorporated. Intrinsic neural fields outperform competing methods in all settings. Additionally, they add flexibility, especially in settings with deformable objects due to the intrinsic nature of our approach.

Acknowledgements. We thank Florian Hofherr, Simon Klenk, Dominik Muhle and Emanuele Rodolà for useful discussions and proofreading. ZL is funded by a KI-Starter grant from the Ministerium für Kultur und Wissenschaft NRW.

References

1. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: IEEE International Conference on Computer Vision (ICCV) (2011)
2. Baatz, H., Granskog, J., Papas, M., Rousselle, F., Novák, J.: Nerf-tex: Neural reflectance field textures. In: Eurographics Symposium on Rendering (EGSR) (2021)
3. Basri, R., Galun, M., Geifman, A., Jacobs, D.W., Kasten, Y., Kritchman, S.: Frequency bias in neural networks for input of non-uniform density. In: International Conference on Machine Learning (ICML) (2020)
4. Benbarka, N., Höfer, T., ul Moqet Riaz, H., Zell, A.: Seeing implicit neural representations as fourier series. In: IEEE Winter Conference of Applications on Computer Vision (WACV) (2022)
5. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.M., Cremers, D.: Anisotropic diffusion descriptors. In: Computer Graphics Forum (CGF). vol. 35 (2016)
6. Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerf: Neural reflectance decomposition from image collections. In: IEEE International Conference on Computer Vision (ICCV) (2021)
7. Boss, M., Jampani, V., Braun, R., Liu, C., Barron, J.T., Lensch, H.P.A.: Neural-pil: Neural pre-integrated lighting for reflectance decomposition. CoRR **abs/2110.14373** (2021)
8. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer Science & Business Media (2008)
9. Burghard, O., Dieckmann, A., Klein, R.: Embedding shapes with green’s functions for global shape matching. Computers & Graphics **68C** (2017)
10. Chen, Z., Nobuhara, S., Nishino, K.: Invertible neural BRDF for object inverse rendering. In: European Conference on Computer Vision (ECCV) (2020)
11. Chibane, J., Pons-Moll, G.: Implicit feature networks for texture completion from partial 3d data. In: European Conference on Computer Vision (ECCV) SHARP Workshop (2020)
12. Crane, K., Weischedel, C., Wardetzky, M.: The heat method for distance computation. Communications of the ACM **60** (2017)
13. Dai, F., Xu, Y.: Approximation theory and harmonic analysis on spheres and balls. Springer (2013)
14. Dwivedi, V.P., Bresson, X.: A generalization of transformer networks to graphs. In: AAAI Workshop on Deep Learning on Graphs: Methods and Applications (2021)
15. Eisenberger, M., Löhner, Z., Cremers, D.: Smooth shells: Multi-scale shape registration with functional maps. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
16. Gargan, D., Neelamkavil, F.: Approximating reflectance functions using neural networks. In: Eurographics Workshop on Rendering Techniques (1998)
17. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) (1997)
18. Hertz, A., Perel, O., Giryas, R., Sorkine-Hornung, O., Cohen-Or, D.: Mesh draping: Parametrization-free neural mesh transfer. CoRR **abs/2110.05433** (2021)
19. Hertz, A., Perel, O., Giryas, R., Sorkine-Hornung, O., Cohen-Or, D.: Sape: Spatially-adaptive progressive encoding for neural optimization. In: Conference on Neural Information Processing Systems (NeurIPS) (2021)

20. Jacot, A., Hongler, C., Gabriel, F.: Neural tangent kernel: Convergence and generalization in neural networks. In: Conference on Neural Information Processing Systems (NeurIPS) (2018)
21. Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* **32** (2013)
22. Lee, J., Jin, K.H.: Local texture estimator for implicit representation function. *CoRR* [abs/2111.08918](#) (2021)
23. Liu, X., Donate, A., Jemison, M., Mio, W.: Kernel functions for robust 3d surface registration with spectral embeddings. In: International Conference on Pattern Recognition (ICPR) (2008)
24. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics (TOG)* **34** (2015)
25. Marin, R., Cosmo, L., Melzi, S., Rampini, A., Rodolá, E.: Spectral geometry in practice. *3DV Tutorial* (2021)
26. Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. *CoRR* [abs/2104.03960](#) (2021)
27. Meronen, L., Trapp, M., Solin, A.: Periodic activation functions induce stationarity. *CoRR* [abs/2110.13572](#) (2021)
28. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. *CoRR* [abs/2112.03221](#) (2021)
29. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision (ECCV) (2020)
30. Morreale, L., Aigerman, N., Kim, V.G., Mitra, N.J.: Neural surface maps. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
31. Novak, R., Xiao, L., Hron, J., Lee, J., Alemi, A.A., Sohl-Dickstein, J., Schoenholz, S.S.: Neural tangents: Fast and easy infinite neural networks in python. In: International Conference on Learning Representations (ICLR) (2020)
32. Oechsle, M., Mescheder, L.M., Niemeyer, M., Strauss, T., Geiger, A.: Texture fields: Learning texture representations in function space. In: IEEE International Conference on Computer Vision (ICCV) (2019)
33. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* **31** (2012)
34. Palafox, P., Bozic, A., Thies, J., Nießner, M., Dai, A.: Neural parametric models for 3d deformable shapes. In: IEEE International Conference on Computer Vision (ICCV) (2021)
35. Parlett, B.N.: The symmetric eigenvalue problem. *Siam* (1998)
36. Peng, L.W., Shamsuddin, S.M.H.: 3d object reconstruction and representation using neural networks. In: International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia (GRAPHITE) (2004)
37. Piperakis, E., Kumazawa, I.: Affine transformations of 3d objects represented with neural networks. In: IEEE International Conference on 3-D Digital Imaging and Modeling (2001)
38. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F.A., Bengio, Y., Courville, A.C.: On the spectral bias of neural networks. In: International Conference on Machine Learning (ICML) (2019)
39. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: Conference on Neural Information Processing Systems (NeurIPS) (2007)

40. Ramasinghe, S., Lucey, S.: Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. *CoRR* **abs/2111.15135** (2021)
41. Ramasinghe, S., Lucey, S.: Learning positional embeddings for coordinate-mlps. *CoRR* **abs/2112.11577** (2021)
42. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: *Symposium on Geometry Processing (SGP)* (2007)
43. Salti, S., Tombari, F., Di Stefano, L.: Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* (2014)
44. Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)* **XX** (20XX)
45. Sharp, N., Crane, K.: A laplacian for nonmanifold triangle meshes. *Computer Graphics Forum* **39** (2020)
46. Sharp, N., Soliman, Y., Crane, K.: The vector heat method. *ACM Transactions on Graphics (TOG)* **38** (2019)
47. Singh, A., Sha, J., Narayan, K.S., Achim, T., Abbeel, P.: Bigbird: A large-scale 3d database of object instances. In: *IEEE International Conference on Robotics and Automation (ICRA)* (2014)
48. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: *Conference on Neural Information Processing Systems (NeurIPS)* (2020)
49. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: *Symposium on Geometry Processing (SGP)* (2007)
50. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Symposium on Geometry Processing (SGP)* (2009)
51. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: *Conference on Neural Information Processing Systems (NeurIPS)* (2020)
52. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P.P., Treitsch, E., Wang, Y., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., Simon, T., Theobalt, C., Nießner, M., Barron, J.T., Wetzstein, G., Zollhöfer, M., Golyanik, V.: Advances in neural rendering. *CoRR* **abs/2111.05849** (2021)
53. Vaxman, A., Ben-Chen, M., Gotsman, C.: A multi-resolution approach to heat kernels on discrete surfaces. *ACM Transactions on Graphics (TOG)* **29** (2010)
54. Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Cremers, D.: Efficient deformable shape correspondence via kernel matching. In: *International Conference on 3D Vision (3DV)* (2017)
55. Wang, P., Liu, Y., Yang, Y., Tong, X.: Spline positional encoding for learning 3d implicit signed distance fields. In: *International Joint Conference on Artificial Intelligence (IJCAI)* (2021)
56. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13** (2004)
57. Wimbauer, F., Wu, S., Rupprecht, C.: De-rendering 3d objects in the wild. *CoRR* **abs/2201.02279** (2022)
58. Xiang, F., Xu, Z., Hasan, M., Hold-Geoffroy, Y., Sunkavalli, K., Su, H.: Neutex: Neural texture mapping for volumetric neural rendering. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021)

59. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond (2021)
60. Yifan, W., Rahmann, L., Sorkine-hornung, O.: Geometry-consistent neural shape representation with implicit displacement fields. In: International Conference on Learning Representations (ICLR) (2022)
61. Yüce, G., Ortiz-Jiménez, G., Besbinar, B., Frossard, P.: A structured dictionary perspective on implicit neural representations. CoRR **abs/2112.01917** (2021)
62. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
63. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
64. Zheng, J., Ramasinghe, S., Lucey, S.: Rethinking positional encoding. CoRR **abs/2107.02561** (2021)