

InfoNeRF: Ray Entropy Minimization for Few-Shot Neural Volume Rendering

Mijeong Kim¹ Seonguk Seo¹ Bohyung Han^{1,2}
¹ECE & ¹ASRI & ^{1,2}IPAI, Seoul National University
 {mijeong.kim, seonguk, bhhan}@snu.ac.kr




Figure 1. Qualitative comparison of InfoNeRF (ours) with other NeRF-based models [11, 20, 37] on the *Lego*, *Materials*, and *Ficus* scenes of Realistic Synthetic 360° dataset in 4-view settings. Existing works often suffer from noise (b), color distortion (c), or blur effect (d), while InfoNeRF (ours) achieves outstanding quality of rendered images with only a few input views. The last column (f) visualizes depth maps estimated by InfoNeRF (ours), which provide clear boundaries and fine details of the objects.

Abstract

We present an information-theoretic regularization technique for few-shot novel view synthesis based on neural implicit representation. The proposed approach minimizes potential reconstruction inconsistency that happens due to insufficient viewpoints by imposing the entropy constraint of the density in each ray. In addition, to alleviate the potential degenerate issue when all training images are acquired from almost redundant viewpoints, we further incorporate the spatial smoothness constraint into the estimated images by restricting information gains from additional rays with slightly different viewpoints. The main idea of our algorithm is to make reconstructed scenes compact along individual rays and consistent across rays in the neighborhood.

The proposed regularizers can be plugged into most of existing neural volume rendering techniques based on NeRF in a straightforward way. Despite its simplicity, we achieve consistently improved performance compared to existing neural view synthesis methods by large margins on multiple standard benchmarks. Our codes and models are available in the project website¹.

1. Introduction

Understanding 3D structure of a natural scene is a critical step for various high-level computer vision applications including object recognition, photorealistic rendering, au-

¹<http://cvlab.snu.ac.kr/research/InfoNeRF>

tonomous driving, virtual reality, and many others. Recent advance of deep learning capacitates high-fidelity 3D reconstruction and recognition, but learning with 3D data is inherently more difficult than its counterpart based on 2D images due to unstructured nature of data format, high memory requirement, and lack of principled architectures. Hence, many researchers investigate the standard models with appropriate training algorithms and the methods for reducing their computational costs, and attempt to solve various challenging tasks.

Novel view synthesis based on neural implicit representations is one of the 3D learning tasks that draws a lot of attention these days since Neural Radiance Field (NeRF) [20] has been introduced. NeRF delivers accurate 3D reconstruction results without explicit modeling of 3D scene structures, but the requirement of many images captured from multiple calibrated cameras hampers the applicability of the method. Therefore, several recent approaches aim to reduce the high computational cost and alleviate the constraints related to datasets [7, 11, 37].

In this line of research, we explore the few-shot prior-free novel view synthesis task, where only a limited number of training images are accessible and other prior information, such as object categories and semantic structures of target scenes, are unavailable. There exist several prior works for this task, but they either work barely with few examples [11] or require narrow baseline assumption to find correspondences using an external module [7]. Other approaches rely on prior knowledge of scenes such as object classes or features. For example, PixelNeRF [37] takes advantage of the features extracted from seen images to compensate for missing information in unseen views while [14, 23] focus on a particular object class, *e.g.*, human, in novel view synthesis.

We address the fundamental drawbacks of existing few-shot novel view synthesis methods: inconsistent reconstruction, which generates noise, blur, or artifacts in rendered images, and overfitting to seen views, which leads to degenerate or trivial solutions. The proposed approach, referred to as InfoNeRF, alleviates the reconstruction inconsistency by imposing the sparsity on the estimated scene, which is achieved by entropy minimization in each ray. The overfitting issues are handled by enforcing the smoothness of the reconstruction with respect to viewpoint changes, which is controlled by minimizing information gains from a pair of slightly different viewpoints. Figure 1 illustrates the outstanding quality of rendered images and depth maps estimated by our model, which delineates clear object boundaries and fine structures using only 4 input views with wide baselines.

Overall, the main contributions and benefits of our algorithm are summarized as follows:

- We propose a novel information-theoretic approach,

InfoNeRF, for the regularization of the neural implicit representations for volume rendering. Our method points out key drawbacks of the existing few-shot novel view synthesis techniques, and introduces two effective regularization schemes, ray entropy minimization and ray information gain reduction.

- Since InfoNeRF is a generic regularization technique and does not require any other external data structures, *e.g.*, voxels or meshes, or additional learnable parameters, it can be applied to various neural volume rendering algorithms with and without scene prior.
- The proposed regularization technique turns out to be effective to alleviate reconstruction inconsistency across multiple views and prevent degenerate solutions by overfitting despite its simplicity. We demonstrate outstanding performance of InfoNeRF on several standard benchmarks for few-shot novel view synthesis.
- To our knowledge, InfoNeRF is the first NeRF variant that performs few-shot novel view synthesis on wide-baseline image datasets without prior information.

2. Related Work

2.1. Novel View Synthesis

Novel view synthesis aims to render realistic images via geometric and photometric understanding of a 3D scene given a set of training images. To address this problem, light fields [15, 29, 35] or image-based rendering [3–6, 27] approaches have been employed traditionally, and the approaches based on deep learning [8, 9, 39, 40] have recently received growing attention. In particular, NeRF [20] achieves photo-realistic rendering results by applying multi-layer perceptrons to differentiable volume rendering successfully. The following works attempt to extend NeRF in various aspects, such as dynamic view synthesis [16], self-calibrated view synthesis [33], real-time rendering [21, 34, 36], relighting [28], and anti-aliasing [1]. Although NeRF-based models have achieved impressive performance, they have a common drawback; they require dense scene sampling, making them difficult to be applied in real-world scenarios. We address the few-shot volume rendering task to improve applicability by reducing the need for many images captured by calibrated cameras.

2.2. Few-shot Novel View Synthesis

To synthesize novel views of a scene given sparse observations, some algorithms estimate depth maps from images since depth is valuable source for view synthesis and 3D reconstruction. Depth information plays a crucial role in depth-guided image interpolation [9, 24], multi-plane image prediction [31], and learned geometry regularization in

complicated scenes [7] for few-shot view synthesis. However, these methods require training images with depth supervision or external depth estimation modules, *e.g.*, multi-view stereo or COLMAP SfM [25], and are susceptible to large projection errors due to incorrect depth predictions.

To overcome the limitation, several approaches exploit multi-view feature semantics by introducing an image encoder for NeRF to estimate color and opacity [14, 23, 37] or achieve semantic consistency between seen images and rendered novel views [11]. These strategies facilitate learning the semantic prior, and allow us to synthesize novel views with only few-shot images. Unlike the aforementioned works, the proposed algorithm does not rely on any prior information or additional pretrained encoders. Our regularization technique is orthogonal to the other methods discussed above and can be integrated into existing few-shot volume rendering approaches straightforwardly.

Some explicit representation methods [17, 36] incorporate sparsity constraints for neural volume rendering, which may also be useful for few-shot novel-view synthesis, although they do not directly address the task. Yu *et al.* [36] adopt a variation of the octree structure with a sparsity prior loss to remove a subset of nodes in the tree corresponding to unobserved regions. On the other hand, Lombardi *et al.* [17] conduct ray marching through voxel grids for volume rendering and regularize the total variation of voxel opacities by enforcing sparse spatial gradient. Our method also employ sparsity via entropy constraints, but is more generic than the explicit methods because our algorithm does not require external data structures and suffer from memory bound to store explicit representations.

3. Preliminaries: NeRF

Neural Radiance Fields (NeRF) [20] is a novel framework to represent a 3D scene with a neural implicit function, where a neural network $f(\cdot, \cdot)$, typically given by an MLP, maps a 3D point $\mathbf{x} = (x, y, z)$ and a unit viewing direction $\mathbf{d} = (\theta, \phi)$ to a volume density σ as well as an emitted RGB color $\mathbf{c} = (r, g, b)$. Following the classical volume rendering theory [19], the rendered RGB color of a target pixel is obtained by integrating colors and densities along a ray. In practice, since the output value (\mathbf{c}, σ) of all continuous points on a ray is not observable, a subset of points are sampled and the rendered color of the ray is approximated by using the quadrature rule as follows:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

where \mathbf{r} denotes a ray, N is the number of samples, and δ_i is the distance between the i^{th} point and its adjacent sample. Note that T_i indicates the accumulated transmittance along

the ray until the i^{th} point, which is given by

$$T_i = \sum_{j=1}^{i-1} \sigma_j \delta_j \quad (2)$$

The points on the ray are sampled in a two-stage hierarchical manner to increase rendering efficiency. In the first stage, the points are sampled uniformly while, in the second stage, the importance sampling is performed based on the density estimated in the first stage. Since all processes are fully differentiable, the neural networks encoding radiance fields are optimized with the following objective:

$$\mathcal{L}_{\text{RGB}} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} \left\| C(\mathbf{r}) - \hat{C}(\mathbf{r}) \right\|_2^2 \quad (3)$$

where \mathcal{R} denotes a set of rays. Note that the positional encoding is also employed before the MLP to map an input coordinate (\mathbf{x}, \mathbf{d}) onto a higher dimensional space, which is helpful to represent high-frequency scenes.

Although NeRF [20] achieves outstanding photorealistic view synthesis results, it requires a lot of images densely captured by calibrated cameras in general. To alleviate this data acquisition issue, we propose a simple yet effective few-shot novel view synthesis approach.

4. Proposed Method

Our approach learns a robust neural volume rendering model based only on a few input images without any prior knowledge about a scene. We focus on how to alleviate the reconstruction inconsistency given by lack of input views and the degeneracy induced by overfitting. This section discusses the proposed information-theoretic regularizations imposed on NeRF-based models to achieve our goals.

4.1. Motivation

Due to the small number of views available, few-shot 3D reconstruction and volume rendering are inherently prone to result in noisy estimations and degenerate solutions. For example, Figure 2a shows that NeRF fails to reconstruct the 3D scene accurately, mainly due to insufficient viewpoints. In addition, the learned model is severely overfit to seen images and converges to a degenerate solution, especially when the viewpoints of training images are similar to each other. Consequently, the rendering fails with only a slight change of view as illustrated in Figure 2b.

4.2. Regularization by Ray Entropy Minimization

To alleviate the reconstruction inconsistency, we impose the sparsity constraint on the reconstructed scene, which is achieved by minimizing the entropy of each ray density function using additional regularization terms. This constraint is reasonable because only a small subset of sampled



(a) Rendered RGB and depth images estimated by NeRF [20] on the Realistic Synthetic 360° dataset in a 4-view setting.

(b) Rendered RGB images for two slightly different viewpoints (left: seen, right: unseen) estimated by NeRF [20] on the DTU dataset in a 3-view setting.

Figure 2. Typical drawbacks of NeRF in the few-shot volume rendering. We are motivated by the following two observations and alleviate the limitations by adding two regularizers in this work. (a) Few-shot NeRF suffers from significant noise in 3D depth estimation, which leads to noisy rendering. (b) Few-shot NeRF provides completely different rendering results between slightly different viewpoints because it extremely overfits to seen input views.

points along a ray hit objects or background in a scene and the rest of points are likely to observe noise.

Ray density Before discussing the ray entropy, we first define the normalized ray density denoted by $p(\mathbf{r})$ as follows:

$$p(\mathbf{r}_i) = \frac{\alpha_i}{\sum_j \alpha_j} = \frac{1 - \exp(-\sigma_i \delta_i)}{\sum_j 1 - \exp(-\sigma_j \delta_j)}, \quad (4)$$

where \mathbf{r}_i ($i = 1, \dots, N$) is a sampled point in a ray, σ_i is the observed density at \mathbf{r}_i , δ_i is a sampling interval around \mathbf{r}_i , and $\alpha_i \equiv 1 - \exp(-\sigma_i \delta_i)$ is the opacity at \mathbf{r}_i . The points on the ray are sampled in a two-stage manner, which are drawn from the uniform distribution followed by from the distribution of opacity as in [20]. Note that we actually utilize the opacity α_i to compute $p(\mathbf{r}_i)$ in (4), instead of the density σ_i , to consider the irregular sampling interval δ_i .

Ray entropy Following Shannon Entropy [26], we define the entropy of a discrete ray density function given by

$$H(\mathbf{r}) = - \sum_{i=1}^N p(\mathbf{r}_i) \log p(\mathbf{r}_i). \quad (5)$$

Because σ and δ values are already calculated to conduct volume rendering procedure in (1), the computation of ray entropy incurs only negligible additional cost.

Disregarding non-hitting rays One issue in ray entropy minimization is that some rays are enforced to have low entropy though they do not hit any objects in the scene. To prevent the potential artifacts induced by this issue, we simply disregard the rays with low density for the entropy minimization. Formally, we employ a mask variable $M(\cdot)$ to indicate the rays that have sufficient observations of the scene,

which is based on the opacity as follows:

$$M(\mathbf{r}) = \begin{cases} 1 & \text{if } Q(\mathbf{r}) > \epsilon \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where

$$Q(\mathbf{r}) = \sum_{i=1}^N 1 - \exp(-\sigma_i \delta_i) \quad (7)$$

denotes the cumulative ray density.

Ray entropy loss Based on the ray entropy computed in (5), our ray entropy minimization loss is defined as follows:

$$\mathcal{L}_{\text{entropy}} = \frac{1}{|\mathcal{R}_s| + |\mathcal{R}_u|} \sum_{\mathbf{r} \in \mathcal{R}_s \cup \mathcal{R}_u} M(\mathbf{r}) \odot H(\mathbf{r}), \quad (8)$$

where \mathcal{R}_s denotes a set of rays from training images, \mathcal{R}_u denotes a set of rays from randomly sampled unseen images, and \odot indicates element-wise multiplication. Note that, NeRF-based models are unavailable to use rays from the unseen images due to the lack of their ground-truths of pixel colors, while our model can utilize them since entropy regularization does not require the ground-truths. We observe that it is beneficial to utilize the rays even from the unobserved viewpoint for better scene reconstruction.

Comparison with existing methods There exist only a few prior works that impose constraints on scene representations or models, but their objectives are different from ours [17, 36], *e.g.*, focusing on improving reconstruction quality and/or achieving real-time processing without considering few-shot training scenarios. In addition, since they rely on 3D volume entropy based on voxel representations, they have to draw a large number of samples to estimate 3D density or occupancy map, resulting in heavy computational cost in terms of both space and time complexities.

On the other hand, InfoNeRF employs the entropy minimization along a ray via 1D sampling and consequently, it runs very efficiently compared to the methods based on 3D volume entropy.

4.3. Regularization by Information Gain Reduction

According to our observations, when the training images have sufficiently diverse viewpoints, the proposed entropy regularization is very helpful to improve the quality of both rendered images and 3D depth estimation in the few-shot setting. However, if all the training images have similar viewpoints to each other, the models are prone to overfit to seen images and fail to generalize to unseen views. This is probably because the lack of diverse observations makes the trained models find degenerate and trivial solutions.

To alleviate the aforementioned limitations, we introduce an additional regularization term to assure the consistent density distribution across rays in the neighborhood. Given an observed ray \mathbf{r} , we sample another ray with a slightly different viewpoint, denoted by $\tilde{\mathbf{r}}$, and minimize the KL-divergence between the density functions of the two rays. The motivation of this objective is to make the observations from two similar viewpoints sufficiently consistent so that the model is generalized to a nearby viewpoint, which is achieved by enforcing smoothness to reconstruction results over spatial view perturbations.

The regularization loss for information gain reduction is given by

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(P(\mathbf{r})||P(\tilde{\mathbf{r}})) = \sum_{i=1}^N p(\mathbf{r}_i) \log \frac{p(\mathbf{r}_i)}{p(\tilde{\mathbf{r}}_i)}, \quad (9)$$

where $\tilde{\mathbf{r}}_i$ is a sampled point for observation in ray $\tilde{\mathbf{r}}$. In our implementation, we obtain $\tilde{\mathbf{r}}$ by slightly rotating the camera pose of \mathbf{r} in the range from -5° to 5° .

4.4. Overall Objective

The total loss function to train a neural implicit model for few-shot neural volume rendering is given by

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RGB}} + \lambda_1 \mathcal{L}_{\text{entropy}} + \lambda_2 \mathcal{L}_{\text{KL}}, \quad (10)$$

where λ_1 and λ_2 are balancing terms for our regularization terms. As mentioned in (3), the reconstruction loss, denoted by \mathcal{L}_{RGB} , is given by

$$\mathcal{L}_{\text{RGB}} = \frac{1}{|\mathcal{R}_s|} \sum_{\mathbf{r} \in \mathcal{R}_s} \left\| C(\mathbf{r}) - \hat{C}(\mathbf{r}) \right\|_2^2, \quad (11)$$

which uses only a set of rays from training images with pixel-level ground-truth, unlike $\mathcal{L}_{\text{entropy}}$ and \mathcal{L}_{KL} which utilize the rays even from unobserved viewpoints.

5. Experiments

We demonstrate the effectiveness of the proposed approach, referred to as InfoNeRF, on the standard benchmarks. This section also discusses the characteristics of our algorithm based on the experiment results.

5.1. Datasets

We describe the details of three benchmarks employed to evaluate our algorithm, which include the Realistic Synthetic 360° [20], ZJU-MoCap [23], and DTU [12] datasets.

Realistic Synthetic 360° This benchmark is common for neural volume rendering, which contains 8 synthetic scenes with view-dependent light transport effects. Each scene has an object at the center and 400 rendered images from inward-facing virtual cameras with different viewpoints. For few-shot training, we randomly sample 4 viewpoints out of 100 training images in each scene, and use the 200 testing images for evaluation.

ZJU-MoCap This dataset consists of multi-view videos capturing human motion from 23 calibrated cameras. Following [23], we sample 4 uniformly distributed viewpoints to construct a training set and use the remaining images for testing.

DTU MVS Dataset (DTU) The images in this dataset contain complex and real-world scenes captured by the calibrated cameras in controlled environments. All the collected images have similar viewpoints and face only one side of a scene. We conduct experiments on 15 scenes, where we optimize the model with 3 images out of 49 views while testing with the remaining 46 views.

5.2. Implementation and Evaluation

Implementation details Our implementation is based on PyTorch [22]. We use Adam optimizer [13] with the initial learning rate of 5×10^{-4} , which decays exponentially by a factor of 10 at every 250,000 iterations. The balancing term for \mathcal{L}_{KL} is decayed by a factor of 2 at every 5,000 iterations. We set the number of rays from seen and unobserved views, denoted respectively by $|\mathcal{R}_s|$ and $|\mathcal{R}_u|$, identically to 1024, and our experiments are conducted with a single NVIDIA Titan XP GPU.

Metrics We evaluate the novel view rendering quality based on the standard image quality metrics, peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) [32]. We also use perceptual metrics, learned perceptual image patch similarity (LPIPS) [38], Fréchet inception distance (FID) [10], and kernel inception distance (KID) [2]. LPIPS estimates normalized features distance between image pair while FID and KID compute the distance in Inception representations [30] between two sets of images.

Table 1. Experimental results of few-shot novel view synthesis on the Realistic Synthetic 360° dataset in the 4-view setting. Our approach outperforms all other existing methods by significant margins in all image quality metrics. The asterisk (*) denotes that the model is pretrained on an external training dataset with dense input views and finet-tuned on this dataset with 4 input views. We run all experimental five times with different viewpoint samples and the same hyperparameters, and compute the average scores and their standard deviations.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	KID \downarrow
NeRF, 100 views	31.01	0.947	0.081	42.83	0.002
PixelNeRF* [37]	16.09 \pm 0.78	0.738 \pm 0.012	0.390 \pm 0.030	265.25 \pm 6.73	0.127 \pm 0.006
NeRF [20]	15.93 \pm 1.06	0.780 \pm 0.014	0.320 \pm 0.049	215.16 \pm 2.32	0.074 \pm 0.012
DietNeRF [11]	16.06 \pm 1.13	0.793 \pm 0.019	0.306 \pm 0.050	197.02 \pm 12.87	0.065 \pm 0.004
InfoNeRF (ours)	18.65\pm0.18	0.811\pm0.008	0.230\pm0.008	181.47\pm4.97	0.062\pm0.004

Table 2. Average PSNRs and standard deviations of individual scenes on the Realistic Synthetic 360° dataset in the 4-view setting.

Method	Lego	Chair	Drums	Ficus	Hotdog	Materials	Mic	Ship	Avg.
NeRF, 100 views	32.54	33.00	25.01	30.13	36.18	29.62	32.91	28.65	31.01
PixelNeRF* [37]	15.14 \pm 0.75	18.87 \pm 1.38	15.10 \pm 0.63	16.60 \pm 0.70	19.37 \pm 1.78	12.31 \pm 1.02	16.35 \pm 0.97	14.96 \pm 0.75	16.09 \pm 0.78
NeRF [20]	15.61 \pm 4.53	18.57 \pm 1.64	12.50 \pm 0.98	16.37 \pm 2.24	19.64 \pm 2.26	15.65 \pm 4.16	14.78 \pm 2.37	14.30 \pm 4.04	15.93 \pm 1.06
DietNeRF [11]	17.13 \pm 4.77	19.37 \pm 3.12	13.74 \pm 1.55	15.76 \pm 3.56	18.24 \pm 5.28	15.00 \pm 5.18	17.71 \pm 1.55	11.51 \pm 4.27	16.06 \pm 1.13
InfoNeRF (ours)	18.92\pm0.51	20.06\pm1.11	14.33\pm0.62	19.41\pm0.07	21.30\pm2.31	18.34\pm0.88	18.55\pm1.71	18.27\pm0.71	18.65\pm0.18

5.3. Results

5.3.1 Realistic Synthetic 360°

We compare our approach with NeRF [20], DietNeRF [11], and PixelNeRF [37] on Realistic Synthetic 360° dataset. NeRF, DietNeRF, and InfoNeRF (ours) are trained with randomly sampled 4 views from scratch. Unlike others, PixelNeRF is pretrained on the DTU [12] dataset with dense input views, and we fine-tune the model with 4 sampled views to handle the domain shift issue between the two datasets.

Table 1 presents overall quantitative results, where InfoNeRF consistently outperforms the baseline algorithms in terms of all metrics with considerable margins while having lower standard deviations. Table 2 breaks down the PSNU scores into 8 individual scenes, where InfoNeRF obviously achieves significant gains for all scenes. Refer to our supplementary document for the results of SSIM and LPIPS, which have the same tendencies.

Figure 1 demonstrates the qualitative results on the novel viewpoints, where InfoNeRF shows outstanding quality in the rendered images compared to all the compared methods. As illustrated in Figure 1(f), the quality of the depth maps estimated by InfoNeRF looks impressive while we notice that all the compared algorithms often fail to reconstruct 3D structures accurately and DietNeRF even has color distortion due to its high-level semantic consistency loss.

5.3.2 ZJU-MoCap

For the ZJU-MoCap dataset, InfoNeRF is evaluated in comparison with NeRF [20], Neural Volume (NV) [17], and Neural Body (NB) [23], where all algorithms are trained with 4 images. Note that, since NB employs a pretrained

Table 3. Quantitative comparison on the ZJU-MoCap dataset in the 4-view setting. NB [23] has the geometric prior by exploiting the pretrained human body model (SMPL).

Method	Prior	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NB [23]	✓	24.18	0.888	0.182
NeRF [20]		20.19	0.794	0.309
NV [17]		21.74	0.827	0.253
InfoNeRF (ours)		22.88	0.838	0.242

human body model denoted by SMPL [18] as its prior, the performance of NB can be regarded as the upper-bound of all other methods.

Table 3 summarizes the experimental results on the ZJU-MoCap dataset, where InfoNeRF achieves the best performance among the methods without using the prior in terms of all the tested metrics. Figure 3 demonstrates the qualitative results of all the compared methods, and the reconstruction result given by InfoNeRF is particularly accurate.

5.3.3 DTU MVS Dataset (DTU)

Contrary to the other two datasets, DTU has substantially different characteristics because the images in each scene have similar viewpoints. PixelNeRF takes advantage of this property and learns the scene-agnostic model successfully while the original NeRF exhibits poor generalization performance in this dataset with few-shot learning.

We compare our algorithms with NeRF and PixelNeRF in this dataset. We train InfoNeRF and NeRF from scratch without exploiting any scene prior, so it is not possible to reconstruct invisible parts of the scene. Therefore, the naïve evaluation of the algorithms without scene prior is not de-




Figure 3. Qualitative comparison on the ZJU-MoCap dataset in 4-view setting. We visualize the rendering results of prior-free algorithms (b-d) including ours, and prior-based algorithm (e). While existing prior-free algorithms (b-c) often suffer from inconsistent reconstruction and missing parts of the human body, InfoNeRF manages to render most of the human body comparable to prior-based algorithm (e).

Table 4. Quantitative comparison on the DTU dataset in 3-view setting. PixelNeRF [37] has the dataset prior by pretraining on other scenes of DTU with dense input views.

Method	Prior	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
PixelNeRF [37]	\checkmark	19.55	0.724	0.286
NeRF [20]		8.50	0.426	0.611
InfoNeRF (ours)		11.23	0.445	0.543

sirable and we use the mask corresponding to the visible parts of each scene for the evaluation performance. Table 4 presents the experimental results on the DTU dataset, where our algorithm achieves outstanding performance compared to NeRF. Note that PixelNeRF achieves the highest performance because it exploits the dataset prior by pretraining on the training split of DTU with dense input views.

5.4. Analysis

Effect of unseen view sampling To verify the effectiveness of sampling from unseen viewpoints for entropy minimization, we run our algorithm by varying the number of rays from unseen views on the *Chair* scene of Realistic Synthetic 360° with the images reduced to half. Table 5 presents that increasing the number of rays from the unseen views achieves gradual improvement by alleviating the reconstruction inconsistency but its benefit is saturated when the number of rays is larger than 1,024. Figure 4 visualizes the benefit of sampling additional rays for unseen views; the use of the rays is helpful for noise reduction in this example.

Benefit of regularization We analyze the impact of the proposed regularization schemes, ray entropy minimization loss, $\mathcal{L}_{\text{entropy}}$, and ray information gain reduction loss, \mathcal{L}_{KL} . Table 6 shows the ablative results of InfoNeRF on the DTU dataset. The entropy minimization loss successfully improves PSNR, but the SSIM value gets worse on this dataset. This is because the models are prone to overfit to the seen images when all the training images have similar viewpoints to each other. However, thanks to our information gain re-

Table 5. Impact of the number of rays sampled from unseen viewpoints on the *Chair* scene of Realistic Synthetic 360° in a 4-view setting. We fix the number of rays for seen views to 1,024 and vary the number of rays for unseen views. Bold and underline indicate the first and second place among the results, respectively.

# of seen rays	# of unseen rays	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
1024	0	20.14	0.834	0.225
	256	20.97	0.844	0.197
	512	21.11	0.851	0.188
	1024	21.37	<u>0.853</u>	<u>0.185</u>
	2048	<u>21.33</u>	0.855	0.167




Figure 4. Benefit of sampling rays for unseen viewpoints. We visualize the rendered images on the *Chair* scene of Realistic Synthetic 360° dataset, where the noise in the background area is removed completely after adding the rays from unseen viewpoints.

duction loss that enforces smoothness to reconstruction results over spatial view perturbations, our full model helps alleviate the overfitting issue and prevent degenerate solutions. Note that, in the other two datasets with substantial viewpoint variations, the entropy minimization loss works well while the information gain reduction loss makes minor contribution in general. Figure 5 visualizes the rendering quality of InfoNeRF in comparison to its ablative models on the DTU dataset. Although InfoNeRF without the information gain reduction loss tends to generate crisp images, there exists a lot of noise in the outputs and inconsistency in the depth maps. On the other hand, our full algorithm




Figure 5. Qualitative comparisons of our method with its ablative models on the DTU dataset in the 3-view setting. We visualize the image synthesis (left) and the depth estimation (right) results for each algorithm. While NeRF suffers from noises, blurs, and artifacts to reconstruct 3D structure, our two loss terms, $\mathcal{L}_{\text{entropy}}$ and \mathcal{L}_{KL} , contribute to outstanding rendering quality and fine depth estimation results.

Table 6. Ablative results of our regularization schemes on the DTU dataset in a 3-view setting.

Method	$\mathcal{L}_{\text{entropy}}$	\mathcal{L}_{KL}	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRF			8.50	0.426	0.611
InfoNeRF w/o $\mathcal{L}_{\text{entropy}}$		✓	8.91	0.439	0.581
InfoNeRF w/o \mathcal{L}_{KL}	✓		10.54	0.418	0.561
InfoNeRF	✓	✓	11.23	0.445	0.543




Figure 6. PSNR with respect to the number of views for training on the Realistic Synthetic 360° dataset.

manages to reconstruct the overall shape and geometry of the scene more accurately. Refer to Table C of the supplementary document for more detailed ablation results on the Realistic Synthetic 360° dataset.

Robustness to the number of views Figure 6 illustrates performance of InfoNeRF by varying the number of views for training on the Realistic Synthetic 360°. Compared to NeRF, InfoNeRF illustrates improved results in terms of all metrics until 8 views, but its merit is saturated as the number of views increases. This is partly because the uncertainty of the reconstructed scene decreases as the number of training views increases, weakening the importance of entropy regularization. See Table D of our supplementary document for SSIM and LPIPS results.

Integration into PixelNeRF To demonstrate the generality of our method, we incorporate the proposed regularization method to PixelNeRF [37], and refer to this version of our model as InfoPixelNeRF. Figure 7 illustrates qualitative



Figure 7. Qualitative comparison on the *Drums* scene in the Realistic Synthetic 360° dataset. InfoPixelNeRF (right) reduces blur significantly compared to PixelNeRF (left), which demonstrates the effectiveness of our regularization schemes.

comparisons between both models, where InfoPixelNeRF reduces blur in the rendered image significantly. Table E of our supplementary file presents detailed comparisons.

6. Conclusion

We proposed an information-theoretic regularization technique for few-shot novel view synthesis. Existing few-shot view synthesis methods suffer from inconsistent reconstruction, which often generates noise, blur, or artifacts in rendered images, and overfitting to seen views, which leads to degenerate solutions. To address these issues, we introduced two effective regularization schemes, ray entropy minimization and ray information gain reduction. Despite its simplicity, the proposed method turns out to be effective to alleviate reconstruction inconsistency across views. We demonstrated outstanding performance of our method on multiple standard benchmarks, and also conducted a detailed analysis of our approach via extensive analysis.

Acknowledgments This work was supported in part by the IITP grants [2021-0-02068, AI Innovation Hub; 2021-0-01343, Interdisciplinary Program in AI (SNU)] and the Bio & Medical Technology Development Program of NRF [2021M3A9E4080782] funded by the Korean government (MSIT) and by Samsung Advanced Institute of Technology.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, 2021. 2
- [2] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *ICLR*, 2018. 5
- [3] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *ACM SIGGRAPH*, 2001. 2
- [4] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. In *TOG*, 2013. 2
- [5] Gaurav Chaurasia, Olga Sorkine, and George Drettakis. Silhouette-aware warping for image-based rendering. In *Computer Graphics Forum*, volume 30, 2011. 2
- [6] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. 2
- [7] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021. 2, 3
- [8] John Flynn, Michael Broxton, Paul Debevec, Matthew Duvall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *CVPR*, 2019. 2
- [9] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *CVPR*, 2016. 2
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 5
- [11] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *ICCV*, 2021. 1, 2, 3, 6
- [12] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014. 5, 6
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [14] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. In *NeurIPS*, 2021. 2, 3
- [15] Marc Levoy and Pat Hanrahan. Light field rendering. In *ACM SIGGRAPH*, 1996. 2
- [16] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2
- [17] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.* 3, 4, 6, 7
- [18] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graph.*, 2015. 6
- [19] Nelson Max. Optical models for direct volume rendering. *IEEE TVCG*, 1995. 3
- [20] 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. 1, 2, 3, 4, 5, 6, 7, 8
- [21] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of neural radiance fields using depth oracle networks. In *EGSR*, 2021. 2
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. 2019. 5
- [23] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2, 3, 5, 6, 7
- [24] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2
- [25] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 3
- [26] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 1948. 4
- [27] Sudipta Sinha, Drew Steedly, and Rick Szeliski. Piecewise planar stereo for image-based rendering. 2009. 2
- [28] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021. 2
- [29] Pratul P Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, and Ren Ng. Learning to synthesize a 4d rgbd light field from a single image. In *ICCV*, 2017. 2
- [30] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 5
- [31] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *CVPR*, 2020. 2
- [32] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 5
- [33] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 2
- [34] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *CVPR*, 2021. 2

- [35] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *ACM SIGGRAPH*, 2000. [2](#)
- [36] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. [2](#), [3](#), [4](#)
- [37] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [5](#)
- [39] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph.*, 2018. [2](#)
- [40] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *ECCV*, 2016. [2](#)