

# Active Scout: Multi-Target Tracking Using Neural Radiance Fields in Dense Urban Environments

Christopher D. Hsu<sup>1,2</sup> and Pratik Chaudhari<sup>2</sup>

**Abstract**—We study pursuit-evasion games in highly occluded urban environments, e.g. tall buildings in a city, where a scout (quadrotor) tracks multiple dynamic targets on the ground. We show that we can build a neural radiance field (NeRF) representation of the city—online—using RGB and depth images from different vantage points. This representation is used to calculate the information gain to both explore unknown parts of the city and track the targets—thereby giving a completely first-principles approach to actively tracking dynamic targets. We demonstrate, using a custom-built simulator using Open Street Maps data of Philadelphia and New York City, that we can explore and locate 20 stationary targets within 300 steps. This is slower than a greedy baseline, which does not use active perception. But for dynamic targets that actively hide behind occlusions, we show that our approach maintains, at worst, a tracking error of 200m; the greedy baseline can have a tracking error as large as 600m. We observe a number of interesting properties in the scout’s policies, e.g., it switches its attention to track a different target periodically, as the quality of the NeRF representation improves over time, the scout also becomes better in terms of target tracking.

## I. INTRODUCTION

Consider a game of cops and robbers in which a quadrotor scout (cop) must search for and track robbers within a city, see Figure 1. Robbers actively avoid the scout by hiding in blind spots, or unknown parts of the environment. In this paper, we ask: what is the next best view for the scout to maximize its information of the targets’ locations? We focus on three specific aspects: (1) If the scout does not have a map of the scene, how does it explore to build one on the fly? (2) How should the scout trade off between learning about the environment and tracking targets? (3) How should targets use blind spots created by occlusions to hide from the scout?

The contributions of this work are as follows. (a) Neural radiance fields (NeRFs) can be trained online to represent the history of color and depth images observed; they should be used to synthesize future views given a sample of future poses. As NeRFs are not probabilistic models, we show that we can build an ensemble from bootstrapped data and calculate a variance over color and depth. (b) Bayes filters can represent a history of detected target locations and we show how to incorporate them into the NeRF representation. (c) The computation of mutual information provides a seamless way of integrating exploration and tracking objectives. With

a ranking of sample poses, we can select a pose and perform a dynamically feasible quadrotor trajectory to the selected waypoint. Finally, (d) we provide a policy for an active target that utilizes knowledge of scout’s history of observations and moves to locations that appear occluded to the scout.

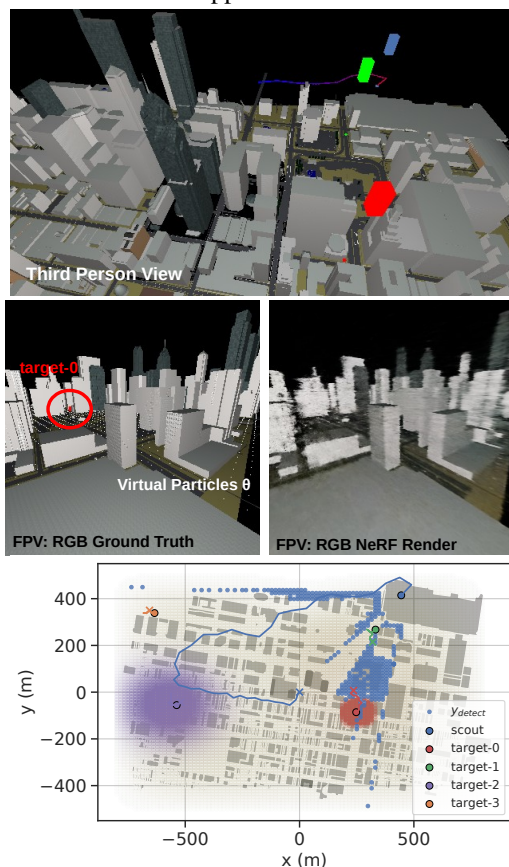


Fig. 1: We present a snapshot of the scout (blue) in a Philadelphia scene with 4 targets. **Top:** third person view of the scene with hovering agent beacons. The red-blue line denotes the scout’s trajectory history. **Middle:** On the left is the scout’s first person view (minus the labels) where it can see a red ‘target-0’ and virtually projected ground particles  $\theta$  (white dots) that help update the target filter. On the right we show the synthesized NeRF rendering of that view after training. **Bottom:** 2D map showing the building footprints (grey), the scout and its current observation of the virtual particles  $y_{detect}$  (blue scatter), and the targets. Each target has an associated filter (opacity denotes weight). In this snapshot, the red target’s filter is updated due to being observed. The purple target was viewed in the past and its uncertainty spreads over time due to the motion model. The orange and green targets have not been viewed so their prior is still rather uniform over the scene (orange+green=brown).

<sup>1</sup>DEVCOM Army Research Laboratory [christopher.d.hsu.civ@army.mil](mailto:christopher.d.hsu.civ@army.mil)

<sup>2</sup>Department of Electrical & Systems Engineering and General Robotics, Automation, Sensing and Perception (GRASP) Laboratory at the University of Pennsylvania. [chs8@seas.upenn.edu](mailto:chs8@seas.upenn.edu), [pratikac@seas.upenn.edu](mailto:pratikac@seas.upenn.edu)

Distribution Statement A. Approved for public release: distribution is unlimited.

## II. RELATED WORK

Inspired by autonomous information gathering problems [1], the prosperous line of work called active perception [2] developed the central tenet that an active perceiver should take control actions that lead to informative observations and use this data in a tight feedback loop to select the next set of controls. The problem discussed here embodies active perception and is rooted in adaptive sampling for which the goal is to choose the best option from a set of samples that minimize prediction uncertainty or maximization of some information gain [3], in comparison to offline non-adaptive where the environment is static and the plan is computed offline [4]. Whether adaptive or non-adaptive, mutual information satisfies submodularity [5] which shows that selecting sequential sensors locations in a greedy fashion has a sensing quality that is provably close to the optimal sensing quality. This property holds for static scenes and has a flavor of the multi-agent sensor problem [6]. It is unclear how this property holds for temporally dynamic scenes or processes of interest. Even so, other works have found success employing a greedy (myopic control) approach to dynamic multi-target problems [7, 8]. On the other hand, non-myopic control solutions have had success in dynamic information gathering problems [9].

In the aforementioned works, the environment for which the robots are deployed in are simplistic with simple measurements models, i.e. bearing or range measurements with noise. In the case of urban environments or occlusions, not much has been done, but works generally consider occlusions within the control optimization scheme as obstacles [10, 11] and these obstacles are often not intrusive to the tracking of the targets. There is work that considers urban environments [12] and they incorporate the occlusions in their particle filter update. However, their sensor is stationary and uses simplistic measurement models. We hypothesize that for an active perceiver in a complex environment with large occlusions (in our case due to tall buildings in a cityscape) more complex measurement and map representations are needed to test the efficacy of these works. In our past work [13], we argued that a neural radiance field (NeRF) [14] is well-suited for active perception tasks for its ability to summarize multi-modal information, e.g. photometric and geometric, in a consistent fashion and synthesize new views to be able to calculate information-based objectives such as predictive information [15]. Uncertainty quantification for next best view selection with radiance fields has also seen success on constrained scenes [16]. On larger scenes NeRFs have been productive with the caveat of well constructed adjustments such as training on progressively different scales of data while also expanding the NeRF concurrently [17], or decomposing the scene into multiple NeRFs [18]. In this work we use a fixed neural network size and online data which trades off high quality resolution reconstruction for a smaller memory footprint and speed of training. In fact, the focus of this

work is the use of NeRFs as a mode to balance exploration and exploitation in target tracking in an unknown complex environment.

However, NeRFs are only one part of the equation. Here we seek to solve the multi-target tracking problem. Past works have used all types of filters to localize targets but predominantly, kalman filters and their variants [19], particle filters [3, 7], and probability hypothesis density (PHD) filters [8, 20]. Single agent tracking with these methods is straightforward but the difficulty lies with multi-target tracking. Generally one would make multiple copies of the filter of choice to track each individual target. However, in the works that use the PHD filter, they forgo the assumption of data association, i.e. given a target the user can update the correct associated filter. In past works, this is a valid thought process to consider as range and bearing measurements of the targets lose vital information. In the context of this work, we employ a perception system, e.g. an rgbd (color and depth) camera to perform measurements. We assume that with our perceptual system we can distinguish between multiple targets and update the associated filter [21].

Finally we want to mention another popular line of work in multi-target tracking: when you have a team of agents. Past works have employed graph neural networks [22], joint or decentralized estimation over the information filter [3, 9, 20], or multi-agent reinforcement learning algorithms [23]. In this work, multi-agent teaming is not the focus. We recount the submodularity property of mutual information which states that the greedy selection of locations of subsequently added sensors is near optimal.

## III. PROBLEM FORMULATION

TABLE I: Key quantities in the text.

$\xi$	scene
$\theta_t^{(i)}$	location of the $i^{\text{th}}$ target at time $t$
$x_t \in \text{SE}(3)$	location of the scout at time $t$
$x_{\text{past}} \equiv x_{0:t}$	past locations of the scout
$x_{\text{future}} \equiv x_{t+\Delta t}$	future location of the scout
$y^{(i)} = (y_{\text{rgb}}, y_{\text{depth}}, y_{\text{detect}}^{(i)})$	RGB, depth images of the scene and detections of the $i^{\text{th}}$ target
$y_{\text{future}} \equiv y_{t+\Delta t}^{(i)}$	future observation
$\mathbf{H}(\cdot), \mathbf{I}(\cdot)$	Shannon entropy and mutual information

Table I is a summary of the the key quantities that will be introduced in the text that follows. Denote the location of the scout (a quadrotor) by  $x_t \in$  the special Euclidean group  $\text{SE}(3)$  and its dynamics by  $\dot{x} = f(x_t, u_t)$  where  $u_t$  denotes the control input; we will elaborate upon the dynamics model later. Locations of the  $m$  mobile ground targets are  $\theta^{(i)} \in \mathbb{R}^2$ . We will assume that the scout can localize itself, e.g. with GPS, i.e.,  $x_t$  is known perfectly. The scout receives RGB and depth images from the scene  $\xi$  and, when the targets are not occluded, it can detect them in the images. Let  $y_t^{(i)} = (y_{\text{rgb}}, y_{\text{depth}}, y_{\text{detect}}^{(i)})$  denote the observation received at time  $t$  corresponding to the  $i^{\text{th}}$  target; it consists of RGB

and depth images from the scout’s location and detections of the target in these images.

The scout searches and tracks the targets. We developed an active perception objective for such problems in [13] and argued that an agent performing active perception should maximize the mutual information that past observations contain about future ones. Future observations are, of course, unavailable, and therefore such an agent should have the ability to synthesize new observations, i.e., a generative model. This representation would ideally be constructed incrementally using past observations. We set

$$u_{\text{future}} \in \underset{\varphi}{\operatorname{argmax}} \mathbb{I}(y_{\text{future}}; y_{\text{past}} | x_{\text{future}}). \quad (1)$$

where  $\varphi \equiv p(x_{\text{future}} | y_{\text{past}})$  is the probability distribution of the next scout location over which we are optimizing, and it depends upon the control  $u_{\text{future}}$ . Mutual information between two random variables  $\mathbb{I}(y; x)$  is defined as

$$\begin{aligned} \mathbb{I}(y; x) &= \int dy dx p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \\ &= \mathbb{H}(y) - \mathbb{H}(y | x), \end{aligned}$$

where  $\mathbb{H}$  is the Shannon entropy. It is equal to the Kullback-Liebler (KL) divergence  $\mathbb{KL}(p(y | x), p(y))$  averaged over all possible realizations of  $x$ . In our case, the mutual information characterizes the discrepancy between the scout’s future observations given its future locations and the scout’s observations given its past locations/observations. The scout takes control actions that maximize this discrepancy, i.e., it maximizes the information gain to take control actions that provide new information about the scene and the targets. To calculate the mutual information, the scout must also be able to sample future observations  $y_{\text{future}}$  given past ones (this means, both how images from the scene  $\xi$  will look and where the targets might be detected in these images). We will discuss how to do this in the following sections.

#### IV. METHODOLOGY

We have three components to the observation  $y$ : images from the static scene  $\xi$  and detections of the dynamic targets  $\theta$ . In this section, we will discuss first how to represent the scene when it is unknown to the scout via neural radiance fields (NeRFs). Next, we will describe how we use a Bayes filter to represent and maintain an estimate of target locations. Given these two representations, we will show how to calculate the most informative next location of the scout. Roughly speaking, the quadrotor first samples future observations  $y_{\text{future}}$  from  $p(y_{\text{future}} | x_{\text{future}}, \xi, \theta_t)$  and calculates the view  $x_{\text{future}}$  that maximizes the information gain. At each step, it updates the Bayes filter and trains the NeRF using new observations  $y_t \sim p(y_t | x_t, \xi, \theta_t)$ . See the system architecture in the Figure 2 schematic. We will now focus on calculating

$$\mathbb{I}(y_{\text{future}}; y_{\text{past}}).$$

There are a few components to this: the mutual information corresponding to RGB and depth images and the mutual information corresponding to the target detections. These are

discussed in the following sections.




Fig. 2: A schematic of the Active Scout system architecture.

##### A. NeRF representation of the scene $\xi$

If the map is unknown, the scout must first build a representation of the scene such that we can synthesize observations corresponding to future states  $x_{\text{future}}$ . We utilize the recent advances in NeRFs [24, 25] to build our representation. We will represent the scene  $\xi : x \rightarrow (c, \sigma)$  as a neural radiance field which takes as input a pose  $x \equiv (R, T) \in \text{SE}(3)$  and outputs color  $c \in \mathbb{R}^3$  and density  $\sigma \in \mathbb{R}_+$ . Each location  $x$  in the NeRF is parameterized using a positional encoding scheme called a multi-resolution hash map which feeds into a multi-layered perceptron (MLP) with 2 layers and 128 neurons per layer, see [25] for more details. We train the NeRF on the fly (using a few iterations of stochastic gradient descent after each time-step) using RGB images  $y_{\text{rgb}}$  and depth images  $y_{\text{depth}}$  with ground-truth pose from the quadrotor as it flies through the city.

The power of the NeRF representation is that it can be used to synthesize images from new viewpoints that the scout might not have seen before. The volume rendering equation lies at the heart of this capability; it is also useful to understand how the NeRF is trained. Assume a pinhole model for the camera where rays emanate from the focus at  $T$ . A point in the distance  $d \in \mathbb{R}$  from the focus along this ray has orientation  $\delta R R$  which can be written as  $x(d) = \delta R R(d, 0, 0)^T + T$ . The additional rotation  $\delta R$  corresponds to all rays that lie in the field of view of the camera. The volume rendering equation samples points along this ray while querying the NeRF for color  $c(x)$  and density  $\sigma(x)$ . The transmittance  $p(d) = \exp(-\int_{d_0}^d \sigma(x(s)) ds)$  is the probability that a ray travels for an additional distance  $d$  from the image image (which is at distance  $d_0$ ) without encountering a solid object. Therefore,  $p(d)\sigma(x(d))$  is the probability that the ray stops at  $d$ . Each rendered pixel has

$$\begin{aligned} \text{color} : y_{\text{rgb}} &= \int_{d_0}^d ds p(s) \sigma(x(s)) c(x(s)), \\ \text{depth} : y_{\text{depth}} &= \int_{d_0}^d ds \sigma(x(s)) s. \end{aligned} \quad (2)$$

These integrals are implemented using quadrature [14] and techniques like [24] can be used to speed up the rendering process by skipping known free space.

We train the NeRF using images, depths, and their corresponding viewpoints collected online from a simulator that renders Open Street Maps data [26], see Figure 3. For each

image and its viewpoint, we query the MLP for the color and density at different points along the ray. The rendered color and depth of each pixel are compared to their ground-truth values to calculate the loss  $\ell = \lambda_1 \ell_{\text{rgb}} + \lambda_2 \ell_{\text{depth}}$ ; we use the  $\ell_1$  loss for both terms. Stochastic gradient descent (SGD) is used to optimize the parameters of the MLP using this objective. We tuned the hyper-parameters  $\lambda_1, \lambda_2$  such that the two terms have approximately equal magnitude during training. As the quadrotor flies through the scene, we expand the training dataset incrementally: adding new images and continuously performing SGD to update the NeRF. For each mini-batch, half the images are sampled from recent observations and the other half are sampled uniformly randomly from past observations [27].




Fig. 3: We show an example of two scenes: Philadelphia (top) and NYC StuyTown (bottom). For each scene we show the ground truth view RGB (left) compared against the NeRF rendering of RGB (middle) and depth (right). These RGB renderings have a PSNR  $\approx 23$ .

We seek to calculate the mutual information  $I(y_{\text{future,rgb}}; y_{\text{past}})$  or  $I(y_{\text{future,depth}}; y_{\text{past}})$  for RGB and depth images respectively. The scene  $\xi$  is a sufficient statistic of the past observations  $y_{\text{past}}$ . So we need to calculate  $p(\xi | y_{\text{past}})$  the probability distribution over the unknown scene and  $p(y_{\text{future,rgb}} | \xi)$  or  $p(y_{\text{future,depth}} | \xi)$ , the probability of scene given candidate future locations of the scout. In [13], we show that we can calculate a distribution over scenes using bootstrapped versions of the training dataset to build an ensemble of NeRFs that together represents  $p(\xi | y_{\text{past}})$ . We use two MLPs  $\{\xi_k\}_{k=1}^2$  to set it to be  $\delta_{\xi_1}(\xi)/2 + \delta_{\xi_2}(\xi)/2$ , where  $\delta$  denotes the Dirac delta distribution. NeRFs are not probabilistic models and therefore we cannot directly compute the likelihood of the scene. However, since the integrals of (2) are just an expectation, we can adjust them to calculate a variance for color:  $\text{var}(y_{\text{rgb}}) = \int_{d_0}^d ds p(s) \sigma(x(s)) (c(x(s)) - y_{\text{rgb}})^2$  and a similar expression for depth  $\text{var}(y_{\text{depth}})$ . If we assume that color and depth have a Gaussian distribution then we can calculate quantities like  $p(y_{\text{future,rgb}} | \xi)$ .

In unbounded scenes such as the city, some rays extend to infinity. This is problematic for our application because the

scout may resort to exploring the sky overhead to maximize the information gain, i.e., the depth infinity, so there is always a mismatch between the volume density  $\sigma$  predicted by the NeRF and the true volume density. We use a technique from [28] to resolve this issue. The authors argue that most rays in the NeRF eventually hit some solid surface. One can therefore model the occupancy  $y_{\text{occ}}$  along a ray as a Bernoulli random variable where the ray hits an obstacle with probability  $1 - p(d_{\text{max}})$  and goes off to infinity with probability  $p(d_{\text{max}})$ . For us, this is effectively an additional observation from the NeRF that depends on the volume density  $\sigma$ . We can calculate  $p(y_{\text{occ}} | x_{t+\Delta t})$  which is probability over the Bernoulli distribution and also add an additional term to the mutual information objective. This term reduces the wasteful exploration that the scout performs to gain information about the open sky.

### B. Bayes Filter to estimate target locations

We represent the probability distribution of the location of the  $i^{\text{th}}$  ground target using  $N$  particles

$$p(\theta_t^{(i)} | y_{\text{past}}) = \sum_{k=1}^N w_{t,k}^{(i)} \delta_{\tilde{\theta}_{t,k}^{(i)}}(\theta_t^{(i)})$$

where each  $\tilde{\theta}_{t,k}^{(i)} \in \mathbb{R}^2$ . In our problem, it is convenient to set up the particles densely on a fixed grid on the ground and set the weights of particles inside buildings (for a known map) to be zero. Updating the probability distribution after each time-step therefore corresponds to implementing a Bayes filter (rather than a particle filter). Note that the choice of the Bayes filter is merely for convenience; we could have also implemented a particle filter for this problem.

a) *Motion Model:* Depending upon the experimental setting, targets will either be stationary, or actively hide from the scout in the blind spots created by the buildings in the city. We compute the latter using a Dijkstra’s algorithm for the target, this is described in the next section. In




Fig. 4: Target noise distribution model.

addition to this, we assume that the scout does not know the true motion model of the targets. If the target is located at the origin, the scout assumes that the probability of it moving to a nearby location is given by the probability distribution in the adjoining figure. It is important to choose this noise distribution carefully.<sup>1</sup> The dynamics update for the Bayes filter corresponds to a convolution of the particle weights with the adjoining kernel, see Figure 4.

<sup>1</sup>Consider the situation when the scout chooses a viewpoint that has line-of-sight of the particle with a high weight. If the noise distribution were Gaussian (symmetric around the origin), if the target moved away, and the scout did not obtain a detection, the posterior  $p(\theta_t)$  would spread in all directions. In theory, this is not an issue, but in practice, such noise leads to a large variance in the posterior. This chosen noise distribution models a target that stays in the vicinity of the origin with a large probability but escapes in the direction of the four corners (as opposed to an arbitrary direction).

b) *Measurement Model*: Given a new observation  $y_{\text{detect}}$  at time  $t$  (target detections) we can update the Bayes filter as  $p(\theta | y_{\text{past}}) \propto p(\theta | y_{0:t-1}) p(y_t | \theta)$ . We assume that when the target is in the field of view of the camera, the scout can detect it with a probability 0.95, i.e., with Bernoulli noise. When the target is hidden due to occlusions caused by the buildings, the scout cannot see the target. We assume that the scout can perfectly distinguish targets from each other and therefore the probability of incorrect data association is zero.

c) *Calculating information gain for target tracking*: The posterior over the targets  $p(\theta_t | y_{\text{past}})$  is a sufficient statistic of the past observations  $y_{\text{past}}$  for detection. To calculate  $I(y_{\text{future,detect}}; y_{\text{past}})$  we need to calculate the probability distribution of the target given some candidate future location of the scout. Given the ground-truth map, the statistic  $p(\theta_t | y_{\text{past}})$ , and under the assumption that the target is stationary, it is straightforward to calculate  $p(y_{\text{future,detect}} | x_{\text{future}}, \theta_t)$ ; this is shown pictorially in Figure 1 (middle left). If we do not have the ground-truth map, we use the underlying voxel grid from the NeRF [24, 25] and the probability  $p(y_{\text{occ}} | x_{\text{future}}, \xi)$  to trace a ray from the camera to each particle of the Bayes filter. If any voxel along this ray has a NeRF volume density above a threshold, the particle is unobservable from that pose. Whether we have the map or not, like we described above, the likelihood  $p(y_{\text{future,detect}} | x_{\text{future}}, \theta_t)$  is a Bernoulli random variable with parameter 0.95. Therefore, we can calculate the information gain  $I(y_{\text{future,detect}}; y_{\text{past}})$  for each target.<sup>2</sup>

### C. Controlling the trajectories of the scout

We use the standard differentially flat [29] dynamical model of a quadrotor where we are able to recover all other parts of the state and control inputs just from the four flat outputs: 3D Euclidean position and yaw. With a flat system, we can design trajectories that satisfy initial and final boundary conditions easily, e.g. any polynomial that fits these conditions, up to control constraints, is a dynamically feasible trajectory. Additionally, we include an independent fifth state, pitch, and altogether, waypoints for the quadrotor are in 5-dimensions.

At each time-step, we sample a set of putative future states  $x_{\text{future}}$  in free space (straightforward with the ground-truth map, voxel grid underlying the NeRF is used otherwise). For each waypoint, we sample future observations  $y_{\text{future}}$  to calculate mutual information. We make a key simplifying assumption:

$$I(y_{\text{future}}; y_{\text{past}}) = I(y_{\text{future,rgb}}; y_{\text{past}}) + I(y_{\text{future,depth}}; y_{\text{past}}) + \lambda \sum_{i=1}^m I(y_{\text{future,detect}}^{(i)}; y_{\text{past}}). \quad (3)$$

<sup>2</sup>This calculation uses a simplistic dynamics model of the target (described above). This is a pragmatic choice. In principle, we could use a more complicated dynamics model, e.g., that targets hide in blind spots. But then calculating  $p(y_{\text{future,detect}} | y_{\text{past}})$  is quite difficult; it would require us to run a different update using a particle filter for each putative scout location  $x_{\text{future}}$ .

Here, the first term corresponds to the information gain for the scene (as if there were no targets, split between independent terms for RGB, depth, and occupancy) and the second term corresponds to information gain for the targets (as if the scene were known). Note that the probability density of the target locations in our Bayes filter is certainly a function of the scene (e.g., observations respect occlusions, targets cannot enter buildings, etc.). This decomposition allows us to calculate mutual information without worrying about calculating the joint distribution of the scene and targets. The hyper-parameter  $\lambda = 10$  enables the scout to trade-off between target tracking and learning the scene (which helps target tracking in the long-term even if it forgoes near-term tracking performance).

We calculate 10 different scout distributions  $\varphi \equiv p(x_{\text{future}} | y_{\text{past}})$  for (1); each of these distributions is represented by 10 particles centered around some waypoint in 3D space. Calculating the mutual information objective is an expensive calculation because it involves many different queries of the NeRF; depending upon the application one could use fewer waypoints. Instead of an argmax in (1) we experimented with a more stochastic policy where waypoints are chosen using multinomial sampling; those with larger I are more likely to be chosen. This scheme breaks the greedy formulation that can cause the scout to be stuck in local minima. Scout trajectories between successive waypoints  $x_t \rightarrow x_{t+\Delta t}$  are calculated using Dijkstra’s algorithm combined with rotortpy [30] to solve a quadratic optimization problem that parameterizes the flat outputs using a 7<sup>th</sup> order polynomial to minimize the integral of the squared snap. Since the camera pitch is independent of the quadrotor dynamics, we linearly interpolate the pitch along this trajectory. We found it helpful to perform an additional  $2\pi$  yaw rotation with some modulation of the pitch at the end of the trajectory; this gives the scout extra information to train the NeRF as well as a larger potential for spotting targets.

## V. SIMULATION EXPERIMENTS

We built a simulator using Open Street Maps [26] that can render the scene (buildings, locations of the targets etc.) using OpenGL. We focus on two specific maps for our simulation experiments. The first is of Center City Philadelphia, see Figure 1, where we set the coordinate origin to be at latitude: 39.9517 and longitude: -75.1671. The map is normalized such that each unit in each direction is roughly 1m. Similarly, our second map is of the apartment complex StuyTown in NYC, see Figure 3 (bottom), with the coordinate origin set to be at latitude: 40.7327 and longitude: -73.9771. The Center City Philadelphia map has taller buildings with irregular heights creating more intricate occlusions than the NYC map that has shorter buildings and more space in between. We bound the altitude the scout can travel such that in the Philadelphia scene it can only fly up to an altitude of 150m whereas in the NYC scene it can fly up to 100m. We do this because if the scout is able to fly to an unbounded

height, it is easy to observe the entire map from a single vantage point. From this simulator, given a pose  $\in \text{SE}(3)$  we are given RGB and depth images.

The scout has a camera with a field of view of 90 degrees and receives RGB and depth images of size  $320 \times 320$ . Each experiment begins with the scout at the origin and collecting observations by first increasing the altitude to some designated maximum; it then performs a  $2\pi$  yaw rotation with some random perturbations to the pitch to fit an initial model of the scene with 30 images, 1 image per step. We train the NeRF on these initial images for 4,000 training iterations before the experiment begins, and train for 4,000 more iterations after each waypoint is reached. Targets are randomly initialized on the ground plane in free space.

Our figure of merit is the mean squared error (RMSE) between the scout’s estimate of the target (mean of the posterior in Bayes filter) and the target’s true position. To highlight targets that are currently being observed, we plot the minimum RMSE as a high opacity plot with the color indicating which target is contributing to this value. In comparison, we plot the maximum RMSE as a low opacity plot and the corresponding target color to represent the neglected target. We also summarize these quantities in Table II.

#### A. Scout Policies

We evaluate variations of scout policies over 2 target policies (stationary and active) in 2 scenes. Each experiment will begin with the scout executing an initialization phase for which they will fly to their maximum height and do a  $2\pi$  yaw rotation scan of the scene. After that, the scout will perform 40 planning steps with each step having 30 control steps. The 4 methods below describe the scout’s policies and information strategy. We will have:

- **GTmap+MAP:** ground truth map + greedy follower (MAP: Maximum A Posteriori),
- **GTmap+MI:** ground truth map + mutual information,
- and **NeRF+MI:** NeRF + mutual information.

The following experiments will evaluate whether neural radiance fields (NeRF) are a valid replacement for a ground truth map (GTmap) given some control policy utilizing the map and target representations. In the NeRF+MI experiments, we will train a NeRF from data collected on the fly and execute mutual information (MI) based control from that representation, (3), and MI over the Bayes filter. In comparison, GTmap+MI will use the ground truth map and MI calculated over the Bayes filter, i.e.  $I = \sum_{i=1}^m \mathbb{I}(y_{\text{future,detect}}^{(i)}; y_{\text{past}})$ .

As our control baselines, we will use the ground truth map plus a greedy control policy that takes control actions that maximize a posteriori (MAP) over the Bayes filter. The MAP policy can be thought as choosing the pose that gives the maximum expected value over the detected target locations

given the future poses such that

$$u_{\text{future}} \in \underset{\varphi}{\text{argmax}} \mathbb{E}_{\theta} [y_{\text{future,detect}}^{(i)}; y_{\text{past}} | x_{\text{future}}],$$

where  $\varphi \equiv p(x_{\text{future}} | y_{\text{past}})$ .

#### B. Target Policies

For each of the scout policies as described above, we evaluate them on 2 different types of control policies: stationary targets and active targets. Stationary targets will be our exploration baseline for these methods. On the other hand, active targets are dynamic and are deliberate in the locations they select to go to as they are aware of the scout’s actions and views.

1) *Stationary Targets (exploration task):* In the first experiment, see Figure 5, we test the three methods against 20 stationary targets in the Philadelphia scene. We observe that all methods are able to localize all the targets within some time. Given a uniform prior over target locations, it can be seen that GTmap+MAP is the quickest at finding all the targets. MI based policies take longer to find all of the targets but they seem to be more thorough in exploring the map. For the observant reader, they will notice that the RMSE increases for some targets for some time. This happens because we still apply the motion model to bring about some uncertainty in the filter. We will see later that when the targets are dynamic, the greedy policy is suboptimal.

This experiment also shows that our method does not need to know how many targets are in the scene a priori since we have a camera sensor that can provide good target identification. In all our experiments, we add a new filter for each new target  $i$  that is found such that there is always 1 extra filter that has not been assigned a target. By always maintaining one extra filter, the scout always has a small opportunity to select a view that promotes exploration. When a new target is found the extra filter is assigned and a new one with a uniform prior is created. We see that in Figure 5 for each method, all targets have been identified and the filter for each target has converged to the true locations.

2) *Actively hiding targets:* Targets with the active policy have knowledge of the scout and maintain a history of graph nodes that the scout has seen during online execution, i.e they have access to  $y_{\text{detect}}$ . During each iteration, an active target will randomly select a graph node from the list of particles the scout has not seen and move there via a path found by Dijkstra’s. We observe that targets that follow this policy will actively hide behind buildings, i.e. more occluded locations, locations that are harder for the scout to see. This is because over time as the scout sees more parts of the map, the parts that have not been seen must be places of the map that the scout must be more deliberate to be able to view. We reset this observed particles buffer at regular intervals (every 10 planning steps) so that targets can return to old graph nodes. This type of reset forces the scout to return to previously seen locations for the sole reason of observing the targets.

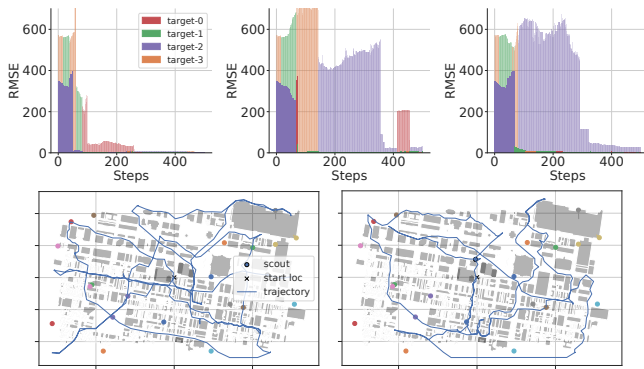


Fig. 5: **Top:** In the Philadelphia scene, we test the three methods GTmap+MAP (left), GTmap+MI (middle), and NeRF:4k+MI (right) on **20 stationary** targets randomly initialized within the scene (the legend shows an example of the first 4 targets colors) and show the RMSE plot over control steps. **Bottom:** We show the scout’s trajectory (blue line) using the GTmap+MAP (left) and NeRF:4k+MI (right) and the locations of the 20 targets in the 2D plot of Philadelphia.

In Table II, we summarize the experiments performed in this paper. We measure the root mean squared error (RMSE) in meters (m) of the targets’ true positions compared to the estimated positions. In the table we report the tracking error mean, minimum, maximum, and corresponding standard deviation which is computed by averaging the RMSE of each target over all targets, over the entire length of the experiment, and over 3 arbitrary seeds (72, 80, 88).

City	Method	Tracking Error (TE) Mean (m)	TE Min (m)	TE Max (m)
Philly	GTmap+MAP	190.668 ± 46.439	1.326 ± 0.802	603.155 ± 33.963
Philly	GTmap+MI	124.671 ± 20.088	0.583 ± 0.283	525.107 ± 92.670
Philly	NeRF:4k+MI	133.818 ± 24.080	0.893 ± 0.293	563.084 ± 61.766
Philly	NeRF:2k+MI	153.711 ± 29.144	1.125 ± 0.423	547.610 ± 90.076
Philly	offlineNeRF+MI	117.800 ± 21.931	1.013 ± 0.569	550.639 ± 16.408
NYC	GTmap+MAP	163.216 ± 18.319	0.677 ± 0.204	572.044 ± 38.461
NYC	GTmap+MI	137.121 ± 12.754	0.976 ± 0.263	550.491 ± 50.295
NYC	NeRF:4k+MI	145.057 ± 34.356	1.492 ± 1.252	586.810 ± 39.409
NYC	NeRF:2k+MI	144.363 ± 29.417	1.573 ± 0.934	594.730 ± 49.646
NYC	offlineNeRF+MI	167.537 ± 24.989	0.781 ± 0.330	570.791 ± 65.133

TABLE II: We provide a table of experiments that will be discussed in Section V-B.2 where the root mean squared error (RMSE) tracking error is averaged over the 4 targets and 3 seeds across the trajectory. We observe that NeRF based policies are similarly performant to ground truth but GTmap+MI is the most performant.

We find that GTmap+MI is the most performant and outperforms GTmap+MAP, a very greedy policy that does not do a good job switching between targets to minimize the overall tracking error. We also observe that NeRF-based policies and their variants with mutual information perform admirably against their ground truth counterparts. In this paper we provide 3 variants to the NeRF+MI policy including NeRF:4k, NeRF:2k, and an offline NeRF. NeRF:4k+MI is the standard setup for which we train the NeRF for 4,000 training steps in between each planning step. NeRF:2k+MI is similar however it is only trained for 2,000 steps for every planning step. Finally, offlineNeRF+MI is an experiment in which the NeRF is *pretrained* on the scene from images collected with the scout following mutual information for 40 planning steps

each with 4,000 training steps. When that phase is complete, the experiment begins and follows the procedure of the two previous NeRF+MI policies where it continues to train the NeRF online but for only 2,000 training steps per planning step. From the metrics in Table II, we observe that more training steps generally leads to better tracking performance. It is important to note that although the experiments are quite stochastic we still see these intuitive trends emerge. We surmise that with more runs, GTmap+MI will overall continue to be the most performant and offlineNeRF+MI will be close behind. These observations leads to our conclusion that NeRFs are a valid replacement to ground truth maps.

In Figure 6, we plot the RMSE of the scout’s estimate of the target over control steps in the Philadelphia scene (top row) and the NYC StuyTown scene (bottom row). After the initialization phase (first 30 control steps) in both scenes, there is a low minimum RMSE as at least 1 target has been seen. This is indicated by the high opacity plot and the color denotes which target contributes to that value, i.e. the red target has the smallest RMSE of the 4 and is plotted. On the other hand, the orange target has the largest RMSE and is plotted in orange with light opacity. In both scenes and all methods, by approximately step 200, all targets have been spotted at least once. From there on out the targets are more actively hiding and is the cause of error for the scout for the rest of the episode. We observe that the NeRF:4k+MI (right) experiments have similar trends to that of GTmap+MI (middle) where different targets are observed over time (varying colors of the full opacity plots) while allowing some targets to remain undetected for some time (light opacity plots). We can see with the variation in colors that over time as targets have not been observed for some time they will gain in RMSE, but upon observation that RMSE will become small and a different target will contribute to the large RMSE.




Fig. 6: In these plots the 4 targets follow the **active** policy: targets move to locations on the map that the scout has not seen. **Top:** In the **Philadelphia** scene, we compare 3 types of scout policies (left: GTmap+MAP, middle: GTmap+MI, right: NeRF:4k+MI). The plots display maximum (light opacity) and minimum (full opacity) mean squared error (RMSE) for the estimated target location against the true target position (seed 88). The color plotted shows which target is contributing to the maximum or minimum error. **Bottom:** Similar to the plots in the top row, the bottom row of plots are the 3 scout policies in the **NYC StuyTown** scene tracking active targets.

Finally, compare the MI based policies to the MAP based policies in Figure 6. In the Philadelphia scene (top), the greedy MAP agent with the ground truth map (left), does a poor job of observing the orange target, allowing its RMSE to explode. We attribute this to the scout greedily checking locations with a large posterior, i.e. places the filter indicates the target is likely to be. Since the targets move to places that are hard to find for the scout, the greedy scout does a bad job of checking hard to view areas, therefore missing hard to find targets.

## VI. CONCLUSION

In this work we study the pursuit-evasion game in which a scout (quadrotor) must track multiple active ground targets in a large urban scene. We showed that even if we do not have a map, we can train a neural radiance field (NeRF) representation of the scene online and still perform admirably against baselines that use the ground truth map. We saw that the NeRF provides a sufficient representation of past observations and that building an ensemble of NeRFs gives us way to calculate probabilistic information. Furthermore, in order to track targets we use a Bayes filter to represent target locations and we demonstrated how the NeRFs' voxel grid can be used to incorporate this filter. The efficacy of our method provides support that mutual information can methodically combine exploration and tracking objectives.

Our NeRF is trained with RGB and depth images collected from our OpenGL based simulator that renders Open Street Maps data. We found that reconstruction quality of the scene is important to tracking tasks and further improvements to building the NeRF for larger scenes should improve performance [17, 18]. Similarly, in the future we would like to relax the assumption of having ground truth depth by utilizing monocular depth estimation models such as Marigold [31].

## VII. ACKNOWLEDGEMENTS

We would like to thank Bethany Allik, Nathan Schomer, and Franklin Shedleski from ARL for the insightful conversations on this work.

## REFERENCES

- [1] A. Wald, "Sequential Tests of Statistical Hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117 – 186, 1945.
- [2] R. Bajcsy, Y. Aloimonos, and J. K. Tsotsos, "Revisiting active perception," 2016.
- [3] G. Hoffmann and C. Tomlin, "Mobile sensor network control using mutual information methods and particle filters," *Automatic Control, IEEE Transactions on*, vol. 55, pp. 32 – 47, 02 2010.
- [4] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser, "Efficient informative sensing using multiple robots," *Journal of Artificial Intelligence Research*, vol. 34, pp. 707–755, 2009.
- [5] A. Krause and C. Guestrin, "Submodularity and its applications in optimized information gathering," *ACM Trans. Intell. Syst. Technol.*, vol. 2, jul 2011.
- [6] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [7] J. R. Spletzer and C. J. Taylor, "Dynamic sensor planning and control for optimally tracking targets," *The International Journal of Robotics Research*, vol. 22, no. 1, pp. 7–20, 2003.
- [8] P. Dames, P. Tokekar, and V. Kumar, "Detecting, localizing, and tracking an unknown number of moving targets using a team of mobile robots," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1540–1553, 2017.
- [9] B. Schlotfeldt, D. Thakur, N. Atanasov, V. Kumar, and G. J. Pappas, "Anytime planning for decentralized multirobot active information gathering," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1025–1032, 2018.
- [10] K. Hausman, G. Kahn, S. Patil, J. Majjler, K. Goldberg, P. Abbeel, and G. S. Sukhatme, "Occlusion-aware multi-robot 3d tracking," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1863–1870, 2016.
- [11] F. Vanegas, J. Roberts, and F. Gonzalez, "Uav tracking of mobile target in occluded, cluttered and gps-denied environments," in *2018 IEEE Aerospace Conference*, pp. 1–7, 2018.
- [12] C. Berry and D. J. Bucci, "Obstructed target tracking in urban environments," 2019.
- [13] S. He, C. D. Hsu, D. Ong, Y. S. Shao, and P. Chaudhari, "Active perception using neural radiance fields," 2023.
- [14] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," 2020.
- [15] W. Bialek, I. Nemenman, and N. Tishby, "Predictability, complexity and learning," 2001.
- [16] W. Jiang, B. Lei, and K. Daniilidis, "Fisherrf: Active view selection and uncertainty quantification for radiance fields using fisher information," 2023.
- [17] Y. Xiangli, L. Xu, X. Pan, N. Zhao, A. Rao, C. Theobalt, B. Dai, and D. Lin, "Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering," 2023.
- [18] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. P. Srinivasan, J. T. Barron, and H. Kretschmar, "Block-nerf: Scalable large scene neural view synthesis," 2022.
- [19] B. Charrow, "Information-theoretic active perception for multi-robot teams," 01 2015.
- [20] A. Banerjee and J. Schneider, "Decentralized multi-agent active search and tracking when targets outnumber agents," 2024.
- [21] P. Tokekar, V. Isler, and A. Franchi, "Multi-target visual tracking with aerial robots," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3067–3072, 2014.
- [22] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," 2022.
- [23] C. D. Hsu, H. Jeong, G. J. Pappas, and P. Chaudhari, "Scalable reinforcement learning policies for multi-agent control," 2021.
- [24] R. Li, M. Tancik, and A. Kanazawa, "Nerfacc: A general nerf acceleration toolbox," 2023.
- [25] T. Majjler, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics*, vol. 41, p. 1–15, July 2022.
- [26] OpenStreetMap contributors, "Planet dump retrieved from <https://planet.osm.org> ." <https://www.openstreetmap.org>, 2017.
- [27] J. Yu, J. E. Low, K. Nagami, and M. Schwager, "Nerfbridge: Bringing real-time, online neural radiance field training to robotics," 2023.
- [28] N. Sajjnderhauf, J. Abou-Chakra, and D. Miller, "Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields," 2022.
- [29] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *2011 IEEE International Conference on Robotics and Automation*, pp. 2520–2525, 2011.
- [30] S. Folk, J. Paulos, and V. Kumar, "Rotorpy: A python-based multirotor simulator with aerodynamics for education and research," *arXiv preprint arXiv:2306.04485*, 2023.
- [31] B. Ke, A. Obukhov, S. Huang, N. Metzger, R. C. Daudt, and K. Schindler, "Repurposing diffusion-based image generators for monocular depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.