
Implicit-ARAP: Efficient Handle-Guided Neural Field Deformation via Local Patch Meshing

Daniele Baieri

University of Milano-Bicocca
daniele.baieri@unimib.it

Filippo Maggioli

Pegaso University
filippo.maggioli@unipegaso.it

Emanuele Rodolà

Sapienza University of Rome
rodola@di.uniroma1.it

Simone Melzi

University of Milano-Bicocca
simone.melzi@unimib.it

Zorah Löhner

University of Bonn, Lamarr Institute
laehner@uni-bonn.de

Abstract

Neural fields have emerged as a powerful representation for 3D geometry, enabling compact and continuous modeling of complex shapes. Despite their expressive power, manipulating neural fields in a controlled and accurate manner – particularly under spatial constraints – remains an open challenge, as existing approaches struggle to balance surface quality, robustness, and efficiency. We address this by introducing a novel method for handle-guided neural field deformation, which leverages discrete local surface representations to optimize the As-Rigid-As-Possible deformation energy. To this end, we propose the local patch mesh representation, which discretizes level sets of a neural signed distance field by projecting and deforming flat mesh patches guided solely by the SDF and its gradient. We conduct a comprehensive evaluation showing that our method consistently outperforms baselines in deformation quality, robustness, and computational efficiency. We also present experiments that motivate our choice of discretization over marching cubes. By bridging classical geometry processing and neural representations through local patch meshing, our work enables scalable, high-quality deformation of neural fields and paves the way for extending other geometric tasks to neural domains.

1 Introduction

Implicit representations—where a surface is defined not explicitly but, for example, as the zero level set of a signed distance function—have long been used in computer graphics. However, they have recently gained renewed attention, particularly due to advances in neural rendering techniques such as NeRF [36]. Neural fields provide a compact representation for implicit representations in the weights of a neural network, and offer several advantages: they support flexible topology, avoid predefined discretization, and integrate naturally with gradient-based optimization methods. These properties make neural fields well-suited for reconstruction tasks. As this representation becomes more widespread, the demand for tools that enable direct analysis and manipulation of implicit surfaces continues to grow.

The traditional representation used in geometry processing applications is the polygonal mesh, in which the surface is explicitly modelled by a collection of connected polygons. This representation




Figure 1: Given a neural surface and deformation handles, we employ our local patch mesh representation to optimize a neural deformation field via the ARAP loss. Our method introduces significant improvements over previous state of the art while employing 0.75% of the time. We show the normalized change in edge lengths (before/after deformation) as a colormap on the deformed surfaces.

provides a high level of interpretability and allows for fine-grained manipulation by artists. Additionally, local surface properties can be easily computed because the neighborhood information is explicitly encoded. Due to its wide adoption and easy manipulation, a number of editing methods have been developed for explicit representations: one of them is optimizing the as-rigid-as-possible (ARAP) energy to deform an object while satisfying user-defined handle constraints and at the same time preserving the surface geometry (in terms of its edges) to a reasonable extent. This yields natural deformations and has been widely adapted in various applications [7, 37, 23].

Evaluating energies such as ARAP directly on implicit surfaces is challenging: properties like edge lengths cannot be computed without applying a discretization scheme, and equivalent properties do not necessarily exist for isosurfaces. Alternating between both representations is certainly possible [35], but it is expensive and inherits the sensitivity to discretization choices [57]. To counter this, we propose a new patch-based meshing for implicit representations which is 1) efficient to compute even at high resolution, 2) not sensitive to the specific choice of discretization, and 3) can be used on all isosurfaces to cover the complete geometric information of the neural field. Along with local patch meshing, our work proposes multiple other significant contributions:

- Extending the ARAP estimation of the thin-shell energy to neural representations;
- Providing an efficient method for handle-guided deformation of neural fields (see Figure 1);
- Introducing a robust and efficient alternative for deforming high-resolution meshes.

2 Related work

Handle-Based Deformations Mesh-based shape editing methods have a long history in geometry processing due to the wide acceptance and explicit representation of meshes [58, 6]. Among these, methods which use handles to indicate the preferred deformation are intuitive for human users to understand and generate. The as-rigid-as-possible (ARAP) energy [51] has become popular due to its straight-forward interpretation and easy optimization while satisfying handle constraints. However, while not complex, its global nature still prevents processing of high-resolution shapes and is sensitive to the mesh discretization. The second aspect was overcome in SR-ARAP [26] with a smoothed and rotation-enhanced ARAP version. While the rigid version is widely used in variety of applications [7, 37, 23], there exist similar formulations other energies related to shape deformation, for example conformal [42, 55] and elastic deformation energies [9]. Instead of using a pre-defined physical energy, it is also possible to learn a set of handles from a collection of shapes [29, 41]. However, these require a suitable set of training data and are restricted to the space of deformations learned initially.

The idea of parametrizing shape deformations with neural networks has been previously explored in recent research: two contributions which are particularly close to our work are Neural Jacobian Fields [2] and Neural Shape Deformation Priors [53]. NFJ is a mesh morphing (*i.e.*, transfer of pose between two given shapes) model which can be trained to inject several different priors, including ARAP deformations. NSDP, on the other hand, is a data-driven handle-based mesh deformation

model. While it allows for very fast deformation of low/medium-resolution meshes, it is limited to the deformation prior and the handle set observed in the training data, limitations overcome by Implicit-ARAP.

While all these methods are designed for polygonal meshes, there are various types of shape representations for which other types of deformation approaches are better suited. One such example which is especially relevant to our work is spatial deformation tasks, which apply to solid models [46, 47], radial basis functions [5] or cages for character animation [24].

Neural Fields Editing Energies like ARAP act directly on surface deformations, making them naturally suited to explicit representations such as triangle meshes. In contrast, editing implicit surfaces is more challenging due to their global coupling – local changes to the surface can affect distant regions of the field. Even locating the surface in space can be non-trivial without certain assumptions. The recent popularity of NeRF-like methods [36] has renewed interest in integrating deformations into implicit frameworks. Dynamic NeRFs have been implemented by introducing a time parameter [45] or optimizing deformation fields [8]. Other works modify the MLP weights of neural fields to generate [17] or edit [3] shapes directly, though this lacks efficient formulations of explicit energies like ARAP and leads to slow optimization. Mehta et al. [35] alternate between explicit and implicit representations to enable deformations, but require costly conversions at each step. Novello et al. [39] supports diverse deformations but lacks handle-based constraints. Neural fields with spatial features allow shape editing through feature-space interpolation [20, 1]. Similarly to our approach, Esturo et al. [18] optimize all isosurfaces simultaneously using divergence-free fields – an approach also applicable beyond implicit functions [15]. Recent work [14] shows cage-based deformations can be modeled with neural fields by learning mappings to barycentric coordinates. Text-driven editing of implicit or hybrid representations (e.g., Gaussian splatting) has also been explored [30, 56, 19, 10, 40].

The method closest to ours, due to Yang et al. [57], optimizes deformation losses sampled from implicit representations but is highly inefficient, requiring several hours. In contrast, we leverage full neural field information efficiently by using small, randomly placed discrete patches. Our use of mesh-based ARAP objectives is supported by [11], which highlights issues with higher-order neural field derivatives, and [28], which advocates finite differences over analytical gradients.

Other efforts to simplify local SDF representations include enhancing marching cubes with Bézier patches [54] and composing surfaces from implicit primitives [31]. However, unlike our approach, these methods do not yield local representations that are both explicit and structurally simple.

3 Method

Our method employs a local patch model (Sec. 3.1) to sample the surface of a neural signed distance field, which is then deformed according to an ARAP-like energy (Sec. 3.2) with an efficient optimization scheme (Sec. 3.3).

3.1 Local patch meshing

Given a 3D surface $\mathcal{S}$ represented implicitly by a neural signed distance field $f_\theta: \mathbb{R}^3 \rightarrow \mathbb{R}$, we require a discrete local representation in order to compute the ARAP energy induced by a given deformation field $d: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. We achieve this by generating local patches for *multiple* isosurfaces of f_θ . The procedure is separated into sampling and projection:

Sampling. We start by sampling k points $V = \{v_j\}_{j=1}^k$ from a disk with radius ρ on the 2D plane (including the origin) and computing the planar Delaunay triangulation. There are several possible distributions to sample the 2D disk: we describe and visualize some possible options in the supplementary material. However, our experiments showed that the particular choice has little influence on our method, as we show in Figure 9.

Projection. Once we obtain the 2D patch $\mathcal{P} = (V, F)$, we can discretize each isosurface by projecting the patch and fitting it to the local geometry. To that end, we approximate a uniform sampling of n points $O = \{o_i\}_{i=1}^n$ from the zero level set $\mathcal{S}_0$ using the rejection/projection algorithm proposed by Yang et al. [57]. Additionally, we sample points from other level sets by sampling

the 3D interval $[-1; 1]^3$ uniformly at random. We project a patch for each sampled point as $V_i = \{o_i + R_i(t(v_j))\}_{j=1}^k$, where t maps 2D points to 3D as $t: (x, y) \mapsto (x, y, 0)$, o_i are the sampled origin points, and R_i are the rotations aligning patch and surface normals at each o_i , which we obtain

as the normalized gradient of f_θ (first column of Figure 2). Then, we map each patch vertex onto level set $l = f_\theta(o_i)$ via the SDF closest surface point formula [57, 12] which maps any 3D point to its nearest neighbor on level set l as

$$p' = p - (f(p) - l) \frac{\nabla f(p)}{\|\nabla f(p)\|}. \quad (1)$$

In practice, when f is a neural SDF, this formula has to be applied recursively to yield accurate results (*i.e.* points whose signed distance is approximately l). Note that the radius ρ of the tangent disk specifies a measure of locality which is highly dependent on the underlying surface $\mathcal{S}$. Figure 2 shows examples of successful projection and radius overestimation.

By sampling origin points from multiple isosurfaces of the signed distance field f_θ , we enable representing any local region of the neural field with a patch mesh. We describe how this structure can be applied to shape deformation in Section 3.3. Moreover, we compare it to the classic SDF meshing algorithm Marching Cubes [32, 33] in terms of its benefits for handle-guided deformation in Section 4.2.




Figure 2: Examples of patch projection in a favourable case (top two rows) and with overestimated radius.

3.2 Deformation model

Following previous work in neural field deformation [8, 39, 45, 57], we represent our deformation as a continuous function of the embedding space $d: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. We model it via a MLP $g_\phi: \mathbb{R}^3 \rightarrow \text{SO}(3) \times \mathbb{R}^3$ mapping 3D coordinates to roto-translations, with parameters ϕ . We use R_ϕ and t_ϕ to refer to the rotation and the translation fields separately. The output layer predicts a 6D vector, where the first three components are interpreted as Euler angles and converted to a rotation matrix. We define the complete deformation as $d_\phi(x) = (R_\phi(x) \cdot x) + t_\phi(x)$.

3.3 Optimization

We optimize our model similarly to Yang et al. [57]; we summarize the process in the supplementary material. The goal is to optimize for target handle positions while regularizing the computed deformation to have some desired properties. Given a set of handles as with source-target position pairs $H = \{(s_i, t_i)\}_{i=1}^h$ (where $s_i = t_i$ in the case of static handles), we fit it via a simple MSE loss $L_{\text{handle}} = \frac{1}{h} \sum_{i=1}^h \|d_\phi(s_i) - t_i\|^2$. The key part of our loss function is the As-Rigid-As-Possible (ARAP) energy introduced by Sorkine and Alexa [51]. While this formulation was previously adopted as a regularizer for generative neural models [16, 23], our work is the first to employ it in the setting of implicit geometry processing. We aim to ensure that our map d_ϕ deforms the level sets of f_θ as-rigidly-as-possible. We evaluate this in a Monte-Carlo fashion, by sampling a set of points $\{x_k\}_{k=1}^n$, where $m \leq n$ points are sampled uniformly from the zero level set of f_θ and $n - m$ points




Figure 3: Loss curves for a single Implicit-ARAP training (Figure 9, first column). We show separate handle losses for static and moving handles for visualization purposes.

uniformly from the bounded volume $[-1; 1]^3$. A surface patch is computed at each of these points via the algorithm presented in Section 3.1, yielding a patch-based representation $\{(V_k, F)\}_{k=1}^n$. Then, the ARAP regularization is computed as

$$L_{\text{arap}} = \frac{1}{n} \sum_{k=1}^n \sum_{(v_i, v_j) \in E_k} w_{i,j} \|(d_\phi(v_i) - d_\phi(v_j)) - (R_\phi(v_i) \cdot (v_i - v_j))\|^2. \quad (2)$$

Where E_k are the mesh edges for a patch mesh (V_k, F) and $w_{i,j}$ are the cotangent Laplacian edge weights. By optimizing the deformed edges (left hand side of the difference) to be as close as possible to the rotational part (right hand side), we effectively mitigate the action of the translation field. The original ARAP formulation only used vertex-wise rotations, as handles were fit as a pre-processing step via Laplacian smoothing of the handle function over the surface. This operation is non-trivial for implicit surfaces, therefore we include a translation field, which allows any given handles set to be fit. The network g_ϕ is optimized with ADAM [25] steps until convergence of the loss function $L = \lambda_1 L_{\text{handle}} + \lambda_2 L_{\text{arap}}$. The entire procedure for computing the L_{arap} loss, which we described in this section, is repeated at each iteration, including all handle points as part of the surface sample. Nonetheless, we have observed convergence to be extremely quick, typically in the order of a few hundreds of iterations, as showed in Figure 3. In our experiments, we usually trained our model for a total of 1000 steps.




Figure 4: Ablation study for our method. We compare mean and standard deviation of the deformed SDF gradient norm for our method against: NFGP, our method with zero level set patches only, and our method without ARAP loss. The best results, in terms of both visual quality and preservation of the SDF field, are obtained by our full pipeline, proving the benefits of its individual components.

4 Experiments

The training procedure described in Section 3.3 can be applied seamlessly for both neural field and mesh deformation, by simply changing the MLP architecture (further details may be found in Section 6.2.1). We provide results for neural fields here, while we refer to Section 7.2 for an evaluation of high-resolution mesh deformation.

Data and baselines We employ two datasets in our evaluation: the first one (TFD) is obtained by designing a set of hand-crafted deformation experiments using mesh data from Thingi10k [59] and the Stanford 3D scanning repository. The second one is the DeFAUST dataset introduced in [34]. We use this data to evaluate the performance of our method against several baselines. In our comparison, we consider the original As-Rigid-As-Possible deformation method proposed by Sorkine and Alexa [51] (ARAP). We also include the spokes and rims variant introduced by Chao et al. [9] (Elastic), who propose the inclusion of an elasticity model in the ARAP optimization, as well as the smooth rotations alternative (SR-ARAP) [26], where the ARAP energy is modified to improve smoothness and volume preservation. The last baseline is the neural field deformation approach introduced by Yang et al. [57], who propose to use the continuous formulation of the thin shell energy at random samplings of the implicit surface to optimize the deformation field, similarly to our method. For this baseline, we will indicate in the following whether the deformation is applied on the SDF field via its inverse (NFGP_{SDF}) or on the original input mesh (NFGP_{Mesh}). Although the Neural Shape Deformation Prior (NSDP) approach proposed by Tang et al. [53] achieves natural looking mesh deformations with interactive performance, it is a data-driven approach that requires the specification of consistent




Figure 5: Qualitative results of our method in comparison to multiple baselines for neural field deformation. Discrete methods [51, 9, 26] are applied by extracting the marching cubes mesh at resolution $R = 512$ and applying a remeshing step, as suggested in [57]. We show true implicit deformations in green, while deformations of the zero-level set meshes are rendered in blue.

handles at training time; consequently, it is not suitable for a more general framework, and cannot be applied to our set of experiments. We conduct a separate comparison on mesh deformation between our method and NSDP in Section 7.3. Implicit-ARAP may be applied in multiple settings:

- Ours_{Mesh}^{MLP}) regular MLP applied to input mesh, trained on its SDF representation
- Ours_{SDF}^{Inv}) invertible MLP applied to input SDF field
- Ours_{SDF}^{MLP}) regular MLP applied to the zero level set mesh of an input SDF

Hardware and hyperparameters All of our experiments were run on a desktop computer with a 12GB NVIDIA RTX4070Ti GPU. Achieving efficiency in this setting allows us to show that our method is suitable for consumer grade hardware (and therefore end-user applications). The hyperparameters we used for the deformation tasks and the SDF fitting are listed in Section 7.1. We sample patch points via the sphere random uniform distribution (see Figure 10).

Metrics To evaluate **meshing accuracy** with respect to some continuous implicit surface, we leverage the SDF f_θ and compare the respective level set value l_i to $f_\theta(p)$ for surface points p in the set of patches $\{\mathcal{P}_i\}_{i=1}^n$:

$$E_{\text{patch}} = \max_{i=1}^n \max_{p \in \mathcal{P}_i} |f_\theta(p) - l_i|. \quad (3)$$

In practice, the innermost max is estimated by evaluating the point-wise error for several points sampled from the triangles of the patch mesh $\mathcal{P}_i = (V_i, F)$. We chose to aggregate via maximum rather than mean, because a single outlier can create severe artifacts in the patch.

We use four metrics to quantitatively evaluate the computed deformations, considering both global and local aspects of the geometry. First, we consider the **percent error in volume and area** of the deformed geometry with respect to the original one. Given a surface S and its deformed version S' , these are computed as

$$E_{\text{vol}} = \frac{|V_S - V_{S'}|}{V_S}, \quad E_{\text{area}} = \frac{|A_S - A_{S'}|}{A_S}, \quad (4)$$

where $V_{\mathcal{X}}$ and $A_{\mathcal{X}}$ indicate volume and surface area of shape $\mathcal{X}$. In order to evaluate the distortion induced on the input geometry, we use two distinct local criteria: **edge lengths and face angle errors**. To obtain consistent values across all experiments, we provide the former as a percentage of the

longest edge in the source mesh. Specifically, the error is computed as

$$\text{EL} = \frac{1}{|E|} \sum_{(u,v) \in E} \frac{||u - v|| - ||d(u) - d(v)||}{\max_{e \in E} ||e||}. \quad (5)$$

For the face angles, we compare the corresponding inner angles of source and deformed triangular faces:

$$\text{FA} = \frac{1}{3|F|} \sum_{f \in F} \sum_{(u,v) \in E(f)} \left| \cos^{-1} \left(\frac{u \cdot v}{||u|| ||v||} \right) - \cos^{-1} \left(\frac{d(u) \cdot d(v)}{||d(u)|| ||d(v)||} \right) \right|. \quad (6)$$

Both metrics require consistent connectivity between vertex sets. To implement them for the neural field pipeline, we extract the zero level set mesh from f_θ and forward-deform its vertices.


Table 1: Average runtime, angle error (FA) and percent errors in volume, area, and edge lengths (EL) for our neural field deformation experiments. Explicit ARAP baselines are run using the zero level set marching cubes mesh as input.

	TFD					DeFAUST				
	Ours	NFGP	ARAP	Elastic	SR-ARAP	Ours	NFGP	ARAP	Elastic	SR-ARAP
Volume	0.00%	6.66%	9.40%	9.01%	5.60%	0.00%	15.15%	17.36%	17.33%	12.09%
Area	0.85%	5.30%	0.29%	0.30%	3.97%	0.63%	20.02%	16.44%	16.41%	8.67%
EL	3.26%	3.34%	0.33%	0.38%	3.44%	2.41%	8.71%	5.72%	5.73%	8.45%
FA	4.541°	3.516°	0.358°	0.379°	4.615°	3.399°	7.577°	1.677°	1.683°	6.833°
Time	2m:48s	14h:26m	9m:51s	9m:44s	10m:27s	5m:36s	14h:26m	8m:11s	8m:38s	7m:06s

4.1 Neural field deformation

The problem of deforming neural fields with handle guidance was first introduced by Yang et al. [57], but the literature is missing follow-up proposals of significant improvements over their work. Berzins et al. [3] hint at neural shape editing as one of the applications of their method, but a complete implementation is not available. Other baselines are obtained by applying explicit methods on the zero level set of a neural field, although these methods do not preserve the neural field information.

Neural Field Pipeline



Architecture Following past literature [57, 38], we define the deformed field as $h(x) = f_\theta(d_\phi^{-1}(x))$. Intuitively, we obtain the SDF value by mapping the query point in the deformed space x to its image x' (s.t. $x = d_\phi(x')$) in the source space, via the inverse of the deformation. Therefore, we employ an invertible MLP architecture based on coordinate splitting, originally proposed by Cai et al. [8]. Since this architecture is derived from the NICE model [13], it retains its volume-preservation property, a notoriously useful prior for deformations (see Table 1). The network is composed of 6 coordinate splitting layers, where the individual coordinate processing blocks are implemented as 3-layer MLPs with Softplus activation, a hidden dimensionality of 256, and 6-frequencies Fourier features encoding [52]. We provide additional details about the invertible MLP architecture in Section 6.2.1.

Results We show averaged quantitative results for our method and baselines in Table 1. For true implicit methods (Ours and NFGP), we compute volume and area directly from the deformed field h . Since the EL and FA metrics require consistent connectivity between source and deformed shape, we compute them by deforming the zero level set mesh of f_θ ; the same holds for the discrete methods ARAP, Elastic and SR-ARAP. We used marching cubes resolution $R = 512$. From the data presented in the table, we observe that our method achieves optimal volume error due to the volume-preserving network architecture. Combining the results in Table 1 with the visualizations in Figure 5, we can appreciate how our method reliably yields plausible results in a fraction of the time required by the baselines, especially NFGP. Moreover, as the results over DeFAUST show, we can easily double the training iterations to achieve even better results in more challenging cases. The visible handle error in some of our outputs stems from our choice of loss balancing weights,

which we set to achieve the best results on average: if a more accurate handle fit is needed, the loss weights can be re-balanced, as we show in Section 7.6. Lastly, Figure 4 shows the benefits of both computing local patches for all level sets and employing the ARAP regularization for our deformation network. We use the deformed SDF gradient norm as a measure of preservation of the signed distance properties. We further discuss this aspect in Section 7.5.

4.2 Local patch meshing

The final section of our experiments is devoted to evaluating our local patch meshing algorithm. We begin by highlighting that **our method should not be considered as a drop-in replacement for marching cubes**: even by sampling a very large number of patches, it is unlikely to cover the entire surface, and a set of largely overlapping patches is not in general a useful representation for the surface. Instead, our method generates discretizations of local surface regions, and we are interested to a) verify how accurately these patches represent the underlying geometry and b) provide indications on how to select useful radius and density values for deformation tasks.




Figure 6: Accuracy evaluation of local patching against marching cubes. We construct the patch meshes in this experiment to match the total vertex count and average edge length of the corresponding marching cubes triangulation. The per-patch vertex count is fixed at 30 for all resolutions.

Reconstruction Ability In Figure 6, we compare marching cubes meshes to local patch meshes constructed with similar vertex count and average edge length. The visualizations show the change in “coarseness” of our representation wrt the resolution, while the graph provides some insights to the accuracy of local patch meshing. The line plot shows the *average* approximation error $\hat{E}_{\text{patch}} = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{p \in \mathcal{P}_i} |f_{\theta}(p) - l_i|$, with the shaded areas covering the entire region between $\hat{E}_{\text{patch}}$ and the maximum error E_{patch} (Equation (3)). For MC, these metrics can be computed by considering the entire mesh as a patch (*i.e.*, $n = 1$ and $\mathcal{P}_1$ is the marching cubes mesh). Despite the marching cubes line hinting at a lower deviation from the mean, our method achieves much lower average error even for coarse patches. This is because the patch vertices are mapped exactly onto the surface, while marching cubes places triangles based on how the surface crosses the sampled voxels; therefore, our error only depends on the overall size of the patch relative to the “flat-ness” of the approximated local surface region.




Figure 7: Variation of deformation error metrics wrt patch radius (left, fixed density = 30) and density (right, fixed radius = 0.03), averaged over TFD dataset.

Approximation Error and Deformation Quality

The inset figure (top right) shows how patch radius and density impact the approximation error (average over TFD shapes). We observe negligible change when varying the patch density, where values above 50 do not result in significant accuracy improvements. This is also the case for deformation error metrics (Figure 7, right). Given that increasing density results in additional cost in deformation time and memory (inset figure, bottom right), a conservative choice appears to be the best one. In our experiments, radius values around 0.01 would usually provide geometrically meaningful patches without significant artifacts (we use absolute units for the radius since input shapes are normalized in the unit cube). On the other hand, while smaller radius values reduce the approximation error, the patches tend to collapse to single points as their radius approaches zero. This results in a less expressive representation of the local geometry, which negatively affects the deformation results as showed in the left plot of Figure 7. Lastly, overly large patches also cause a degradation in performance due to the increase in approximation error. In Section 7.8, we provide visualizations of Implicit-ARAP outputs as the patch radius and density change. The results are visually consistent with the quantitative evaluation presented here.





Figure 8: LPM and MC as underlying triangulation for deformation tasks. Our sampling is both memory efficient and expressive.




Figure 9: Mesh deformation results for multiple patch distributions. The final results are qualitatively identical and the variations in metrics are negligible in most cases. See Figure 6 for the input shape.

Discretization In Figure 8, we present qualitative results of mesh deformation using both patches and marching cubes (MC) as underlying discretizations for ARAP energy. We show three different variations for MC: zero level set only with resolution $R = 256$, zero level set only with resolution $R = 128$, and meshing all level sets with $R = 128$ to make the computed energy similar to that of our method. The last option leads to meshes with very high triangle count especially for higher SDF values, which increases the memory requirement and limits the resolution $R = 128$. Consequently, the optimization converges to a rigid transformation, probably due to the coarse-ness of the level sets'

representation. Using only the zero level set, the model seems unable to properly apply the ARAP prior to the deformation, for resolutions $R = 256$ and $R = 128$. This is likely due to using a single discretization during the optimization process: without seeing multiple possibilities (e.g. random patches), the continuous deformation model may find local optima that “cheat” the ARAP loss without actually resulting in local rigidity. Overall, our patching approach appears more stable, reliable, and efficient. Moreover, the results presented in Figure 9 suggest that the variation in error metrics due to the choice in patch point sampling is negligible.

5 Conclusions

Discussion We presented a novel way to apply as-rigid-as-possible deformations to neural fields which is highly efficient and more flexible and robust than previous work. To this end, we proposed to mesh patches from several isosurfaces of a signed distance field and then compute the energy on those to regularize a deformation field encoded in a neural network. This has important advantages because it detaches the computational complexity from the resolution and allows for regularization that includes properties of the embedding space, e.g., the volume-preservation of our invertible model. The core idea can be applied seamlessly in the context of deforming high resolution meshes and neural fields: in the latter case, we employ an invertible deformation which allows to define the output neural field, at the cost of generality. The combination of these properties – directly inferring the new SDF and general deformation space – is hard to obtain due to the unpredictable possible changes in the SDF from an unconstrained deformation, but it would make for a challenging future work. In the context of mesh deformation, we believe that employing more efficient neural SDF representations provides an interesting direction for future investigation. Nevertheless, we believe our work is a valuable step in the direction of efficient and flexible editing of neural fields, and that our local discretization could be applied to solve more geometric problems in the implicit domain.

Limitations The general framework of our deformation method (“transporting” the input neural field along the inverse deformation) is adopted from previous contributions in the literature. As such, our method inherits the common limitation of lacking exact guarantees on the preservation of any properties of the input field (for instance, we discuss gradient norm preservation in Section 7.5). Additionally, the optimization of the ARAP loss is very dependent on the discretization: while our experiments showed that our patches are a valid choice, it would be beneficial to improve this representation with the goals of a) removing hyperparameters and b) increase surface coverage. Lastly, concerning mesh deformation (see Section 7.2), our method’s constant time scaling is only beneficial for high resolution meshes: further improving the optimization time would allow our method to also target low/medium resolution meshes and improve its generality.

Acknowledgments

This work was partially supported by the Italian Ministry of Education, Universities and Research under the grant *Dipartimenti di Eccellenza 2023-2027* of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca and by the PRIN 2022 project *GEOPRIDE Geometric primitive fitting on 3D data for geometric analysis and 3D shapes*. We acknowledge the support of NVIDIA Corporation with the RTX A5000 GPUs granted through the Academic Hardware Grant Program to the University of Milano-Bicocca for the project *Learned representations for implicit binary operations on real-world 2D-3D data*. This work is further supported by the MUR FIS2 grant n. FIS-2023-00942 “NEXUS” (cup B53C25001030001), and by Sapienza University of Rome via the Seed of ERC grant “MINT.AI” (cup B83C25001040001).

References

- [1] Jad Abou-Chakra, Feras Dayoub, and Niko Sünderhauf. Particlenerf: A particle-based encoding for online neural radiance fields. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5963–5972, 2024. doi: 10.1109/WACV57701.2024.00587.
- [2] Noam Aigerman, Kunal Gupta, Vladimir G. Kim, Siddhartha Chaudhuri, Jun Saito, and Thibault Groueix. Neural jacobian fields: learning intrinsic mappings of arbitrary meshes. *ACM*

- Trans. Graph.*, 41(4), July 2022. ISSN 0730-0301. doi: 10.1145/3528223.3530141. URL <https://doi.org/10.1145/3528223.3530141>.
- [3] Arturs Berzins, Moritz Ibing, and Leif Kobbelt. Neural implicit shape editing using boundary sensitivity, 2023.
 - [4] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014. doi: 10.1109/CVPR.2014.491.
 - [5] Mario Botsch and Leif Kobbelt. Real-time shape editing using radial basis functions. *Computer Graphics Forum*, 24(3):611–621, 2005. doi: <https://doi.org/10.1111/j.1467-8659.2005.00886.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2005.00886.x>.
 - [6] Mario Botsch, Robert W. Sumner, Mark Pauly, and Markus Gross. Deformation transfer for detail-preserving surface editing. In *Vision, modeling, and visualization 2006 : proceedings, November 22-24, 2006, Aachen, Germany*, pages 357 – 364, Berlin, 2006. Akademische Verlagsgesellschaft. ISBN 1-58603-688-2.
 - [7] Aljaz Bozic, Pablo Palafox, Michael Zollhöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. *Advances in Neural Information Processing Systems*, 33:18727–18737, 2020.
 - [8] Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, and Juyong Zhang. Neural surface reconstruction of dynamic scenes with monocular rgb-d camera. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
 - [9] Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. A simple geometric model for elastic deformations. *ACM Trans. Graph.*, 29(4), jul 2010. ISSN 0730-0301. doi: 10.1145/1778765.1778775. URL <https://doi.org/10.1145/1778765.1778775>.
 - [10] Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting, 2023.
 - [11] Aditya Chetan, Guandao Yang, Zichen Wang, Steve Marschner, and Bharath Hariharan. Accurate differential operators for hybrid neural fields, 2023. URL <https://arxiv.org/abs/2312.05984>.
 - [12] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, December 2020.
 - [13] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: non-linear independent components estimation. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings*, 2015. URL <http://arxiv.org/abs/1410.8516>.
 - [14] Ana Dodik, Oded Stein, Vincent Sitzmann, and Justin Solomon. Variational barycentric coordinates. *ACM Trans. Graph.*, 42(6), December 2023. ISSN 0730-0301. doi: 10.1145/3618403. URL <https://doi.org/10.1145/3618403>.
 - [15] Marvin Eisenberger, Zorah Löhner, and Daniel Cremers. Divergence-free shape correspondence by deformation. *Computer Graphics Forum*, 38(5):1–12, 2019. doi: <https://doi.org/10.1111/cgf.13785>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13785>.
 - [16] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuomorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7473–7483, 2021.

- [17] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *International Conference on Computer Vision (ICCV)*, 2023.
- [18] Janick Martinez Esturo, Christian Rössl, and Holger Theisel. Continuous Deformations of Implicit Surfaces. In Reinhard Koch, Andreas Kolb, and Christof Rezk-Salama, editors, *Vision, Modeling, and Visualization (2010)*. The Eurographics Association, 2010. ISBN 978-3-905673-79-1. doi: 10.2312/PE/VMV/VMV10/219-226.
- [19] Jiemin Fang, Junjie Wang, Xiaopeng Zhang, Lingxi Xie, and Qi Tian. Gaussianeditor: Editing 3d gaussians delicately with text instructions. In *CVPR*, 2024.
- [20] Jun Gao, Wenzheng Chen, Tommy Xiang, Clement Fuji Tsang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Learning deformable tetrahedral meshes for 3d reconstruction. In *Advances In Neural Information Processing Systems*, 2020.
- [21] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems 2020*, pages 3569–3579. 2020.
- [22] John C. Hart. Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12:527–545, 1996. URL <https://api.semanticscholar.org/CorpusID:20278>.
- [23] Qixing Huang, Xiangru Huang, Bo Sun, Zaiwei Zhang, Junfeng Jiang, and Chandrajit Bajaj. Arapreg: An as-rigid-as possible regularization loss for learning deformable shape generators. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5795–5805, Los Alamitos, CA, USA, oct 2021. IEEE Computer Society. doi: 10.1109/ICCV48922.2021.00576. URL <https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00576>.
- [24] Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. Harmonic coordinates for character articulation. *ACM Trans. Graph.*, 26(3):71–es, July 2007. ISSN 0730-0301. doi: 10.1145/1276377.1276466. URL <https://doi.org/10.1145/1276377.1276466>.
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <https://api.semanticscholar.org/CorpusID:6628106>.
- [26] Zohar Levi and Craig Gotsman. Smooth rotation enhanced as-rigid-as-possible mesh animation. *IEEE Transactions on Visualization and Computer Graphics*, 21(2):264–277, 2015. doi: 10.1109/TVCG.2014.2359463.
- [27] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [28] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [29] Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. Deepmetahandles: Learning deformation meta-handles of 3d meshes with biharmonic coordinates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12–21, June 2021.
- [30] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [31] W. Liu, Y. Wu, S. Ruan, and G. S. Chirikjian. Marching-primitives: Shape abstraction from signed distance function. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8771–8780, Los Alamitos, CA, USA, jun 2023. IEEE Computer Society. doi: 10.1109/CVPR52729.2023.00847. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR52729.2023.00847>.

- [32] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 21(4):163–169, aug 1987. ISSN 0097-8930. doi: 10.1145/37402.37422. URL <https://doi.org/10.1145/37402.37422>.
- [33] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’87, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery. ISBN 0897912276. doi: 10.1145/37401.37422. URL <https://doi.org/10.1145/37401.37422>.
- [34] Filippo Maggioli, Daniele Baieri, Zorah Löhner, and Simone Melzi. Sshade: a framework for scalable shape deformation via local representations, 2024. URL <https://arxiv.org/abs/2409.17961>.
- [35] Ishit Mehta, Manmohan Chandraker, and Ravi Ramamoorthi. A level set theory for neural implicit evolution under explicit flows. In *European Conference on Computer Vision (ECCV)*, 2022.
- [36] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [37] Yuichi Nagata and Shinji Imahori. Creation of dihedral escher-like tilings based on as-rigid-as-possible deformation. *ACM Transactions on Graphics*, 43(2):1–18, 2024.
- [38] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019.
- [39] Tiago Novello, Vinícius da Silva, Guilherme Schardong, Luiz Schirmer, Hélio Lopes, and Luiz Velho. Neural implicit surface evolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14279–14289, October 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Novello_Neural_Implicit_Surface_Evolution_ICCV_2023_paper.html.
- [40] Francesco Palandra, Andrea Sanchietti, Daniele Baieri, and Emanuele Rodolà. Gsedite: Efficient text-guided editing of 3d objects via gaussian splatting, 2024. URL <https://arxiv.org/abs/2403.05154>.
- [41] Xiaoyu Pan, Jiaming Mai, Xinwei Jiang, Dongxue Tang, Jingxiang Li, Tianjia Shao, Kun Zhou, Xiaogang Jin, and Dinesh Manocha. Predicting loose-fitting garment deformations using bone-driven motion networks. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022.
- [42] Nikolas Paries, Patrick Degener, and Reinhard Klein. Simple and efficient mesh editing with consistent local frames. In *15th Pacific Conference on Computer Graphics and Applications (PG’07)*, pages 461–464, 2007. doi: 10.1109/PG.2007.43.
- [43] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [45] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [46] Thomas W. Sederberg and Scott R. Parry. Free-form deformation of solid geometric models. *SIGGRAPH Comput. Graph.*, 20(4):151–160, August 1986. ISSN 0097-8930. doi: 10.1145/15886.15903. URL <https://doi.org/10.1145/15886.15903>.

- [47] Thomas W. Sederberg and Scott R. Parry. Free-form deformation of solid geometric models. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '86*, page 151–160, New York, NY, USA, 1986. Association for Computing Machinery. ISBN 0897911962. doi: 10.1145/15922.15903. URL <https://doi.org/10.1145/15922.15903>.
- [48] Dario Seyb, Alec Jacobson, Derek Nowrouzezahrai, and Wojciech Jarosz. Non-linear sphere tracing for rendering deformed signed distance fields. *ACM Trans. Graph.*, 38(6), November 2019. ISSN 0730-0301. doi: 10.1145/3355089.3356502. URL <https://doi.org/10.1145/3355089.3356502>.
- [49] Nicholas Sharp et al. Polyscope, 2019. www.polyscope.run.
- [50] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.
- [51] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 109–116, 2007.
- [52] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- [53] Jiapeng Tang, Markhasin Lev, Wang Bi, Thies Justus, and Matthias Nießner. Neural shape deformation priors. In *Advances in Neural Information Processing Systems*, 2022.
- [54] Holger Theisel. Exact isosurfaces for marching cubes. *Computer Graphics Forum*, 21(1): 19–32, 2002. doi: <https://doi.org/10.1111/1467-8659.00563>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.00563>.
- [55] Amir Vaxman, Christian Müller, and Ofir Weber. Conformal mesh deformations with möbius transformations. *ACM Transactions on Graphics (TOG)*, 34:1 – 11, 2015.
- [56] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3825–3834, 2021. URL <https://api.semanticscholar.org/CorpusID:245006265>.
- [57] Guandao Yang, Serge Belongie, Bharath Hariharan, and Vladlen Koltun. Geometry processing with neural fields. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 22483–22497. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/bd686fd640be98efaae0091fa301e613-Paper.pdf.
- [58] Yizhou Yu, Kun Zhou, Dong Xu, Xiaohan Shi, Hujun Bao, Baining Guo, and Heung-Yeung Shum. Mesh editing with poisson-based gradient field manipulation. *ACM Trans. Graph.*, 23(3):644–651, aug 2004.
- [59] Qingnan Zhou and Alec Jacobson. Thingi10k: A dataset of 10, 000 3d-printing models. *CoRR*, abs/1605.04797, 2016. URL <http://arxiv.org/abs/1605.04797>.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [\[Yes\]](#)

Justification: Our experiments prove the efficiency of Implicit-ARAP. Additionally, our method appears robust to drastic deformations, always providing natural-looking results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are discussed in the Conclusions section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We did not provide theoretical results in our paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We provide a complete and detailed description of our algorithm across the main and supplementary material, including components of our method adapted from previous literature.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: The code and part of our data (two TFD shapes) are provided in anonymized format as part of our supplementary materials.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All the necessary details are provided. Additionally, we presented results discussing our choices for method-specific hyperparameters (patch radius and density).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The specific case of handle-guided shape deformation does not support large-scale analysis, as also attested by past contributions. The core reason is the lack of established datasets.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Our method is designed to run in end-user applications, and we provided various evaluations of its time/memory performance on consumer-grade hardware (specified in the main paper).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We reviewed the Code of Ethics and acknowledge our paper's conformity.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[No\]](#)

Justification: While empowering 3D artists with faster and more robust tools for their work may be seen as a positive societal impact, we did not feel it was relevant enough to be discussed in the paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our experiments employ a small set of CC-0 3D models found on various public repositories. In our opinion, their redistribution poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The data we employed was obtained from DeFAUST, Thingi10k and the Stanford 3D scanning repository, which were all properly cited in the paper. We provided license names for each object in the supplementary material.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: Our submission includes a complete code release with instructions and minimal samples of data necessary to test the proposed method.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Contents

1	Introduction	1
2	Related work	2
3	Method	3
3.1	Local patch meshing	3
3.2	Deformation model	4
3.3	Optimization	4
4	Experiments	5
4.1	Neural field deformation	7
4.2	Local patch meshing	8
5	Conclusions	10
6	Additional details	23
6.1	Local patch meshing	23
6.2	Model	23
6.2.1	Network architectures	23
6.3	Training procedure	24
6.4	Implementation	25
7	Additional experiments	25
7.1	Hyperparameters	25
7.2	High resolution mesh deformation	26
7.3	Comparison to NSDP	28
7.4	Ground truth evaluation	28
7.5	SDF preservation	29
7.6	Loss balancing	29
7.7	Matching ARAP resolution	30
7.8	Correlation of deformation results and patch parameters	30
8	Data license	31

6 Additional details

6.1 Local patch meshing




Figure 10: Different distributions for sampling the 2D disk (100 pts per patch).

We visualize four possible types of discretization in Figure 10, such as uniform sampling in polar coordinates, normalized normal sampling, linear (deterministic) sampling, and combining uniform coordinates with normal radius. Even though the patches are visibly different both in terms of point distribution and triangle appearance, the specific choice has little influence on our method, as we show in Figure 9.

Algorithm 1 presents the rejection/projection algorithm to sample the zero level set of an SDF introduced by Yang et al. [57]. In our implementation, several parts of the algorithm are parallelized for efficiency: for instance, the inner loop samples a large number of points x at once, retaining and projecting those with absolute distance $< \tau$. By performing enough sampling attempts in a single iteration, the algorithm can frequently terminate in a single step (*i.e.*, by retaining at least N points among those that were sampled).

Algorithm 1 SDF zero level set sampling.

```

1: procedure REJECTPROJECTSAMPLING( $f_\theta, N, \tau, t_{\max}$ )
2:    $S \leftarrow \emptyset$ 
3:   while  $|S| < N$  do
4:      $x \sim \mathcal{U}([-1; 1]^3)$  ▷ Rejection step
5:     while  $|f_\theta(x)| > \tau$  do ▷ Sample  $x$  from bounded 3D domain
6:        $x \sim \mathcal{U}([-1; 1]^3)$  ▷ Ensure close to surface
7:     end while
8:     end while
9:     for  $t = 1 \rightarrow t_{\max}$  do ▷ Projection step
10:       $x \leftarrow x - f_\theta(x) \frac{\nabla f_\theta(x)}{\|\nabla f_\theta(x)\|}$  ▷ Iterate closest surface point
11:    end for
12:     $S \leftarrow S \cup \{x\}$ 
13:  end while
14:  return  $S$ 
15: end procedure

```

6.2 Model

6.2.1 Network architectures

Shape model To represent the input shape internally to our deformation algorithm, we adopt a neural SDF model proposed in previous literature [50, 52]. This model is suitable to our application due to its efficiency on consumer-grade hardware and robustness with respect to the input geometry. The SDF is represented via a MLP network with 8 layers, a hidden dimensionality of 256, a residual connection at the fourth layer, 6-frequencies Fourier features encoding, and Softplus activation. This network is optimized via eikonal training, originally proposed by Gropp et al. [21], which employs the following four losses:

$$L_{\text{zero}} = \mathbb{E}_{x \in \mathcal{S}_0} |f_\theta(x)|, \quad (7)$$

$$L_{\text{eikonal}} = \mathbb{E}_{x \in \mathbb{R}^3} \left\| \|\nabla f_\theta(x)\| - 1 \right\|^2, \quad (8)$$

$$L_{\text{normals}} = \mathbb{E}_{x \in \mathcal{S}_0} \left(1 - \frac{\nabla f_\theta(x) \cdot \mathbf{n}(x)}{\|\nabla f_\theta(x)\| \|\mathbf{n}(x)\|} \right), \quad (9)$$

$$L_{\text{penalty}} = \mathbb{E}_{x \in \mathbb{R}^3} \exp(-\alpha |f_\theta(x)|). \quad (10)$$

Defined on a discretization of the represented surface $\mathcal{S}_0$. Intuitively, these respectively constrain the network to: **1)** vanish on surface points (sampled from the input mesh triangles) **2)** have unitary norm of gradient **3)** have gradient aligned with surface normals (indicated by $\mathbf{n}(x)$ for surface point x), and **4)** have minimal zero level set, to avoid artifacts due to under-determination. We list the values for loss weights and α which we employed in our implementation in Table 3. The Adam optimizer runs for a total of 10000 steps and uses a starting learning rate of 10^{-4} and a scheduler which halves it at steps 1000, 2000, and 5000.

Deformation model. For the invertible network employed in the neural field deformation pipeline, we use 6 coordinate splitting layers, where the individual coordinate processing blocks are implemented as 3-layer MLPs with Softplus activation, a hidden dimensionality of 256, and 6-frequencies Fourier features encoding [52]. Each of these layers splits and combines coordinates according to the layer index i , by selecting the “focus” coordinate $w = p_{i \bmod 3}$, where $p = (x, y, z)$ is the input vector. Each layer predicts a translation of the “focus” coordinate and a 2D roto-translation of the two others, which we progressively aggregate to obtain $R_\phi(x)$ and $t_\phi(x)$. The layer architecture for this model is visualized in Figure 11. Contrarily to the Lipschitz-continuous MLP used in NFGP [57], this architecture allows for an analytic expression of its inverse and thus is more efficient, as it does not require fixed point iterations for inversion.

For the mesh deformation pipeline, we use a standard MLP composed of 8 linear layers with a hidden dimensionality of 256 and Softplus activation. We apply Fourier features encoding with 6 frequencies at the input layer and a residual connection at the 4th layer. For both networks, we adopt the neural deformation initialization scheme of [43, 8], which allows the initial state of the model to predict the identity transformation without causing symmetry breaking failures or numerical instability interacting with the Adam optimizer (which are common when the weight matrices are entirely initialized to zero).

6.3 Training procedure

Algorithm 2 shows all the high-level operations performed in a single step of Implicit-ARAP training. First, a single flat patch is constructed via the sampling procedure described in Section 3.1. Then, we construct a set of origin points for surface patches: we include 1) the handle sources, 2) a quasi-uniform sampling of the zero level set (see Algorithm 1) and 3) a uniform sampling of the bounded embedding space. The for loop in lines 7–10 aligns the flat patch at each origin and deforms it according to the local geometry. This operation is performed in parallel for all patches to avoid a bottleneck. Finally, the patch mesh is constructed as the union of all surface patches, and its vertices are deformed via d_ϕ . The computation of the ARAP and handle losses is described in Section 3.3.




Figure 11: Diagram for the forward (top) and inverse (bottom) passes of the invertible MLP architecture we employ in our neural field deformation pipeline.




Figure 12: Results for method and baselines for a non-bijective deformation, which invertible networks (Inv, NFGP) cannot represent.

Algorithm 2 Implicit-ARAP training loop.

```
1: procedure TRAIN( $T, \lambda_1, \lambda_2, n, m, \rho, k, f_\theta, d_\phi, \{(s_i, t_i)\}_{i=1}^h$ )
2:   for  $t \leftarrow 1$  to  $T$  do
3:      $V, F \leftarrow \text{DISKSAMPLE}(\rho, k)$  ▷ Section 3.1, “Sampling”
4:      $O \leftarrow \{s_i\}_{i=1}^h$  ▷ Handle sources as origins
5:      $O \leftarrow O \cup \text{REJECTIONSAMPLE}(f_\theta, m - h)$ 
6:      $O \leftarrow O \cup \text{UNIFORMSAMPLE}([-1; 1]^3, n - m)$ 
7:     for  $j \leftarrow 1$  to  $n$  do ▷ Section 3.1, “Projection”
8:        $V_j \leftarrow \text{ALIGN}(V, o_j, f_\theta)$ 
9:        $V_j \leftarrow \text{PROJECT}(V_j, f_\theta)$ 
10:    end for
11:     $\mathcal{P} \leftarrow \{(V_j, F)\}_{j=1}^n$ 
12:     $\mathcal{P}' \leftarrow \{(d_\phi(V_j), F)\}_{j=1}^n$  ▷ Section 3.2
13:     $L_{\text{arap}} \leftarrow \text{ARAPLOSS}(\mathcal{P}, \mathcal{P}')$  ▷ Section 3.3
14:     $L_{\text{handle}} \leftarrow \frac{1}{h} \sum_{i=1}^h \|d_\phi(s_i) - t_i\|^2$  ▷ Section 3.3
15:     $L \leftarrow \lambda_1 L_{\text{handle}} + \lambda_2 L_{\text{arap}}$ 
16:     $\text{OPTIMIZE}(L; \phi)$  ▷ Compute gradients, ADAM step
17:  end for
18: end procedure
```

6.4 Implementation

We implemented our algorithm in Python, relying on PyTorch [44] for neural network primitives, linear algebra and automatic differentiation. In addition, we used Polyscope [49] for visualization, extending its GUI with functionalities for easy point picking, which we used to design deformation experiments. While our viewer renders a 3D triangle mesh extracted with marching cubes for the sake of efficiency, the points selected on the shape by the user are mapped exactly onto the implicit surface via iterations of the SDF closest point equation (Equation (1)), allowing to select an arbitrary set of accurate handles. For a given set of points, the user can then specify an affine transformation and save both the resulting handle transforms and the original positions. Our codebase is available at this url.

7 Additional experiments

7.1 Hyperparameters

Where unspecified, all of our deformation experiments were run using the hyperparameters showed in Table 2. To train the neural SDFs, we used the architecture described in Section 6.2.1 with the hyperparameters listed in Table 3.

Table 2: Hyperparameter values for the deformation procedures. k and ρ refer to patch density and radius, respectively, while λ_i are the loss balancing weights. n and m define the number of sampled patches per training step (see Algorithm 2).

Data	Mesh		Field	
	TFD	DeFAUST	TFD	DeFAUST
Steps	1000	2000	1000	2000
LR	10^{-3}	10^{-3}	10^{-3}	10^{-3}
λ_1	10^3	10^3	10^3	10^3
λ_2	10^1	10^4	10^3	10^4
k	30	30	30	30
ρ	0.03	0.01	0.03	0.01
n	2000	2000	2000	2000
m	1000	1000	1000	1000

Table 3: Hyperparameter values for the SDF fitting procedure.

α	λ_{zero}	λ_{eikonal}	λ_{normals}	λ_{penalty}
100	3000	100	50	3000





Figure 13: Qualitative results of our method in comparison to multiple baselines for mesh deformation, on TFD examples.

7.2 High resolution mesh deformation

Our pipeline can be applied for classic mesh deformation, where the trained network is used to deform the vertices of an input mesh (internally converted to a neural SDF for optimization).

Architecture For the mesh deformation pipeline, we do not require invertibility as we are simply interested in the forward deformation of a set of points in the zero level set of f_θ . Therefore, we use a standard MLP composed of 8 linear layers with a hidden dimensionality of 256 and Softplus activation. We apply Fourier features encoding with 6 frequencies at the input layer and a residual connection at the 4th layer. As mentioned in the main manuscript, the network is again initialized to predict the identity transformation [43, 8]. This architecture is $\sim 1.53 \times$ more efficient to query than the invertible one, resulting in faster runtime for the deformation procedure (excluding the internal SDF fitting). Figure 12 is also related to this section, as it shows how the standard MLP we use in this task may represent deformations such that bijections between source and deformed shape are not possible (*e.g.* changes of topology).

Results Well-established explicit methods [51, 9, 26] are very efficient for most use-cases, but they do not scale to meshes with millions of vertices due to super-linear time and memory costs. By representing the input geometry implicitly in the weights of a neural network, and only computing Laplacian edge weights for small triangle patches, our method achieves runtime and VRAM usage independent of the input size.



As highlighted in Table 4, Implicit-ARAP’s runtime is comparable to that of the explicit baselines. However, the time scaling of our method is constant, therefore it is expected to remain efficient even at resolutions higher than those employed in our evaluation. Moreover, most ($\sim 72\%$) of our runtime is spent fitting a neural SDF to the input mesh: employing a faster procedure for this step could greatly improve our performance. The remainder of Table 4 summarizes the performance of our method: Implicit-ARAP yields a clear improvement over SR-ARAP and NFGP. However, ARAP and Elastic achieve much better metrics than our method for area, EL and FA. Since these two methods provide the best solution (in a least-squares sense) for local rigidity of the given input mesh under the handle constraints, this is to be expected. However, including useful priors such as smoothness (SR-ARAP) immediately aligns the scale of error values to ours and NFGP.

Combining these results with the qualitative evaluation we present in Figures 13 and 14 provides a clear picture of the accuracy, robustness and efficiency of our method, which consistently yields accurate results with minimal artifacts. We point out that the results of ARAP baselines for the cubes experiment (Figure 13) are correct, as the 2nd and 4th cubes are unconstrained and may be mapped arbitrarily. Implicit-ARAP and NFGP, on the other hand, exploit the spectral bias of neural networks to propagate the handle maps smoothly over the whole domain. However, NFGP deforms the individual cubes into trapezoids more evidently than our method. The last advantage in using an implicit method like ours lies in its independence from the quality of the input shape’s connectivity: for example, the CGAL implementation of explicit baselines failed to run using the buddha mesh (see Figure 8).




Figure 14: Qualitative results of our method in comparison to multiple baselines for mesh deformation, on DeFAUST examples.

Table 4: Average runtime, angle error (FA) and percent errors in volume, area, and edge lengths (EL) for our high-resolution mesh deformation experiments. The time for our method and NFGP includes 4m:44s required for the initial neural SDF fit, meaning our deformation phase runs in 1m:49s on average for 1000 steps.

	TFD					DeFAUST				
	Ours	NFGP	ARAP	Elastic	SR-ARAP	Ours	NFGP	ARAP	Elastic	SR-ARAP
Volume	4.27%	8.41%	11.82%	11.42%	8.24%	0.99%	15.02%	12.41%	12.25%	11.69%
Area	2.83%	7.04%	0.23%	0.25%	3.46%	0.82%	20.25%	0.45%	0.45%	5.36%
EL	0.76%	0.87%	0.12%	0.12%	0.89%	1.06%	8.71%	0.37%	0.38%	3.39%
FA	3.448°	4.178°	0.486°	0.495°	4.743°	1.517°	8.859°	0.517°	0.523°	5.628°
Time	6m:33s	14h:31m	8m:06s	8m:00s	9m:07s	8m:22s	14h:31m	7m:15s	8m:27s	7m:11s

7.3 Comparison to NSDP

In this section, we compare Implicit-ARAP to NSDP [53], a recent contribution in neural methods for 3D mesh deformation. We highlight the main differences between the two methods in Table 6. Differently from Implicit-ARAP, NSDP is data-driven, therefore to ensure fairness we carry out this evaluation on the original test set employed by the authors, the DeformingThings4D dataset [27]. The results are available in Table 5. We observe that Implicit-ARAP, leveraging the local rigidity prior, achieves a clear advantage for all deformation quality metrics. However, on the low/medium resolution data available in DT4D, NSDP can compute deformations faster than Implicit-ARAP, by pre-training the model on large amounts of data. This comes at the cost of generality, as NSDP requires that the handles employed at inference time be “semantically” the same as those used during training (*e.g.*, limb extremities of a humanoid character).

Table 5: Comparison of Implicit-ARAP and NSDP [53] performance on DT4D [27].

	Volume	Area	EL	FA	Time	VRAM
Ours	0.58%	2.04%	2.03%	4.728°	6m:33s	1.3GB
NSDP	7.80%	5.76%	3.59%	6.603°	0m:18s	2.1GB

Table 6: Highlights of differences between our method and Neural Shape Deformation Priors [53].

	Implicit-ARAP	NSDP
<i>Training</i>	Trained for each input	Pretrained on large dataset
<i>Handle set</i>	Any	Defined at training time
<i>Evaluation data resolution</i>	Medium/high (50k-500k)	Medium/low (10k-50k)
<i>Deformation prior</i>	Local rigidity (ARAP)	Linear blend skinning (data prior)
<i>Inference-time optimization</i>	Yes	No

7.4 Ground truth evaluation

Throughout our evaluation, we ran deformation experiments which were obtained from 3D shape pairs (*i.e.*, shape pairs from FAUST [4], animations from DT4D [27]). In these cases, one may leverage the “target” shape information to evaluate how closely the algorithm approximates the deformation prior showed in the data, by computing the difference between the deformed shape and the target shape [53]. This provides an additional criterion of evaluation for deformation algorithms, which while not an “objective” metric of correctness, can provide additional insight into the behavior of a particular method. The results are available in Table 7.

Table 7: Evaluation of Implicit-ARAP’s ability to approximate target shapes, when available for our deformation experiments. For reference, NSDP averaged $7.79 \cdot 10^{-4}$ on DT4D.

DeFAUST	DT4D
$3.54 \cdot 10^{-3}$	$3.44 \cdot 10^{-3}$

7.5 SDF preservation

We are interested in evaluating to what degree Implicit-ARAP deformations preserve the signed distance properties of the input field. Like past contributions, our method lacks theoretical guarantee of exact preservation, and obtaining such a guarantee is an open problem. In Figure 4, we provide an ablation experiment showing that our formulation is beneficial unitary gradient norm preservation. We complete that result with Table 8, which presents the average of the same data over all our experiments.

To give a practical measure of how well the SDF properties are preserved, we provide sphere-traced [22] renders of SDFs deformed with Implicit-ARAP in Figure 15. We used the sphere tracing implementation from Polyscope [49]. We note that, while ad-hoc rendering methods for deformed fields such as [48] could be used, the scope of this work allows us to be satisfied with these basic results. The results show that the deformed fields are rendered without significant artifacts.

Table 8: SDF gradient norm preservation, averaged over all our experiments. Only true implicit methods are evaluated. The mean and STD of the deformed field gradient norm are evaluated over a dense uniform sampling (1M points) of $[-1; 1]^3$. Recall that for exact SDFs $\|\nabla f\| = 1$.

	Mean $\ \nabla f_\theta \circ d_\phi^{-1}\ $	STD $\ \nabla f_\theta \circ d_\phi^{-1}\ $
Ours	1.016	0.146
NFGP	1.029	0.290
Input	0.999	0.077



Figure 15: Sphere-traced renders of SDFs deformed with Implicit-ARAP. The deformation inputs may be found in Figures 5 and 13.




Figure 16: Alternative visualization of the ablation study presented in Figure 4.

7.6 Loss balancing

As we mentioned in the main manuscript, in our neural representation setting we cannot ensure a 100% accurate handle fit, as we have to optimize the input handles via gradient descent, balancing their contribution with that of the ARAP loss. As a result, especially in difficult cases with substantial deformations, the optimization may favor a better preservation of local rigidity and sacrifice some

measure of handle accuracy. In Figure 17, we show that it is possible to re-balance the loss term and achieve accurate handle fit, at the cost of a less natural-looking deformation.




Figure 17: Decreasing the weight of the ARAP loss (λ_2) allows to balance handle satisfaction and deformation quality. λ_1 is set to 1000 as for all other experiments.

7.7 Matching ARAP resolution

Furthermore, we explore the connection between our method and the original ARAP through the results in Figure 18. ARAP results are generally sensitive to the input mesh resolution: different discretizations of the same shape will result in outputs which, while technically correct, exhibit different properties. For low resolution meshes, local rigidity will be optimized between large triangles, often resulting in behaviours typical of skeleton-based linear blend skinning deformations. On the other hand, high resolution meshes will exhibit local rigidity in more localized regions. This is apparent from the top row of Figure 18.

To determine if Implicit-ARAP would exhibit the same phenomenon, we ran our method varying radius and density so that the average edge length of our training patches would approximate that of the ARAP input meshes. Our results show that this is not the case, as the variations in our outputs are insubstantial.




Figure 18: Comparison of our method to ARAP with matching resolutions (showed as a percentage of full resolution). The training patches are constructed with radius 0.03 (avoiding degenerate patches) and manually selecting the density k which best approximates the average edge length of the corresponding ARAP input mesh (which we obtained by decimating the sculpture mesh at multiple resolutions). We also show the total runtime t (without SDF fit in our case).

7.8 Correlation of deformation results and patch parameters

Lastly, Figure 19 completes the discussion about correlation of patch parameters and deformation performance Section 4.2. As we anticipated from the quantitative evaluation, we observe that small patches (radius below 0.01) result in artifacts due to the insufficient influence of the ARAP energy in

the optimized deformation. Viceversa, when the patches are too large, artifacts arise from interaction of geodesically distant regions (*e.g.*, the front left paw with the body, or the neck with the spine). Finally, coherently with the results in Figure 7, variations in patch density do not have a significant impact on visual quality.




Figure 19: Qualitative results for Implicit-ARAP, varying patch density and radius.

8 Data license

For our experiments, we gathered a small collections of 3D shapes from Thingi10k [59] and the Stanford 3D scanning repository. Then, we designed various deformation experiments for each of these shapes. Finally, we included a set of deformation experiments automatically generated from the FAUST [4] dataset. In Table 9, we list the license name of each of these assets.

Asset	Reference	License
DeFAUST shapes	Figure 14	FAUST License
Armadillo (Stanford)	Figures 12, 13 and 15	CC BY-NC 4.0
Dragon (Stanford)	Figures 1, 2, 5 and 19	CC BY-NC 4.0
Buddha (Stanford)	Figures 5, 8 and 15	CC BY-NC 4.0
Piranha Plant (Thingi10k)	Figures 13 and 15	CC BY-NC 4.0
Hand (Thingi10k)	Figure 2	CC BY-NC-SA 4.0
Sculpture (Thingi10k)	Figures 6, 9, 13 and 18	CC BY-NC 4.0
Cat (Thingi10k)	Figure 4	CC BY-NC-ND 4.0
Octopus (Thingi10k)	Figures 5 and 17	CC BY-SA 4.0

Table 9: References in the paper and license for all 3D meshes we employed in our experiments.