InfoNorm: Mutual Information Shaping of Normals for Sparse-View Reconstruction

Xulong Wang^{1,2*} Siyan Dong^{3*} Youyi Zheng^{1†} Yanchao Yang^{3,4†}

¹ State Key Lab of CAD&CG, Zhejiang University ² Chohotech co. ltd.
 ³ Institute of Data Science, The University of Hong Kong
 ⁴ Department of Electrical and Electronic Engineering, The University of Hong Kong

Abstract. 3D surface reconstruction from multi-view images is essential for scene understanding and interaction. However, complex indoor scenes pose challenges such as ambiguity due to limited observations. Recent implicit surface representations, such as Neural Radiance Fields (NeRFs) and signed distance functions (SDFs), employ various geometric priors to resolve the lack of observed information. Nevertheless, their performance heavily depends on the quality of the pre-trained geometry estimation models. To ease such dependence, we propose regularizing the geometric modeling by explicitly encouraging the mutual information among surface normals of highly correlated scene points. In this way, the geometry learning process is modulated by the second-order correlations from noisy (first-order) geometric priors, thus eliminating the bias due to poor generalization. Additionally, we introduce a simple yet effective scheme that utilizes semantic and geometric features to identify correlated points, enhancing their mutual information accordingly. The proposed technique can serve as a plugin for SDF-based neural surface representations. Our experiments demonstrate the effectiveness of the proposed in improving the surface reconstruction quality of major states of the arts. Our code is available at: https://github.com/Muliphein/InfoNorm.

Keywords: 3D from multi-view images \cdot surface reconstruction

1 Introduction

3D surface reconstruction from multi-view images is an essential task in the computer vision and graphics community, with practical applications including content creation for virtual reality and robot-scene interaction. However, indoor scenes present challenges due to the large and complex scenes captured from sparse viewpoints. These challenges encompass issues such as occlusion and ambiguity arising from limited observations. As a result, traditional Multi-View Stereo (MVS) methods [28,36], which require substantial overlap among images, may not produce satisfactory results in these scenarios.

^{*} Equal contributions (wangxulong@zju.edu.cn, siyan3d@hku.hk).

[†] Corresponding authors.

² This work was done during the author's internship at Chohotech Co. ltd..



Fig. 1: 3D scene reconstruction from sparse views on Replica [26] (first row) and Scan-Net++ [38] (second row). By enforcing the mutual information between the normals of highly correlated scene points, the proposed method can effectively enhance the reconstruction quality of the baselines (VolSDF [37] and GeoNeuS [8]).

Recently, implicit scene representations, e.g., Neural Radiance Fields (NeRFs) [21], have been proposed to encode 3D scenes into a set of neural network parameters, with subsequent works enhancing the representational capacity for geometry to enable surface extraction. The key idea involves incorporating signed distance functions (SDFs) [23] into the learned radiance field. Although NeRFs with SDF produce high-quality reconstructions of simple scenes with sufficient images, they still underperform in large and complex scenes captured from sparse viewpoints. More recent approaches, such as NeuRIS [29] and MonoSDF [40], propose using monocular geometric priors to guide the learning of the radiance field. However, they can be affected by inevitable errors and noise in the monocular geometry estimation modules.

We explore geometric regularization in a second-order perspective. We argue that both the geometry itself and the correlation of geometry among scene regions are crucial for surface reconstruction, especially when dealing with varying indoor scenes under sparse views with noisy geometry estimates. The correlation is expressed as mutual information between scene entities under random perturbations of the network weights. More explicitly, we propose to enforce the mutual information between the normals of two points in the scene that are considered geometrically correlated, where the mutual information can be efficiently computed by measuring the cosine similarity between the gradients of the normals with respect to the perturbed weights. Further, we leverage a combination of pretrained semantic and geometric features [4,7,12] to identify the correlated scene regions that endorse high mutual information. This results in an easy-to-use geometric shaping¹ technique that can be applied to any SDF-based neural radiance fields. To thoroughly evaluate its effectiveness, we test the proposed with

¹ We use the term "shaping" to refer to the process of optimizing network parameters to enforce consistencies in the scene geometry via aligning the gradients.

multiple state-of-the-art neural surface reconstruction models. The experimental results confirm the effectiveness of the proposed shaping technique, showing improved surface reconstruction quality (e.g., Fig. 1) on challenging scenes with limited views. To summarize:

- We propose to leverage second-order geometric correlation in the form of mutual information between normals to regularize the surface reconstruction of SDF-based NeRF representations.
- We develop a pipeline that efficiently enforces the mutual information of the normals of points, deemed geometrically correlated with pre-trained multimodal features, to improve the reconstruction quality of the scene.
- We verify the effectiveness of the proposed second-order regularization technique across a broad spectrum of baselines and demonstrate its usefulness as an easy-to-use plugin in improving 3D surface modeling.

2 Related Work

3D Reconstruction with neural radiance fields (NeRFs). NeRF [21] is originally designed for novel view synthesis, with its variants serving as scene representations [15, 32, 33, 35, 39] for tasks related to editing [2, 31, 44], semantic segmentation [5, 14, 16, 34], 3D shape generation [20, 27], and so on. We mainly focus on neural surface representations that enable the extraction of high-quality 3D meshes from the trained NeRFs.

Early works such as NeuS [30] and VolSDF [37] combine signed distance functions (SDFs) with volume rendering to extract high-quality surfaces. Recently, SparseNeuS [17] introduces a cascaded geometry reasoning framework that generalizes well to novel scenes. GeoNeuS [8] uses surface points from structurefrom-motion and neighbor-view patches to reconstruct surfaces, while NeuDA [3] employs a deformable anchor to adaptively encode geometric details. NeAT [19] broadens the scope of reconstruction to arbitrary topology, no longer confining to watertight surfaces. Moreover, Neuralangelo [13] uses a hash grid for position encoding and implements a coarse-to-fine optimization with numerical gradients. Despite their promising results, most of these methods need sufficient training views, and their performance declines when dealing with large and complex scenes under sparse viewpoints.

Incorporating priors into NeRFs. One can augment scene representations by introducing various priors to NeRFs. VDN-NeRF [43] normalizes the spatial feature to align with monocular features. LERF [11] introduces a new branch whose outputs align with DINO [4] and CLIP [24] features. Semantic-NeRF [42] encodes semantic information for segmentation tasks, while JacobiNeRF [34] shapes the color gradients to encode the mutual information in terms of semantic similarity from DINO [4] features for label propagation.

For surface reconstruction, GeoNeuS [8] incorporates patch similarity from neighboring images and sparse surface points derived from structure-from-motion

[25], while I²-SDF [44] combines the rendering equation with VolSDF [37] for better novel views and geometry. To cope with complex indoor scenes, NeuRIS [29] incorporates the normal priors as additional geometric supervision. Further, MonoSDF [40] uses both depth and normal to improve the surface reconstruction quality. However, these methods highly depend on the quality of the supplemented geometric information. In contrast, our method leverages the secondorder information in geometric prior, rather than directly using the prior as a first-order supervision signal, which makes it more robust to errors and noise. There are also methods utilizing strong assumptions derived for human-made scenes. For instance, ManhattanSDF [10] works under the assumption that all captured scenes adhere to the Manhattan-world concept, thus, predicting semantic segmentation to identify walls and floors. While these methods primarily focus on planar scenes, the proposed technique can work for general scenes under noisy correlation information.

3 Method

We start with a discussion on the preliminaries in Sec. 3.1 to provide the background on how we can learn from posed images an implicit scene representation (NeRF), improve 3D scene reconstruction via the combination of NeRF and SDF representations, and employ mutual information as a constraint in the NeRF training process. Following this, in Sec. 3.2, we delve into how mutual information can be used as constraints on NeRF's implied surfaces, so that we can learn better geometries for higher-quality 3D reconstruction of the scene. Finally, we elaborate on how we can employ off-the-shelf semantic and geometry features to assist the encoding of mutual information through the proposed geometric shaping framework.

3.1 Preliminaries

Neural radiance fields. NeRF [21] learns an implicit scene representation from a set of posed images. For a specific pixel in an image, NeRF samples a list of 3D points, $\{x|x_t = \mathbf{o} + t\mathbf{v}, t \in [t_{near}, t_{far}]\}$, from the camera center \mathbf{o} along the viewing direction \mathbf{v} within a bounded space. Specifically, it employs a Multilayer Perceptron (MLP) to encode the volume density and color of each sampled point. We denote all the parameters of NeRF as θ , the density function as σ , and the color function as c. The pixel's rendered color $\hat{\mathbf{C}}$ from a NeRF is calculated by the discrete volume rendering:

$$\hat{\mathbf{C}}(\mathbf{o}, \mathbf{v}; \theta) = \sum_{i=1}^{N} \omega(\mathbf{o}, \mathbf{v}, t_i; \theta) c(x_{t_i}, \mathbf{v}; \theta),$$
where $\omega(\mathbf{o}, \mathbf{v}, t_i; \theta) = \prod_{j=1}^{i-1} (1 - \alpha_j) \alpha_i.$
(1)

Here, $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$, where σ_i is the abbreviation of $\sigma(x_{t_i}; \theta)$ and $\delta_i = t_{i+1} - t_i$. A photometric loss is utilized to minimize the discrepancy between the rendered color and its ground truth:

$$L_C = ||\mathbf{C} - \hat{\mathbf{C}}||_1, \tag{2}$$

so that the parameters θ can be optimized to fit the scene. Despite the ability to handle photorealistic novel view synthesis, NeRFs suffer from noisy and unrealistic surface extraction when it comes to 3D reconstruction, as discussed in the literature [30,37].

NeRFs with signed distance functions. To produce high-quality surfaces, it is beneficial to use signed distance functions (SDFs) [6, 23] to represent geometry. Consequently, the surface can be reconstructed by extracting the zero-level set [18] of the SDF values. The combination of NeRF and SDF can be achieved by deriving the NeRF's densities from the SDF, or defining Laplace's cumulative distribution function. Accordingly, α_i in Eq. 1 can be represented by a function Φ which converts SDF values to volume densities as

$$\alpha_i = \Phi(x, f(x; \theta), i). \tag{3}$$

Generally, NeRFs with an SDF achieve improved surface reconstruction quality, where the key to maintaining the characteristics of SDF is an eikonal loss [9]:

$$L_E = \frac{1}{N} \sum_{i=1}^{N} (||\nabla f(x_i)||_2 - 1)^2.$$
(4)

Mutual information in NeRFs. Mutual information (MI) quantifies the statistical dependence between two random variables, providing an estimate of how much information they share. It has been introduced for semantic segmentation and motion modeling tasks [34, 41]. In this paper, We study the mutual information under perturbations of the NeRF parameters, and focus more on the geometry properties. Let $I(p_i)$ and $I(p_j)$ denote two pixels derived from the NeRF rendering process \mathbb{F} . And we assume that a subset of the NeRF weights, denoted by θ^D , is perturbed by a random noise vector $\mathbf{n} \in \mathbb{R}^D$ sampled from a uniform distribution on the sphere \mathbb{S}^{D-1} by a small step $\gamma \ll 1.0$. The random variables corresponding to the perturbed pixels can be written through a Taylor expansion:

$$\hat{I}(p_i) = I(p_i) + \gamma \mathbf{n} \frac{\partial \mathbb{F}_i}{\partial \theta^D},$$
$$\hat{I}(p_j) = I(p_j) + \gamma \mathbf{n} \frac{\partial \mathbb{F}_j}{\partial \theta^D}.$$

We can now characterize the mutual information between $\hat{I}(p_i)$ and $\hat{I}(p_j)$ under the joint probability distribution $\mathbb{P}(\hat{I}(p_i), \hat{I}(p_j))$. As derived in JacobiNeRF, the

mutual information can be computed as:

$$\begin{split} \mathbb{I}(\hat{I}(p_i), \hat{I}(p_j)) &= \mathbb{H}(\hat{I}(p_i)) - \mathbb{H}(\hat{I}(p_i) | \hat{I}(p_j)) \\ &= \log(\frac{1}{\sqrt{1 - \cos^2 \beta}}) + \text{const.} \end{split}$$

Here, β denotes the angle between $\partial \mathbb{F}_i / \partial \theta^D$ and $\partial \mathbb{F}_i / \partial \theta^D$. It is observed that:

$$\mathbb{I}(\hat{I}(p_i), \hat{I}(p_j)) \propto ||\cos\beta||.$$
(5)

Therefore, to enhance the MI between two random pixels induced by NeRF perturbations, one can enforce the absolute value of the cosine similarity. That is, if two pixels p_i , p_j come with high correlation (e.g., from the same object in the scene), their gradients regarding the perturbed parameters should be in the same or opposite direction.

3.2 Mutual Information Shaping of Geometry in SDF-based NeRFs

While the previous approach focuses on enforcing mutual information for image synthesis, we explore a mutual information shaping of the geometries for quality surface reconstruction in SDF-based NeRFs.

Specifically, we represent a scene via a NeRF backbone with an SDF head and a color head, as shown in Fig. 2. Note that the method to be proposed is capable of being applied to any network similar to NeRF and SDF representations. Given that the two heads partially share parameters, we denote all the parameters as θ , and use $f(x;\theta)$ and $c(x, \mathbf{v};\theta)$ to represent the SDF and color branches, respectively. To achieve mutual information shaping of the geometry, we focus on manipulating the parameters of $f(x;\theta)$, while keeping $c(x, \mathbf{v};\theta)$ unchanged. Below, we first explain why and how we compute mutual information with surface normals derived from SDF, i.e., $f(x;\theta)$, as a geometry-aware constraint. Next, we present a simple yet effective strategy that extracts regions with high correlation, which need to be encoded into the weights using the mutual information shaping technique. Finally, we elaborate on our loss function and training process to reconstruct the scene from posed images.

Mutual information among surface normals. We explicitly encode mutual information into the function $f(x; \theta)$. However, instead of computing mutual information between the density values, we proceed by manipulating the surface normal vectors that can be derived from f. Specifically, for a 3D point x on a surface of the scene, its normal vector $N(x; \theta)$ can be calculated as the partial derivative of the 3D coordinate:

$$N(x;\theta) = \frac{\partial f(x;\theta)}{\partial x}.$$
(6)

There are several reasons to manipulate surface normal vectors but not SDF values: 1) Shaping the SDF values with respect to a subset θ^D of the parameters may not effectively impose the constraints on all parameters. However, the



Fig. 2: An overview of the pipeline, where we apply mutual information shaping on the geometric branch to enforce consistencies that help enhance the surface reconstruction. Specifically, the model consists of the NeRF backbone, an SDF head, and a color head. The predicted density and color are supervised by the classic eikonal loss L_E and the photometric reconstruction loss L_C , respectively. We encode geometry-aware mutual information into a subset of the parameters θ^D of f(x) to constrain the density field learning for better surface quality, which is achieved by the proposed mutual information loss L_M computed on top of the estimated surface normal.

calculation of the surface normal includes the entire set of parameters in f, thus a more global regularity. 2) Moreover, shaping f with the mutual information between its values may induce side effects as density could be distorted to satisfy the constraints. However, shaping the normal $\partial f/\partial x$ leave more space for the SDF values to cope with the scene reconstruction task, thus minimizing potential negative effects.

According to Eq. 5, we need to calculate the partial gradients $\partial N/\partial \theta^D$ to compute the mutual information. When it comes to a vector rather than a scalar, the gradients consist of three components corresponding to the **x**, **y**, and **z** axes:

$$\frac{\partial N(x;\theta)}{\partial \theta^D} = \frac{\partial N_{\mathbf{x}}(x;\theta)}{\partial \theta^D}, \frac{\partial N_{\mathbf{y}}(x;\theta)}{\partial \theta^D}, \frac{\partial N_{\mathbf{z}}(x;\theta)}{\partial \theta^D}.$$
(7)

For ease of computation, we concatenate the gradients into a vector, denoted as ∂N for simplicity:

$$\partial N = \operatorname{Concat}(\frac{\partial N_{\mathbf{x}}}{\partial \theta^{D}}, \frac{\partial N_{\mathbf{y}}}{\partial \theta^{D}}, \frac{\partial N_{\mathbf{z}}}{\partial \theta^{D}}).$$
(8)

Since a 2D pixel is accumulated from 3D, we aggregate the normal-based information along a viewing ray as in Eq. 1 as:

$$\mathbb{N}(p_i) = \sum_{x \in \text{ray}} \omega(t_i; \theta) \partial N.$$
(9)

Finally, we obtain that the mutual information between two normals is:

$$\mathbb{I}(\hat{N}(p_i), \hat{N}(p_j)) \propto ||\frac{\mathbb{N}(p_i) \cdot \mathbb{N}(p_j)}{|\mathbb{N}(p_i)||\mathbb{N}(p_j)|}||.$$
(10)

Please refer to our supplementary material for a detailed derivation. From Eq. 10, we observe that the mutual information in surface normal vectors correlates



Fig. 3: An example of positive samples from different features. Given an anchor pixel (marked by a blue circle in the first image) on the wall-like cabinet, semantic features like DINO often correlate both the cabinet and the ceiling (as shown by the red area in the second image). Meanwhile, geometric features such as (noisy) monocular normals can not distinguish between parallel planes that are not connected (the third image). By combining semantic and geometric features, we can obtain positive samples of the anchor pixel with better geometric consistency, thus, high mutual correlation.

positively with the cosine similarity of their gradients to the network parameters. This conclusion is similar to that of Eq. 5. As a result, we can now identify the pixels that are of high correlations, and increase the aggregated mutual information during the NeRF training process.

Extracting correlated scene regions from image pixels. To extract scene regions that need to be shaped to have high mutual information, we present a simple yet effective approach with off-the-shelf visual features. Given the input images, we extract pixel-wise semantic and geometric features to identify correlated scene regions. Following previous research [34], we apply DINO [4] as the semantic feature extractor to obtain a high-level understanding of the scene regions. Mean-while, as we shape mutual information on top of the surface normal, we employ a monocular normal estimator [7] as the geometric feature extractor for a detailed low-level perception of the 3D points in the scene.

DINO features can group different parts of an object, like the seat and back of a chair, which share similar semantics. However, the surface normal vectors from these parts have limited correlation, thus should not be constrained. Conversely, using only geometric features can wrongly suggest high correlation between unrelated areas, like the seat and the floor. As these areas aren't closely connected in 3D space, there is no reason to have better surface reconstruction when enhancing their mutual information. Similar phenomena on wall-like structures can be observed in Fig. 3, where we show that by combining semantic and geometric features together we can better identify scene regions that are spatially connected and have similar normal directions.

More explicitly, we combine semantic and geometric features with an intersection operator. It is simple yet effective. The process of identifying correlated scene regions (in normal) is implemented as positive pixel pair selection from input images. Given a randomly selected pixel from an image, we first extract similar pixels from neighboring images using semantic and geometric features, respectively. This extraction is implemented by applying two thresholds, β_S and β_G , to filter the cosine similarity for semantic and geometric features. We then feed the intersection of the top similar pixels from each feature type into NeRF's training process as positive samples, allowing us to shape them with high mutual information. In order to leverage the contrastive learning framework to encourage mutual information, we randomly select pixels not involved in the aforementioned process to serve as negative samples (independent pixels that are not correlated to the anchor pixel).

Loss functions for training. The aforementioned contrastive learning process is performed with a tailored InfoNCE loss [22]:

$$L_M = -\log \frac{\sum \exp(||\cos(\partial \mathbb{N}_i, \partial \mathbb{N}_{i+1})||)}{\sum_j \exp(||\cos(\partial \mathbb{N}_i, \partial \mathbb{N}_j)||)},\tag{11}$$

where i, i+, and j denote the anchor, positive, and negative pixels, respectively. This loss function effectively enhances mutual information (according to Eq. 10) among correlated pixels, while maintaining low mutual information among independent pixels. Finally, to train the NeRF backbone and retain the characteristics of the SDF, we adopt the basic photometric loss (Eq. 2) and the eikonal loss (Eq. 4). The latter two are derived from single images, while the former can be calculated from either single images or a set of neighboring images. The full loss function is defined as the weighted sum of three loss terms:

$$L = L_C + \lambda_E L_E + \lambda_M L_M. \tag{12}$$

With the above, all network parameters θ can be trained end-to-end.

4 Experiments

In this section, we evaluate the effectiveness of our method through experiments conducted on multiple datasets with various baselines. First, we introduce the two datasets used to conduct the experiments in Sec. 4.1. Next, we detail the implementation of the proposed method in Sec. 4.2, where we adapt a spectrum of state-of-the-art models for validating the effectiveness of the geometric shaping. Then, we report the comparisons between the baselines and our models in Sec. 4.3. Lastly, we conduct a series of ablation studies in Sec. 4.4, to verify the necessity of each component in our method.

4.1 Datasets

We evaluate the proposed method using two public datasets from the literature: ScanNet++ [38] and the Replica [26] dataset. For the ScanNet++ dataset, we choose 8 representative scenes, while for the Replica dataset, we select 4 representative scenes. The scenes from both datasets represent a variety of indoor settings, including offices, bedrooms, and so on. As our focus is on scene reconstruction from sparse views, we uniformly sample 50-85 images from the official (and Replica from Semantic-NeRF [42]) scanning sequences for each scene to use as input in our evaluation. For detailed statistics of each scene, please refer to our supplementary material.

Evaluation metrics. Following previous works [13, 30, 40], we evaluate the quality of surface reconstruction using two metrics: Chamfer distance and F-score. Chamfer distance is a measure that evaluates the difference between reconstructed and ground-truth points sets. F-score is a statistical metric that integrates precision and recall, providing a more comprehensive understanding of the reconstruction quality.

4.2 Implementation Details

As the proposed geometric shaping can be easily integrated into any network similar to NeRF with SDF representations, we select several state-of-the-art models as our baselines and apply the proposed to them for demonstrating the effectiveness of our approach. We first choose NeuS [30] and VolSDF [37], the two representative pioneers for NeRFs with SDFs. We also apply our method to the following works: GeoNeuS [8], I²-SDF [44], NeuRIS [29], MonoSDF [40], and Neuralangelo [13]. In addition to the basic loss functions for NeRF and SDF, GeoNeuS incorporates patch similarity from neighboring views and surface points derived from structure-from-motion. I²-SDF integrates the rendering equation with VolSDF. We disregard its normal and depth loss terms. NeuRIS utilizes monocular normal estimation as a supervision on the geometry, and discards the low-quality normal information that lacks sufficient confidence through multi-view patching. MonoSDF employs both depth and normal estimation, for which we set the decay of the normal and depth at 30k iterations to prevent the method from degeneration. Neuralangelo utilizes a multi-resolution hash grid and numerical gradient. We set the hash encoding dictionary size to 20, and the encoding dimension to 4. For all the base models mentioned above, we enhance them by incorporating our proposed mutual information loss during training. The resulting methods are respectively labeled as NeuS⁺, VolSDF⁺, GeoNeuS⁺, and so on.

4.3 Comparisons

Results on the ScanNet++ dataset. The quantitative comparisons are reported in Tab. 1. We also visualize the surface reconstruction results in Fig. 4 for qualitative comparisons. We can observe that by applying the proposed mutual information shaping on the geometry, the surface reconstruction quality can be improved across all baselines, as measured by mean Chamfer distance, F-score, and the visuals.

NeuS and VolSDF tend to produce rough surfaces and excessively smoothed boundaries. However, with our geometry-aware mutual information, the surfaces become smoother, while the boundaries are more crispy. Compared to GeoNeuS, a method that uses patch-similarity among neighboring views and requires sufficient textures, our method provides significant improvements. Applying the proposed technique to I^2 -SDF can also enhance its reconstruction quality.

As for more recent works, we find that NeuRIS is somewhat sensitive to the quality of normal estimation. It can either be affected by noise in the normal

Table 1: Quantitative results on the ScanNet++ dataset. Each cell includes the original baseline number in black, followed by the improvement using our mutual information. Positive and negative improvements are marked in green and orange, respectively. The last column shows the average improvement over all scenes, validating the effectiveness of the proposed shaping for better geometric reconstruction (despite a small potion of noise on a few scenes).

		0a7c	0a18	6ee2	7b64	56a0	9460	a08d	e0ab	Mean
Chamfer (m)↓	NeuS [30] ⁺	0.08-0.04	0.05-0.03	0.03+0.01	0.06-0.04	0.01-0.00	0.24-0.14	0.04-0.03	0.02-0.00	0.07-0.03
	VolSDF [37] ⁺	0.90-0.83	0.31 - 0.29	0.04 + 0.02	0.10 - 0.03	0.03 - 0.01	0.14 + 0.03	0.39-0.38	0.04-0.00	0.24-0.19
	GeoNeuS [8] ⁺	0.03-0.02	1.03 - 1.02	0.28 - 0.27	0.05-0.03	0.02 - 0.01	0.22-0.20	0.01 - 0.01	0.26 - 0.25	0.24-0.23
	I^2 -SDF $[44]^+$	0.27-0.24	0.22 - 0.03	0.04 - 0.01	0.47 - 0.45	0.03-0.02	0.13 + 0.05	0.03+0.02	0.30 - 0.07	0.19-0.09
	NeuRIS [29] ⁺	0.01-0.00	0.01-0.00	0.03 + 0.02	0.07 - 0.04	0.04 - 0.03	0.41 - 0.21	0.03-0.01	0.03-0.00	0.08-0.03
	MonoSDF [40] ⁺	0.01-0.00	0.01 + 0.02	0.02 - 0.01	0.02 - 0.00	0.01 - 0.00	0.05 - 0.02	0.01 - 0.00	0.01 - 0.00	0.02-0.00
	Neuralangelo [13] ⁺	0.17-0.01	0.84 - 0.66	0.52 - 0.33	0.05 - 0.01	0.04 + 0.02	4.76 - 2.85	0.33 + 0.05	0.12 + 0.03	0.85 - 0.48
F-score \uparrow	NeuS [30] ⁺	0.49+0.06	0.48 + 0.10	0.67 + 0.00	0.58 + 0.16	0.81-0.01	0.58-0.02	0.67 + 0.17	0.75 + 0.09	0.63+0.07
	VolSDF [37] ⁺	0.45 + 0.35	0.47 + 0.27	0.67 + 0.00	$0.55 {+} 0.05$	0.71 + 0.06	0.52 + 0.18	0.48 + 0.35	$0.70 {+} 0.02$	0.57 + 0.16
	GeoNeuS [8] ⁺	0.69 + 0.20	0.47 + 0.44	$0.52 {+} 0.34$	0.70 + 0.06	0.87 + 0.03	0.68 + 0.23	$0.75 {+} 0.14$	$0.37 {+} 0.51$	0.63 + 0.24
	I^2 -SDF $[44]^+$	0.33 + 0.52	0.28 + 0.00	$0.58 {+} 0.15$	$0.41 {+} 0.43$	$0.75 {+} 0.10$	$0.65 {+} 0.07$	$0.65 {+} 0.07$	$0.40 {+} 0.25$	0.51 + 0.20
	NeuRIS [29] ⁺	0.72 + 0.03	$0.68 {+} 0.01$	$0.65 {+} 0.01$	$0.69 {+} 0.06$	$0.85 {+} 0.01$	$0.24 {+} 0.17$	$0.64 {+} 0.13$	$0.69 {+} 0.09$	0.65 + 0.06
	MonoSDF [40] ⁺	0.92-0.01	0.88 <mark>-0.01</mark>	$0.81 {+} 0.04$	$0.79 {+} 0.06$	$0.89 {+} 0.00$	$0.88 {+} 0.05$	$0.91 {+} 0.00$	0.89 <mark>-0.01</mark>	0.87 + 0.01
	Neuralangelo $[13]^+$	0.42 + 0.05	0.14 + 0.30	$0.35 {+} 0.12$	$0.63 {+} 0.05$	0.71 - 0.01	0.17 + 0.22	$0.29 {+} 0.06$	0.51 <mark>-0.01</mark>	0.40 + 0.10

Table 2: Quantitative results on the Replica dataset.

	office0	office1	room0	room1	Mean
NeuS [30] ⁺	0.005 + 0.001	0.023-0.012	0.068-0.038	0.130-0.087	0.057-0.035
G VolSDF [37] ⁺	0.004 + 0.006	0.016 - 0.004	0.080 - 0.071	0.034 - 0.025	0.033-0.023
^H GeoNeuS [8] ⁺	0.023-0.021	0.014 -0.009	0.040 -0.035	0.002 -0.000	0.020-0.017
夏 I ² -SDF [44] ⁺	0.032 - 0.009	0.056 - 0.033	-	-	0.044-0.021
H NeuRIS [29] ⁺	0.014 - 0.006	3.975 - 3.963	0.025 - 0.017	0.015 + 0.002	1.007-0.996
뒨 MonoSDF [40] ⁺	0.003-0.000	0.005 - 0.001	0.004 - 0.000	0.004 - 0.001	0.004-0.001
Neuralangelo [13] ⁺	0.024 - 0.018	0.030-0.019	0.146 - 0.012	0.082 - 0.047	0.070-0.024
NeuS [30] ⁺	0.92 + 0.00	0.80 + 0.01	0.78 <mark>-0.03</mark>	$0.72 {+} 0.08$	0.80+0.02
VolSDF [37] ⁺	0.94-0.04	$0.78 {+} 0.01$	$0.78 {+} 0.16$	$0.82 {+} 0.07$	0.83 + 0.05
_ϕ GeoNeuS [8] ⁺	$0.89 {+} 0.08$	0.91 + 0.04	0.90 + 0.09	0.98 + 0.01	0.920 + 0.05
9 I ² -SDF [44] ⁺	$0.72 {+} 0.09$	$0.56 {+} 0.20$	-	-	0.64 + 0.14
⁹ _L NeuRIS [29] ⁺	0.79 + 0.06	$0.01 {+} 0.68$	$0.89 {+} 0.05$	$0.77 {+} 0.03$	0.61 + 0.20
MonoSDF [40] ⁺	0.96 + 0.00	0.90 + 0.02	0.98 + 0.00	0.95 + 0.02	0.95 + 0.01
Neuralangelo [13] ⁺	$0.89 {+} 0.07$	0.49 + 0.36	$0.65 {+} 0.15$	$0.71 {+} 0.18$	0.69 + 0.19

or completely disregard the normal information due to low confidence. Particularly, when images are sparse, the regions covered by fewer viewpoints tend to lose supervision. Our mutual information scheme does not depend on the absolute values of the surface normal, but on the similarity among these values. This makes it more robust to noise and less likely to discard low-quality but useful normal information. Consequently, our method enhances the reconstruction quality of NeuRIS. MonoSDF achieves the best overall performance among the baselines because it uses both depth and normal estimations as geometry supervision. However, our method still offers improvements. Notably, it produces smoother planes and clearer boundaries. This indicates that not only the supervision of geometric cues is important, but also the correlation between the geometric entities matters.



Fig. 4: Visual results on the ScanNet++ dataset. Each row shows a comparison with a different baseline, from top to bottom, NeuS [30]⁺, VolSDF [37]⁺, GeoNeuS [8]⁺, I²-SDF [44]⁺, NeuRIS [29]⁺, MonoSDF [40]⁺, and Neuralangelo [13]⁺. Red boxes are overlaid to help the contrast.



Fig. 5: Reconstruction results after 10K iterations in the training process with (NeuS⁺) and without (NeuS) the proposed geometric shaping technique.

Last but not least, the experiments on Neuralangelo demonstrate the potential of our method in working with hash grid encoding and numerical gradient. This highlights our method's versatility and ease of use, which can be broadly applied in various reconstruction architectures.

Results on the Replica dataset. In Tab. 2 we report the numerical results on the Replica dataset, with respect to NeuS [30], VolSDF [37], GeoNeuS [8], I²-SDF [44], NeuRIS [29], MonoSDF [40], and Neuralangelo [13]. As observed, the proposed geometric shaping technique still provides positive improvements when incorporated into these baselines, aligning with the conclusion drawn from the ScanNet++ data. Qualitative examples can also be found in Fig. 1.

4.4 Analyses

Ablation studies. We select three representative scenes from the ScanNet++ dataset and conduct ablation studies to evaluate the effectiveness of our method designs. We conduct ablation studies using two top-performing baselines: NeuRIS and MonoSDF.

The results are reported in Tab. 3. Instead of using the surface normal formula to inject mutual information, we first perform a basic mutual information shaping separately in the color and SDF values. These two variations are labeled as base⁺ (color)

Table 3: Ablation studies on three scenes from ScanNet++. The best and second-best results are marked in **bold** and <u>underlined</u>.

	Chamfer (m)↓				F-score ↑			
	6ee2	7b64	9460	Mean	6 ee 2	7b64	9460	Mean
NeuRIS [30]	0.029	0.070	0.405	0.168	0.65	0.69	0.24	0.53
NeuRIS ⁺ (color)	0.023	0.038	0.315	0.125	0.72	0.68	<u>0.34</u>	0.58
NeuRIS ⁺ (SDF)	0.127	0.053	0.360	0.180	0.47	0.61	0.30	0.46
NeuRIS ⁺ w/o normal	0.036	0.025	0.338	0.133	0.66	<u>0.75</u>	0.32	0.58
NeuRIS ⁺ w/o DINO	0.061	0.033	0.381	0.158	0.63	0.72	0.32	0.56
NeuRIS ⁺ (Full)	0.044	$\underline{0.033}$	0.198	0.092	<u>0.66</u>	0.76	0.41	0.61
MonoSDF [40]	0.020	0.016	0.046	0.028	0.81	0.79	0.88	0.83
MonoSDF ⁺ (color)	0.024	0.019	0.033	<u>0.025</u>	0.79	0.78	0.87	0.81
MonoSDF ⁺ (SDF)	0.588	0.029	3.443	1.353	0.38	0.65	0.21	0.42
MonoSDF ⁺ w/o normal	0.128	0.023	0.066	0.072	0.68	0.73	0.80	0.74
MonoSDF ⁺ w/o DINO	0.064	0.016	0.022	0.034	0.78	<u>0.83</u>	<u>0.93</u>	<u>0.85</u>
MonoSDF ⁺ (Full)	0.014	0.020	$\underline{0.023}$	0.019	0.85	0.85	0.93	0.87

and base⁺ (SDF), respectively. Interestingly, we find that, even though color is not directly related to geometry, by shaping it, we can still obtain improvements on the geometry in some cases. Moreover, shaping directly on the SDF values

severely degrade the reconstruction quality, which evidences that the shaping on SDF values could incur a competition with the goal to faithfully learn the geometry within the neural radiance field. We further examine the effectiveness of combining semantic and geometric features to identify geometrically highly correlated regions. As demonstrated in Fig. 3, using only one type of the features can lead to incorrect correlations. Tab. 3 also shows that eliminating any one of the features can decrease the performance, confirming the effectiveness of our positive-negative sample selection mechanism levering the multimodal features.

Time and memory consumption. Our method provides benefits for 3D reconstruction with limited additional effort. Specifically, it only increases the optimization time and memory usage during the training process, without affecting the model architecture and final storage. The time and memory consumption are reported in Tab. 4. It is observed that the overhead on the training time highly depends on the architecture used in each basline. For example, the train-

Fable 4: Time and memory consump

	Training Time (h)	VRAM (GB)
NeuS [30] ⁺	$3.00 {+} 0.37$	$7.89 {+} 1.68$
VolSDF [37] ⁺	$3.18 {+} 1.96$	$16.36 {+} 4.68$
GeoNeuS [8] ⁺	$3.26 {+} 0.16$	$5.72 {+} 3.56$
I ² -SDF [44] ⁺	3.47 + 2.84	$10.08 {+} 8.09$
NeuRIS [29] ⁺	$3.70 {+} 0.20$	$11.73 {+} 4.12$
MonoSDF [40] ⁺	$2.65 {+} 1.32$	$12.83 {+} 9.42$
Neuralangelo [13] ⁺	$3.12 {+} 2.06$	$16.44 {+} 1.00$

ing time increase on NeuRIS is marginal while more significant on VolSDF. These studies provide comprehensive information on the trade-offs between training time and performance gain. However, we do find that baselines with the proposed shaping technique can achieve a better performance (e.g., early geometries from NeuS and NeuS⁺ shown in Fig. 5) at the same training steps due to the effective regularization by mutual information.

5 Conclusion

We explore the usage of the second-order geometric correlations as a regularization and introduce a mutual information shaping technique of the normals for better surface reconstruction quality. With a simple yet effective correlated point selection mechanism that does not depend on accurate geometry information, we can perform a contrastive learning that maximizes the detected correlations. In return, we can observe a boost in the geometric quality of the reconstructed surfaces in state-of-the-art neural scene representations.

Extensive evaluations on various baselines and datasets show the effectiveness of our method, which improves 3D reconstruction quality without resorting to a more sophisticated surface modeling. The overhead caused by the shaping of the normals varies with different neural architectures of the scene representation, e.g., some are marginal, and some are significant, which we deem as a limitation of our current pipeline. We believe that an acceleration in computation can be achieved using a coarse-to-fine strategy, which we leave as a future work. Acknowledgments. We thank the anonymous reviewers for their valuable feedback. This work is supported by the Early Career Scheme of the Research Grants Council (grant # 27207224), the HKU-100 Award, and in part by NSF China (No. 62172363). Siyan Dong would also like to thank the support from HKU Musketeers Foundation Institute of Data Science for the Postdoctoral Research Fellowship.

References

- Bae, G., Budvytis, I., Cipolla, R.: Estimating and exploiting the aleatoric uncertainty in surface normal estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13137–13146 (2021)
- Bao, C., Zhang, Y., Yang, B., Fan, T., Yang, Z., Bao, H., Zhang, G., Cui, Z.: Sine: Semantic-driven image-based nerf editing with prior-guided editing field. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20919–20929 (2023)
- Cai, B., Huang, J., Jia, R., Lv, C., Fu, H.: Neuda: Neural deformable anchor for high-fidelity implicit surface reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8476–8485 (2023)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
- Chen, Y., Wu, Q., Zheng, C., Cham, T.J., Cai, J.: Sem2nerf: Converting single-view semantic masks to neural radiance fields. In: European Conference on Computer Vision. pp. 730–748. Springer (2022)
- Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques. pp. 303–312 (1996)
- Do, T., Vuong, K., Roumeliotis, S.I., Park, H.S.: Surface normal estimation of tilted images via spatial rectifier. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16. pp. 265– 280. Springer (2020)
- Fu, Q., Xu, Q., Ong, Y.S., Tao, W.: Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. Advances in Neural Information Processing Systems 35, 3403–3416 (2022)
- Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. arXiv preprint arXiv:2002.10099 (2020)
- Guo, H., Peng, S., Lin, H., Wang, Q., Zhang, G., Bao, H., Zhou, X.: Neural 3d scene reconstruction with the manhattan-world assumption. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5511– 5520 (2022)
- Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: Lerf: Language embedded radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19729–19739 (2023)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026 (2023)

- 16 X. Wang & S. Dong et al.
- Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8456– 8465 (2023)
- Liu, F., Zhang, C., Zheng, Y., Duan, Y.: Semantic ray: Learning a generalizable semantic field with cross-reprojection attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17386–17396 (2023)
- Liu, Y.L., Gao, C., Meuleman, A., Tseng, H.Y., Saraf, A., Kim, C., Chuang, Y.Y., Kopf, J., Huang, J.B.: Robust dynamic radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13–23 (2023)
- Liu, Z., Milano, F., Frey, J., Siegwart, R., Blum, H., Cadena, C.: Unsupervised continual semantic adaptation through neural rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3031– 3040 (2023)
- Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In: European Conference on Computer Vision. pp. 210–227. Springer (2022)
- Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: Seminal graphics: pioneering efforts that shaped the field, pp. 347–353 (1998)
- Meng, X., Chen, W., Yang, B.: Neat: Learning neural implicit surfaces with arbitrary topologies from multi-view images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 248–258 (2023)
- Metzer, G., Richardson, E., Patashnik, O., Giryes, R., Cohen-Or, D.: Latent-nerf for shape-guided generation of 3d shapes and textures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12663– 12673 (2023)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165– 174 (2019)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- 26. Straub, J., Whelan, T., Ma, L., Chen, Y., Wijmans, E., Green, S., Engel, J.J., Mur-Artal, R., Ren, C., Verma, S., Clarkson, A., Yan, M., Budge, B., Yan, Y., Pan, X., Yon, J., Zou, Y., Leon, K., Carter, N., Briales, J., Gillingham, T., Mueggler, E., Pesqueira, L., Savva, M., Batra, D., Strasdat, H.M., Nardi, R.D., Goesele, M., Lovegrove, S., Newcombe, R.: The replica dataset: A digital replica of indoor spaces (2019)

- Tertikas, K., Despoina, P., Pan, B., Park, J.J., Uy, M.A., Emiris, I., Avrithis, Y., Guibas, L.: Partnerf: Generating part-aware editable 3d shapes without 3d supervision. arXiv preprint arXiv:2303.09554 (2023)
- Wang, F., Galliani, S., Vogel, C., Speciale, P., Pollefeys, M.: Patchmatchnet: Learned multi-view patchmatch stereo. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14194–14203 (2021)
- Wang, J., Wang, P., Long, X., Theobalt, C., Komura, T., Liu, L., Wang, W.: Neuris: Neural reconstruction of indoor scenes using normal priors. In: European Conference on Computer Vision. pp. 139–155. Springer (2022)
- Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689 (2021)
- Weder, S., Garcia-Hernando, G., Monszpart, A., Pollefeys, M., Brostow, G.J., Firman, M., Vicente, S.: Removing objects from neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16528–16538 (2023)
- 32. Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., Lin, D.: Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In: European conference on computer vision. pp. 106–122. Springer (2022)
- 33. Xu, L., Xiangli, Y., Peng, S., Pan, X., Zhao, N., Theobalt, C., Dai, B., Lin, D.: Grid-guided neural radiance fields for large urban scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8296– 8306 (2023)
- 34. Xu, X., Yang, Y., Mo, K., Pan, B., Yi, L., Guibas, L.: Jacobinerf: Nerf shaping with mutual information gradients. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16498–16507 (2023)
- Yan, Z., Li, C., Lee, G.H.: Nerf-ds: Neural radiance fields for dynamic specular objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8285–8295 (2023)
- Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018)
- Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. Advances in Neural Information Processing Systems 34, 4805–4815 (2021)
- Yeshwanth, C., Liu, Y.C., Nießner, M., Dai, A.: Scannet++: A high-fidelity dataset of 3d indoor scenes. In: Proceedings of the International Conference on Computer Vision (ICCV) (2023)
- Yu, H., Julin, J., Milacski, Z.A., Niinuma, K., Jeni, L.A.: Dylin: Making light field networks dynamic. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12397–12406 (2023)
- Yu, Z., Peng, S., Niemeyer, M., Sattler, T., Geiger, A.: Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. Advances in neural information processing systems 35, 25018–25032 (2022)
- Zhang, Y., Yang, G., Guibas, L., Yang, Y.: Infogaussian: Structure-aware dynamic gaussians through lightweight information shaping. arXiv preprint arXiv:2406.05897 (2024)
- 42. Zhi, S., Laidlow, T., Leutenegger, S., Davison, A.J.: In-place scene labelling and understanding with implicit scene representation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15838–15847 (2021)

- 18 X. Wang & S. Dong et al.
- Zhu, B., Yang, Y., Wang, X., Zheng, Y., Guibas, L.: Vdn-nerf: Resolving shaperadiance ambiguity via view-dependence normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 35–45 (2023)
- 44. Zhu, J., Huo, Y., Ye, Q., Luan, F., Li, J., Xi, D., Wang, L., Tang, R., Hua, W., Bao, H., Wang, R.: I²-sdf: Intrinsic indoor scene reconstruction and editing via raytracing in neural sdfs. In: CVPR (2023)