

Nuvo: Neural UV Mapping for Unruly 3D Representations

Pratul P. Srinivasan, Stephan J. Garbin, Dor Verbin, Jonathan T. Barron, and Ben Mildenhall

Google

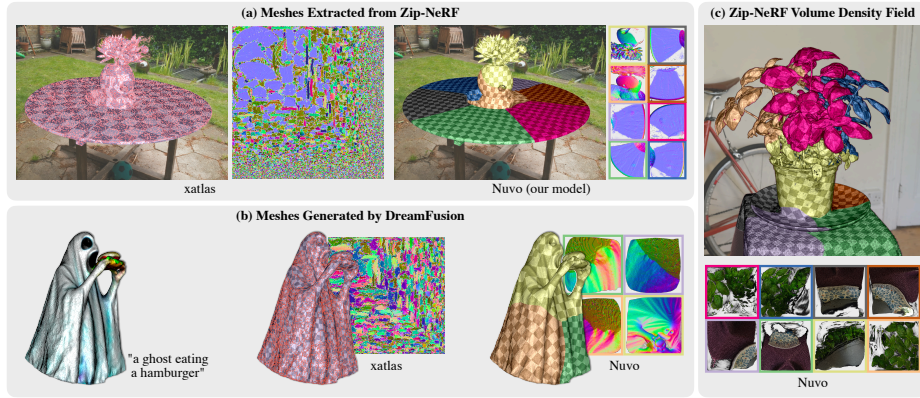


Fig. 1: Nuvo is a technique for UV mapping geometry produced by state-of-the-art 3D reconstruction and generation models such as Neural Radiance Fields (NeRFs) [22]. When applied to such geometry, existing UV mapping algorithms like xatlas [45] produce fragmented texture atlases (as shown by the chart boundaries marked in orange) that are unusable for appearance editing. Nuvo produces high-quality editable UV mappings for these 3D models, and is robust to challenging input geometry such as (a) meshes extracted from trained NeRF models and (b) meshes generated by text-to-3D models such as DreamFusion [27]. Nuvo can even operate directly on (c) non-surface geometry, *e.g.* NeRF volume density, without requiring a triangulated mesh.

Abstract. Existing UV mapping algorithms are designed to operate on well-behaved meshes, instead of the geometry representations produced by state-of-the-art 3D reconstruction and generation techniques. As such, applying these methods to the volume densities recovered by neural radiance fields and related techniques (or meshes triangulated from such fields) results in texture atlases that are too fragmented to be useful for tasks such as view synthesis or appearance editing. We present a UV mapping method designed to operate on geometry produced by 3D reconstruction and generation techniques. Instead of computing a mapping defined on a mesh’s vertices, our method Nuvo uses a neural field to represent a continuous UV mapping, and optimizes it to be a valid and well-behaved mapping for just the set of visible points, *i.e.* only points that affect the scene’s appearance. We show that our model is robust to the challenges posed by ill-behaved geometry, and that it produces editable UV mappings that can represent detailed appearance.

1 Introduction

Surface parameterization (“UV mapping”) is the process of flattening a 3D surface onto a plane, and it is a core component of 3D content creation pipelines that enables representing and editing detailed appearance on surface geometry. For complex, real-world meshes, this usually necessitates finding a sequence of cuts such that distortion of the mapping can be minimal. If those cuts result in multiple disconnected components that get packaged into one texture, this is commonly referred to as a “texture atlas”.

Existing UV mapping algorithms are generally designed to work with well-behaved meshes, such as those created by specialized 3D artists. However, an increasing amount of 3D content does not fall into this category: State-of-the-art methods for reconstructing and generating 3D representations from images or text are based on Neural Radiance Fields (NeRFs), which represent geometry as volumetric fields instead of meshes [22]. Level sets of volume density are generally not smooth, and triangulating these level sets using techniques such as marching cubes [20] produces meshes with multiple connected components, holes that connect to internal “hidden” geometry, and many small “bumpy” triangles. More formally, such meshes are typically not manifold or locally smooth, and frequently contain a large number of connected components. Existing UV mapping methods either cannot operate on such meshes or produce heavily fragmented UV atlases that complicate downstream applications. For instance, appearance editing in texture space becomes difficult, and optimization through differentiable rendering is complicated by discontinuities in the surface parameterization which could cause instabilities or necessitate the use of very large texture representations.

We present an approach, which we call Nuvo, that addresses these issues by using neural fields to directly optimize a UV mapping that satisfies the myriad requirements of a well-behaved surface parameterization. Our method simply requires a representation of scene geometry that allows for sampling visible 3D points, and we optimize Nuvo from scratch for each scene by minimizing a set of losses that encourage Nuvo to represent a well-behaved mapping for observed points. Because our method uses point sampling on the surface as its fundamental operation, it can be applied to any implicit surface representation, as well as polygonal meshes, without strong limitations on manifoldness, connectivity, or smoothness. For instance, Nuvo can generate texture atlases directly from NeRF’s volume density representation of geometry and it can be applied to extracted meshes while remaining agnostic to the connectivity of the underlying mesh. Because Nuvo’s UV mapping representation is not tied to any underlying mesh, it does not suffer from the chart fragmentation issues that can arise when texturing non-smooth meshes with many small triangles.

We test Nuvo on a variety of 3D geometry representations: well-behaved meshes, volume density fields reconstructed by NeRF, meshes extracted from NeRF’s volume density fields, and meshes produced by text-to-3D generative models. As illustrated in Figure 1, Nuvo produces texture atlases that are high-

quality with low chart fragmentation, and can therefore be used to represent and edit detailed surface appearance.

2 Related Work

Preliminaries Surfaces reconstructed by neural field representations can have vastly different topology depending on the scene. As such, triangulated meshes of these representations do not have guarantees on manifoldness, uniformity of triangle areas, smoothness, number of connected components, or genus. Since only developable surfaces (such as a cylinder) allow for isometric mappings to the plane without cuts, our method is designed to optimise both segmentation of the surface, *i.e.* cuts, and the UV mapping, *i.e.* the parameterization, simultaneously.

A mapping is isometric if it preserves both lengths and angles, and conformal if it only preserves angles. Connected components in UV space are referred to as charts, and their collective is called an atlas [21]. Packing charts into an atlas is typically done as a post-processing step, but because Nuvo produces square charts (similarly to Carr *et al.* [4]) which are trivial to pack, we will not discuss packing in detail. Please refer to reviews by Sheffer *et al.* [34] and Floater *et al.* [11] for further mesh parameterization details.

Mesh Cutting & Parameterization Mesh parameterization is a long-standing problem [5, 40]. However, only a few existing methods address the same problem setting as we do by simultaneously optimizing both cuts and UV mappings.

Many methods instead start with a user-specific boundary and optimize a UV mapping for that specific cut. Least Squares Conformal Mapping (LSCM) [18] is a widely-used technique that produces conformal maps for a pre-cut mesh. LSCM and related techniques construct and solve a system of equations that is a function of the mesh and its connectivity. This strategy ties these methods to the given mesh topology, and imposes strict requirements like manifoldness. In contrast, Nuvo never explicitly builds and solves such systems and instead relies on sampling points on the surface.

Sorkine *et al.* [36] propose the first method to simultaneously optimize cuts and surface parameterization by starting from seed triangles and iteratively adding elements until a distortion bound is reached. Geometry Images [14] and Multi-Chart Geometry Images [32] encode a 3D triangle mesh on a regularly sampled grid, which allows the geometry to be stored using image compression techniques and facilitates surface re-meshing. Rectangular Multi-Chart Geometry Images improves these algorithms to ensure one-to-one texel assignment across chart boundaries while forming rectangular charts that can be packaged easily [4], like our method. Their cutting and parameterization algorithm is similar to Sorkine *et al.* [36], but with a different objective function. Unlike Nuvo, these methods require manifold triangle meshes and were developed for well-behaved surfaces with significantly lower complexity than those we address.

AutoCuts [28] and OptCuts [19] generate parameterizations of 3D surfaces that have both minimal cut length and low distortion by alternating the optimization of the mapping distortion (continuous) and cut locus (discrete). While

AutoCuts is designed for an interactive workflow, OptCuts operates fully automatically. While OptCuts method can produce mappings with low distortion and simple boundaries on simpler meshes, it is unable to run on high-complexity meshes extracted from neural fields, even after manual cleanup.

A related line of work focuses on improving existing UV parameterizations. Smith and Schaefer [35] minimize symmetric Dirichlet energy to improve the distortion of an existing parameterization, which can be very efficient [30] and is more robust than prior work such as LSCM [18]. Several recent methods optimize a pre-existing UV mapping to minimise a reconstruction loss using differentiable rendering [17, 38]. However, these techniques must be provided with an initial surface parameterization. Variational Surface Cutting [33] investigates the theoretical problem of optimizing a surface cut to minimize distortion in the continuous setting, and show that this can be done without explicitly parameterizing the surface to compute distortion.

Neural 3D Representations & Parameterization State-of-the-art methods for 3D reconstruction and generation are based on neural fields, which typically represent 3D geometry as some variation of a volumetric density field [2, 26, 31, 44] parameterized by some combination of MLPs with positional encoding [39] and feature grids [12, 24, 37]. Some recent works focus on finding 2D parameterizations of 3D signals without using explicit triangle meshes. For example, Layered Neural Atlases [16] convert videos into a set of 2D atlases where an MLP maps from a 3D coordinate in the video to a 2D atlas containing color and transparency. AtlasNet [13], FoldingNet [43], and follow-up methods [3, 29] use learned mappings from 2D to 3D for generating and reconstructing geometry of objects. Other recent works [6, 7] learn UV mapping-based models of appearance for the specialized cases of dynamic human models and 3D human faces. Neural Surface Maps [23] use MLPs to encode single-chart surface mappings for geometry processing problems including establishing correspondences between surfaces. Williams *et al.* [41] train sets of MLPs to map from 2D to 3D for the goal of reconstructing surfaces from point clouds.

Our work builds upon NeuTex [42], which proposes a joint optimization framework for appearance and surface parameterization. Similar to Nuvo, NeuTex uses a cycle-consistency loss to encourage invertible mappings from 3D to 2D. However, NeuTex and related followups [9, 10, 46] cannot represent atlases with multiple charts, and are therefore not suited to represent mappings for general scenes with multiple connected components.

3 Method

Given a representation of scene geometry that allows for sampling visible points in the scene (*e.g.* a NeRF or a mesh), Nuvo generates a UV mapping that partitions the scene geometry into n charts using $2n + 1$ MLPs (see Figure 2):

1. One “chart assignment” MLP $c : \mathbb{R}^3 \rightarrow \Delta^{n-1}$ assigns points on the surface to charts by mapping from a 3D point to a probability mass function (PMF) of a categorical distribution over the n charts.




Fig. 2: Overview of Nuvo. Our model uses a neural field to represent a scene’s UV mapping. Our “chart assignment” MLP $c(\cdot)$ outputs probabilities of a categorical distribution over charts for any surface point \mathbf{x} , our “texture coordinate” MLPs $t_i(\cdot)$ map from 3D points \mathbf{x} to 2D UV coordinates \mathbf{u} , and our “surface coordinate” MLPs $s_i(\cdot)$ map from 2D UV coordinates to 3D points on the surface. Here we visualize Nuvo’s mappings for charts $i = 1$ and $i = 3$ in an atlas consisting of $n = 4$ charts.


2. A set of n “texture coordinate” MLPs $\{t_i : \mathbb{R}^3 \rightarrow [0, 1]^2\}$, each of which describes chart i ’s UV mapping from a 3D point to the corresponding 2D “UV” texture coordinate (bounded to lie between 0 and 1).
3. A set of n “surface coordinate” MLPs $\{s_i : [0, 1]^2 \rightarrow \mathbb{R}^3\}$, each of which describes the inverse of each t_i by mapping from a 2D texture coordinate to a 3D point.

3.1 Losses

We optimize Nuvo by minimizing a weighted sum of losses that encourage the recovery of a well-behaved mapping for the scene. The losses are averaged over batches of either random 3D points $\mathbf{x} \in \mathcal{G}$ sampled from the input scene geometry, or random 2D points $\mathbf{u} \in \mathcal{T}$ distributed uniformly in texture space $[0, 1]^2$.

Mapping bijectivity Each chart’s texture coordinate mapping from 3D points on the surface to 2D points on the plane should be approximately bijective, *i.e.* both injective (multiple 3D points should not map to the same texture coordinate) and surjective (the entire texture space should be “used” by the mapping). Injectivity is important, as we require the ability to make separate edits to the appearance of all distinct surface points. However, since our method is designed to be applied to 3D volumetric representations such as NeRFs (and not just 2D surfaces embedded in 3D), it cannot be strictly injective. Instead, we encourage the mappings to be invertible but allow them to “collapse” for regions of the volume when doing so does not hinder representing appearance. Surjectivity is desirable but does not need to strictly hold, as a small amount of unused texture space is tolerable for our purposes. We encourage bijectivity

Fig. 3: Visualization of our bijectivity and chart assignment entropy losses. These losses together encourage Nuvo to partition the scene into a set of approximately bijective chart mappings. $\mathcal{L}_{\text{entropy}}$ maximizes chart i ’s probability for 3D points $s_i(\mathbf{u})$ that are mapped to from 2D points in chart i . $\mathcal{L}_{3 \rightarrow 2 \rightarrow 3}$ and $\mathcal{L}_{2 \rightarrow 3 \rightarrow 2}$ regularize each chart’s mapping to be bijective by encouraging t_i and s_i to be each other’s inverse for all 3D surface points \mathbf{x} and all 2D UV points \mathbf{u} .



by minimizing two cycle consistency losses. The first loss minimizes the squared distance travelled by 3D points \mathbf{x} after being mapped to 2D and back:

$$\mathcal{L}_{3 \rightarrow 2 \rightarrow 3} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \sum_i c(\mathbf{x})[i] \cdot \|s_i(t_i(\mathbf{x})) - \mathbf{x}\|_2^2. \quad (1)$$

We weight the loss by chart probabilities $c(\cdot)[i]$ (the PMF value corresponding to chart i predicted by the chart assignment MLP $c(\cdot)$) such that the cycle consistency loss for a given point \mathbf{x} under chart i is proportional to the current probability estimate that point \mathbf{x} belongs to chart i . Our second cycle consistency loss minimizes the squared distance travelled by 2D points \mathbf{u} after being mapped to 3D and back:

$$\mathcal{L}_{2 \rightarrow 3 \rightarrow 2} = \frac{1}{|\mathcal{T}|} \sum \sum_i \|t_i(s_i(\mathbf{u})) - \mathbf{u}\|_2^2. \quad (2)$$

Note that both both cycle consistency losses are necessary to encourage a bijective mapping: $\mathcal{L}_{3 \rightarrow 2 \rightarrow 3}$ encourages chart mappings to be injective (one-to-one), but it does not penalize mappings that only map to a subregion of texture space, while $\mathcal{L}_{2 \rightarrow 3 \rightarrow 2}$ encourages the entirety of texture space to be covered by an invertible mapping, but it does not penalize mappings that are degenerate for parts of the 3D scene.

Chart assignment entropy Intuitively, if we sample a 2D point \mathbf{u} in chart i , the corresponding 3D point mapped to by surface coordinate MLP s_i should have a chart assignment PMF that is close to a one-hot distribution where the value for chart i is 1 and the rest are 0. We can therefore encourage the chart assignment MLP to confidently partition the 3D scene by minimizing the cross-

entropy of the output PMFs with these one-hot distributions:

$$\mathcal{L}_{\text{entropy}} = -\frac{1}{|\mathcal{T}|} \sum_{\mathbf{u} \in \mathcal{T}} \sum_i \log c(s_i(\mathbf{u}))[i]. \quad (3)$$

Figure 3 illustrates how our bijective and chart assignment losses together encourage Nuvo to represent a texture atlas that confidently partitions the scene into a set of n charts.

Surface coordinate We encourage the surface coordinate MLPs to approximate the input geometry by minimizing the symmetric Chamfer distance between random 2D texture points \mathbf{u} mapped to 3D by surface coordinate networks and random 3D points \mathbf{x} sampled from scene geometry:

$$\begin{aligned} \mathcal{L}_{\text{surface}} = & \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \min_{\mathbf{x}' \in \bigcup_i s_i(\mathcal{T})} \|\mathbf{x}' - \mathbf{x}\|_2^2 \\ & + \frac{1}{|\bigcup_i s_i(\mathcal{T})|} \sum_{\mathbf{x}' \in \bigcup_i s_i(\mathcal{T})} \min_{\mathbf{x}} \|\mathbf{x}' - \mathbf{x}\|_2^2, \end{aligned} \quad (4)$$

where $\bigcup_i s_i(\mathcal{T})$ is the union of the 3D points mapped to by all surface coordinate MLPs from random 2D UV points $\mathbf{u} \in \mathcal{T}$. This loss is meant to encourage the entire UV domain to be mapped to by points on the 3D geometry. Without ensuring that all 2D texture points map back to the actual 3D surface, it is possible for the MLPs to only represent injective mappings to a subregion of UV space and satisfy $\mathcal{L}_{2 \rightarrow 3 \rightarrow 2}$ with injective mappings from 3D regions not on the surface to the remainder of UV space.


Chart assignment clustering We penalize texture atlas fragmentation using a clustering loss that minimizes the distance from each 3D point to the centroids of the points assigned to each chart. We weight the loss by the chart assignment probabilities such that distances from a point \mathbf{x} to the centroid of chart i is penalized proportionally to the probability that point \mathbf{x} belongs to chart i :

$$\mathcal{L}_{\text{cluster}} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \sum_i c(\mathbf{x})[i] \cdot \left\| \left(\frac{\sum_{\mathbf{x}' \in \mathcal{G}} c(\mathbf{x}')[i] \cdot \mathbf{x}'}{\sum_{\mathbf{x}' \in \mathcal{G}} c(\mathbf{x}')[i]} - \mathbf{x} \right) \right\|_2^2, \quad (5)$$

Distortion To enable texture editing and to encourage an even allocation of texture resolution over the scene, we want our chart mappings to have low distortion. In particular, we would like the chart mappings to be conformal (angle-preserving), and we would like the amount of stretching or shrinking in the mapping to be uniform, as satisfying these properties ensures that a texture wrapped onto the surface resembles the original texture image.

To this end we impose regularizers on the texture coordinate mappings' differentials, as illustrated in Figure 4. The differential of mapping t_i at point \mathbf{x} , $Dt_i(\mathbf{x})$, is a linear mapping of vectors in the 3D surface's tangent space at \mathbf{x} to 2D UV vectors \mathbf{u} . If a mapping is conformal, applying the differential to orthogonal tangent space vectors should result in orthogonal vectors. If a mapping has

Fig. 4: Visualization of our distortion loss. The differential of mapping t_i , $Dt_i(\mathbf{x})$, transforms tangent vectors on the surface to vectors in UV space. $\mathcal{L}_{\text{conformal}}$ encourages the mappings to preserve angles by penalizing the cosine between the differential $Dt_i(\mathbf{x})$ applied to orthogonal tangent vectors at \mathbf{x} and $\mathcal{L}_{\text{stretch}}$ encourages the mappings to have uniform stretch by penalizing variation over the scene of the areas of parallelograms spanned by the differentials of orthogonal tangent vectors.



uniform stretch, the degree to which $Dt_i(\mathbf{x})$ stretches 3D tangent vectors should be constant for all \mathbf{x} .

We generate random orthogonal tangent space unit vectors, $r_{\mathbf{x}}$ and $v_{\mathbf{x}}$ using the normal at each point \mathbf{x} , and then transform the endpoints of $\epsilon p_{\mathbf{x}}$ and $\epsilon q_{\mathbf{x}}$ ($\epsilon = 10^{-2}$ in all experiments) to obtain UV vectors $Dt_i(\epsilon p_{\mathbf{x}})$ and $Dt_i(\epsilon q_{\mathbf{x}})$ (note that we drop the dependence of the differential on \mathbf{x} for readability). These vectors approximate the mappings’ differentials and let us define a distortion loss that is the sum of conformal and uniform stretch regularizers: $\mathcal{L}_{\text{distortion}} = \mathcal{L}_{\text{conformal}} + \mathcal{L}_{\text{stretch}}$. Our conformal regularizer is the squared cosine between the transformed vectors, weighted by the probability that \mathbf{x} belongs to chart i :

$$\mathcal{L}_{\text{conformal}} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \sum_i c(\mathbf{x})[i] \cdot \left(\frac{Dt_i(\epsilon p_{\mathbf{x}}) \cdot Dt_i(\epsilon q_{\mathbf{x}})}{\|Dt_i(\epsilon p_{\mathbf{x}})\| \|Dt_i(\epsilon q_{\mathbf{x}})\|} \right)^2. \quad (6)$$

Our uniform stretch regularizer penalizes the squared difference between the mapping’s stretch (the area of the spanned parallelogram) and a single scalar “average stretch” scalar parameter σ that we optimize for in each scene:

$$\mathcal{L}_{\text{stretch}} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \sum_i c(\mathbf{x})[i] \cdot \left\| \|Dt_i(\epsilon p_{\mathbf{x}}) \times Dt_i(\epsilon q_{\mathbf{x}})\|_2 - \sigma \right\|_2^2. \quad (7)$$

This is functionally equivalent to simply minimizing the variance of the stretch over the scene, but we found that alternative to have the undesirable effect of collapsing the texture coordinate MLPs to map all points to the same texture coordinate, while our loss results in stable optimization.

Texture optimization We find that optimizing Nuvo’s mappings to be usable for representing surface normals helps encourage injectivity. To this end we impose a penalty on the difference between surface normals optimized in UV space and the true surface normals:

$$\mathcal{L}_{\text{texture}} = \frac{1}{|\mathcal{G}|} \sum_{\mathbf{x} \in \mathcal{G}} \sum_i c(\mathbf{x})[i] \cdot \|N_i(t_i(\mathbf{x})) - n(\mathbf{x})\|_2^2, \quad (8)$$




Fig. 5: Visualization of mappings for varying number of charts n . Nuvo produces editable UV mappings across a wide range of user-specified number of charts ($n = 4, 8, 16$ shown here).

where $N_i(\mathbf{u})$ is the value of a pixel grid for chart i at 2D points \mathbf{u} (using bilinear interpolation) and $n(\mathbf{x})$ is the surface normal at \mathbf{x} . We jointly optimize pixel grid values in N_i alongside the mapping MLPs to minimize this loss.

4 Experiments

Datasets We compare UV mappings across four datasets: **bunny** (72K vertices) and **lion** (750K vertices) are well-behaved meshes with a single connected component and smooth manifold geometry. **garden vase** (1.4M vertices) and **amigurumi** (2.2M vertices) are meshes extracted by marching cubes from Zip-NeRF [2] reconstructions. These two scenes are captured in the “Mip-NeRF 360” style [1] with roughly 185 images per dataset, where one eighth of these are reserved for testing. The meshes of these two scenes have non-smooth geometry, thousands of separate connected components, and non-manifold edges.

When optimizing UV mappings for well-behaved meshes, we sample random 3D points \mathbf{x} uniformly distributed on the surface. When optimizing UV mappings for meshes extracted from NeRF reconstructions, we sample random camera rays that view the mesh, and use their intersection points as \mathbf{x} .

4.1 Nuvo Represents Detailed Appearance

We first validate Nuvo’s ability to produce UV mappings that effectively represent detailed appearance for a given geometry. As our goal is to compute UV mappings for geometry produced by NeRF and other view synthesis approaches, we evaluate the mappings produced by Nuvo and our baselines by measuring how useful they are for view synthesis.

We conduct this evaluation on the **garden vase** and **amigurumi** datasets. Starting with an optimized NeRF (we use the state-of-the-art Zip-NeRF [2]), we first extract a mesh using marching cubes [20], and then compute a UV mapping using our algorithm. Next, we define an appearance model consisting

	gardenvase		amigurumi	
	PSNR Params		PSNR Params	
BakedSDF [44]	22.55	37.8M	28.63	59.4M
Ours, $n = 2$	23.11	9.1M	28.63	11.5M
Ours, $n = 2$ (b)	23.09	10.1M	28.47	14.1M
Ours, $n = 8$	23.48	9.1M	28.81	11.5M
Ours, $n = 8$ (b)	22.98	10.1M	28.42	14.1M
Ours, $n = 32$	23.72	9.1M	28.92	11.5M
Ours, $n = 32$ (b)	23.41	10.1M	28.67	14.1M

Table 1: View synthesis comparison. Nuvo’s UV mappings effectively represent detailed surface appearance for view synthesis. We optimize the view-dependent appearance model used in BakedSDF [44], but in UV atlas space instead of on mesh vertices. We evaluate Nuvo’s view synthesis performance on two scenes from the Mip-NeRF 360 dataset [1] using $n = 2, 8, 32$ charts, and fix memory usage by using a texture resolution of $256\sqrt{2/n} \times 256\sqrt{2/n}$ for each chart. Additionally, we evaluate “baked” versions of Nuvo (indicated with “b”) by precomputing and storing the MLP-predicted UV coordinates on mesh vertices to show that Nuvo’s MLPs do not need to be kept after optimization, and that our UV maps can be used in standard graphics pipelines.

of a diffuse color and three spherical Gaussian view-dependent color lobes (the same appearance model used in BakedSDF [44]) on a 2D grid in UV space, and optimize this representation to best reproduce the training images.

Table 1 compares Nuvo’s view synthesis results to those of BakedSDF, which uses the same view-dependent appearance model defined directly on the extracted mesh’s vertices. We can see that optimizing our UV mappings for view synthesis performs similarly or even better than optimizing the same appearance model directly on mesh vertices, and we are able to achieve this performance while using less memory than BakedSDF as we do not allocate memory to triangles that do not directly influence view synthesis.

To further demonstrate Nuvo’s usefulness in standard graphics pipelines, we “bake” the optimized UV coordinates onto the mesh as vertex attributes (we select the chart with the maximum probability for each vertex). This can be thought of as a piecewise linear approximation of the optimized UV mappings. The “(b)” rows in Table 1 demonstrate that even though our UV mappings are optimized as continuous MLP-parameterized functions, baking them onto a mesh incurs a minimal decrease in performance.

4.2 Nuvo Produces Editable Mappings

Baselines We compare UV mappings produced by our model (using $n = 1, 2, 8$, and 32 charts) to:

NeuTex [42], which optimises a single-chart parameterization with a bijective consistency loss ($\mathcal{L}_{3 \rightarrow 2 \rightarrow 3}$) and no penalty on distortion;




Fig. 6: UV mapping and appearance editing of the **amigurumi** mesh extracted from Zip-NeRF volume densities. (a) Nuvo is able to recover a high-quality editable mapping that supports realistic appearance editing, while (b) NeuTex [42]’s mappings exhibit significant distortion (red circles), and (c) xatlas’ [45] and (d) Blender’s [8] atlases are too fragmented for appearance editing. Chart boundaries are marked in orange.

xatlas [45], an open-source library used throughout the 3D content creation industry based on work such as LSCM [18] and multi-chart geometry images [32]. Xatlas has been used in many recent reconstruction methods [15, 25];

Blender Smart UV [8], a robust industrial tool which we found to behave similarly to “automatic UV” methods in proprietary and commercial applications;

OptCuts [19] alternately minimizes distortion and cut length, and is the most recent academic method we found for obtaining a UV atlas.

We use the official implementations of xatlas, OptCuts, and Blender 2.93, all with default settings. We were unable to run OptCuts on our more complex scenes as it failed to find a suitable initial UV embedding, even after manually simplifying the meshes by removing non-manifold geometry and multiple connected components.

Metrics We compare UV mappings using two metrics: “Boundary” and “Editability” (higher is better for both), computed over a test set of randomly-sampled camera viewpoints for each dataset. “Boundary” quantifies texture atlas fragmentation by measuring the fraction of rendered pixels that do not lie on chart boundaries. “Editability” quantifies a user’s ability to wrap arbitrary texture on the geometry. First, we compute the UV coordinates for the vertices of the triangle intersected by each camera ray (using the “baked” versions of our MLP-parameterized mappings for fair comparison). Next, we measure each triangle’s editability as the average of angular distortion and area distortion metrics of the linear mapping implied by the UV coordinates of the three vertices. The image’s total editability score is the average of each pixel’s triangle editability masked by chart boundaries; pixels corresponding to chart boundaries are not considered editable. Please refer to the supplementary materials for a full detailed definition of these metrics.

	bunny		lion		gardenvase		amigurumi	
	Boundary	Editability	Boundary	Editability	Boundary	Editability	Boundary	Editability
Ours, $n = 1$	0.987	0.946	0.988	0.914	0.986	0.682	0.983	0.836
Ours, $n = 2$	0.982	0.937	0.977	0.890	0.974	0.859	0.983	0.813
Ours, $n = 8$	0.970	0.923	0.964	0.882	0.933	0.780	0.966	0.809
Ours, $n = 32$	0.924	0.860	0.903	0.811	0.823	0.704	0.927	0.784
NeuTex [42]	0.990	0.777	0.987	0.652	0.989	0.516	0.994	0.659
xatlas [45]	0.926	0.924	0.692	0.689	0.457	0.457	0.249	0.248
Blender Smart UV [8]	0.876	0.837	0.680	0.657	0.778	0.764	0.603	0.581
OptCuts [19]	0.987	0.954	0.970	0.939	—	—	—	—

Table 2: Editability comparison. Nuvo is competitive with the state-of-the-art OptCuts method specifically designed for well-behaved meshes (**bunny** and **lion**) and Nuvo significantly outperforms all baselines on the challenging meshes extracted from Zip-NeRF (**gardenvase** and **amigurumi**), where OptCuts is unable to compute any UV mapping. The “Boundary” metric quantifies the texture atlas fragmentation (higher is better *i.e.* less fragmented), and the “Editability” metric quantifies a user’s ability to wrap arbitrary new textures on the geometry (higher is better, best method is highlighted in red and second best method is highlighted in orange), which requires that the mapping has both low fragmentation and low distortion.

4.3 Ablation Studies

The loss ablations in Table 3 demonstrate that our full model achieves the best tradeoff of minimizing texture fragmentation (“Boundary”), minimizing distortion (“Stretch” and “Conformal”), and utilizing all of UV space (“UV Efficiency”). Figure 5 visualizes Nuvo’s results when varying the number of charts, Figure 8 visualizes the effect of our texture optimization loss, and Figure 9 visualizes ablations of the distortion and clustering losses.

	Bound.	Stretch	Conf.	UV Eff.	Avg.
Ours, $n = 4$	0.976	0.976	0.940	0.846	0.935
w/o $\mathcal{L}_{3 \rightarrow 2 \rightarrow 3}$	0.973	0.974	0.898	0.617	0.865
w/o $\mathcal{L}_{2 \rightarrow 3 \rightarrow 2}$	0.974	0.990	0.914	0.247	0.782
w/o $\mathcal{L}_{\text{entropy}}$	0.976	0.988	0.946	0.456	0.841
w/o $\mathcal{L}_{\text{surface}}$	0.982	0.987	0.948	0.381	0.824
w/o $\mathcal{L}_{\text{cluster}}$	0.953	0.979	0.947	0.785	0.916
w/o $\mathcal{L}_{\text{distortion}}$	0.974	0.983	0.588	0.965	0.878
w/o $\mathcal{L}_{\text{texture}}$	0.968	0.981	0.937	0.848	0.933

Table 3: Loss ablations for bunny using $n = 4$. Removing either of the bijectivity, chart assignment entropy, surface coordinate, or chart assignment clustering losses results in mappings that only use a subregion of UV space (low “UV Efficiency”). Ablating the distortion regularization results in a mapping with much worse conformal distortion. Removing the texture optimization loss does not significantly affect metrics, but can result in non-invertible mappings, as shown in Figure 8.




Fig. 7: UV mapping comparisons on the lion mesh. (a) Nuvo produces a mapping with a simple boundary and low overall distortion. (b) NeuTex’s [42] mapping is distorted and challenging to edit. (c) xatlas and (d) Blender produce atlases with significant fragmentation that preclude appearance editing. (e) OptCuts [19] recovers a high-quality mapping for this well-behaved mesh, but is unable to run on the more challenging meshes extracted from NeRF. Chart boundaries are marked in orange.

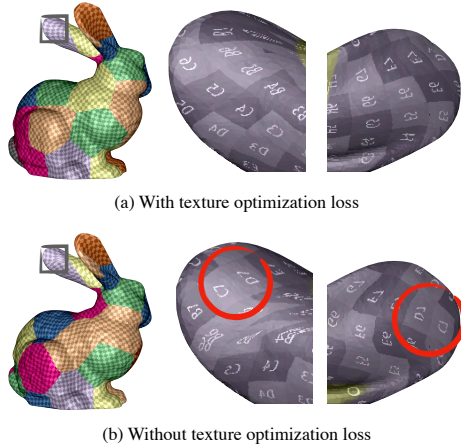


Fig. 8: Texture optimization loss $\mathcal{L}_{\text{texture}}$ encourages the optimized mappings to be injective (one-to-one) by penalizing UV mappings that cannot accurately represent surface normals. Inspecting the optimized mappings for the bunny’s left ear (grey box) from either side, we see that an ablation trained without texture optimization maps both sides of the ear to the same coordinates in UV space; the “C7” and “D7” UV coordinates (red circles) are visible from both sides of the ear. Our complete model trained with texture optimization does not have this degeneracy.




Fig. 9: Loss ablations visualization. (a) Our complete model recovers a high-quality UV mapping for the bunny model. (b) Removing our distortion loss $\mathcal{L}_{\text{distortion}}$ results in a texture atlas with non-uniform warping that hinders editing in 2D texture space. (c) Removing our chart assignment clustering loss $\mathcal{L}_{\text{cluster}}$ leads to fragmented and irregularly-shaped charts, which also makes appearance editing cumbersome.

4.4 Discussion

Table 2 demonstrates that Nuvo’s UV mappings are competitive with OptCuts’ on simpler well-behaved meshes, such as `lion` shown in Figure 7, and significantly better than all baselines for the more challenging cases of geometry extracted from NeRF models, such as `amigurumi` shown in Figure 6. These results are in line with our expectations. OptCuts effectively minimizes both boundary length and distortion, but its optimization procedure is expensive and does not apply to the unstructured non-manifold geometry produced by NeRF. `xatlas`’ and Blender’s strategy of starting with many cuts and attempting to merge mappings is not able to produce simple boundaries for non-smooth surfaces, resulting in heavily-fragmented atlases. Moreover, Nuvo is more efficient than these alternatives and only requires 20-40 minutes to optimize a UV mapping. OptCuts takes 2 hours for `bunny` and 35 hours for `lion`. `xatlas` and Blender SmartUV take about 10 minutes for `bunny`, close to an hour for `lion`, and over a day for the `gardenvase` and `amigurumi` NeRF meshes.

Limitations One of Nuvo’s strengths is that it optimizes a UV mapping with point samples instead of explicitly parameterizing the mapping over an entire mesh. However, this makes it harder to absolutely guarantee that it is bijective or that it globally minimizes distortion for a given boundary. Finally, while Nuvo is able to automatically optimize editable UV mappings for challenging geometry, it currently lacks interactive mapping capabilities such as allowing users to specify cut locations or regions for which they would particularly like to minimize distortion. We think that extending Nuvo to address these deficiencies would be fruitful directions for future work.

5 Conclusion

We have presented Nuvo, a method that produces editable UV atlases without severe fragmentation and distortion, even for challenging geometry created by 3D reconstruction and generation techniques. By focusing on visible surfaces and parameterizing mappings using neural fields instead of directly on mesh vertices, Nuvo can handle scenes with complexity beyond the capability of prior approaches. We believe that this work opens up numerous possibilities for creative and artistic editing of reconstructed 3D content.

References

1. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. *Conference on Computer Vision and Pattern Recognition* (2022) 9, 10
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. *International Conference on Computer Vision* (2023) 4, 9

3. Bednarik, J., Parashar, S., Gundogdu, E., Salzmänn, M., Fua, P.: Shape reconstruction by learning differentiable surface representations. Conference on Computer Vision and Pattern Recognition (2020) 4
4. Carr, N.A., Hoberock, J., Crane, K., Hart, J.C.: Rectangular Multi-Chart Geometry Images. Eurographics Symposium on Geometry Processing (2006) 3
5. Catmull, E.E.: A subdivision algorithm for computer display of curved surfaces. The University of Utah (1974) 3
6. Chen, Y., Wang, X., Chen, X., Zhang, Q., Li, X., Guo, Y., Wang, J., Wang, F.: Uv volumes for real-time rendering of editable free-view human performance. Conference on Computer Vision and Pattern Recognition (2023) 4
7. Chen, Z., Yin, K., Fidler, S.: Auv-net: Learning aligned uv maps for texture transfer and synthesis. Conference on Computer Vision and Pattern Recognition (2022) 4
8. Community, B.O.: Blender - a 3D modelling and rendering package, <http://www.blender.org> 11, 12
9. Das, S., Ma, K., Shu, Z., Samaras, D.: Learning an isometric surface parameterization for texture unwrapping. European Conference on Computer Vision (2022) 4
10. Deprelle, T., Groueix, T., Aigerman, N., Kim, V.G., Aubry, M.: Learning joint surface atlases. ECCV Workshop on Learning to Generate 3D Shapes and Scenes (2022) 4
11. Floater, M.S., Hormann, K.: Surface parameterization: a tutorial and survey. Advances in Multiresolution for Geometric Modelling (2005) 3
12. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Conference on Computer Vision and Pattern Recognition (2022) 4
13. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: AtlasNet: A Papier-Mâché approach to learning 3d surface generation. Conference on Computer Vision and Pattern Recognition (2018) 4
14. Gu, X., Gortler, S.J., Hoppe, H.: Geometry images. SIGGRAPH (2002) 3
15. Hasselgren, J., Hofmann, N., Munkberg, J.: Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. Conference on Neural Information Processing Systems (2022) 11
16. Kasten, Y., Ofri, D., Wang, O., Dekel, T.: Layered neural atlases for consistent video editing. ACM Transactions on Graphics (2021) 4
17. Knodt, J., Pan, Z., Wu, K., Gao, X.: Joint UV Optimization and Texture Baking. ACM Transactions on Graphics (2023) 4
18. Lévy, B., Petitjean, S., Ray, N., Maillot, J.: Least squares conformal maps for automatic texture atlas generation. ACM Transactions on Graphics (2002) 3, 4, 11
19. Li, M., Kaufman, D.M., Kim, V.G., Solomon, J., Sheffer, A.: OptCuts: Joint Optimization of Surface Cuts and Parameterization. ACM Transactions on Graphics (2018) 3, 11, 12, 13
20. Lorensen, W.E., Cline, H.E.: Marching Cubes: A High Resolution 3D Surface Construction Algorithm. SIGGRAPH (1987) 2, 9
21. Maillot, J., Yahia, H., Verroust, A.: Interactive texture mapping. SIGGRAPH (1993) 3
22. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. European Conference on Computer Vision (2020) 1, 2
23. Morreale, L., Aigerman, N., Kim, V.G., Mitra, N.J.: Neural surface maps. In: Conference on Computer Vision and Pattern Recognition (2021) 4

24. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* (2022) 4
25. Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. *Conference on Computer Vision and Pattern Recognition* (2022) 11
26. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. *Conference on Computer Vision and Pattern Recognition* (2019) 4
27. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3D using 2D Diffusion. *International Conference on Learning Representations* (2022) 1
28. Poranne, R., Tarini, M., Huber, S., Panozzo, D., Sorkine-Hornung, O.: AutoCuts: Simultaneous Distortion and Cut Optimization for UV Mapping. *ACM Transactions on Graphics* (2017) 3
29. Poursaeed, O., Fishes, M., Aigerman, N., Kim, V.G.: Coupling explicit and implicit surface representations for generative 3d modeling. *European Conference on Computer Vision* (2019) 4
30. Rabinovich, M., Poranne, R., Panozzo, D., Sorkine-Hornung, O.: Scalable locally injective mappings. *ACM Transactions on Graphics* (2017) 4
31. Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P.P., Mildenhall, B., Geiger, A., Barron, J.T., Hedman, P.: MERF: Memory-Efficient Radiance Fields for Real-Time View Synthesis in Unbounded Scenes. *ACM Transactions on Graphics* (2023) 4
32. Sander, P.V., Wood, Z.J., Gortler, S.J., Snyder, J., Hoppe, H.: Multi-chart geometry images. *Eurographics Symposium on Geometry Processing* (2003) 3, 11
33. Sharp, N., Crane, K.: Variational surface cutting. *ACM Transactions on Graphics* (2018) 4
34. Sheffer, A., Praun, E., Rose, K.: Mesh Parameterization Methods and Their Applications (2006) 3
35. Smith, J., Schaefer, S.: Bijective parameterization with free boundaries. *ACM Transactions on Graphics* (2015) 4
36. Sorkine, O., Cohen-Or, D., Goldenthal, R., Lischinski, D.: Bounded-distortion piecewise mesh parameterization. *IEEE Visualization* (2002) 3
37. Sun, C., Sun, M., Chen, H.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: *Conference on Computer Vision and Pattern Recognition* (2022) 4
38. Sun, H., Wang, S., Wu, W., Jin, Y., Bao, H., Huang, J.: Efficient texture parameterization driven by perceptual-loss-on-screen. *Computer Graphics Forum* (2022) 4
39. 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. *Conference on Neural Information Processing Systems* (2020) 4
40. Tutte, W.T.: How to draw a graph. *Proceedings of the London Mathematical Society* (1963) 3
41. Williams, F., Schneider, T., Silva, C., Zorin, D., Bruna, J., Panozzo, D.: Deep geometric prior for surface reconstruction. *Conference on Computer Vision and Pattern Recognition* (2019) 4
42. Xiang, F., Xu, Z., Hasan, M., Hold-Geoffroy, Y., Sunkavalli, K., Su, H.: NeuTex: Neural Texture Mapping for Volumetric Neural Rendering. *Conference on Computer Vision and Pattern Recognition* (2021) 4, 10, 11, 12, 13

- 43. Yang, Y., Feng, C., Shen, Y., Tian, D.: Foldingnet: Point cloud auto-encoder via deep grid deformation. Conference on Computer Vision and Pattern Recognition (2018) 4
- 44. Yariv, L., Hedman, P., Reiser, C., Verbin, D., Srinivasan, P.P., Szeliski, R., Barron, J.T., Mildenhall, B.: BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. ACM Transactions on Graphics (2023) 4, 10
- 45. Young, J.: xatlas, <https://github.com/jpcy/xatlas> 1, 11, 12
- 46. Zhan, F., Liu, L., Kortylewski, A., , Theobalt, C.: General neural gauge fields. In: International Conference on Learning Representations (2023) 4