

Self-Aligning Depth-regularized Radiance Fields for Asynchronous RGB-D Sequences

Yuxin Huang¹ Andong Yang² Yuantao Chen^{1,3} Runyi Yang^{1,4}
 Zhenxin Zhu^{1,5} Chao Hou^{1,6} Hao Zhao^{1,†} Guyue Zhou¹
¹AIR, Tsinghua University

²University of Chinese Academy of Sciences ³Chinese University of Hong Kong (Shenzhen)

⁴Imperial College London ⁵Beihang University ⁶The University of Hong Kong

572142867olive@gmail.com, zhaohao@air.tsinghua.edu.cn

Abstract

It has been shown that learning radiance fields with depth rendering and depth supervision can effectively promote the quality and convergence of view synthesis. However, this paradigm requires input RGB-D sequences to be synchronized. In the UAV city modeling scenario, there exists asynchrony between RGB images and depth images due to the different frequencies of the solid-state LiDAR and RGB sensors. To synthesize high-quality views in such a scenario, we propose a novel time-pose function, which is an implicit network that maps timestamps to $SE(3)$ elements. To train this function, we also design a joint optimization scheme to jointly learn the large-scale depth-regularized radiance fields and the time-pose function. Furthermore, we propose a large synthetic dataset with diverse controlled mismatches and ground truth to evaluate this new problem setting systematically. The proposed approach has been evaluated on both datasets and in a real drone. To evaluate the impact of view density, each algorithm was tested on three different trajectories with different view densities. Compared to state-of-the-art baseline methods, the proposed approach reduces reconstruction error by 35.26% in city modeling scenarios. Our code is available at github.com/saythe17/AsyncNeRF.

1. Introduction

Radiance field learning has made significant breakthroughs in view synthesis, augmented reality, autonomous navigation, and computational photography etc [2, 22]. Incorporating depth rendering and depth supervision into radiance fields has been demonstrated as a helpful regulariza-

tion technique in several recent studies [21, 35, 36]. However, this technique has not yet been successfully introduced into radiance field learning from UAV (Unmanned Aerial Vehicle) images. A closer look at the aforementioned works indicates that they presume synchronized RGB and depth signals, which is hard to guarantee in UAV vision. Due to the power supply and payload limitations of drones, sensors such as LiDAR have to operate at a lower frequency, leading to the problem of asynchronous sensor data. In some cases where RGB data is used for visual odometry, the low output frequency of the pose estimation can even result in inaccurate depth image poses. Therefore, we study the *problem* of learning depth-regularized radiance fields from asynchronous RGB-D sequences.

As a recap, the canonical radiance field [16] learns a neural network parameterized by θ that represents a 3D scene from input images I and their intrinsic/extrinsic parameters \mathcal{T}_I . To alleviate the reliance on \mathcal{T}_I , some works [8, 10, 32] aim to resolve a different problem that self-calibrates \mathcal{T}_I . In other words, they jointly learn θ and \mathcal{T}_I from input images I . Similarly, the *formulation* considered here is to learn scene representation θ , camera parameters \mathcal{T}_I and \mathcal{T}_D from inputs RGB images I and depth images D .

We utilize an important prior specific to this problem: RGB-D frames are sampled from the same physical trajectory. As shown in Fig. 1-a/b, \mathcal{T}_I and \mathcal{T}_D can be considered as samples from a function that maps timestamps to $SE(3)$ elements. We name this function as **time-pose function** and model it with a neural network parameterized by ϕ . As such, we address the *problem* with a *new formulation* that learns scene representation θ and the time-pose function ϕ from inputs RGB images I and depths D . Both θ and ϕ are implicit neural representation networks (or coordinate-based networks) that allow fully differentiable training.

We also propose an effective learning scheme. In the **first** stage, we fit the time-pose function ϕ using one modal-

[†]Corresponding author: Hao Zhao.

*This work was sponsored by Tsinghua-Toyota Joint Research Fund (20223930097).

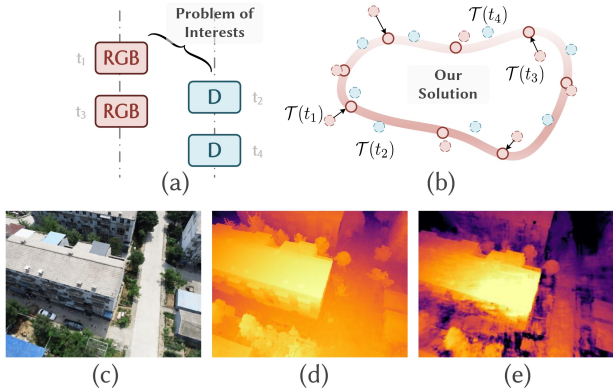


Figure 1. To learn a depth-regularized radiance field using (a) asynchronous RGB-D sequences, we propose a (b) time-pose function to map from timestamp to camera pose. For a (c) novel view, our method can render a more accurate (d) depth map than (e) Mega-NeRF.

ity (e.g., RGB images) and infer the poses of the other using a balanced pose regression loss and a speed regularization term. **Secondly**, we bootstrap a large-scale radiance field θ based upon Mega-NeRF [30] using the outputs of the trained time-pose function. Thanks to the time-pose function, depth regularization can be imposed here in spite of RGB-D misalignment. **Thirdly**, since the cascade of two fully differentiable implicit representation networks, we jointly optimize the 3D scene representation θ and compensate for pose errors by updating ϕ .

Due to limited available datasets, we contribute a synthetic dataset (named AUS) for systematic evaluation, generating realistic drone trajectories of varying difficulty levels from six large-scale 3D scenes. Simple trajectories are heuristically designed in a zig-zag pattern, while complex ones are generated using manual control signals in simulation. We also control the mismatch between RGB-D sequences with various protocols to cover a wide range of real-world scenarios the algorithm may encounter. Through comprehensive experiments, we show the proposed method outperforms baseline methods without regularization and our design choices contribute positively to performance. To further verify the effectiveness of the proposed method, we present a real-world evaluation using asynchronous data on the drone. Our depth rendering results (on unseen viewpoint) are shown in Fig. 1-d, which outperforms Mega-NeRF, shown in Fig. 1-e, across multiple metrics. This success is credited to the usage of depth regularization, which was made possible by our novel algorithm.

To summarize, we have the following contributions: (1) We formalize the new *problem* of learning depth-regularized radiance fields from asynchronous RGB-D sequences, which is rooted in many UAV applications. (2) We identify an essential domain-specific prior in this problem: RGB-D frames are sampled from the same underlying

trajectory. We instantiate this prior into a novel time-pose function and develop a cascaded, fully differentiable implicit representation network. (3) To systematically evaluate the task, we contribute a photo-realistically rendered synthetic dataset that simulates different types of mismatch. (4) Through benchmarking on this new dataset and real-world asynchronous RGB-D sequences, we demonstrate that our method can promote performance over baselines.

2. Related Works

Large-scale Radiance Fields. Neural Radiance Field (NeRF) [16] has shown impressive results in neural reconstruction and rendering. Some follow-up works [12, 27, 42] have significantly improved the rendering quality and training efficiency. With the development of deep learning frameworks, Large Language Models (LLMs) have exhibited exceptional performance across various tasks [46], leading to methods [19, 33] that integrate LLMs with NeRF to enhance scene understanding. However, its capacity to model large-scale unbounded 3D scenes is limited. Several strategies [15, 28, 30, 36, 39, 48] have been proposed to address this limitation, with a common principle of dividing large scenes into blocks or decomposing the scene into multiple levels. Block-NeRF [28] clusters images by dividing the whole scene according to street blocks. Mega-NeRF [30] utilizes a clustering algorithm that partitions sampled 3D points into different NeRF submodules. BungeeNeRF [36] trains NeRFs using a growing model of residual blocks with predefined multiple scales of data. Switch-NeRF [15] uses a gating network to jointly learn scene decomposition and NeRFs without any priors on geometry. However, these prior works fail to leverage the rich geometric information in depth images for effective regularization.

Depth-regularized Radiance Fields. Volumetric rendering requires extensive samples and sufficient views to effectively differentiate between empty space and opaque surfaces. Depth maps can serve as geometric cues, providing regularization constraints and sampling prior, which accelerates NeRF’s convergence towards the correct geometry. DS-NeRF [3] enhances this process using depth supervision from 3D point clouds, estimated by structure-from-motion, and a specific loss for rendered ray termination distribution. Mono-SDF [41] and Dense-Depth Prior [22] further supplement this with a pretrained dense monocular depth estimator for less-observed and textureless areas. To adapt NeRF for outdoor scenarios, URF [21] rasterizes a pre-built LiDAR point cloud map to generate dense depth images and alleviates floating elements by penalizing floaters in the free space. Moreover, S-NeRF [38] completes depth on sparse LiDAR point clouds using a confidence map, effectively handling street-view scenes with limited perspectives. However, those methods are not readily applicable to UAV-captured images due to the lack of suitable synchronized

sensors for long ranges.

Broader UAV Vision and Synchronization. Like autonomous driving, UAV vision is drawing increasing attention due to its unique characteristics. Broader UAV vision covers many topics like counting [7, 34], trajectory forecasting [18], intention prediction [37], object tracking [17], physics understanding [43], next-best-view prediction [6], 3D reconstruction [45], and calibration [20, 47]. Sensor synchronization is challenging for UAV vision (and other settings) and several works address the problem from an algorithmic perspective. One possibility is to adopt tailored hardware designs or software protocols [1] to synchronize all the devices. Another branch of sensor-agnostic methods utilizes temporal priors by using Sum-of-Gaussians [4] or parametric interpolation functions [40]. These methods considered the temporal-spatial relationships but are sub-optimal since they fail to take advantage of all available sensory observations (e.g. images or depths).

3. Preliminaries

Problem & Challenge. Our goal is to learn a neural radiance field parameterized by θ for large-scale scene representation from UAV images as done in prior works [30, 36]. However, these works fail to leverage depth supervision, which is known [3, 22] as useful for training floater-less NeRFs. To our knowledge, there are no easily accessible synchronized RGB-D sensor suites for **large-scale** outdoor scenes, and synchronizing them according to timestamp cannot fully address the misalignment issue. Instead of using expensive hardware, we take an algorithmic perspective.

Pose Representation. We use a translation vector $\hat{\mathbf{x}} \in \mathbb{R}^3$ and a rotation matrix $\hat{\mathbf{q}} \in \mathbf{SO}(3)$ as pose representation.

Input & output. There are some prior works on large-scale scene modeling using aerial images [5, 6, 30, 36]. In this study, we assume an input RGB-D stream captured by the drone: a set of RGB camera images $\{I^{(i)}\}_{i=1}^{N_I}$ and a set of depth maps $\{D^{(j)}\}_{j=1}^{N_D}$ shown in Fig. 1-a and we aim to recover the spatiotemporal transformations between them. Without loss of generality, we assume a set of camera poses $\{\mathcal{T}_I^{(i)}\}_{i=1}^{N_I}$ for color images are obtained by a Structure from Motion (SfM) algorithm. The neural scene representation parameterized by θ outputs an image \hat{I} as well as a depth map \hat{D} at a given perspective camera pose \mathcal{T}_I and \mathcal{T}_D .

4. Method

We introduce in Section 4.1 the details of learning an implicit time-pose function. In Section 4.2, we describe our neural scene representation networks and the bootstrapping strategy. In Section 4.3, we adopt depth supervision and jointly train the time-pose function with RGB-D pairs. In Section 4.4, we summarize the overall pipeline of our

method, for clarity.

4.1. Time-Pose Function

We represent the camera trajectory as an implicit time-pose function ϕ whose input is a timestamp, and whose output is a 6-DoF pose.

Definition. Specifically, the time-pose function can be represented as $\phi : t \rightarrow \hat{\mathcal{T}} = [\hat{\mathbf{x}}, \hat{\mathbf{q}}]$, where t is the timestamp of capture, and $\hat{\mathcal{T}}$ is the estimated pose represented by a translation vector $\hat{\mathbf{x}} \in \mathbb{R}^3$ and a rotation matrix $\hat{\mathbf{q}} \in \mathbf{SO}(3)$.

We represent the time-pose function using a neural network. Fig. 3 depicts the network structure, comprising a compact 1-D multi-resolution hash grid $\{\mathcal{G}^{(l)}\}_{l=1}^L$ and an MLP decoder.

Hash encoding. For the queried time-stamp t_i , the hash encoding of $\lfloor t_i \rfloor - 1$, $\lfloor t_i \rfloor$, and $\lfloor t_i \rfloor + 1$ are extracted in each layer $\mathcal{G}^{(l)}$ of the multi-resolution hash grid. We perform quadratic interpolation on the extracted hash feature:

$$\begin{aligned} \mathcal{V}_i^l = & \frac{1}{2}(t_i - \lfloor t_i \rfloor)(t_i - \lfloor t_i \rfloor - 1)\mathcal{G}^{(l)}(\lfloor t_i \rfloor - 1) \\ & - (t_i - \lfloor t_i \rfloor + 1)(t_i - \lfloor t_i \rfloor - 1)\mathcal{G}^{(l)}(\lfloor t_i \rfloor) \\ & + \frac{1}{2}(t_i - \lfloor t_i \rfloor + 1)(t_i - \lfloor t_i \rfloor)\mathcal{G}^{(l)}(\lfloor t_i \rfloor + 1). \end{aligned} \quad (1)$$

The interpolated feature are then concatenated together $\mathcal{V}_i = \text{concat}\{\mathcal{V}_i^l\}_{l=1}^L$.

Decoding. After obtaining the interpolated feature vector, an MLP with two separated decoder heads is used to predict the output translation \hat{x}_i and rotation \hat{q}_i vectors respectively. Specifically, we use a 10-layer MLP with 1024 dimensions in each layer, including a skip connection that concatenates the timestamp input to the 5-th layer which is shown in Fig. 3. The forward pass can be expressed in the following equations:

$$\mathcal{V}_i = \text{concat}\{\text{interp}(\text{h}(t; \pi_l), \mathcal{G}^{(l)})_{l=1}^L\}, \quad (2)$$

$$\hat{\mathcal{T}}_i = [\hat{x}_i, \hat{q}_i] = [l_{\text{trans}}(\mathcal{F}_{\text{MLP}}(\mathcal{V}_i)), l_{\text{rot}}(\mathcal{F}_{\text{MLP}}(\mathcal{V}_i))], \quad (3)$$

where interp denotes interpolation, h is the hash function parameterized by π_l . \mathcal{F}_{MLP} , l_{trans} , l_{rot} are the MLP networks and the decoder heads, with Φ_{MLP} , Φ_{trans} , Φ_{rot} representing their parameters.

Justification. The reason why we choose this architecture is as follows: The time-pose function is a coordinate-based function that may contain coarse and fine-level feature components [29], which is shown by the ground truth trajectory in Fig. 5, and this architecture allows us to sample the hash encodings from each grid layer with different resolutions and perform quadratic interpolation on the extracted encodings to obtain a feature vector \mathcal{V}_i when querying a specific timestamp t that is in the range of all timestamps.

RGB-D Transformation. Using the time-pose function, we can predict the camera pose for the depth images. Since




Figure 2. **Method Pipeline.** The time-pose function is modeled using a 1-D multi-resolution hash grid with direct and speed losses. After bootstrapping the scene representation networks with pure RGB signals, the predicted depth sensor poses are used for jointly optimizing the NeRFs’ parameters θ . At each timestamp (t_i from RGB sequence or t_j from depth sequence), only one modality of sensor signals is provided, thus only one loss term is activated (shown on the right).




Figure 3. Our implementation of the Time-Pose Function with a multi-resolution hash grid. Blue and orange networks are of different resolution.

both the depth maps and the RGB images are collected by the same drone, they cover the same spatial-temporal footprints except for the difference in the placement of the two sensors on the aircraft. For every depth frame, we first predict the RGB camera pose using the capture timestamps of the depth sensor with the time-pose function and then transform the predicted RGB camera pose to the depth sensor pose with a pre-calibrated pose transformation $\mathcal{T}_{I \rightarrow D}$ between sensors.

Training. To train the Time-Pose Function, we propose the following objective function:

$$\mathcal{L} = \lambda_{\text{trans}} \mathcal{L}_{\text{trans}} + \lambda_{\text{rot}} \mathcal{L}_{\text{rot}} + \lambda_{\text{speed}} \mathcal{L}_{\text{speed}}, \quad (4)$$

where $\mathcal{L}_{\text{trans}}$, \mathcal{L}_{rot} , $\mathcal{L}_{\text{speed}}$ are translation, rotation and speed losses respectively as shown in the left panel of Fig. 2. and λ_{trans} , λ_{rot} , λ_{speed} are the weighting parameters. Note that λ_{trans} and λ_{rot} are automatically adjusted as explained in a later paragraph.

To optimize the translation and the rotation vectors, we minimize the mean square error (MSE) between the estimated and ground-truth camera poses:

$$\mathcal{L}_{\text{trans}} = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{x}_i)^2, \quad \mathcal{L}_{\text{rot}} = \frac{1}{n} \sum_{i=1}^n (q_i - \hat{q}_i)^2. \quad (5)$$

Uncertainty Balancing. Since \hat{x} and \hat{q} are in different units, the scaling factor λ_{trans} and λ_{rot} play an important role in balancing the losses. To prevent translation and rotation from negatively influencing each other in training and to tap into possible mutual facilitation, we make the weighting factors learnable by using homoscedastic uncertainty [9] as $\mathcal{L}_{\sigma} = \mathcal{L}_{\text{trans}} \exp(-\hat{s}_{\text{trans}}) + \hat{s}_{\text{trans}} + \mathcal{L}_{\text{rot}} \exp(-\hat{s}_{\text{rot}}) + \hat{s}_{\text{rot}}$, where \hat{s} are learnable parameters, thus the loss terms are balanced during training course¹.

Speed. Observing that the time-pose function is essentially a function of translational displacement and angular displacement with respect to time, we can use the linear speed at time t_i to supervise the gradient of the network output, with regard to the input vectors. Since the linear speed variation is small and the angular speed variation is relatively larger in the scenes captured by the drone, only the average linear speed is used to supervise the neural network and the latter is not supervised in our method:

$$\begin{aligned} \mathcal{L}_{\text{speed}} &= \text{MSE}(v(t_i), \hat{v}(t_i)) \\ &= \frac{1}{n} \sum_{i=1}^n \left(v(t_i) - \frac{\partial \hat{x}}{\partial t}(t_i) \right)^2 \\ v(t_i) &= \frac{\partial x}{\partial t} \Big|_{t=t_i} \approx \frac{x_i - x_{i-1}}{t_i - t_{i-1}} \end{aligned} \quad (6)$$

4.2. Bootstrapping Large-scale Neural Radiance Fields

Partition. In this part, we introduce our proposed scene representation shown in the right half of Fig. 2. Due to the limited capacity of MLPs, we follow Mega-NeRF [30] and partition the scene map into a series of equal-sized blocks in

¹Manual selection of weights requires laborious tuning, but comparable performance can be achieved.

terms of spatial scope, and each block learns its individual scene representation with an implicit field. In this stage, we optimize the scene representation with pure RGB data. Specifically, the radiance field is denoted as $\{f_{\text{NeRF}}^{(i)}\}_{i=1}^{N_x \times N_y}$, where N_x, N_y denotes the spatial grid size. Each implicit function represents a geographic region with $\mathbf{r}_i^{\text{centroid}}$ as its centroid. The k -th scene model can be written as:

$$f_{\text{NeRF}}^{(k)}(\mathbf{o}_{\text{pts}}, \mathbf{d}) \rightarrow (c, \sigma), \quad (7)$$

where $k = \arg \min_i \|\mathbf{r}_{\text{pts}} - \mathbf{r}_i^{\text{centroid}}\|_2$ and γ is the positional encoding function.

Rendering. For view synthesis, we adopt volume rendering techniques to synthesize color image \hat{I} and depth map \hat{D} . Specifically, we sample a set of points for each emitted camera ray in a coarse-to-fine manner [16] and accumulate the radiance and the distance along the corresponding ray to calculate the rendered color \hat{I} and depth \hat{D} . To obtain the radiance of a spatial point \mathbf{r}_{pts} , we use the nearest scene model for prediction. A set of per-image appearance embedding [14] is also optimized simultaneously in the training.

$$\begin{aligned} \hat{I}(\mathbf{o}, \mathbf{d}) &= \int_{\text{near}}^{\text{far}} T(t) \sigma^{(k)}(\mathbf{r}(t)) \cdot c^{(k)}(\mathbf{r}(t), \mathbf{d}) dt, \\ \hat{D}(\mathbf{o}, \mathbf{d}) &= \int_{\text{near}}^{\text{far}} T(t) \sigma^{(k)}(\mathbf{r}(t)) \cdot t dt, \end{aligned} \quad (8)$$

where \mathbf{o} and \mathbf{d} denote the starting point and orientation of the sampled ray, $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ represents the sampled point coordinates in the world space, and $T(t) = \exp\left(-\int_{\text{near}}^t \sigma^{(k)}(\mathbf{r}(s)) ds\right)$ is the accumulated transmittance, where σ and c denotes opacity and view-dependent color for a given position. We optimize the scene representation model with only the photometric error as $\mathcal{L}_{\text{bootstrap}} = \text{MSE}(I, \hat{I})$. We conduct ablation study and observe that this bootstrapping is critical to the challenging third stage which jointly learns θ and ϕ using asynchronous RGB-D data.

4.3. Joint Optimization

While the time-pose function learns a good initialization from the RGB sequence, there are still errors to be compensated. In this section, we describe how we perform simultaneous mapping and pose optimization, which compensates for the initial error of the time-pose function.

Formulation. We jointly optimize the inaccurate camera poses and the implicit maps: when fitting parameters $\Theta_{\text{NeRF}}^{(k)}$ of the scene representation, the estimated depth camera poses $\hat{\mathcal{T}}_D^{(j)} \in \text{SE}(3)$ will be simultaneously optimized on the manifold:

$$\theta, \{\hat{\mathcal{T}}_D\} = \underset{\theta, \mathcal{T} \in \text{SE}(3)}{\text{argmin}} \mathcal{L}(\{I^{(i)}\}, \{D^{(j)}\} | \theta, \{\hat{\mathcal{T}}_D\}), \quad (9)$$

where \mathcal{L} is the objective function we demonstrate in the next paragraph, and the gradients on \mathcal{T} can be further backpropagated to update ϕ .

Loss. To train the implicit representation to obtain photo-realistic RGB rendering maps and accurate depth map estimation, we update the mapping losses as:

$$\begin{aligned} \mathcal{L} &= \lambda_{\text{color}} \sum_i \text{MSE}(I^{(i)}, \hat{I}^{(i)}) \\ &+ \lambda_{\text{depth}}(\alpha) \sum_j \text{MSE}(D^{(j)}, \hat{D}^{(j)}), \end{aligned} \quad (10)$$

where λ_{color} and $\lambda_{\text{depth}}(\alpha)$ are weighting hyper-parameters for color and depth loss, in which the depth loss weight starts to grow from zero gradually with the training process. To compensate for the error from the time-pose function extracted poses, we jointly optimize two implicit representation networks thanks to the end-to-end differentiable nature.

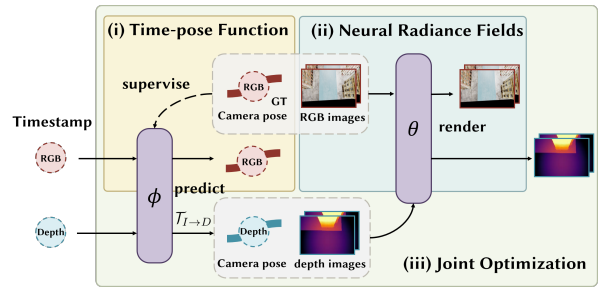


Figure 4. **Three-step Optimization.** (i) A time-pose function ϕ is trained to predict camera poses from timestamps; (ii) The neural radiance field parameterized by θ is bootstrapped with pure RGB losses; (iii) Both of the parameters θ, ϕ are jointly optimized with RGB-D supervision.

4.4. Pipeline Summary

We propose a 3-step optimization as shown in Fig. 4. First, since the time-pose relationship for the RGB captures is given, a time-pose function on the RGB sequence can be trained. Then, to train the neural radiance field, we bootstrap the network with pure RGB supervision. To further enable training with RGB-D supervision, we use the previously trained time-pose function and a pre-calibrated pose transformation $\mathcal{T}_{I \rightarrow D}$ to estimate the corresponding depth camera poses $\{\mathcal{T}_D^{(j)}\}$ of the depth timestamps $\{t_D^{(j)}\}$. Since both of the networks are differentiable, we jointly optimize the networks in an end-to-end manner in the third stage.

5. Experiments

In Section 5.1, we introduce how we designed the AUS dataset. In Section 5.2, we conduct qualitative and quantitative evaluations of our proposed methods and compared them with baseline methods and demonstrate the effective-

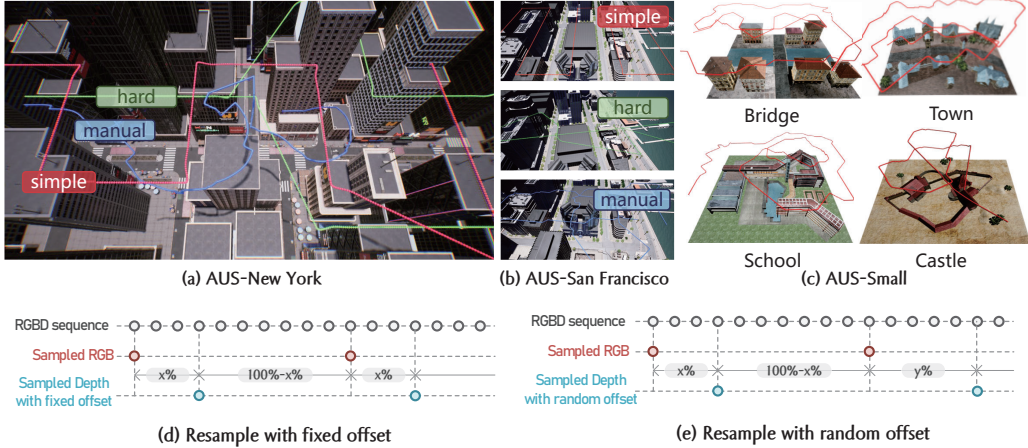


Figure 5. We propose a photo-realistically rendered dataset named Asynchronous Urban Scene (AUS) for evaluation. (a/b) are large-scale city scenes designed according to New York and San Francisco while (c) contains (relatively) small-scale scenes provided by Urban-Scene3D. Drone trajectories of different difficulty levels are visualized in (a-c). On these trajectories, we first capture an RGB-D sequence with an enough high framerate. Then we exploit two resampling strategies: fixed offset (d) and random offset (e). x equals 30 in (d) for every RGB-D pair. x equals 30 while y equals 50 in (e).

ness of our 3-step optimization pipeline. In Section 5.3, we perform some ablation experiments.

5.1. Asynchronous Urban Scene (AUS) Dataset

Dataset Collection. Our Asynchronous Urban Scene (AUS) dataset as illustrated in Fig. 5 is generated using Airsim [24], a simulator plug-in for Unreal Engine. With 3D city models loaded in Unreal Engine, the simulator can output photorealistic and high-resolution RGB images with synchronized depth images (resampled later) according to the a drone trajectory and a capture framerate. We choose Airsim as it strikes a good balance between rendering quality and dynamics modeling flexibility.

3D City Scene Models. To generate the AUS dataset, we exploit a total of six scene models, covering two large-scale ones shown in Fig. 5-a/b and four (relatively) small-scale ones shown in Fig. 5-c. The former uses the New York and San Francisco city scenes provided by Kirill Sibirakov [25], in which AUS-NewYork covers a $250 \times 150m^2$ area with many detailed buildings and AUS-SanFrancisco consists of a $500 \times 250m^2$ area near the Golden Gate Bridge. The latter uses four model files provided in the Urban-Scene3D dataset [11]. As such, at the scene level, AUS features a good coverage of both large-scale modern cities and smaller cultural heritage sites.

Trajectory Generation. Trajectory complexity matters for our problem. In many real-world applications, photographers may manually control drones to capture a city. To build a meaningful and comprehensive benchmark, we use three types of trajectories: a trivial Zig-Zag trajectory named simple, a more complex randomly generated trajectory named hard, and a most complex manually controlled trajectory named manual in Fig. 5. In AUS-Small,

we provide manually controlled trajectories, since the scene are relatively small and using the former two trajectory strategies leads to an unrealistically large overlap between frames.

Mismatch Resampling. We first sample synchronized RGB-D sequences in the simulator at a high frequency (50fps) then re-sample RGB and depth images with various offsets to create asynchronous RGB-D sequences. As shown in Fig. 5-d/e, we exploit two settings for the AUS dataset. In Fig. 5-d, every RGB-D pair is resampled according to a fixed offset denoted by percent x . For example, we sample the RGB image at 5fps or say every 10 frames and $x = 30$ means every depth image is 3 frames later than the RGB counterpart. In Fig. 5-e, the offset between an RGB-D pair is randomly selected, simulating a challenging real-world asynchronous sequence. Offset ablation will be shown later.

5.2. Comparison to State-of-the-art

We evaluate our proposed method against NeRF-W [14], city-scale Mega-NeRF [30], linear interpolation and SIREN [26]. The latter two are variants of our 1D hash-function. We present the quantitative results in Table. 2 and qualitative results in Fig. 6. NeRF-W is the baseline from which we borrow the aforementioned idea of per-image appearance embedding and Mega-NeRF is a large-scale scene modeling framework which our network is built upon. Linear interpolation and SIREN [26] are two alternative approaches that can estimate camera pose from timestamps.

RGB-D View Synthesis We present the RGB-D view synthesis results qualitatively in Fig. 6 in which our method synthesize photo-realistic images and accurate depth maps, while baseline methods fail at predicting reliable depth




Figure 6. **Qualitative Results.** Due to the large size of the dataset, we have selected a few representative samples to showcase our results. It can be seen that our method can render photo-realistic novel views (left) and the best depth estimation results (right).

Table 1. Performance of the time-pose function. For large cities, results for different trajectories are separated by slashes. For small scenes, results for different scenes are separated by slashes.

Scene	Time-Pose Function	
	Rotation ($^{\circ}$)	Translation (m)
NY Full	0.66 / 0.59 / 3.70	1.84 / 1.12 / 0.46
SF Full	0.17 / 0.67 / 0.65	1.34 / 1.45 / 0.94
Small	1.51 / 0.68 / 0.70 / 1.05	0.95 / 1.35 / 0.89 / 0.38
Mean	1.04	1.07
Scene	Joint Optimization	
NY Full	0.13 / 0.09 / 1.47	0.34 / 0.56 / 0.20
SF Full	0.05 / 0.41 / 0.02	0.32 / 1.09 / 0.66
Small	0.49 / 0.36 / 0.68 / 0.38	0.57 / 0.85 / 0.56 / 0.12
Mean	0.41 (-0.63)	0.53 (-0.54)
Scene	PoseNet [13]	
NY Full	0.08 / 0.10 / 1.91	0.18 / 0.12 / 1.50
SF Full	0.04 / 0.49 / 0.36	0.15 / 1.13 / 0.98
Small	0.40 / 0.44 / 0.60 / 0.62	0.45 / 0.77 / 0.50 / 0.02
Mean	0.50	0.64

maps (e.g., in Town, the depth values around glasses are obviously inaccurate). For quantitative comparison (Table. 2), the standard metrics for novel view synthesis and depth estimation are used for evaluation. For RGB view synthesis, image quality metrics including PSNR, SSIM [31], and the VGG implementation of LPIPS [44] are used. The Peak Signal-to-Noise Ratio (PSNR) is in inverse ratio to the mean square error. Our method achieves a higher PSNR than other methods, with an improvement of 1.11 dB over Mega-NeRF. In addition to PSNR, the Structural Similarity Index (SSIM) is designed to model image distortion as a combination of three factors: structure, luminance, and contrast. As

it can be seen, the SSIM value obtained by our method is the highest among all images, reaching 0.8206. The Learned Perceptual Image Patch Similarity (LPIPS) metric focuses more on perceptual factors that are relevant to human vision, allowing for better capturing of perceptual differences between images. The LPIPS score typically ranges from 0 to 1, where a smaller value indicates higher perceptual quality of the image. Our method has achieved the minimum LPIPS score. Therefore, the results indicate that our method achieves a higher image quality.

For depth estimation, Root Mean Square Error (RMSE), Root Mean Square Logarithmic Error (RMSE log), $\delta_{\text{thr}} = \frac{1}{MN} \sum_{i,j} (\max(\frac{f_{ij}}{g_{ij}}, \frac{g_{ij}}{f_{ij}}) < 1.25^{\text{thr}})$ are used, where f and g are $M * N$ depth images. Our method achieves a significant reduction in RMSE by 35.26 and a decrease in RMSE log by 0.3332. Additionally, it demonstrates an improvement in the $\delta_{1,2,3}$ metric, with an average enhancement around 12%.

Depth Pose Estimation We evaluate the performance of our time-pose function (Table. 1) to localize depth sensor poses. As shown quantitatively, our method can achieve an average pose error of $1.04m$ and 1.07° in the first stage. After joint optimization in the third stage, our method cuts half the errors to $0.53m$ and 0.41° . The results show that the time-pose function learns an accurate implicit trajectory from the RGB sequence that can estimate accurate poses for depth frames. By further tuning the time-pose function jointly with the scene representation network, the accuracy of the predicted depth sensor poses can be improved. Furthermore, we conduct comparative experiments with the PoseNet [13] method, and the results show that the our method performs better than PoseNet [13].

Real-world Evaluation. In the real-world experiments, we use the DJI M300 UAV equipped with a high-definition RGB camera and LiDAR to collect real data, where the RGB camera collects images at the frame rate of 30fps and the LiDAR collects depth information at 240Hz. The poses of the RGB images are provided by COLMAP [23]. The fixed transformations between sensors are provided by the producer. A qualitative comparison with Mega-NeRF is provided in Fig. 1. In terms of depth rendering, the figure shows that our method outperforms Mega-NeRF in generating better depth maps. In terms of quantitative results, our method achieves an average pose error of $2.53m$ and 2.87° , while Mega-NeRF has a higher average pose error of $3.86m$ and 4.03° .

5.3. Ablation Studies

Joint Optimization for Pose Error Compensation. To demonstrate the importance of rectifying erroneous poses of depth images in asynchronous RGB-D sequences using the time-pose function, we train a Mega-NeRF [30] with depth supervision but disabled the joint optimization stage.

Table 2. Evaluation results on the AUS dataset. For the NY and SF, we report the mean performances on all sequences (Simple / Hard / Manual). For the Bridge / Town / School / Castle scenes, we integrate their results and presented them as the "Small".

Scene	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	RMSE \downarrow	RMSE log \downarrow	$\delta_1(\%) \uparrow$	$\delta_2(\%) \uparrow$	$\delta_3(\%) \uparrow$
NY Mean	NeRF-W	23.32	0.8105	0.2249	17.40	0.2630	80.17	90.11	94.72
	Mega-NeRF	23.53	0.8375	0.1920	23.99	0.2943	80.11	88.77	92.87
	Linear Interp.	21.12	0.8034	0.3065	27.14	0.4643	64.24	83.12	88.72
	SIREN	20.78	0.8012	0.3543	32.45	0.5012	58.97	81.50	87.80
	Ours	24.33	0.8346	0.1833	6.15	0.0816	94.85	98.23	99.22
SF Mean	NeRF-W	19.21	0.6610	0.3632	24.93	0.1877	81.81	91.54	96.93
	Mega-NeRF	20.53	0.7334	0.2619	23.56	0.1713	88.58	94.74	96.83
	Linear Interp.	18.01	0.6514	0.4576	30.84	0.2501	72.38	87.38	95.32
	SIREN	18.22	0.6423	0.4462	28.49	0.2379	74.01	88.08	95.61
	Ours	22.14	0.7930	0.2620	7.64	0.0789	96.34	98.80	99.70
Small	NeRF-W	22.61	0.6855	0.3384	99.90	1.1233	59.03	67.27	68.45
	Mega-NeRF	26.58	0.8193	0.2250	92.00	1.0326	70.13	75.08	76.29
	Linear Interp.	21.34	0.6586	0.3575	128.44	1.1585	44.12	55.93	60.27
	SIREN	22.01	0.6345	0.3571	132.76	1.1298	47.94	58.34	64.83
	Ours	27.53	0.8343	0.2218	20.00	0.3382	92.79	96.56	98.18
Mean	NeRF-W	21.71	0.7190	0.3088	47.41	0.5247	73.67	82.97	86.70
	Mega-NeRF	23.55	0.7967	0.2263	46.52	0.4994	79.61	86.20	88.66
	Linear Interp.	20.16	0.7012	0.3739	62.14	0.6243	60.25	75.48	81.44
	SIREN	20.34	0.6926	0.3859	64.56	0.6230	60.31	75.97	82.75
	Ours	24.66	0.8206	0.2224	11.26	0.1662	94.66	97.86	99.03
		(+1.11)	(+0.0239)	(-0.0039)	(-35.26)	(-0.3332)	(+15.05)	(+11.66)	(+10.37)

Table 3. Ablation Study on the joint optimization stage.

Scene	Ours		w/o depth input		w/o joint optimization	
	PSNR \uparrow	RMSE \downarrow	PSNR \uparrow	RMSE \downarrow	PSNR \uparrow	RMSE \downarrow
NY Mean	24.24	5.93	24.03	42.15	19.70	15.94
SF Mean	22.70	7.26	20.00	32.17	19.07	11.39
Bridge	29.06	26.55	27.98	120.41	22.35	96.16
Town	25.32	15.61	24.69	129.50	20.14	81.99
School	26.51	21.19	25.57	63.10	21.91	42.74
Castle	28.22	16.66	28.06	54.99	23.23	38.90
Mean	26.01	15.53	25.01	73.72	21.07	47.85

Table 4. Ablation study on the Time-Pose Function.

	Rotation($^\circ$)	Translation(m)	PSNR	RMSE
RGB Init.	0.72	1.38	20.41	12.04
Ours full	0.09	0.56	24.33	6.15

From the evaluation results (Table. 3), we observe its substantial impact on the rendering quality (PSNR for RGB and RMSE for depth), which shows that jointly optimizing the time-pose function and the scene representation significantly helps reduce geometric error. We also provide the setting *without depth input*, which shows better PSNR than *w/o joint optimization* but much worse RMSE than it. This indicates that the improvement in metrics brought by the joint optimization is not due to the addition of depth image.

Time-pose function. We perform an ablation study with time-pose function (Table. 4). As a comparison, we use the pose from RGB images as the initial value of the depth camera pose and perform joint optimization in sequence NewYork-Hard. This represents a typical case of

asynchronous sensor data, where the depth image and RGB image are misaligned in both position and time. The results show that our method achieves better results by using the initial values generated by the time-pose function rather than directly using the transformed poses from the RGB sequence. The RGB poses are effective initial values when the drone motion is small, such as small rotations or slow movements. However, their performance deteriorates when the motion in the sequence becomes larger.

6. Conclusion

In this paper, we present a method to learn depth-supervised neural radiance fields from asynchronous RGB-D sequences. We leverage an important prior that the sensors cover the same spatial-temporal footprints and propose to utilize this prior with an implicit time-pose function. With a 3-staged optimization pipeline, our method calibrates the RGB-D poses and trains a large-scale implicit scene representation. Our experiments on a newly proposed large-scale dataset show that our method can effectively register depth camera poses and learns the 3D scene representation for photo-realistic novel view synthesis and accurate depth estimations. The future work includes leveraging the imagery features in the time-pose function and add more scenes into the AUS Dataset.

References

- [1] Sameer Ansari, Neal Wadhwa, Rahul Garg, and Jiawen Chen. Wireless Software Synchronization of Multiple Distributed Cameras. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–9, May 2019. [3](#)
- [2] Haowen Deng, Mai Bui, Nassir Navab, Leonidas Guibas, Slobodan Ilic, and Tolga Birdal. Deep bingham networks: Dealing with uncertainty and ambiguity in pose estimation. *International Journal of Computer Vision*, pages 1–28, 2022. [1](#)
- [3] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. [2, 3](#)
- [4] A. Elhayek, C. Stoll, N. Hasler, K. I. Kim, H.-P. Seidel, and C. Theobalt. Spatio-temporal motion tracking with unsynchronized cameras. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1870–1877, June 2012. ISSN: 1063-6919 CCF: A. [3](#)
- [5] Antoine Guedon, Pascal Monasse, and Vincent Lepetit. Scone: Surface coverage optimization in unknown environments by volumetric integration. In *Advances in Neural Information Processing Systems*, 2022. [3](#)
- [6] Antoine Guédon, Tom Monnier, Pascal Monasse, and Vincent Lepetit. Macarons: Mapping and coverage anticipation with rgb online self-supervision. *arXiv preprint arXiv:2303.03315*, 2023. [3](#)
- [7] Meng-Ru Hsieh, Yen-Liang Lin, and Winston H Hsu. Drone-based object counting by spatially regularized regional proposal network. In *Proceedings of the IEEE international conference on computer vision*, pages 4145–4153, 2017. [3](#)
- [8] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5846–5854, 2021. [1](#)
- [9] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5974–5983, 2017. [4](#)
- [10] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741–5751, 2021. [1](#)
- [11] Liqiang Lin, Yilin Liu, Yue Hu, Xingguang Yan, Ke Xie, and Hui Huang. Capturing, reconstructing, and simulating: the urbanscene3d dataset. In *ECCV*, 2022. [6](#)
- [12] Junchen Liu, Wenbo Hu, Zhuo Yang, Jianteng Chen, Guoliang Wang, Xiaoxue Chen, Yantong Cai, Huan-ang Gao, and Hao Zhao. Rip-nerf: Anti-aliasing radiance fields with ripmap-encoded platonic solids. In *SIGGRAPH'24 Conference Proceedings*, 2024. [2](#)
- [13] Qi Ma, Danda Pani Paudel, Ajad Chhatkuli, and Luc Van Gool. Continuous pose for monocular cameras in neural implicit representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5291–5301, June 2024. [7](#)
- [14] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7206–7215, June 2021. ISSN: 2575-7075. [5, 6](#)
- [15] Zhenxing Mi and Dan Xu. Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In *International Conference on Learning Representations (ICLR)*, 2023. [2](#)
- [16] 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 Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, pages 405–421, Cham, 2020. Springer International Publishing. [1, 2, 5](#)
- [17] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for uav tracking. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 445–461. Springer, 2016. [3](#)
- [18] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th international conference on computer vision*, pages 261–268. IEEE, 2009. [3](#)
- [19] Ziqiao Peng, Wentao Hu, Yue Shi, Xiangyu Zhu, Xiaomei Zhang, Hao Zhao, Jun He, Hongyan Liu, and Zhaoxin Fan. Synctalk: The devil is in the synchronization for talking head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 666–676, 2024. [2](#)
- [20] Amir M Rahimi, Raphael Ruschel, and BS Manjunath. Uav sensor fusion with latent-dynamic conditional random fields in coronal plane estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4527–4534, 2016. [3](#)
- [21] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12932–12942, 2022. [1, 2](#)
- [22] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022. [1, 2, 3](#)
- [23] Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, June 2016. [7](#)
- [24] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation

- for autonomous vehicles. In *Field and Service Robotics: Results of the 11th International Conference*, pages 621–635. Springer, 2018. 6
- [25] Kirill Sibiriyakov. Artstation page <https://www.artstation.com/vegaart>, 2022. 6
- [26] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *arXiv*, 2020. 6
- [27] Xiaowei Song, Jv Zheng, Shiran Yuan, Huan-ang Gao, Jingwei Zhao, Xiang He, Weihao Gu, and Hao Zhao. Sags: Scale-adaptive gaussian splatting for training-free anti-aliasing. *arXiv preprint arXiv:2403.19615*, 2024. 2
- [28] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 2
- [29] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains, June 2020. [arXiv:2006.10739](https://arxiv.org/abs/2006.10739) [cs]. 3
- [30] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-NeRF: Scalable Construction of Large-Scale NeRFs for Virtual Fly-Throughs. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12912–12921, June 2022. ISSN: 2575-7075. 2, 3, 4, 6, 7
- [31] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 7
- [32] 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. 1
- [33] Yuxi Wei, Zi Wang, Yifan Lu, Chenxin Xu, Changxing Liu, Hao Zhao, Siheng Chen, and Yanfeng Wang. Editable scene simulation for autonomous driving via collaborative llm-agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024. 2
- [34] Longyin Wen, Dawei Du, Pengfei Zhu, Qinghua Hu, Qilong Wang, Liefeng Bo, and Siwei Lyu. Detection, tracking, and counting meets drones in crowds: A benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7812–7821, 2021. 3
- [35] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuantao Chen, Runyi Yang, Yuxin Huang, Xiaoyu Ye, Zike Yan, Yongliang Shi, Yiyi Liao, and Hao Zhao. Mars: An instance-aware, modular and realistic simulator for autonomous driving. *CICAI*, 2023. 1
- [36] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In *The European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3
- [37] Dan Xie, Sinisa Todorovic, and Song-Chun Zhu. Inferring” dark matter” and” dark energy” from videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2224–2231, 2013. 3
- [38] Ziyang Xie, Junge Zhang, Wenye Li, Feihu Zhang, and Li Zhang. S-neRF: Neural radiance fields for street views. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [39] Linning Xu, Yuanbo Xiangli, Sida Peng, Xingang Pan, Nanxuan Zhao, Christian Theobalt, Bo Dai, and Dahua Lin. Grid-guided neural radiance fields for large urban scenes, 2023. 2
- [40] Anqi Joyce Yang, Can Cui, Ioan Andrei Barsan, Raquel Urtasun, and Shenlong Wang. Asynchronous Multi-View SLAM. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5669–5676, May 2021. ISSN: 2577-087X. 3
- [41] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [42] Shiran Yuan and Hao Zhao. Slimmerf: Slimmable radiance fields. In *2024 International Conference on 3D Vision (3DV)*, pages 64–74, 2024. 2
- [43] Kuo-Hao Zeng, Roozbeh Mottaghi, Luca Weihs, and Ali Farhadi. Visual reaction: Learning to play catch with your drone. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11573–11582, 2020. 3
- [44] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, June 2018. ISSN: 2575-7075. 7
- [45] Runze Zhang, Siyu Zhu, Tian Fang, and Long Quan. Distributed very large scale bundle adjustment by global camera consensus. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 29–38, 2017. 3
- [46] Zifan Zheng, Yezhaohui Wang, Yuxin Huang, Shichao Song, Mingchuan Yang, Bo Tang, Feiyu Xiong, and Zhiyu Li. Attention heads of large language models: A survey, 2024. 2
- [47] Qiang Zhou, Weize Li, Lihan Jiang, Guoliang Wang, Guyue Zhou, Shanghang Zhang, and Hao Zhao. Pad: A dataset and benchmark for pose-agnostic anomaly detection. *arXiv preprint arXiv:2310.07716*, 2023. 3
- [48] Zhenxin Zhu, Yuantao Chen, Zirui Wu, Chao Hou, Yongliang Shi, Chuxuan Li, Pengfei Li, Hao Zhao, and Guyue Zhou. Latitude: Robotic global localization with truncated dynamic low-pass filter in city-scale nerf. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8326–8332. IEEE, 2023. 2