

MeshFeat: Multi-Resolution Features for Neural Fields on Meshes

Mihir Mahajan^{*1} , Florian Hofferr^{*1,2} , and Daniel Cremers^{1,2} 

¹ Technical University of Munich

² Munich Center for Machine Learning

Abstract. Parametric feature grid encodings have gained significant attention as an encoding approach for neural fields since they allow for much smaller MLPs, which significantly decreases the inference time of the models. In this work, we propose MeshFeat, a parametric feature encoding tailored to meshes, for which we adapt the idea of multi-resolution feature grids from Euclidean space. We start from the structure provided by the given vertex topology and use a mesh simplification algorithm to construct a multi-resolution feature representation directly on the mesh. The approach allows the usage of small MLPs for neural fields on meshes, and we show a significant speed-up compared to previous representations while maintaining comparable reconstruction quality for texture reconstruction and BRDF representation. Given its intrinsic coupling to the vertices, the method is particularly well-suited for representations on deforming meshes, making it a good fit for object animation.

Keywords: Feature Encodings · Multi-Resolution · Meshes

1 Introduction

Capturing and modeling our 3D world realistically and effectively is becoming increasingly relevant for AR, VR, CGI, simulations, and computer games. Classical approaches for scene representation often rely on explicit discrete data structures like voxel grids to store function values like SDF. Alternatively, point clouds, and meshes are used to store geometry or UV-Maps, and discrete texture maps to represent appearance. For all these approaches, the memory footprint is strongly coupled to the resolution, making high-resolution representations often impractical. Since the breakthrough work by Mildenhall et al. [29], the focus has shifted towards neural implicit representations, which use Multilayer Perceptrons (MLP) to represent scenes in a continuous and resolution-free manner.

The key to high-quality reconstructions is to use input encodings to overcome the spectral bias of neural networks [34] and enable the model to learn high-frequency details. Models based on frequency encodings like positional encoding [29, 43] or Fourier features [35, 41] tend to be costly and slow to evaluate since all the information on the scene is stored in the weights of the MLP. This results

^{*} Equal contribution

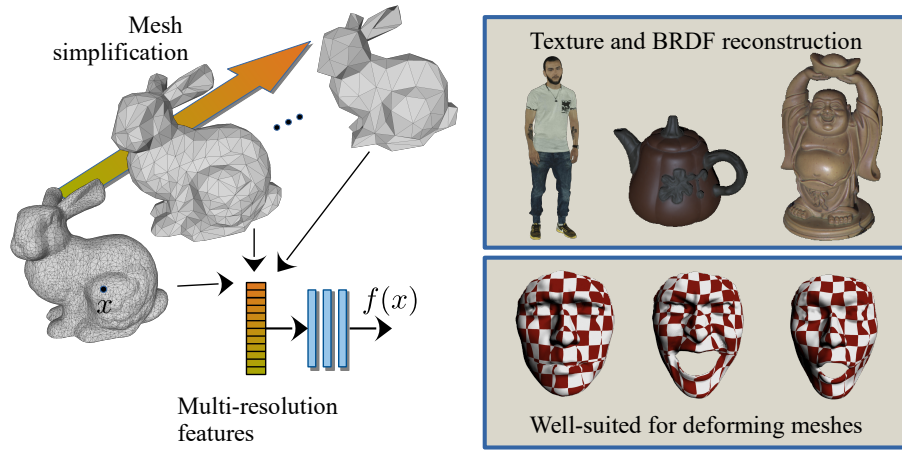


Fig. 1: We present MeshFeat, a parametric encoding strategy for neural fields on meshes. We propose a multi-resolution strategy based on mesh simplification to enhance the efficiency of the encoding. Our approach allows for much smaller MLPs than previous frequency-based encodings, resulting in significantly faster inference times. We evaluate the method for texture reconstruction and BRDF estimation and demonstrate, that the encoding is well suited to represent signals on deforming meshes.

in a strong coupling of different parts of the scene and necessitates large neural networks. On the other hand, the use of multi-resolution feature grids [4, 6, 15, 21, 24, 25, 33, 39, 42] as input encoding has led to a large gain in evaluation speed, making the models real-time capable. The idea of feature grids is to store local information about the scene in spatially distributed features and use only a small MLP to decode this information to the quantity of interest. This decoupling of spatial information from information decoding, combined with the resulting possibility of smaller networks, is the reason for the performance gain for those approaches.

While neural representations have demonstrated impressive performance, numerous 3D computer graphics pipelines continue to utilize meshes as a fundamental data structure. Meshes remain a preferred choice in many applications due to the existence of efficient algorithms, intuitive editing capabilities, and convincing animation possibilities. Consequently, a noteworthy avenue of research explores the integration of mesh-based representations with neural fields, aiming to capitalize on the strengths of both approaches. Texture fields by Oechsle *et al.* [31] embed the points on the mesh into the surrounding Euclidean space and leverages neural fields to regress texture values. NeuTex [47] and other methods [1, 41] follow the same idea, but additionally use frequency encodings to capture higher frequency details and circumvent the spectral bias issue. Unfortunately, these Euclidean embeddings do not take the mesh geometry into account, and for points where Euclidean and geodesic distance varies significantly, these methods can show bleeding artifacts [18]. Intrinsic neural fields by Koestler *et al.* [18] ad-

addresses this by transferring the idea of frequency encodings from the Euclidean space to manifolds by using the eigenfunctions of the Laplace-Beltrami Operator. This makes the encoding intrinsic to the mesh. While they show promising results in reconstruction quality, they require heavy preprocessing to calculate the eigenfunctions. Moreover, they require a large number of eigenfunctions, resulting in a substantial memory footprint of the model. All methods employ frequency-based encodings and, therefore, suffer from the requirement of large networks, hindering efficient inference and rendering.

In this work, we propose MeshFeat, a novel parametric multi-resolution feature encoding for neural fields on meshes, for which we adopt ideas of feature grids from Euclidean space. We leverage the spatial structure given by the mesh and employ a mesh simplification algorithm for our multi-resolution strategy. Similarly to feature grids in Euclidean space, decoupling spatial information and decoding allows for much smaller MLPs and significantly accelerates the inference compared to previous methods. A schematic overview of our approach can be seen in Fig. 1.

We summarize our main contributions as follows:

- We introduce MeshFeat, a parametric feature encoding for neural fields on meshes which includes an effective multi-resolution strategy based on mesh simplification.
- We demonstrate that this feature encoding allows for significantly smaller MLPs compared to previous approaches, leading to a large speed-up in inference time.
- We perform a thorough comparison of our approach to state-of-the-art methods for the tasks of single-object texture reconstruction and BRDF estimation. Moreover, we demonstrate the native applicability of our approach to signal representation on deforming meshes.

Please see our project page at <https://maharajamihir.github.io/MeshFeat/> for the source code.

2 Related Work

Neural Fields Scene representation based on coordinate-based networks has been used extensively in recent years. The idea is to use a simple multi-layer perception (MLP) to map a spatial position to quantities of interest like occupancy [5, 26, 33], signed distance [10, 32], texture [31], density and color for volume rendering [29, 45, 49] and more. In contrast to classical methods for scene representation purely based on explicit and discrete data structures like voxel grids, these implicit approaches are not limited by a discrete resolution resulting in a much more memory-efficient representation. The key to high-fidelity reconstructions with neural fields is to encode the input to a higher dimensional space to overcome the low-frequency bias of neural networks [2, 34].

Input Encodings The different input encoding strategies can be divided into two categories: frequency-based encodings and parametric approaches. Among the most popular choices for the first type are positional encodings, which originate from transformer architectures [43] and have been first introduced in the seminal NeRF paper [29] for neural fields. The idea is to use a sequence of sine and cosine functions to encode the input coordinates. Tancik *et al.* analyze these positional encodings using neural tangent kernel analysis and generalize the idea to the Fourier feature encoding [41]. A similar effect is achieved by SIRENs, which are MLP with periodic activation functions [37]. Hertz *et al.* complement this idea with an adaptive frequency masking scheme [12].

In contrast, parametric encodings do not directly apply an encoding function to the input. Instead, they use the input coordinate as an interpolation point for learnable features stored in more classical data structures like grids [4, 6, 15, 21, 24, 25, 33, 39, 42]. To evaluate the model, the grid features are interpolated at the input coordinate, and the result is used as an input for the MLP. Extending a single-resolution grid to a multi-resolution approach adds different scales of locality to the model. This decoupling of local and global information helps to improve efficiency and convergence properties. As further improvements, the use of hash grids [30] and octrees [40] have been proposed.

Parametric encodings allow for much smaller MLPs than their frequency-based counterparts, making evaluating these models significantly faster. Intuitively, feature encodings contain the information “What is at this position” and the MLP decodes this into quantities of interest. In contrast, frequency-based encodings do not contain learned information, and the MLP needs to store all the information on the scene in its weights. The smaller computational cost for parametric encodings comes at the price of a larger memory footprint.

Input Encodings for Neural Fields on Meshes Encoding strategies specifically for meshes have received much less attention than encodings for neural fields in the Euclidean space. A common strategy, used *e.g.* by Texture Fields [31], is to embed the points on the mesh into the surrounding Euclidean space. This enables the use of established encoding strategies like Fourier features, as done by Text2Mesh [27] or variants of positional encoding, as done by Hertz *et al.* [13].

The drawback of this strategy is that the embedding to the surrounding space makes the encodings *extrinsic* and the properties of the underlying manifold represented by the mesh are not considered, which can lead to artifacts [18]. Koestler *et al.* show, that the equivalent to positional encoding in the Euclidean space are the eigenfunctions of the Laplace-Beltrami Operator on manifolds [18]. Grattarola and Vandergheynst investigate a similar idea [9]. This results in an *intrinsic* encoding that respects the properties of the manifold and enables applications like texture transfer. This approach has been used successfully to model deformation fields on meshes [22, 44] and for scene stylization [14]. While intrinsic encodings are elegant, in practice, many eigenfunctions have to be stored per vertex, resulting in a large model size. Also, the computation of the eigenfunctions is costly since the eigenvalues of a large matrix need to be computed.

All approaches for encodings on meshes discussed so far follow the idea of frequency encodings and thus require large MLPs for high-quality reconstruction, resulting in slow evaluation times. In contrast, we present a parametric encoding on a mesh, enhanced by a multi-resolution strategy. Only very few previous works consider this form of encoding. Yang *et al.* [48] train features on a mesh, which they use as a scaffold for a neural scene representation in 3D. In contrast to our work, their approach does not contain hierarchical features, and moreover, they do not consider a signal restricted to the mesh surface. While Kim *et al.* propose multi-resolution features, they only consider spherical meshes [17]. In contrast, we are not restricted to a certain class of meshes.

Multi-Resolution Approaches on Meshes Multi-resolution approaches for meshes have been used for various tasks. Lee *et al.* use mesh simplification based on vertex removal for adaptive re-meshing [19]. Liu *et al.* investigate mesh simplification for multi-grid solvers on curved surfaces [23]. Jiang *et al.* propose a prismatic shell for meshes that can help to avoid geometric artifacts for mesh simplification algorithms for extreme vertex reductions at the cost of additional computational overhead [16]. MeshCNN uses a learned mesh-pooling for a multi-resolution approach to mesh analysis tasks [11].

3 Method

In this section we present MeshFeat, a parametric encoding strategy for the parametrization of neural fields directly on meshes. We adapt the idea of multi-resolution grids for the Euclidean space to meshes to obtain an efficient encoding. Instead of using a regular voxel grid in Euclidean space, we use the mesh vertices as pre-defined locations to store the feature vectors. We apply a mesh simplification algorithm to obtain different resolutions of the initial mesh for our multi-resolution approach. Similar to feature grids in Euclidean space, our approach allows to use very small MLPs resulting in a fast evaluation time.

3.1 Multi-Resolution Feature Encoding

Mesh Simplification To obtain our multi-resolution feature encoding, we simplify the initial mesh $M = (V, F)$ to multiple resolutions using the quadric error metric decimation by Garland and Heckbert [7, 8]. We denote the sequence of resolutions for the mesh simplification by $(r^{(i)})_i$, where $r^{(i)} \in [0, 1]$ means, that the i -th resolution has $|V^{(i)}| = r^{(i)}|V|$ number of vertices. This yields the sequence of meshes $((V^{(i)}, F^{(i)}))_i$. Moreover, we store the mapping

$$m^{(i)} : V \rightarrow V^{(i)} \quad v \mapsto m^{(i)}(v) = v^{(i)} \quad (1)$$

that assigns a vertex $v \in V$ in the original mesh to the vertex $v^{(i)} \in V^{(i)}$ in resolution i to which v was collapsed in the decimation algorithm. Note that in contrast to other simplification-based multi-resolution schemes, we do not compute a geometric mapping between the resolutions, which reduces the computational overhead. Please see the appendix for further information.

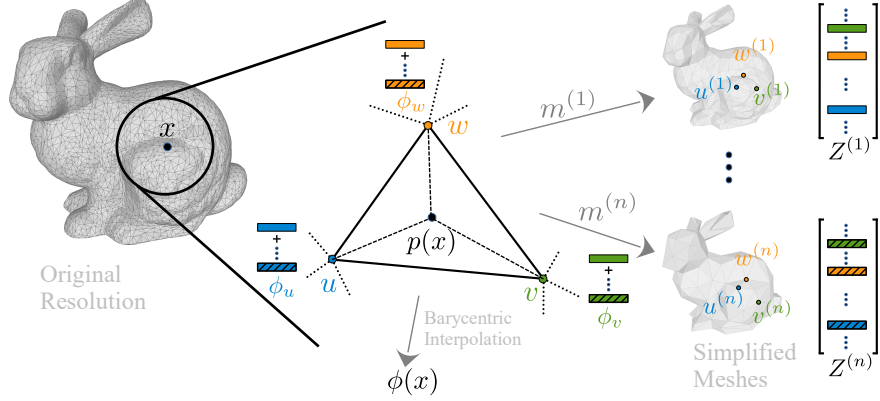


Fig. 2: Overview of our multi-resolution feature approach on the mesh. To get the feature encoding $\phi(x)$ for a point x on the original mesh, we determine the vertices u, v, w of the respective triangle. Using the mappings $m^{(i)}$, we gather the corresponding features from the different resolutions. By summing them, we obtain the features ϕ_u, ϕ_v, ϕ_w at the vertices in the original mesh. We receive the final feature encoding $\phi(x)$ by barycentric interpolation of the features at the vertices.

Multi-Resolution Strategy For each of the resolutions, we use a learnable feature matrix $Z^{(i)} \in \mathbb{R}^{|V^{(i)}| \times d}$, where $d \in \mathbb{N}$ is the feature dimension. We denote the feature vector of vertex $v^{(i)} \in V^{(i)}$ by $Z^{(i)}(v^{(i)}) \in \mathbb{R}^d$.

In contrast to parametric encodings on Euclidean grids, we do not interpolate the features for each resolution but rather accumulate the features from all resolutions on the original resolution and interpolate the results only once. The reason is that a point on the mesh for a certain resolution does not have a direct meaning on the meshes of the other resolutions. The Euclidean embedding of the point, for example, does not even need to be on the other meshes. Therefore, we use the mappings $m^{(i)}$ to “pull” all features of the coarser resolutions to the finest resolution (which is where we want to evaluate the function on the mesh). We obtain the combined feature vector ϕ_v for a vertex $v \in V$ in the original mesh by summing the feature vectors from all resolutions

$$\phi_v = \sum_i Z^{(i)}(m^{(i)}(v)). \quad (2)$$

See Fig. 2 for an overview of our multi-resolution feature pipeline.

Feature Interpolation To obtain the feature vector for an arbitrary point x on the mesh, we compute the barycentric coordinates $p(x) = [\lambda_1, \lambda_2, \lambda_3]^\top$ within the respective triangle of the original mesh. Let $v_1, v_2, v_3 \in V$ be the vertices of this triangle. Then the encoding $\phi(x)$ for the point x reads

$$\phi(x) = \sum_i \lambda_i \phi_{v_i} = [\phi_{v_1}, \phi_{v_2}, \phi_{v_3}] p(x). \quad (3)$$

3.2 Feature Regularization

Depending on the training setup, it can happen that the features for some vertices will never receive a training signal. For example, for training from images and very fine meshes, some triangles might not be hit during the ray mesh intersection leading to missing supervision for vertices in that region. We use a regularization of the features based on the mesh Laplacian to avoid the resulting artifacts. We denote the mesh Laplacian of the mesh in the original resolution by $L \in \mathbb{R}^{|V| \times |V|}$. To remove the dependency on the mesh scale and avoid the necessity to tune the regularization weight to each mesh individually, we normalize L by its spectral norm, *i.e.* we consider $\hat{L} = L / \|L\|_2$. For the regularization, we accumulate the features for all vertices of the original resolution in the matrix $\Phi \in \mathbb{R}^{|V| \times d}$ and sum the absolute values of the normalized Laplacian applied to the features, *i.e.*

$$\mathcal{L}_{\text{reg}} = \sum_{i,j} |(\hat{L}\Phi)_{i,j}|. \quad (4)$$

This can be seen as the sum over the 1-norms of the Laplacian applied to the individual entries of the latent codes. This loss term penalizes large variations of neighboring feature values since the Laplacian computes the deviation from the local average. Still, due to the 1-norm, we allow for sparse larger changes. Note that while we compute the loss term only for the original resolution, the regularization affects all resolutions since we apply the Laplacian to the accumulated features. In practice, we use the robust Laplacian by Sharp and Crane [36].

3.3 Model Architecture and Training Details

To decode the feature vectors into the final function value, we employ a standard MLP with ReLU activation functions. For the experiments in this paper, we use a hidden dimension of 32 with 2 hidden layers. Also, we employ a sigmoid output-nonlinearity since all of our reconstructed signals are in the range $[0, 1]$. To prevent overfitting, a weak L2 regularization (with weight 10^{-5}) proved to be helpful.

We found that a total of 4 resolutions $r^{(i)} \in \{1, 0.1, 0.05, 0.01\}$ for the mesh simplification achieved optimal results for all experiments. Note, that for $r^{(1)} = 1$, the mapping $m^{(1)}$ yields the identity function, *i.e.* we use the original mesh as our finest resolution. We initialize the feature matrices using a normal distribution with $\sigma = 5 \cdot 10^{-4}$.

For training, we use a batch size of $b = 8000$. We train for 1000 epochs for the texture reconstruction task and for 500 epochs for the BRDF estimation. We found that different learning rates for the MLP parameters and the features are crucial. We use $\text{lr}_\theta = 2 \cdot 10^{-4}$ for the weights of the neural network and $\text{lr}_Z = 5 \cdot 10^{-3}$ for our latent codes. We use a factor $\lambda_{\text{reg}} = 1.5 \cdot 10^{-6}$ to balance the regularization loss with the data loss.

4 Experiments

In this section, we present a detailed evaluation of our proposed approach. First, we evaluate our method for single-object texture reconstruction from multi-view images. We compare against state-of-the-art methods [18, 31, 47]. Moreover, we perform an ablation study to evaluate the significance of our modeling choices. As a second application of our method, we consider the reconstruction of a parametric BRDF from multi-view images of a single object with calibrated lighting. We use our approach to estimate the parameters of the well-known Disney BRDF [3] spatially varying on the mesh. Furthermore, we demonstrate that our method is well-suited to represent quantities on deforming meshes due to the tight coupling with the mesh. We refer to the supplementary material for additional experiments and further training information.

Concurrent Models We compare our encoding strategy to state-of-the-art work for neural fields on meshes with different encoding strategies. We use Texture Fields (TF) [31] as a method with an *extrinsic* encoding strategy. Consistent with the experiments done by Koestler *et al.* [18], we augment the method with random Fourier features (RFF). Moreover, we compare against Intrinsic Neural Fields (INF) [18], which uses an intrinsic and frequency-based encoding based on the eigenfunctions of the Laplace-Beltrami operator. Both methods use an MLP with 6 hidden layers of width 128. To keep the experiments consistent with previous work [18], we also include a modified version of NeuTex [47] in our experiments on texture reconstruction, even though the method was originally designed for geometry estimation along with texture. NeuTex follows a UV-mapping-based approach and learns separate networks for geometry, UV-mapping and texture. Following Koestler *et al.* [18], we allow it to take advantage of the given geometry. We refer the reader to their work for details on this modification. Finally, as a reference, we include a non-neural baseline, for which we learn color values directly on the vertices of the original mesh and use barycentric interpolation to obtain the values on the triangles. Again, we use our regularization term to account for unsupervised vertices. All methods are adapted into the same pipeline for a fair comparison.

4.1 Texture Reconstruction from Multi-View Images

We follow the experimental setup of Intrinsic Neural Fields [18], initially proposed by Oechsle et al. [31] and use the same dataset. The input to all methods is a set of five 512x512 pixel posed images alongside their respective intrinsic and extrinsic camera matrices. Additionally, the triangle mesh of the object is given. In the preprocessing stage, we compute the ray-mesh intersection for each pixel to obtain the corresponding point on the mesh. We use the Euclidean embedding of this point as input for TF and use barycentric coordinates in the respective triangle for the other methods. Additionally, for INF, the values of the eigenfunctions of the Laplace-Beltrami operator (LBO) must be computed at the

Table 1: Texture reconstruction from multi-view images. Our multi-resolution feature encoding significantly improves inference speed compared to state-of-the-art neural methods while maintaining similar reconstruction quality. Note that DSSIM and LPIPS are scaled by 100. The trade-off for the speed-up is a slight increase in the number of parameters; however, we still have a significantly smaller model compared to INF, which generally produces the best results. While the non-neural reference baseline has the fastest evaluation time, it shows significantly decreased reconstruction quality, particularly for the coarser mesh. For our method, using the regularizer shows a substantial improvement in reconstruction quality. While for coarser meshes like the *cat*, we obtain better reconstruction results for an increase in feature dimension ($d \uparrow$) to $d = 10$, we did observe slight overfitting for finer meshes like the *human*.

Object	Method	PSNR \uparrow	DSSIM \downarrow	LPIPS \downarrow	# Params. \downarrow	Speedup \uparrow
human $ V = 129k$	NeuTex [47]	27.32	0.549	0.954	793k	1.0x
	TF+RFF [31, 41]	32.10	0.232	0.423	331k 🏆	1.96x
	INF [18]	32.46 🏆	0.215	0.390 🏆	133130k	3.06x 🏆
	Ours ($\lambda_{reg} = 0$)	31.25	0.323	0.510	604k	13.49x
	Ours ($d \uparrow$)	32.46	0.203 🏆	0.410	1058k	12.08x
	Ours ($d = 4$)	32.51 🏆	0.202 🏆	0.400 🏆	604k 🏆	13.49x 🏆
	Non-neural (ref.)	32.01	0.225	0.432	391k	28.32x
cat $ V = 33k$	NeuTex	31.56	0.338	0.336	793k	1.0x
	TF+RFF	34.33	0.162 🏆	0.247 🏆	331k 🏆	1.96x
	INF	34.76 🏆	0.166 🏆	0.202 🏆	36430k	3.07x
	Ours ($\lambda_{reg} = 0$)	33.27	0.305	0.453	166k	13.33x
	Ours ($d \uparrow$)	34.65 🏆	0.191	0.295	411k	11.94x 🏆
	Ours ($d = 4$)	34.23	0.238	0.349	166k 🏆	13.33x 🏆
	Non-neural (ref.)	33.01	0.400	0.775	123k	32.52x

mesh vertices. We would like to point out that this computation can take *hours*, depending on the mesh size, while the other methods do not require any preprocessing. For the training, we employ an L1 loss on the color values. We render 200 images for evaluation and report PSNR, DSSIM³ [46] and LPIPS [50]. All experiments were done on an Nvidia A10G Tensor Core with 24GiB of memory.

The results in Tab. 1 and Fig. 3 show that our model achieves results on par with the state-of-the-art while showing a notable speed-up compared to previous approaches.

Reconstruction Quality Tab. 1 shows that our method yields quantitative results that are on par with the other methods for the reconstruction quality. As can be seen in Fig. 3, we even achieve a noticeable qualitative improvement for some regions, while other regions are slightly worse. We found that areas of high quality often correspond with a finer mesh in that area, which is to be expected since, for those regions, more features are available.

³ For better readability we use the Structural Dissimilarity: $DSSIM = (1 - SSIM)/2$

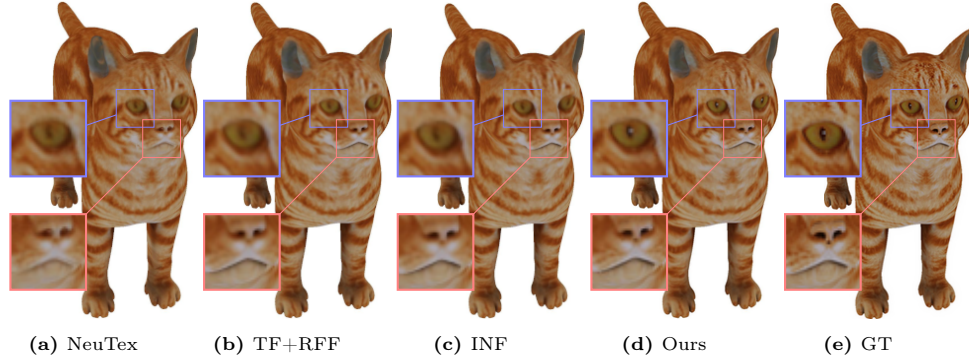


Fig. 3: Qualitative results for texture reconstruction from multi-view images on the *cat*. Our method enables high-quality reconstructions, matching state-of-the-art methods in visual fidelity while offering a significant speedup. Because the baseline methods are based on frequency encodings, they lead to an over-smoothing of intricate details around the eye, which only our method can capture. Furthermore, NeuTex is unable to capture spatially fast changing color like on the mouth of the *cat* and shows distortions inside the ear.

Inference Time To benchmark the inference speed, we measure the evaluation time of the encoding stage and the MLP for 2^{15} points for each model. Note that this does not include the ray mesh intersection or the computation of the eigenfunctions. To reduce variance, we average the results over 300 measurements of the same batch. Tab. 1 shows a significant speed-up of our method compared to the other methods, which we attribute to the much smaller MLP size. Also, it is important to note that our encoding strategy is computationally lightweight since it involves a single matrix multiplication with low dimensions due to the small feature size. As expected, the non-neural baseline is the fastest.

Number of Parameters For the number of parameters, we count all parameters that are necessary to evaluate the model. This includes learnable components like the weights of the network and the features, as well as non-learnable ones like the Fourier feature matrix and the LBO eigenfunctions values. We observe that the trade-off for the evaluation speed of our method is a slightly increased number of parameters⁴ due to the feature matrices. This observation is similar to feature-grid-based methods in the Euclidean space.

4.2 Ablation Study

Multi-Resolution Features We accumulate information distributed over multiple resolutions of our mesh to enable different scales of locality. This way, coarse resolutions can capture more global information, and fine resolutions can act as

⁴ As a reference: a 512x512 3-channel color image consists of over 786k pixel values.

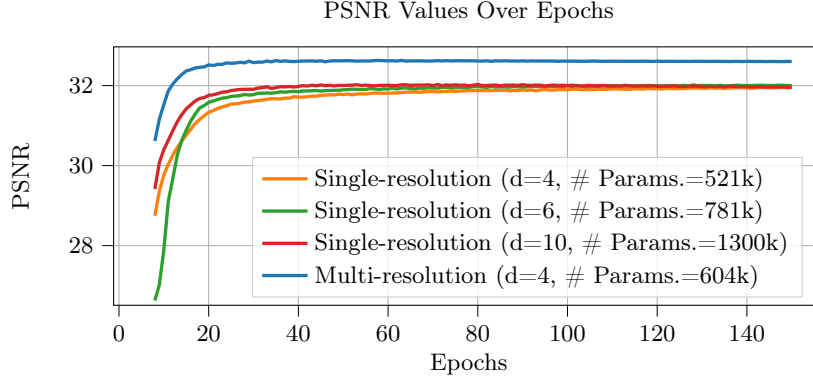


Fig. 4: Validation PSNR over training time in epochs. Our multi-resolution feature encoding leads to higher reconstruction quality despite having fewer parameters than a single-resolution encoding. For the latter, we use the finest resolution and arrange parameters the same way as in the multi-resolution approach. While a higher feature dimension d leads to faster convergence in the single-resolution setting, it does not improve the reconstruction quality despite the increased number of parameters.

correcting terms for local changes. Fig. 4 shows the effectivity of our multi-resolution approach and demonstrates that a single-resolution setup struggles to achieve high visual fidelity, even if the number of parameters exceeds the multi-resolution approach. It confirms that our multi-resolution approach allows an effective sharing of global information on different scales and enables reconstructions of higher quality while maintaining rapid convergence.

Regularization Frequency encodings require the MLP to map encoded spatial information to function values. Due to strong coupling of spatial information in a single network, frequency encodings can interpolate for regions, where training data is sparse. For our parametric encoding approach, working with fine meshes and sparse training data can lead to unsupervised latent codes, creating visual artifacts in our reconstructions. Our latent code regularization (Eq. (4)) uses the mesh Laplacian to circumvent this issue. Fig. 5 shows how our regularization reduces these artifacts significantly. The positive effect of our regularization scheme is also underlined quantitatively in Tab. 1.

4.3 Deforming Meshes

For many real-world applications of meshes that include animations, it is crucial that the representation of *e.g.* texture can also be used under deformations of the mesh. Since our features are stored at the mesh vertices and are therefore intrinsic to the mesh, they are unaffected by deformations. Consequently, our method supports mesh deformations natively without any additional computational overhead. While the same is true for other intrinsic encodings [18],

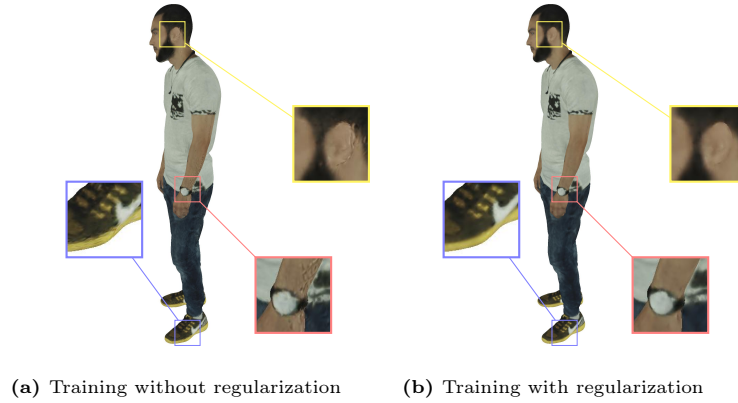


Fig. 5: Qualitative results of our method for texture reconstruction with and without the regularization based on the mesh Laplacian. The results on the left, without the regularization, show visual artifacts around the ear, and on the shoe and arm. This is a direct result of sparse training data resulting in untrained feature vectors. Our regularization acts as a smoothing term that enables feature information to be diffused to the unsupervised areas, significantly reducing the noise.




Fig. 6: Qualitative evaluation of texture represented by our method under mesh deformation. We train our network for a checker texture on the reference mesh (*left*) and render images for various deformations. Since the mesh topology remains unchanged under deformations, and our features are intrinsic to the mesh, our representation is unaffected by the deformation and can be evaluated similarly to the reference configuration. The results show that the representation is consistent under deformations and produces no visible artifacts.

extrinsic methods like [31, 41] need to map the intersection point in the deformed configuration back to the reference configuration, which adds a slight computa-

tional overhead. In Fig. 6, we show qualitatively that a texture represented by our method is consistent when the mesh is deformed and exhibits no noticeable artifacts. For the experiments, we use a checker texture on two meshes from the data provided by Sumner and Popovic [38], which we learn directly on the respective reference mesh and evaluate on some of its deformations.

4.4 BRDF Reconstruction from Images with Calibrated Lighting

As a second application of our method, we consider the estimation of a spatially varying *bidirectional reflectance distribution function* (BRDF) on a mesh. The BRDF $f(x, l, v)$ describes how much of the irradiance incident from the light direction $l \in \mathbb{S}^2$ is reflected in the viewing direction $v \in \mathbb{S}^2$. To obtain the total outgoing radiance $L_o(x, v)$ at position x in the viewing direction v , the incoming irradiance $L_i(x, l)$ needs to be integrated over the hemisphere \mathbb{H} , which is known as the rendering equation

$$L_o(x, v) = \int_{\mathbb{H}} f(x, l, v) L_i(x, l) \cos \theta_l \, dl. \quad (5)$$

For the estimation of the BRDF, we consider the special case of a single, directional light. As a result, the irradiance L_i is independent of the position x , and the rendering equation Eq. (5) reduces to a single evaluation for the direction l of the directional light. The simplified equation reads

$$L_o(x, v) = f(x, l, v) L_i \mathbb{I}_s(x, l) \cos \theta_l, \quad (6)$$

where we have introduced the indicator function $\mathbb{I}_s(x, l)$ to account for cast shadows.

Several models have been proposed to parametrize the BRDF. In this work, we choose the isotropic variant of the state-of-the-art Disney BRDF, which uses 12 parameters with values in the unit interval [3]. We compare against texture fields and intrinsic neural fields and modify the methods to predict these parameters spatially varying on the mesh by adjusting the output dimension accordingly. For Texture Fields, we found that an exponential learning rate scheduler with $\gamma = 0.9$ is necessary to prevent overfitting. All other methods remain unchanged. We employ a batch size of $b = 2^{14}$ color values under a given light direction for the training of all methods.

We use the established DiLiGenT-MV real-world dataset for our experiments [20]. It contains HDR images of 5 objects with complex reflective behavior. The images are taken from 20 calibrated views, captured under 96 calibrated directional lights each. Moreover, the dataset contains meshes of the objects. For training, we use 10 views with 30 lights each and evaluate the models on 20 lights for each of the remaining 10 views. We follow [28] and combine an L1 loss with a gamma correction function applied to the color values. Since the dataset contains HDR images, very bright regions would otherwise dominate the loss, and information from darker regions would be suppressed. We refer to the supplement for more details. The results in Tab. 2 and Fig. 7 show, that




Fig. 7: Qualitative Results of the BRDF reconstruction for the *cow* of the DiLiGenT-MV real-world dataset. The renderings show, that our method is able to reconstruct sharp boundaries in the material and yields results that are practically indistinguishable from the baseline methods, while being almost an order of magnitude faster.

Table 2: Results of the BRDF reconstruction. All metrics are averaged over the experiments for the 5 objects of the DiLiGenT-MV real-world dataset. Note that DSSIM and LPIPS are scaled by 100. Our method archives again similar reconstruction quality to the other methods while maintaining the large speed-up observed previously.

Method	PSNR \uparrow	DSSIM \downarrow	LPIPS \downarrow	# Params. \downarrow	Speedup \uparrow
TF+RFF	42.13	0.6718	1.50 🏆	332k 🏆	1.00x
INF	42.21 🏆	0.666 🏆	1.53 🥈	204930k	1.08x 🥈
Ours	42.17 🥈	0.6700 🥈	1.60	930k 🥈	7.58x 🏆

for BRDF estimation, we achieve similar reconstruction quality to the other methods while maintaining the significant speed-up in evaluation time. For more qualitative results, we refer to the supplement.

5 Conclusion

We have introduced MeshFeat as a multi-resolution parametric feature encoding for neural fields on meshes. The approach adapts the idea of parametric feature grids from the Euclidean space to meshes. Compared to existing approaches, this enables us to work with considerably smaller MLPs, leading to a significant speed-up. Our proposed multi-resolution approach is based on mesh simplification and enables the effective sharing of global feature information. This allows for higher reconstruction quality than a single-resolution approach while simultaneously requiring a smaller feature dimension. We have demonstrated the computational efficiency of our encoding for texture and BRDF representation on meshes, where we achieve reconstruction results that are on par with competing methods based on frequency-based encodings with significantly longer evaluation times. Additionally, we showcased our method’s native applicability to signals on deforming meshes. We identify a more texture-adaptive multi-resolution feature approach as a promising direction for future research. We hope our work provides valuable insights applicable to various areas, including virtual and augmented reality, animations, and computer graphics engines.

Acknowledgments. We thank Christian Koke for helpful discussions. This work was supported by ERC Advanced Grant SIMULACRON.

References

1. Baatz, H., Granskog, J., Papas, M., Rousselle, F., Novák, J.: Nerf-tex: Neural reflectance field textures. In: Comput. Graph. Forum (2022)
2. 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: ICML (2020)
3. Burley, B., Studios, W.D.A.: Physically-based shading at disney. In: Acm Siggraph (2012)
4. Chabra, R., Lenssen, J.E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., Newcombe, R.: Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In: ECCV (2020)
5. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: CVPR (2019). <https://doi.org/10.1109/CVPR.2019.00609>
6. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022)
7. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: SIGGRAPH (1997). <https://doi.org/10.1145/258734.258849>
8. Garland, M., Heckbert, P.S.: Simplifying surfaces with color and texture using quadric error metrics. In: IEEE Visualization Conference (1998). <https://doi.org/10.1109/VISUAL.1998.745312>
9. Grattarola, D., Vanderghenst, P.: Generalised implicit neural representations. In: NeurIPS (2022)
10. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. In: ICML (2020)
11. Hanocka, R., Hertz, A., Fish, N., Giryas, R., Fleishman, S., Cohen-Or, D.: Meshcnn: a network with an edge. ACM TOG (2019). <https://doi.org/10.1145/3306346.3322959>
12. Hertz, A., Perel, O., Giryas, R., Sorkine-Hornung, O., Cohen-Or, D.: SAPE: spatially-adaptive progressive encoding for neural optimization. In: NeurIPS (2021)
13. Hertz, A., Perel, O., Giryas, R., Sorkine-Hornung, O., Cohen-Or, D.: Mesh draping: Parametrization-free neural mesh transfer. Comput. Graph. Forum (2023). <https://doi.org/10.1111/CGF.14721>
14. Hwang, I., Kim, H., Kim, Y.M.: Text2scene: Text-driven indoor scene stylization with part-aware details. In: CVPR (2023). <https://doi.org/10.1109/CVPR52729.2023.00188>
15. Jiang, C.M., Sud, A., Makadia, A., Huang, J., Nießner, M., Funkhouser, T.: Local implicit grid representations for 3d scenes. In: CVPR (2020)
16. Jiang, Z., Schneider, T., Zorin, D., Panozzo, D.: Bijective projection in a shell. ACM TOG (2020). <https://doi.org/10.1145/3414685.3417769>
17. Kim, H., Jang, Y., Lee, J., Ahn, S.: Hybrid neural representations for spherical data. CoRR (2024). <https://doi.org/10.48550/ARXIV.2402.05965>
18. Koestler, L., Grittner, D., Möller, M., Cremers, D., Lähner, Z.: Intrinsic neural fields: Learning functions on manifolds. In: ECCV (2022). https://doi.org/10.1007/978-3-031-20086-1_36
19. Lee, A.W.F., Sweldens, W., Schröder, P., Cowsar, L.C., Dobkin, D.P.: MAPS: multiresolution adaptive parameterization of surfaces. In: SIGGRAPH (1998). <https://doi.org/10.1145/280814.280828>

20. Li, M., Zhou, Z., Wu, Z., Shi, B., Diao, C., Tan, P.: Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE TIP* (2020). <https://doi.org/10.1109/TIP.2020.2968818>
21. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: *CVPR* (2023)
22. Lin, S., Zhou, B., Zheng, Z., Zhang, H., Liu, Y.: Leveraging intrinsic properties for non-rigid garment alignment. In: *ICCV* (2023). <https://doi.org/10.1109/ICCV51070.2023.01332>
23. Liu, H.D., Zhang, J.E., Ben-Chen, M., Jacobson, A.: Surface multigrid via intrinsic prolongation. *ACM TOG* (2021). <https://doi.org/10.1145/3450626.3459768>
24. Liu, L., Gu, J., Lin, K.Z., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. *NeurIPS* (2020)
25. Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: *ICCV* (2021)
26. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: *CVPR* (2019)
27. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. In: *CVPR* (2022). <https://doi.org/10.1109/CVPR52688.2022.01313>
28. Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In: *CVPR* (2022). <https://doi.org/10.1109/CVPR52688.2022.01571>
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: *ECCV* (2020)
30. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG* (2022). <https://doi.org/10.1145/3528223.3530127>
31. Oechsle, M., Mescheder, L.M., Niemeyer, M., Strauss, T., Geiger, A.: Texture fields: Learning texture representations in function space. In: *ICCV* (2019). <https://doi.org/10.1109/ICCV.2019.00463>
32. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: *CVPR* (2019)
33. Peng, S., Niemeyer, M., Mescheder, L.M., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: *ECCV* (2020). https://doi.org/10.1007/978-3-030-58580-8_31
34. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: *ICML* (2019)
35. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. *NIPS* (2007)
36. Sharp, N., Crane, K.: A Laplacian for Nonmanifold Triangle Meshes. *Comput. Graph. Forum* (2020)
37. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: *NeurIPS* (2020)
38. Sumner, R.W., Popović, J.: Deformation transfer for triangle meshes. *ACM TOG* (2004)
39. Sun, C., Sun, M., Chen, H.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: *CVPR* (2022)

40. Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3D shapes. In: CVPR (2021)
41. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. NeurIPS (2020)
42. Tang, D., Dou, M., Lincoln, P., Davidson, P., Guo, K., Taylor, J., Fanello, S., Keskink, C., Kowdle, A., Bouaziz, S., Izadi, S., Tagliasacchi, A.: Real-time compression and streaming of 4d performances. In: ACM TOG (2018)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NIPS (2017)
44. Walker, T., Mariotti, O., Vaxman, A., Bilen, H.: Explicit neural surfaces: Learning continuous geometry with deformation fields. CoRR (2023). <https://doi.org/10.48550/ARXIV.2306.02956>
45. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: NeurIPS (2021)
46. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE TIP (2004). <https://doi.org/10.1109/TIP.2003.819861>
47. Xiang, F., Xu, Z., Hasan, M., Hold-Geoffroy, Y., Sunkavalli, K., Su, H.: Neutex: Neural texture mapping for volumetric neural rendering. In: CVPR (2021). <https://doi.org/10.1109/CVPR46437.2021.00704>
48. Yang, B., Bao, C., Zeng, J., Bao, H., Zhang, Y., Cui, Z., Zhang, G.: Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In: ECCV (2022). https://doi.org/10.1007/978-3-031-19787-1_34
49. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. In: NeurIPS (2021)
50. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018). <https://doi.org/10.1109/CVPR.2018.00068>