

Towards Efficient Neural Scene Graphs by Learning Consistency Fields

Yeji Song¹

ldynx@snu.ac.kr

Chaerin Kong¹

veztylord@snu.ac.kr

Seoyoung Lee²

seoyoung1215@snu.ac.kr

Nojun Kwak*¹

nojunk@snu.ac.kr

Joonseok Lee*^{2,3}

joonseok@snu.ac.kr

¹ Department of Intelligence and Information,
Seoul National University
Seoul, Korea

² Graduate School of Data Science,
Seoul National University
Seoul, Korea

³ Google Research
Mountain View, CA, USA

Abstract

Neural Radiance Fields (NeRF) achieves photo-realistic image rendering from novel views, and the Neural Scene Graphs (NSG) [16] extends it to dynamic scenes (video) with multiple objects. Nevertheless, computationally heavy ray marching for every image frame becomes a huge burden. In this paper, taking advantage of significant redundancy across adjacent frames in videos, we propose a feature-reusing framework. From the first try of naively reusing the NSG features, however, we learn that it is crucial to disentangle object-intrinsic properties consistent across frames from transient ones. Our proposed method, *Consistency-Field-based NSG (CF-NSG)*, reformulates neural radiance fields to additionally consider *consistency fields*. With disentangled representations, CF-NSG takes full advantage of the feature-reusing scheme and performs an extended degree of scene manipulation in a more controllable manner. We empirically verify that CF-NSG greatly improves the inference efficiency by using 85% less queries than NSG without notable degradation in rendering quality. Code will be available at <https://github.com/ldynx/CF-NSG>.



Figure 1: Comparison of the qualitative results from NSG [16] and our CF-NSG. CF-NSG greatly improves the efficiency by using only 15% of the original queries.

1 Introduction

Being a long standing issue in computer vision, various approaches have been proposed for the task of novel view synthesis, including those using point clouds, discretized voxel grids and textured mesh [3, 20, 25]. Neural Radiance Fields (NeRF) [12] represents a static scene by evaluating a volumetric density (transparency) and a view-dependent color. The object’s details can be expressed at a low storage cost via an implicit 5D (spatial location (x, y, z) and viewing direction (θ, ϕ)) function implemented with a multi-layer perceptron (MLP).

A similar idea is applied to understand a video scene. Unlike a still image, a video has an additional temporal-axis, and objects in the scene may be static or dynamic (moving) across time. Neural Scene Graphs (NSG) for dynamic scenes [16] provides a considerable potential on various applications by enabling the understanding of a *complex scene with dynamic multi-objects*, which has been tricky to model. Also, NSG extends the task of novel view synthesis to *novel scene manipulation*, allowing spatial rearrangements of dynamic objects in the scene (e.g., the second row in Fig. 6). However, direct application of NSG is still not practical due to heavy computational overhead at inference. Along with other NeRF variants, NSG also suffers from expensive ray marching for all sampled queries for rendering. Moreover, since NSG repeats this procedure for each image frame independently, the total inference cost scales with the spatio-temporal resolution of the video. Observing this limitation, a natural question arises: how can we efficiently reduce this computational overhead?

We start from an idea that most objects in a video do not significantly change in adjacent frames. In other words, there is much redundancy in the representation of the scene across image frames as visual semantics are largely consistent. Reusing the previously-computed representations, instead of repeating computations for similar queries again, would be a plausible way to significantly improve efficiency.

We conduct a simple experiment to measure redundancy across frames. We divide an object’s bounding box into uniform spatial bins and assign queries (x, y, z) to each bin, storing the estimated color and density values. Fig. 2(a) shows the values across all frames for two selected bins [A] and [B]. We found that the color and density values from the consecutive frames hardly or only slightly change within a limited range. Fig. 2(b) shows the ratio of all bins whose values change less than a threshold ϵ in two consecutive frames. As the range of radiance and density are within $[0, 1]$, the possible threshold for ϵ is also within $[0, 1]$. We note that redundancy is already close to 70% at $\epsilon \approx 0.04$. To facilitate better understanding of ϵ , we generate images with random salt-and-pepper noise of $\epsilon = \{0.04, 0.18\}$ for all pixels. As shown in Fig. 2(c), the image with $\epsilon = 0.04$ noise is almost identical to the original image maintaining high PSNR. The image with $\epsilon = 0.18$ noise (90% redundancy in (b)) still well-preserved the overall scene. This experiment indicates that most of bins share highly similar values across frames. Hence, actively leveraging this redundancy is not only beneficial but in fact necessary to guide our model to learn dynamic scenes more efficiently.

To this end, we propose to identify visual components that are consistent across frames, and reuse them for improved efficiency. Fig. 1 shows the rendered image using our method with 85% less queries without significantly degrading rendering quality. In this paper, we propose Consistency-Field-based NSG (CF-NSG), which effectively reduces redundant computation across frames for more efficient rendering. We first review NSG in Sec. 2, and analyze a critical factor that hinders the feature reusing mechanism in Sec. 3. Then, we introduce our CF-NSG in Sec. 4. We empirically demonstrate that our model greatly improves efficiency with no distinguishable compromise in image quality in Sec. 5.

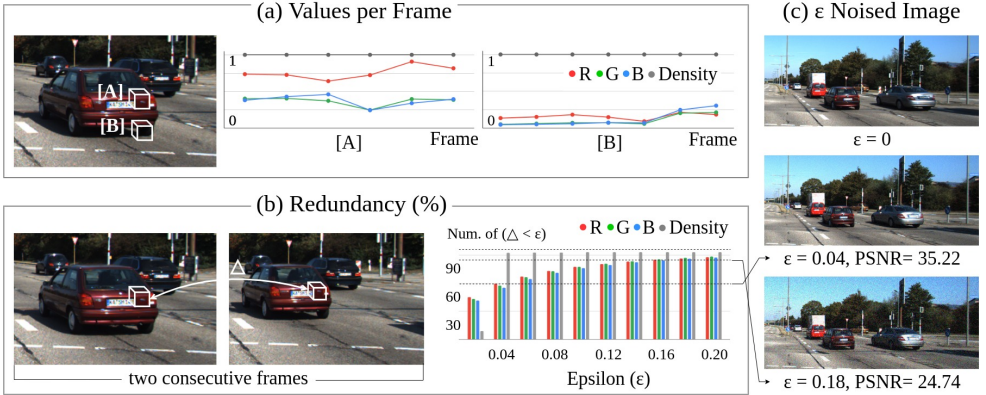


Figure 2: **Visualization of large redundancy in NSG [16].** (a), (b): Most of estimated color and density values are hardly or only slightly change across frames. (c): We add salt-and-pepper noise of ϵ to an image proving a majority of bins in (b) preserve similar values.

2 Preliminaries

Neural Radiance Fields (NeRF) [12] captures 3D representations of static objects or scenes. The implicit representation is encoded into an MLP, which maps a 3D spatial location $\mathbf{x} = (x, y, z)$ to its volumetric density σ , combined with a viewing direction $\mathbf{d} = (\theta, \phi)$ to an emitted color $\mathbf{c} = (r, g, b)$. The color of the pixel $C(\mathbf{r})$ along a camera ray \mathbf{r} is estimated by accumulating the color and transmittance of N sampled points along the ray:

$$\hat{C}(\mathbf{r}) = \sum_{i=0}^{N-1} T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad T_i = \exp\left(-\sum_{j=0}^{i-1} \sigma_j \delta_j\right) \quad (1)$$

where δ_i is the distance between adjacent sampled points. The NeRF network is optimized to reduce the L_2 distance between the estimated colors $\hat{C}(\mathbf{r})$ for a random a batch of rays \mathcal{R} and their ground truth (GT); that is, $\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \|C(\mathbf{r}) - \hat{C}(\mathbf{r})\|_2^2$.

Neural Scene Graphs (NSG) [16] classifies objects in a multi-object dynamic scene (video) according to whether they are moving or static across frames. They define the mappings from global coordinates to each dynamic object’s canonical coordinates using transformations of the object (e.g., translation, rotation and scaling). They also group dynamic objects into classes (e.g., car, pedestrian, cyclist), defining a common canonical coordinates of each class. Meanwhile, all the static objects are grouped as the background. Then, canonical coordinates of the each class as well as the background are represented using separate NeRFs.

The overall process can be expressed by

$$\mathcal{F}_{\theta_{bg}} : (\mathbf{x}, \mathbf{d}) \rightarrow (\mathbf{r}, \mathbf{g}, \mathbf{b}, \sigma), \quad \mathcal{F}_{\theta_c} : (\mathbf{x}_o, \mathbf{d}_o, \mathbf{p}_o, \mathbf{l}_o) \rightarrow (\mathbf{r}, \mathbf{g}, \mathbf{b}, \sigma), \quad (2)$$

where θ_{bg} , θ_c are respective weights for static and dynamic representation models. Here, c denotes each class of dynamic objects, where $c \in \{c_1, \dots, c_n\}$ for n classes. In the dynamic model, the set of dynamic objects belonging to class c (\mathcal{O}_c) shares weights θ_c , while latent vector \mathbf{l}_o is uniquely learned for an individual object $o \in \mathcal{O}_c$. \mathbf{p}_o is the global location of object o in the scene. Each \mathbf{x}_o and \mathbf{d}_o refer to the spatial location and the viewing direction, respectively, where the corresponding color values are desired. Note that $\mathbf{x}_o \in [-1, 1]^3$




Figure 3: **Rendering results with naive reusing of NSG.** We reproduce a frame seen during training with a naive reuse of NSG features, resulting in degradation of image quality.

and \mathbf{d}_o are in the canonical coordinates of object o , different from \mathbf{x} and \mathbf{d} in global coordinates. Outputs (r, g, b, σ) from the dynamic models are mapped from canonical coordinates to global coordinates and integrated simultaneously with the outputs from the static model along the rays, composing a scene similarly as in [14].

3 First Try: A Naive Reuse of NSG Features

NSG extends the realm of NeRF-based applications to videos with multiple dynamic objects. However, computation cost at inference is still prohibitive, as image rendering requires ray marching operations per frame. Thus, we propose a *feature-reusing framework* that stores redundant features across frames during training, then at inference reuses them instead of going through a full forward pass. Assuming the features in the nearby voxels are similar, we create quantized memory bins in the coordinates as follows. For the dynamic models (\mathcal{F}_{θ_c} in Eq. (2)), we quantize the object-specific 3D coordinates (within the object’s bounding box) into spatial bins, similarly as in the experiment in Fig. 2. For the static model ($\mathcal{F}_{\theta_{bg}}$ in Eq. (2)), following NSG, we define radiance fields and grid memory bins on a set of 2D planes instead of volume. If a query is assigned to the bin whose value is expected not to significantly change across frames, we reuse previously computed results from the bin.

A Naive Approach of Reusing NSG features. At the beginning, we use a pretrained NSG model to compute the color (RGB) and density (σ) for the training set, caching them in the memory bins. To estimate the expected change of each bin across frames, we additionally store the gradients of color and density with respect to global location \mathbf{p}_o , viewing direction \mathbf{d}_o and \mathbf{d} , inputs that may change across frames. A bin with a small gradient indicates that values in that bin are most likely not to change across frames. Therefore, we examine whether queries are assigned to the bins whose gradient norm is smaller than a predetermined threshold τ_d at inference. For such queries, we directly retrieve the stored values from the bins. Otherwise, the full MLP path is used to compute new RGB and σ . We apply gradient norm $\partial = \|\partial\omega/\partial\mathbf{p}_o\|_2^2 + \|\partial\omega/\partial\mathbf{d}_o\|_2^2$ for the dynamic models, and $\partial = \|\partial\omega/\partial\mathbf{d}\|_2^2$ for the static model, where $\omega \in \{r, g, b, \sigma\}$. Fig. 5(b) presents an overall naive reusing process.

Fig. 3 illustrates the qualitative results. Fig. 3(a) shows an example of a reproduced dynamic object, a vehicle. It clearly shows that reusing RGB color values generates abnormalities, e.g., blended colors at the rear of the grey car. The naive reusing approach is also imperfect for reproducing a static background. As seen in Fig. 3(b), the street lamp on the right or trees in the back (yellow circled) cannot be pinned down, displaying ghost effects. This indicates the reusing decision was made erroneously, reusing the features that are significantly changes across the frames.

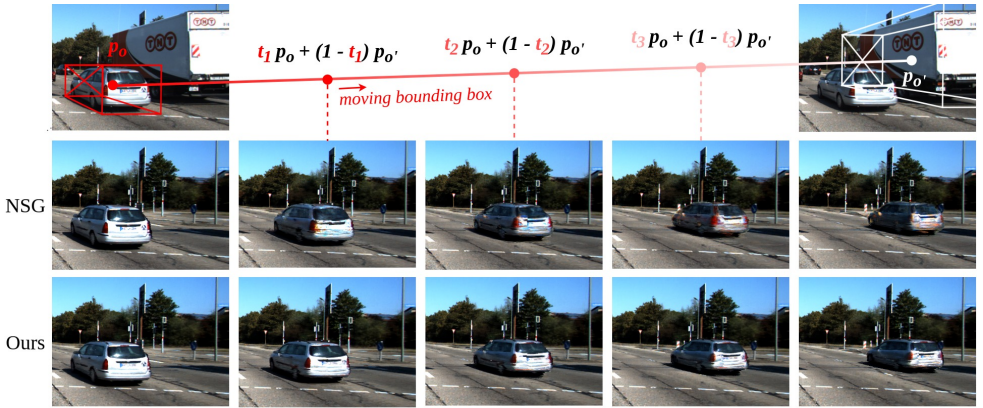


Figure 4: **Translation experiment of a dynamic object.** In NSG, an object’s color and shape of tail light change depending on the object’s global location. Meanwhile, CF-NSG well preserves canonical properties of the object disentangled from environmental factors.

Limitation of the NSG Representations. From the previous example, we notice two inadequacies in naive reusing of NSG that may lead to the disappointing results. First, gradient norm may not provide a convincing criterion for which values in the bins are reusable and advantageous for the rendering quality. Second, RGB color values may not be appropriate for direct reuse. As shown in Fig. 2, RGB and σ hardly change in most frames. Then, what factors do cause these inadequacies? We notice that RGB values change abruptly once in a while due to external factors such as shadow from nearby environment, global illumination, or the change of a dynamic object’s location. Considering features that are independent of these external factors would lead to a more reliable criterion as well as an appropriate reuse.

First, we define two distinct properties that are canonical to the object (*e.g.*, intrinsic color [27] or shape) and those from external factors. The former should be maintained consistent across frames regardless of the object’s current location or viewing direction, while the latter may change depending on the spatio-temporal environment. Here, spatio-temporal consistency incorporates both actual time change or movement of an object (temporal change) and the viewpoint change coming from the use of stereo-cameras on different locations (spatial change). Therefore, we aim to reuse features representing canonical properties.

However, we hypothesize that NSG fails to completely disentangle features intrinsic to each object and those affected by its transient environment. Since RGB color values are produced by both object-internal and external factors, supervision from RGB color (GT pixel values) alone would be unable to instruct about canonical properties of the object. To support our hypothesis, we conduct an experiment shown in Fig. 4. The first row shows two objects, a white car (o) and a truck behind it (o') in the training set, whose bounding boxes p_o and $p_{o'}$ are marked in red and white, respectively. In the second row, we render the white car (o) at intermediate locations between p_o and $p_{o'}$, using the NSG representation of the car (I_o). Even though we only change its physical location, we observe that canonical properties of the object, *e.g.*, its color and shape of the tail lights, are significantly contaminated depending on its location. This result reveals that the object-intrinsic and environmental representations are severely entangled in NSG. That is, the object-intrinsic representation learned by NSG is heavily affected by transient factors, such as the viewing direction or its location.

4 The Proposed Method: Consistency-Field-based NSG

To improve the efficiency of NSG by the feature-reusing framework while maintaining the image quality, our observations from Sec. 3 imply two conditions: 1) a solid criterion should be used to determine which bins are reusable without hurting the rendering quality, and 2) the reused features should be able to properly represent the characteristics that is intrinsic to each object. We propose the Consistency-Field-based NSG (CF-NSG), illustrated in Sec. 4.1 and Fig. 5, satisfying these two conditions and achieving disentangled representations. By the term *consistency*, we refer to the characteristics of the query that strongly show canonical and consistent properties of the object across frames. We call our method *Consistency-Field* since CF-NSG reformulates the radiance fields to additionally measure the consistency of each query across frames. We also discuss how we further boost the efficiency and adapt our final model to practical settings under limited memory in Sec. 4.2.

4.1 Consistency Fields and Reusable Features

Solid Reusability Criterion. We have shown in Sec. 3 that the gradient-norm-based naive reusing cannot retain the quality of rendered images. Therefore, we take a learnable approach to enable our model to establish a stronger criterion. Compared to NSG in Fig. 5(a), our CF-NSG in Fig. 5(c-d) additionally returns *consistency scores* $s \in [0, 1]$ estimating how consistent the query is across frames. We indicate $s_{bg}(\mathbf{x})$ for the static model, and $s_c(\mathbf{x}_o)$ for the dynamic model of object o belonging to class c , respectively. The queries with higher consistency scores than a predetermined threshold τ have a strong tendency to be invariant spatio-temporally, meaning that they are safe to be reused in other frames.

Since there is no ground truth for consistency scores, we give an auxiliary loss at training as follows: both the full feed-forward pass (**black solid line** in Fig. 5(c)) and the reuse pass (**red dashed line**) are activated, producing $\omega_{\text{full}}, \omega_{\text{reuse}} \in \{(r, g, b, \sigma) : r, g, b, \sigma \in [0, 1]\}$, respectively. Then, they are aggregated to

$$\omega_{\text{mixed}} = s \cdot \omega_{\text{reuse}} + (1 - s) \cdot \omega_{\text{full}} \quad (3)$$

A batch of ω_{mixed} are integrated along rays to yield an interpolated image. The color difference between this interpolated and GT images induces the loss, which is backpropagated to update s as well as other parameters. At inference, we choose each query’s path based on s , to reuse (**red dashed line** in Fig. 5(c)), skip (**red solid line**, see Sec. 4.2) or full feed-forward.

Reusable Representations. As observed in Sec. 3, the quality of the rendered images deteriorates as inappropriate features are reused. To overcome this issue, we explicitly learn a *canonical feature* $\mathbf{y} \in \mathbb{R}^l$, where l is a hyperparameter. \mathbf{y}_{bg} indicates that of the static model and is a function of \mathbf{x} . The canonical feature \mathbf{y}_o of a dynamic object o is a function of its latent vector \mathbf{l}_o along with \mathbf{x}_o . These canonical features differ from the intermediate features of NSG. NSG uses intermediate features to map the positional and directional inputs to higher-frequency, whereas our canonical features are for well-representing spatio-temporally consistent properties through the feature-reusing framework. It is these canonical features \mathbf{y} (along with the density σ) that are stored in the corresponding memory bin at training (**blue dashed line** in Fig. 5(c)), contrast to the actual RGB values used in the naive reusing in Sec. 3, Fig. 5(b). Also, reusing is based on consistency score instead of gradient norm ∂ at inference. **Overall Pipeline.** Fig. 5(c-d) illustrate the overall pipeline of our CF-NSG for the dynamic model. CF-NSG learns consistency score s and canonical features \mathbf{y}_o simultaneously, and reuses \mathbf{y}_o depending on s . If the model reuses improper features largely affected by transient

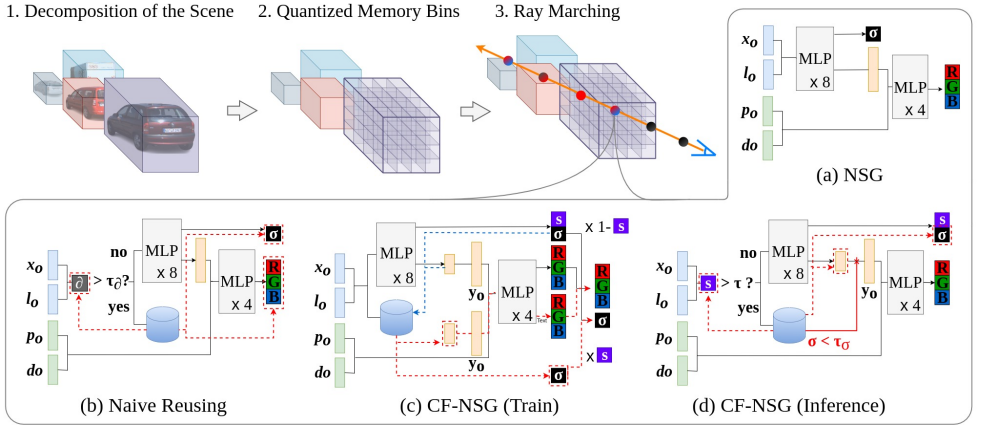


Figure 5: The overall pipeline of CF-NSG (the dynamic model). (a)–(d): Instead of computing the whole queries, CF-NSG utilizes feature-reusing frameworks reducing an amount of full feed-forward pass at inference. (c): CF-NSG learns consistency fields and appropriately reuses features by loss from interpolated outputs (see Eq. 3) and regularizer (see Eq. 4)

factors, L_2 -loss between the rendered image and GT would increase and imposes a penalty. Thus, y_o with a high s are encouraged to be even more object-intrinsic across frames. The static model also follows a similar structure, illustrated more in detail in the Supp. Materials. We also provide pseudocodes for both dynamic and static models in the Supp. Materials.

Training Objective. Our full objective is as follows:

$$\mathcal{L} = \sum_{r \in \mathcal{R}} [\|\hat{C}(r) - C(r)\|_2^2 + \|\hat{C}_{\text{mixed}}(r) - C(r)\|_2^2] + \sum_{x, x_o \in \mathcal{X}} \left[\frac{1}{\|s_{bg}(x)\|^2} + \frac{1}{\|s_c(x_o)\|^2} \right] + \frac{1}{v} \|l_o\|_2^2, \quad (4)$$

where we denote the predicted pixel color from a ray r by full feed-forward rendering as \hat{C} , the one by aggregation in Eq. (3) as \hat{C}_{mixed} , and the reference color as C . The original training objective of NSG consists of the L_2 -loss on the rendered image (the first term in Eq. (4)) and a Gaussian prior on the object latent vector l_o (the last term). To encourage our model to reason about consistency, we introduce additional regularization to the objective function. Since we aggregate ω_{mixed} by a convex combination of full feed-forward outputs and reused outputs, the model might converge to the trivial solution of zero reuse. Thus, we add a regularization term to penalize low reuse scores for s_{bg} and $s_c(x_o)$ (the second term).

Where Does Disentanglement Come from? In NSG, where L_2 -loss between the rendered and GT images is the only supervision, intermediate features do not need to preserve unique information about an object or a scene from l_o , x_o or x , diluted with other inputs quickly. CF-NSG, on the other hand, is simultaneously trained by an additional task, estimating the consistency score s to predict the proper reusability. Since the model produces both the canonical features y and s , y is naturally aware of consistent properties of the corresponding bin. For this reason, while NSG appears to confuse intrinsic vehicle properties and its transient location in Fig. 4, CF-NSG is less affected thanks to better disentanglement between them.

Are Environmental Effects Not Considered? We emphasize that CF-NSG *does not* disregard transient factors and is able to learn environmental properties as well. The second MLP, a shallow 4-layer MLP, receives transient inputs (global location p_o , viewing direction d_o and d) along with y that well-represents object-intrinsic properties, enabling the second MLP to learn transient properties more effectively. This process is not skipped regardless of s .

4.2 Further Improvements

Skipping. We utilize the consistency score s more aggressively to further boost the efficiency. We find that a bin with a high consistency score s and a low density σ is likely to be an empty space without a relevant scene content. We skip such a bin and allow queries to be more densely distributed near the content. NSVF [10] is closely related to our work in that they improve the efficiency of NeRF by defining a set of voxels in the scene and skipping the empty voxels. However, the difference arises from the way of finding the empty space. NSVF prunes the empty voxels based on σ , while our method uses both the s and σ . Relying only on σ requires a finer adjustment of the threshold between an almost transparent car window and a truly empty space. On the other hand, they can be distinguished through s ; the former has a lower s , while the latter has a higher s , since the latter has consistently low density over time. Our method enjoys additional gain in efficiency by skipping more aggressively.

Memory-efficient Implementation. Storing canonical features $\mathbf{y} \in \mathbb{R}^l$ for each spatial bin imposes a considerable memory footprint, making the framework impractical. That is, we are faced with two contradicting desiderata: keeping as little information in memory as possible, while mapping it to a rich feature space. To this end, we apply a learnable factorization inspired by [21, 22]. Specifically, \mathbf{y} is factorized to $\text{flatten}(\mathbf{u}_1 \times \mathbf{u}_2) + \mathbf{z}$, where $\mathbf{z} \in \mathbb{R}^l$, $\mathbf{u}_i \in \mathbb{R}^{m^2}$, and $\sqrt{l} = m^2$. Each \mathbf{u}_i is similarly factorized again by $\mathbf{u}_i = \text{flatten}(\mathbf{v}_{i,1} \times \mathbf{v}_{i,2})$, where $\mathbf{v}_{i,j} \in \mathbb{R}^m$. Here, \mathbf{z} keeps shared information about each component of the scene (*e.g.*, each dynamic object or background) and is shared for all queries belonging to the same component. Since it is stored only once per each component, we keep \mathbf{z} unfactorized to fully enjoy l degree of freedom. On the other hand, the rest, memory-bin-specific components, are aggressively factorized to rank- m for efficient memory footprinting. We use $l = 256, m = 4$. Instead of storing 256D, we store $4\text{D} \times 4$, thereby reducing the memory usage by 93%.

5 Experiment

Implementation Details. We use KITTI [5] and Objectron [11] datasets that provide 3D bounding boxes for objects in a scene. KITTI provides multi-object tracking information captured by stereo-cameras while Objectron contains more diverse types of objects with more drastic camera view changes. We set baselines including NSG [16], NeRF [12] and NeRF with temporal inputs [16]. We also compare with NSVF [10] that improved efficiency of NeRF and D-NeRF [14] that extended NeRF to the dynamic scenes. However, since the former is mainly targeted to the static scenes and the latter is not considering efficiency, they are not directly comparable. Nevertheless, we include them as baselines to show a general tendency. We do not compare with NSFF [9] since the unbounded scene of KITTI induced an unstable training for depth estimation loss. We refer to Supp. Materials for more details.

Comparison with Baselines. Tab. 1 quantitatively compares the quality of images from each implicit neural representation framework, with its computational cost indicated by the number of queries. Impressively, our CF-NSG achieves image quality close to that of NSG by only using 15–53% of queries. Note that the table is sorted by the number of queries used by each method, so for the methods below NSG, NSG scores are considered as the *upper bound*. When we train vanilla NSG with reduced number of queries (‘NSG-reduced’ in Tab. 1), we observe significant performance drop. Fig. 6 compares our method to baselines qualitatively. The first row compares our CF-NSG against baselines on reproduction of a frame seen during training, and the second row compares NSG and CF-NSG on novel scene

Dataset	Method	#Queries	PSNR(\uparrow)	SSIM [\square](\uparrow)	LPIPS [\square](\downarrow)	iOF [\square] $\times 10^6$ (\downarrow)	tLP [\square] $\times 100$ (\downarrow)
KITTI	D-NeRF [\square]	8.72 \times	16.33	0.505	0.418	3.823	4.835
	NeRF [\square]	7.90 \times	20.99	0.621	0.446	2.702	3.840
	NSVF [\square]	6.38 \times	22.95	0.706	0.386	2.831	5.071
	NeRF+time [\square]	1.64 \times	24.86	0.653	0.492	2.272	1.563
	NSG [\square]	1 \times	29.54	0.914	0.171	0.619	0.265
	NSG-reduced [\square]	0.75 \times	24.69	0.702	0.452	1.625	1.990
	CF-NSG (ours)	0.15\times	28.70	0.891	0.204	0.766	0.266
Objectron <i>chair</i>	NSG [\square]	1 \times	29.20	0.864	0.286	0.224	0.411
	NSG-reduced [\square]	0.29 \times	27.65	0.866	0.263	0.245	0.296
	CF-NSG (ours)	0.28\times	28.58	0.878	0.259	0.229	0.354
Objectron <i>camera</i>	NSG [\square]	1 \times	29.17	0.854	0.273	0.414	0.673
	NSG-reduced [\square]	0.85 \times	22.82	0.701	0.454	1.182	1.383
	CF-NSG (ours)	0.53\times	26.01	0.789	0.363	0.471	0.634

Table 1: **Quantitative Comparison on KITTI [5] and Objectron [10].** CF-NSG achieves comparable results with NSG and outperforms other baselines. Methods above NSG use more queries than NSG, while those below NSG use less (thus, NSG is an *upper-bound*).

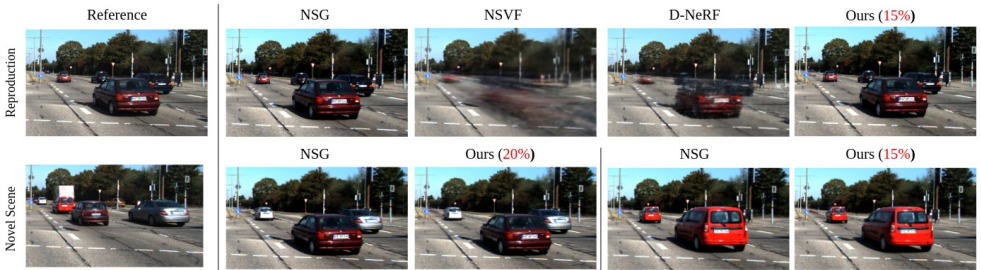


Figure 6: **Qualitative comparisons on KITTI [5].** We compare NSG [16], NSVF [21], D-NeRF [19] and CF-NSG. We note that CF-NSG uses 80–85% fewer number of queries.

manipulation, where we sample dynamic objects in the reference frame and generate a scene in a new arrangement. We can see that the rendered images of NSG and ours are almost indistinguishable from human eyes in spite of significantly less number of (15–20%) queries.

We also show that our method indeed learns disentangled representations regarding consistency. We conduct a similar experiment as in Sec. 3 for CF-NSG, and show the result in the third row of Fig. 4. Our CF-NSG relatively well preserves object-intrinsic properties regardless of the object’s location. Making use of disentangled representations, we also can apply CF-NSG to various scene compositions in a more stable manner (see Supp. Material).

Ablation Study. In Tab. 2, we show the benefit of using score-based skipping and feature-reusing framework progressively. After skipping, 62.7% of queries are left, while 75.8% of them take the reusing pass. We also show our choice of 100 bins per dimension balances the trade-off between the rendered image quality and the memory cost. More ablations on various skipping and memory-efficient implementations are in Supp. Materials.

Trade-off between Computation and Memory Cost. Since our approach stores and reuses consistent components, our efficiency gain is accompanied by additional memory footprint. Tab. 3 represents the rendering quality and additional memory cost as a function of com-

Method	#Queries	PSNR(\uparrow)	#Bins per dim.	PSNR(\uparrow)	Mem.(MB)
NSG[10]	1	29.54	75	28.59	165.73
+ Score-based skipping	0.627	28.80	100	28.70	313.83
+ Feature-Reusing (CF-NSG)	0.152	28.70	125	28.72	503.12

Table 2: **Ablations for each component and bin size.** We validate score-based skipping and feature-reusing frameworks respectively and evaluate effect of bin size.

#Queries	0.15 \times	0.30 \times	0.45 \times	0.60 \times	0.90 \times	1 \times	NSG(1 \times)
FLOPs/frame	8.78×10^{11}	1.66×10^{12}	2.30×10^{12}	2.95×10^{12}	4.50×10^{12}	5.00×10^{12}	5.12×10^{12}
Mem.(MB)	313.8	249.9	184.3	131.92	119.76	0	0
PSNR(\uparrow)	28.70	28.73	28.78	28.78	28.78	28.88	29.54

Table 3: **Trade-off between speed and additional memory usage with CF-NSG.** We reveal relation between additional memory cost and number of queries for full feed-forward pass.

putation cost in the number of queries and FLOPs. We observe a relatively small memory footprint of 300MB leads to large efficiency gains by up to 85% less queries for forward pass, with little drop in the image quality. (The actual speed improvement in FLOPs is about 82.9%, slightly less than 85%, since reusing is still not completely free.) In practice, a user may flexibly choose a proper reusing rate considering the resource budget.

6 Related Work

Recently, the advancement of implicit or neural representations has enabled researchers to achieve photo-realistic views [7, 10, 15]. To suggest a better representation for over-smoothed renderings of the existing methods, Mildenhall *et al.* [12] introduced Neural Radiance Fields (NeRF). However, training and rendering processes based on neural representations often require time-consuming ray marching. To improve efficiency of NeRF-based models, some research has introduced more efficient data structures *e.g.*, caching [9, 6, 13], visual hull [8], sparse voxel [11], view-dependent multiplane image [24] and octree [26].

Another stream of research introduced representations of complex and dynamic scenes [9, 18, 19]. To clarify the difference between static and dynamic scene representations, the latter involves movement in both the camera and the scene. Hence, it is important to consider interactions with the scene such as global illumination to determine the appearance of a dynamic object. Due to this difference, more factors need to be considered to efficiently render the dynamic scenes besides simply modeling each canonical space of dynamic object as an efficient radiance fields for static scenes *i.e.*, directly applying [9, 6, 8, 10, 13, 24, 26].

7 Summary

We propose CF-NSG, a novel framework for representing multi-object dynamic scenes efficiently by utilizing a feature-reusing framework based on consistency-fields. CF-NSG is able to render images using only 15% of the original number of queries with little compromise to the image quality. Also CF-NSG enjoys more extended novel scene manipulation, taking advantage of disentangled representation with respect to spatio-temporal consistency.

Acknowledgement

This work was supported by the New Faculty Startup Fund from Seoul National University and by National Research Foundation (NRF) grant (No. 2021H1D3A2A03038607/15%, 2022R1C1C1010627/15%) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant (No. 2022-0-00264/10%, 2021-0-01778/10%, 2022-0-00320/25%, No.2022-0-00953/25%) funded by the Korea government (MSIT).

References

- [1] Adel Ahmadyan, Liangkai Zhang, Artsiom Ablavatski, Jianing Wei, and Matthias Grundmann. Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7822–7831, 2021.
- [2] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. Learning temporal coherence via self-supervision for GAN-based video generation. *ACM Transactions on Graphics (TOG)*, 39(4):75–1, 2020.
- [3] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3D object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017.
- [4] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. FastNeRF: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021.
- [5] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [6] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021.
- [7] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3D scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020.
- [8] Naruya Kondo, Yuya Ikeda, Andrea Tagliasacchi, Yutaka Matsuo, Yoichi Ochiai, and Shixiang Shane Gu. VaxNeRF: Revisiting the classic for voxel-accelerated neural radiance field. *arXiv preprint arXiv:2111.13112*, 2021.
- [9] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021.
- [10] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *arXiv preprint arXiv:2007.11571*, 2020.

- [11] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.
- [12] 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 *European conference on computer vision*, pages 405–421. Springer, 2020.
- [13] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4): 102:1–102:15, July 2022.
- [14] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021.
- [15] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020.
- [16] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2856–2865, 2021.
- [17] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [18] Keunghong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021.
- [19] 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*, pages 10318–10327, 2021.
- [20] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. DeepVoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [21] Ivan Skorokhodov, Savva Ignatyev, and Mohamed Elhoseiny. Adversarial generation of continuous images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10753–10764, 2021.
- [22] Joseph Suarez. Language modeling with recurrent highway hypernetworks. *Advances in neural information processing systems*, 30, 2017.

- [23] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. IEEE, 2003.
- [24] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. NeX: Real-time view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8534–8543, 2021.
- [25] Xinchun Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3D object reconstruction without 3D supervision. *arXiv preprint arXiv:1612.00814*, 2016.
- [26] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoc-trees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021.
- [27] Mingyang Zhang, Pengli Wang, and Xiaoman Zhang. Vehicle color recognition using deep convolutional neural networks. In *Proceedings of the 2019 International Conference on Artificial Intelligence and Computer Science*, page 236–238. Association for Computing Machinery, 2019.
- [28] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.