A Visual Active Search Framework for Geospatial Exploration

Anindya Sarkar[1]   Michael Lanier[1]   Scott Alfeld[2]   Jiarui Feng[1]   Roman Garnett[1]
Nathan Jacobs[1]   Yevgeniy Vorobeychik[1]

[1]{anindya,lanier.m,feng.jiarui,garnett,jacobsn,yvorobeychik}@wustl.edu,
[2]salfeld@amherst.edu

[1]Department of Computer Science & Engineering, Washington University in St. Louis
[2]Department of Computer Science, Amherst College

Abstract

Many problems can be viewed as forms of geospatial search aided by aerial imagery, with examples ranging from detecting poaching activity to human trafficking. We model this class of problems in a visual active search (VAS) framework, which has three key inputs: (1) an image of the entire search area, which is subdivided into regions, (2) a local search function, which determines whether a previously unseen object class is present in a given region, and (3) a fixed search budget, which limits the number of times the local search function can be evaluated. The goal is to maximize the number of objects found within the search budget. We propose a reinforcement learning approach for VAS that learns a meta-search policy from a collection of fully annotated search tasks. This meta-search policy is then used to dynamically search for a novel target-object class, leveraging the outcome of any previous queries to determine where to query next. Through extensive experiments on several large-scale satellite imagery datasets, we show that the proposed approach significantly outperforms several strong baselines. We also propose novel domain adaptation techniques that improve the policy at decision time when there is a significant domain gap with the training data. Code is publicly available at this link.

1. Introduction

Consider a large national park that hosts endangered animals, which are also in high demand on a black market, creating a major poaching problem. An important strategy in an anti-poaching portfolio is to obtain aerial imagery using drones that helps detect poaching activity, either ongoing, or in the form of traps laid on the ground [1, 2, 3, 9, 8]. The quality of the resulting photographs, however, is generally somewhat poor, making the detection problem extremely difficult. Moreover, park rangers can only inspect relatively few small regions to confirm poaching activity, doing so sequentially. Crucially, inspecting such regions

yields *new ground truth* information about poaching activity that we can use to decide which regions to inspect in the future. We can distill some key generalizable structure
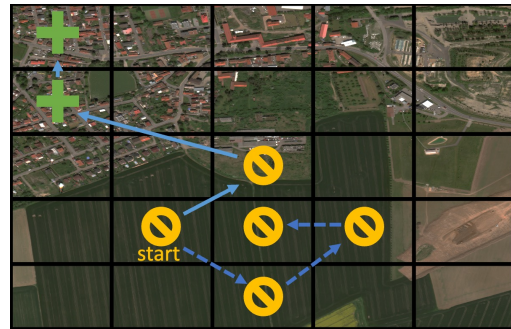

Figure 1: A comparison of a greedy search policy (dashed line) with an active search strategy (solid line) for the *small car* target class. The greedy policy is not able to adapt when a car is not found in the starting cell and needlessly searches many similar cells. The active strategy adapts and explores regions with different visual characteristics, eventually finding the objects of interest.

from this scenario: given a broad area image (often with a relatively low resolution), *sequentially query* small areas within it (e.g., by sending park rangers to the associated regions, on the ground), with each query returning the ground truth, to identify as many target objects (e.g., poachers or traps) as possible. The number of queries we can make is typically limited, for example, by budget or resource constraints. Moreover, query results (e.g., detected poaching activity in a particular region) are *highly informative about the locations of target objects in other regions*, for example, due to spatial correlation. We refer to this general modeling framework as *visual active search (VAS)*. Numerous other scenarios share this broad structure, such as identification of drug or human trafficking sites, broad area search-and-rescue, identifying landmarks, and many others.

A simple solution to the broad-area search problem is to divide the broad area into many grid cells, train a classifier

to predict existence of a target object in each grid cell, and simply explore the top K cells in terms of predicted likelihood of the object being in the cell. We call this the *greedy* policy, which essentially reduces geospatial active search to the familiar *object identification (or detection)* problem. In Figure 1, we offer some intuition about why this idea fails to capture important problem structure. Suppose we look for small cars in an image, starting in the grid marked *start*, which we initially think is the best place to look. The greedy policy being a one-shot predictor of likely grid cells containing target, continues to explore similar regions (marked as $\oslash$). What this approach ignores, as does framing the problem as traditional one-shot object identification, is the fact that both success and failure of past queries are informative due to complex spatial correlation among objects and other patterns in the scene; here, because the car was not found, we proceed to instead explore regions that have somewhat different characteristics. The key to visual active search, therefore, is to learn how to make use of the ground truth information obtained over a sequence of queries to decide where to query next. Additionally, Section 5 of [11] provides a rigorous analysis of the general sub-optimality of greedy policies in active search settings.

Relationship to Active Search and Active Learning *VAS* is closely related to *active search* [11, 10, 15, 13]. Active search is typically concerned with binary classification, and aims to maximize the number of discovered positively-labeled inputs. It uses a function f, which predicts labels of inputs x, as a means to this end, with each query serving the dual-purpose of improving f as well as identifying a positive instance. A central concern in active search, therefore, is achieving a balance of exploration (learning f) and exploitation (identifying target inputs). This consideration is also the key distinction between active search and *active learning* [25], which is concerned *solely* with improving the predictive quality of f. Thus, if we had only a single query, active learning would typically choose x for which prediction is highly uncertain, whereas active search would choose x for which $f(x)$ is most confidently positive. We empirically show that active learning is inappropriate for solving the active search problem.

However, current active search approaches typically lack a pre-search training phase, and are therefore effective in relatively low dimensions and for relatively simple model classes such as k-nearest-neighbors. In *VAS*, in contrast, our goal is to *learn how to search*, that is, to learn how to best use information obtained from previous search queries in choosing the next query. We experimentally demonstrate the advantage of *VAS* over conventional active search below.

Contributions We propose a deep reinforcement learning approach to solve the *VAS* problem. Our key contribution is a novel policy architecture which makes use of a natural representation of search state, in addition to the task im-

age input, which the policy uses to dynamically adapt to the task at hand at decision time, without additional training. Additionally, we consider a variant of *VAS* in which the nature of input tasks at decision time is sufficiently different to warrant test-time adaptation, and propose several adaptation approaches that take advantage of the *VAS* problem structure. We extensively evaluate our proposed approaches to *VAS* on two satellite imagery datasets in comparison with several baseline, including a state-of-the-art approach for a related problem of identifying regions of an image to zoom into [30]. Our results show that our approach significantly outperforms all baselines.

In summary, we make the following contributions:
- We propose *visual active search (VAS)*, a novel visual reasoning model that represents an important class of geospatial exploration problems, such as identifying poaching activities, illegal trafficking, etc.
- We propose a deep reinforcement learning approach for *VAS* that learns how to search for target objects in a broad geospatial area based on aerial imagery.
- We propose two new variants of *test-time adaptation (TTA)* variants of *VAS*: (a) *Online TTA* and (b) *Stepwise TTA*, as well as an improvement of the *FixMatch* state-of-the-art *TTA* method [26].
- We perform extensive experiments on two publicly available satellite imagery datasets, xView and DOTA, in a variety of settings, and demonstrate that proposed approaches significantly outperform all baselines.

2. Related Work

Foveated Processing of Large Images Numerous papers [33, 31, 30, 36, 32, 22, 29, 20, 19, 35] have explored the use of low-resolution imagery to guide the selection of image regions to process at high resolution, including a number using reinforcement learning to this end. Our setting is quite different, as we aim to choose a sequence of regions to query, where each query yields the *true label*, rather than a higher resolution image region, and these labels are important for both guiding further search, and as an end goal.

Reinforcement Learning for Visual Navigation Reinforcement learning has also been extensively used for visual navigation tasks, such as point and object localization [5, 18, 21, 7]. While similar at the high level, these tasks involve learning to decide on a sequence of visual navigation steps based on a local view of the environment and a kinematic model of motion, and commonly do not involve search budget constraints. In our case, in contrast, the full environment is observed initially (perhaps at low resolution), and we sequentially decide which regions to query, and are not limited to a particular kinematic model.

Active Search and Related Problem Settings Garnett et al. [11] first introduced *Active Search (AS)*. Unlike *Active Learning* [24], *AS* aims to discover members of valuable

and rare classes rather than on learning an accurate model. Garnett et al. [11] demonstrated that for any $l > m$, a l-step lookahead policy can be arbitrarily superior than an m-step one, showing that a nonmyopic active search approach can be significantly better than a myopic one-step lookahead. Jiang et al. [15, 14] proposed approaches for efficient non-myopic active search, while Jiang et al. [13] introduced consideration of search cost into the problem.

We note two crucial differences between our setting and the previous works on active search. First, we are the first to consider the problem in the context of vision, where the problem is high-dimensional, while prior techniques rely on a relatively low dimensional feature space. Second, we use reinforcement learning as a means to learn a search policy, in contrast to prior work on active search which aims to design efficient search algorithms.

3. Model

At the center of our task is an aerial image x which is partitioned into N grid cells, $x = (x^{(1)}, x^{(2)}, ..., x^{(N)})$. We can also view x as the disjoint union of these N grid cells, each of which is a sub-image. A subset (possibly empty) of these grid cells contain an instance of the target object. We formalize this by associating each grid cell j with a binary label $y^{(j)} \in \{0, 1\}$, where $y^{(j)} = 1$ iff grid cell j contains the target object. Let $y = (y^{(1)}, y^{(2)}, ..., y^{(N)})$.

We do not know y a priori, but can sequentially query to identify grid cells that contain the target object. Whenever we query a grid cell j, we obtain both the associated label $y^{(j)}$, (i.e., whether it contains the target object) *and* accrue utility if the queried cell actually contains a target object. Our ultimate goal is to find as many target objects as possible through a sequence of such queries given a total query budget constraint $\mathcal{C}$,

Formally, let $c(j, k)$ be the cost of querying grid cell k if we start in grid cell j. For the very first query, we can define a dummy initial grid cell d, so that cost function $c(d, k)$ captures the initial query cost. Let q_t denote a query performed in step t. Our ultimate goal is to solve the following optimization problem:

$$\max_{\{q_t\}} U(x; \{q_t\}) \equiv \sum_t y^{(q_t)}$$
$$\text{s.t.} : \sum_{t \geq 0} c(q_{t-1}, q_t) \leq \mathcal{C}, \tag{1}$$

where $c(q_{-1}, q_0) = c(d, q_0)$.

In order to succeed, we need to use the labels from previously queried cells to decide which cell to query next. This is a conventional setup in active search, where an important means to accomplish such learning is by introducing a model f to predict a mapping between (in our instance) a grid cell and the associated label (whether it contains a target object) [10, 15, 13]. However, in many domains of
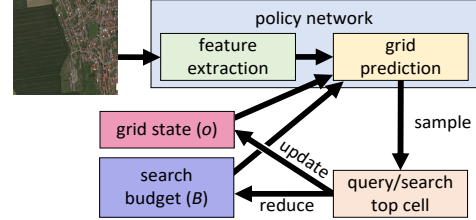

Figure 2: An overview of the *VAS* framework.

interest, such as most visual domains of the kind we consider, the query budget $\mathcal{C}$ and the number of grid cells N are very small compared to the dimension of the input x, far too small to learn a meaningful prediction f. Instead, we suppose that we have a dataset of *tasks* (aerial images) for which we have labeled whether each of the grid cells contains the target object. Let this dataset be denoted by $\mathcal{D} = \{(x_i, y_i)\}$, with each $x_i = (x_i^{(1)}, x_i^{(2)}, ..., x_i^{(N)})$ the task image and $y_i = (y_i^{(1)}, y_i^{(2)}, ..., y_i^{(N)})$ its corresponding grid cell labels. Then, at decision (or inference) time, we observe the task aerial image x, including its partition into the grid cells, and choose queries $\{q_t\}$ sequentially to maximize $U(x; \{q_t\})$.

We consider two variations of the model above. In the first, each instance (x_i, y_i) in the training data $\mathcal{D}$, as well as (x, y) at decision time (when y is unobserved before queries) are generated i.i.d. from the same distribution. In the second variation, while instances (x, y) are still i.i.d. at decision time, their distribution can be different from that of the training data $\mathcal{D}$. The latter variation falls within the broader category of *test-time adaptation (TTA)* settings, but with the special structure pertinent to our model above.

4. Solution Approach

Visual active search over the area defined by x and its constituent grid cells is a dynamic decision problem. As such, we model it as a budget-constrained episodic Markov decision process (MDP), where the search budget $\mathcal{C}$ is defined for each instance x *at decision time*. In this MDP, the actions are simply choices over which grid cell to query next; we denote the set of grids by $A = \{1, ..., N\}$. Since in our model there is never any value to query a grid cell more than once, we restrict actions available at each step to be only grids that have not yet been queried (in principle, this restriction can also be learned). Policy network inputs include: 1) the overall input x, which is crucial in providing the broad perspective on each search problem, 2) outcomes of past search queries o (we detail our representation of this presently), and 3) remaining budget $B \leq \mathcal{C}$. State transition simply updates the remaining budget and adds the outcome of the latest search query to state. Finally, an immediate reward for query a grid cell j is $R(x, o, j) = y^{(j)}$. We represent outcomes of search query history o as fol-
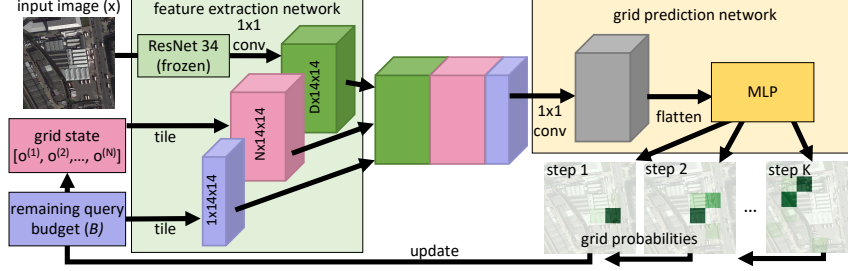
Figure 3: Our VAS policy network architecture, showing the grid probabilities at three different steps.

lows. Each element of o corresponds to a grid cell j, so that $o = (o^{(1)}, \ldots, o^{(N)})$. $o^{(j)} = 0$ if j has not been previously queried. If grid cell j has been previously queried,

$$o^{(j)} \leftarrow \begin{cases} 1, & \text{if } y^{(j)} = 1 \\ -1, & \text{if } y^{(j)} = 0. \end{cases} \qquad (2)$$

Armed with this MDP problem representation, we next describe our proposed deep reinforcement learning approach for learning a search policy that makes use of a dataset $\mathcal{D}$ of past search tasks. Specifically, we use the *REINFORCE* policy gradient algorithm [28] to directly learn a search policy $\psi(x, o, B; \theta)$ where θ are the parameters of the policy that we learn. Specifically, we maximize the following objective function:

$$\nabla J(\theta) = \sum_{i=1}^{M} \sum_{t=1}^{T_i} \mathbb{1}_{\sum_{t \geq 0} c(q_{t-1}, q_t) \leq \mathcal{C}} \nabla \log \psi_\theta(a_t^i | x_i, o_t^i, B_t^i) \mathbb{R}_t^i \qquad (3)$$

Where M is the number of example search task seen during training and $\mathbb{R}_t$ is the discounted cumulative reward defined as $\mathbb{R}_t = \sum_{k=t}^{T} \gamma^{k-t} R_k$ with a discount factor $\gamma \in [0, 1]$.

The output of the search policy ψ is a probability distribution over A, with $\psi_j(x, o, B; \theta)$ the probability that grid cell $j \in A$ is selected by the policy ψ.

In general, ψ will output a positive probability over all possible grid cells $j \in A$. However, in our setting there is no benefit to querying any grid cell $j \in A$ that has previously been queried, i.e., for which $o^{(j)} \neq 0$. Consequently, both at training (when the next decision is generated) and decision time, we restrict consideration only to $j \in A$ with $o^{(j)} = 0$, that is, which have yet to be queried, and simply renormalize the output probabilities of ψ. Formally, we define $\psi'_j(x, o, B; \theta) = 0$ for j with $o^{(j)} \neq 0$, and define

$$\psi'_j(x, o, B; \theta) = \frac{\psi_j(x, o, B; \theta)}{\sum_{k \in A: o^{(k)} = 0} \psi_k(x, o, B; \theta)}.$$

Grid cells j are then samples from ψ'_j at each search step during training. At decision time, on the other hand, we choose the grid cell j with the maximum value of ψ'_j. This approach allows us to simply train the policy network

ψ without concern about feasibility of particular grid cell choices at decision time. In addition, to ensure that the policy is robust to search budget uncertainty, we use randomly generated budgets $\mathcal{C}$ at training time for different task instances. In the case of query costs $c(j, k) = 1$ for all grid cells j, k, each episode has a fixed length $\mathcal{C}$. In general, episodes have no fixed length, and end whenever we exhaust the total cost budget $\mathcal{C}$. The overview of our proposed *VAS* framework is depicted in Figure 2.

Next, we detail the proposed policy network architecture, and subsequently describe an adaptation of our approach when instances at decision time follow a different distribution from those in the training data $\mathcal{D}$, that is, the test-time adaptation (TTA) setting.

4.1. Policy Network Architecture

As shown in Figure 3, the policy network $\psi(x, o, B; \theta)$ is composed of two components: 1) the image feature extraction component $f(x; \phi)$ which maps the aerial image x to a low-dimensional latent feature representation z, and 2) the grid selection component $g(z, o, B; \zeta)$, which combines the latent image representation z with outcome of past search queries o and remaining budget B to produce a probability distribution over grid cells to search in the next time step. Thus, the joint parameters of ψ are $\theta = (\phi, \zeta)$.

We use a frozen ResNet-34 [12], pretrained on ImageNet [16], as the feature extraction component f, followed by a 1×1 convolution layer. We combine this with the budget B and past query information o as follows. We apply the tiling operation in order to convert o into a representation with the same dimensions as the extracted features $z = f(x)$, aiding us to effectively combine latent image feature and auxiliary state feature while preserving the grid specific spatial and query related information. Similarly, we apply tiling to the scalar budget B to transform it to match the size of z and the tiled version of o. Finally, we concatenate the features (z, o, B) along the channels dimension and pass them through the grid prediction network g. This consists of 1×1 convolution to reduce dimensionality, flattening, a small MLP with ReLU activations, and a final output (softmax) that represents the current grid probability. This

yields the full policy network to be trained end to end via REINFORCE: $\psi(x, o, B; \theta) = g(f(x; \phi), o, B; \zeta)$.

4.2. Test-Time Adaptation

A central issue in our model, as in traditional active search, is that tasks faced at decision time may in some respects be novel, unlike tasks faced previously (e.g., represented in the dataset $\mathcal{D}$). We view this issue through the lens of *test-time adaptation (TTA)*, in which predictions are made on data that comes from a different distribution from training data. While a myriad of TTA techniques have been developed, they have focused almost exclusively on *supervised learning* problems, rather than active search settings of the kind we study. Nevertheless, two common techniques can be either directly applied, or adapted, to our setting: 1) *Test-Time Training (TTT)* [27] and 2) *FixMatch* [26].

TTT makes use of a self-supervised objective at both training and prediction time by adding a self-supervised head r as a component of the policy model. The associated self-supervised loss (which is added during training) is a quadratic reconstruction loss $\|x - r(z; \eta)\|$, where $z = f(x; \phi)$ is the latent embedding of the input aerial image x and η the parameters of r. At decision time, a new task image x is used to update policy parameters using just the reconstruction loss before we begin the search. Adaptation of *TTT* to our VAS domain is therefore direct.

The original variant of *FixMatch* uses pseudo-labels at decision time, which are predictions on weakly augmented variants of the input image x (keeping only those which are highly confident), to update model parameters. In our domain, however, we can leverage the fact that we obtain *actual* labels whenever we query regions of the image. We make use of this additional information as follows. Whenever a query j is successful (i.e., $y^{(j)} = 1$), we construct a label vector as the one-hot vector with a 1 in the location of the successful grid cell j. However if $y^{(j)} = 0$, we associate each queried grid cell with a 0, and assign uniform probability distribution over all unqueried grids. We then update model parameters using a cross-entropy loss.

Even as we adapted them, *TTT* and *FixMatch* do not fully take advantage of the rich information obtained at decision time in the *VAS* context as we proceed through each input task: we not only observe the input image x, but also observe query results over time during the search. We therefore propose two new variants of *TTA* which are specific to the *VAS* setting: (a) *Online TTA* and (b) *Stepwise TTA*. In *Online TTA*, we update parameters of the policy network after each task is completed during decision time, which yields for us both the input x and the *observations* o of the search results, which only partially correspond to y, since we have only observed the contents of the previously queried grid cells. Nevertheless, we can simply use this partial information o as a part of the *REINFORCE* pol-

icy gradient update step to update the policy parameters θ. In *Stepwise TTA*, we update the policy network parameters, even during the execution of a particular task, at decision time, once every $m < \mathcal{C}$ steps. The main difference between *Online* and *Stepwise* variations of our *TTA* approaches is consequently the frequency of updates. Note that we can readily compose both of these *TTA* approaches with conventional *TTA* methods, such as *TTT* and *FixMatch*.

5. Experiments

Evaluation Metric We evaluate the proposed approaches in terms of the average number of target objects discovered (we shorten it to **ANT**).

Baselines We compare the proposed VAS policy learning framework with the following baselines:
1. *random search*, where each grid is chosen uniformly at random among those which haven't been explored,
2. *greedy classification*, in which we train a classifier ψ_{gc} to predict whether a particular grid has a target object and search the grids most likely to contain the target until the search budget is exhausted, and
3. *greedy selection*, based on the approach by Uzkent and Ermon [30] which trains a policy ψ_{gs} which yields a probability of zooming into each grid cell j. We select grids according to ψ_{gs} until the budget $\mathcal{C}$ is saturated.
4. *active learning*, in which we randomly select the first grid to query and then choose $\mathcal{C} - 1$ grids using a state-of-the-art active learning approach by Yoo et al. [37].
5. *conventional active search*, an active search method by Jiang et al. [15], using a low-dimensional feature representation for each image grid from the same feature extraction network as in our approach.

Query Costs We consider two ways of generating query costs: (i) $c(i, j) = 1$ for all i, j, where $\mathcal{C}$ is just the number of queries, and (ii) $c(i, j)$ is based on Manhattan distance between i and j. Most of the results we present reflect the second setting; the results for uniform query costs are qualitatively similar and provided in the Supplement.

Datasets We evaluate the proposed approach using two datasets: xView [17] and DOTA [34]. xView is a satellite imagery dataset which consists of large satellite images representing 60 categories, with approximately 3000 pixels in each dimensions. We use 67% and 33% of the large satellite images to train and test the policy network respectively. DOTA is also a satellite imagery dataset. We re-scale the original $\sim 3000 \times 3000$px images to 1200×1200px. Unless otherwise specified, we use $N = 36$ non-overlapping pixel grids each of size 200×200.

5.1. Results on the xView Dataset

We begin by evaluating the proposed approaches on the xView dataset, varying search budgets $\mathcal{C} \in \{25, 50, 75\}$ and

number of grid cells $N \in \{30, 48, 99\}$. We consider two target classes: *small car* and *building*. As the dataset contains variable size images, take random crops of 2500×3000 for $N = 30$, 2400×3200 pixels for $N = 48$, and 2700×3300 for $N = 99$, thereby ensuring equal grid cell sizes.

Table 1: **ANT** comparisons for the *small car* target class on xView.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
random search ($N = 30$)	3.41	3.95	4.52
greedy classification ($N = 30$)	3.91	4.60	4.76
greedy selection [30] ($N = 30$)	3.90	4.63	4.78
active learning [37] ($N = 30$)	3.92	4.58	4.73
conventional active search [15] ($N = 30$)	3.61	4.17	4.70
VAS ($N = 30$)	**4.61**	**7.49**	**9.88**
random search ($N = 48$)	3.20	3.66	4.11
greedy classification ($N = 48$)	3.87	4.29	4.52
greedy selection [30] ($N = 48$)	3.89	4.42	4.53
active learning [37] ($N = 48$)	3.87	4.28	4.51
conventional active search [15] ($N = 48$)	3.26	3.74	4.32
VAS ($N = 48$)	**4.56**	**7.45**	**9.63**
random search ($N = 99$)	1.10	2.15	2.96
greedy classification ($N = 99$)	1.72	2.79	3.36
greedy selection [30] ($N = 99$)	1.78	2.83	3.41
active learning [37] ($N = 99$)	1.69	2.78	3.33
conventional active search [15] ($N = 99$)	1.42	2.31	3.10
VAS ($N = 99$)	**2.72**	**4.42**	**5.78**

Table 2: **ANT** comparisons for the *building* target class on xView.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
random search ($N = 30$)	3.97	4.94	5.39
greedy classification ($N = 30$)	4.69	5.27	5.80
greedy selection [30] ($N = 30$)	4.84	5.33	5.82
active learning [37] ($N = 30$)	4.67	5.24	5.80
conventional active search [15] ($N = 30$)	4.15	5.20	5.51
VAS ($N = 30$)	**5.65**	**9.31**	**12.20**
random search ($N = 48$)	3.47	3.96	4.26
greedy classification ($N = 48$)	3.90	4.43	4.61
greedy selection [30] ($N = 48$)	3.95	4.51	4.67
active learning [37] ($N = 48$)	3.88	4.43	4.60
conventional active search [15] ($N = 48$)	3.70	4.11	4.38
VAS ($N = 48$)	**5.61**	**9.26**	**12.15**
random search ($N = 99$)	1.55	2.99	4.18
greedy classification ($N = 99$)	2.17	3.96	4.84
greedy selection [30] ($N = 99$)	2.29	4.21	5.22
active learning [37] ($N = 99$)	2.17	3.95	4.82
conventional active search [15] ($N = 99$)	1.68	3.10	4.33
VAS ($N = 99$)	**4.29**	**6.91**	**8.98**

The results are presented in Table 1 for the *small car* class and in Table 2 for the *building* class. We see substantial improvements in performance of the proposed *VAS* approach compared to all baselines, ranging from 15–260% improvement relative to the most competitive state-of-the-art approach, *greedy selection*. There are two general consistent trends. First, as the number of grids N increases compared to $\mathcal{C}$ (corresponding to sets of rows in either table), performance of all methods declines, as the task becomes more challenging. However, the decline in perfor-

mance is typically much greater for our baselines than for *VAS*. Second, overall performance improves as $\mathcal{C}$ increases (columns in both tables), and the relative advantage of *VAS* increases, as it is better able to take advantage of the greater budget than the baselines.

Figure 4: Comparison of policies learned using *VAS* (left) and the *greedy selection* baseline method (right).

In Figure 5 we visually illustrate *VAS* search strategy in comparison with the *greedy selection* baseline (the best performing baseline). The plus signs correspond to successful queries, $\oslash$ to unsuccessful queries, and arrows represent query order. This shows that *VAS* quickly learns to take advantage of the visual similarities between grids (after the first several failed queries, the rest are successful), whereas our most competitive baseline—*greedy selection*—fails to take advantage of such information. During the initial search phase, the VAS policy explores different types of grids before exploiting grids it believes to have target objects.

Finally, we perform an ablation study to understand the added value of including remaining budget B as an input in the *VAS* policy network. To this end, we modify the combined feature representation of size $(2N + 1) \times 14 \times 14$, consisting of input and auxiliary state features, each of size $N \times 14 \times 14$, and a single channel of size 14×14 containing the information of remaining search budget, as depicted in Figure 3. We only eliminate the channel from the combined feature representation that contains the information about the number of queries left, resulting in $2N \times 14 \times 14$ size feature map. The resulting policy network is then trained just as the original *VAS* architecture.

Table 3: Comparative **ANT** performance of *VAS without remaining search budget* and *VAS* using *small car* as the target class.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
VAS w/o remaining search budget ($N = 30$)	4.47	7.38	9.62
VAS ($N = 30$)	**4.61**	**7.49**	**9.88**
VAS w/o remaining search budget ($N = 48$)	4.34	7.31	9.49
VAS ($N = 48$)	**4.56**	**7.45**	**9.63**
VAS w/o remaining search budget ($N = 99$)	2.63	4.29	5.69
VAS ($N = 99$)	**2.72**	**4.42**	**5.78**

We compare the performance of the policy without remaining search budget (referred to as *VAS without remaining search budget*) with *VAS* in Table 3. Across all problem sizes and search budgets, we observe a relatively small but

consistent improvement (∼ 1–3%) from using the remaining search budget B as an explicit input to the policy network.

5.2. Results on the DOTA Dataset

Next, we repeat our experiments on the DOTA dataset. We use *large vehicle* and *ship* as our target classes. In both cases, we also report results with non-overlapping pixel grids of size 200×200 and 150×150 ($N = 36$ and $N = 64$, respectively). We again use $\mathcal{C} \in \{25, 50, 75\}$.

Table 4: **ANT** comparisons for the *large vehicle* target class on DOTA.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
random search ($N = 36$)	1.79	3.50	5.10
greedy classification ($N = 36$)	2.64	4.07	5.88
greedy selection [30] ($N = 36$)	2.82	4.21	5.97
active learning [37] ($N = 36$)	2.63	4.06	5.84
conventional active search [15] ($N = 36$)	1.92	3.63	5.34
VAS ($N = 36$)	**4.63**	**6.79**	**8.07**
random search ($N = 64$)	1.48	2.96	3.91
greedy classification ($N = 64$)	2.59	3.77	5.48
greedy selection [30] ($N = 64$)	2.72	4.10	5.77
active learning [37] ($N = 64$)	2.57	3.74	5.47
conventional active search [15] ($N = 64$)	1.64	3.15	4.23
VAS ($N = 64$)	**5.33**	**8.47**	**10.51**

Table 5: **ANT** comparisons for the *ship* target class on the DOTA dataset.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
random search ($N = 36$)	1.73	3.07	4.26
greedy classification ($N = 36$)	2.04	3.65	4.92
greedy selection [30] ($N = 36$)	2.33	3.84	5.01
active learning [37] ($N = 36$)	2.01	3.64	4.91
conventional active search [15] ($N = 36$)	1.86	3.25	4.40
VAS ($N = 36$)	**3.31**	**5.34**	**6.74**
random search ($N = 64$)	1.26	2.33	3.14
greedy classification ($N = 64$)	1.89	3.06	3.75
greedy selection [30] ($N = 64$)	2.07	3.32	4.02
active learning [37] ($N = 64$)	1.87	3.05	3.72
conventional active search [15] ($N = 64$)	1.41	2.48	3.38
VAS ($N = 64$)	**3.58**	**6.38**	**7.83**

The results are presented in Tables 4 and 5, and are broadly consistent with our observations on the xView dataset, with *VAS* outperforming all baselines by ∼ 40 – 80%, with the greatest improvement typically coming with a higher search budget $\mathcal{C}$.

5.3. Visualization of VAS Strategy

In figure 5, we demonstrate the sequential behavior of a pretrained VAS policy during inference. We have shaded the grid such that darker indicates higher probability. The darkest grid at each step is the target to be revealed. In the first row we have VAS searching for *large vehicles*. In step 1, VAS looks at the roof of a building, which looks very similar to a *large vehicle* in the overhead view. Next in step 3, it searches a grid with a large air conditioner which also looks similar to a *large vehicle*. Having viewed these two

confusers, VAS now learns that the rest of the grids with building roofs likely contain no *large vehicles*. It is important to note that this eliminates a large portion of the middle of the image from consideration as it is entirely roof tops. In step 5, it moves to an area which is completely different, a road where it finds a *large vehicle*. VAS now aggressively begins examining grids with roads. In steps 7 through 13 it searches roads discovering *large vehicle*. Finally in step 15 it explores to a parking lot containing a *large vehicle*. In our middle example we have VAS targeting *small cars*. In step 1, VAS targets a road and fails to find a car. In step 3, it searches another road in a different region and finds a car. Having explored regions with prominent major roads it moves to a parking lot in step 5 and finds a car. It now searches a similar parking lot in step 7. Having explored grids with parking lots it goes back to searching minor roads for the duration of its search. VAS does not visit a parking lot in the north east corner, but this parking lot is visually much different from the other two (i.e. it's not rectangular). In our bottom example we have VAS searching for *ships*. In step 1, VAS searches near a harbor. Having found a *ship* it begins exploring similar harbor regions. In step 3 and 5 it searches other parts of the same harbor finding *ships*. In steps 7-9, it searches areas similar to the harbor without *ships*. VAS now learns that *ships* are not likely present in the rest of the dock and explores different regions leaving the rest of the dock unexplored. These three examples demonstrate VAS's tendency for the explore-exploit behavior typical of reinforcement learning algorithms. Additionally, we note that VAS has an ability to eliminate large areas that would otherwise confuse standard greedy approaches.

5.4. Efficacy of Test-Time Adaptation

One of the important features of the visual active search problem is that queries actually allow us to observe partial information about target labels *at inference time*. Here, we evaluate how our approaches to *TTA* that take advantage of this information perform compared to the baseline *VAS without TTA*, as well as state-of-the-art *TTA* baselines discussed in Section 4.2 (where *FixMatch* is adapted to also take advantage of observed labels).

Consider first the case where there is no difference between training and test distribution over classes. As before we consider xView and DOTA for analysis. The results are presented in Figure 6, and show a consistent pattern. The *TTT* approach performs the worst, followed by (out adaptation of) *FixMatch*, which is only slightly better than *TTT*. *Stepwise TTA* outperforms both *TTT* and *FixMatch*, albeit slightly, and *Online TTA* is, somewhat surprisingly, much better than all others (this is surprising since it has a lower frequency of model update compared to *Stepwise TTA*).

Finally, we consider a *TTA* setting in which the domain exhibits a non-trivial distributional shift at inference time.
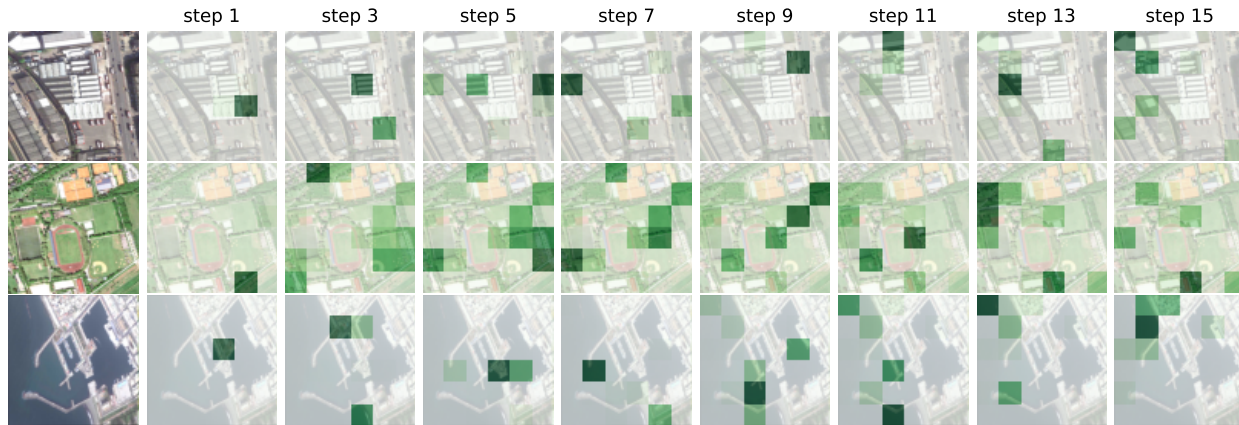
Figure 5: Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS for different target types.
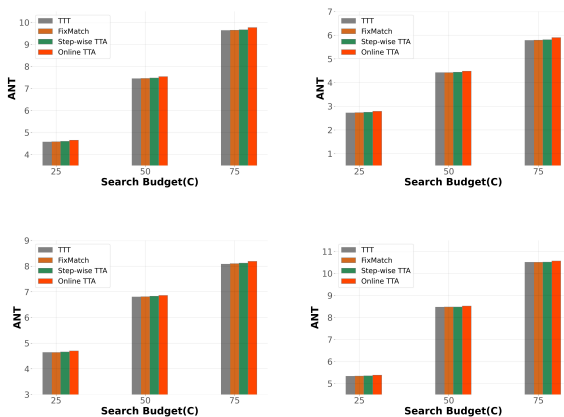

Figure 6: Comparative results of TTA methods on VAS framework. xView (top; *small car* target): (left) $N = 48$, (right) $N = 99$. DOTA (bottom; *large vehicle* target): (left) $N = 36$, (right) $N = 64$.

Table 6: Comparative results on xView dataset with *small car* and *Building* as the target class during training and inference respectively.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
without TTA ($N = 30$)	5.28	8.58	11.42
TTT [27] ($N = 30$)	5.30	8.61	11.45
FixMatch [26] ($N = 30$)	5.31	8.62	11.47
Stepwise TTA ($N = 30$)	5.33	8.64	11.50
Online TTA ($N = 30$)	**5.42**	**8.69**	**11.58**

Table 7: Comparative results on DOTA dataset with *large vehicle* and *ship* as the target class during training and inference respectively.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
without TTA ($N = 36$)	2.69	4.38	5.84
TTT [27] ($N = 36$)	2.70	4.39	5.84
FixMatch [26] ($N = 36$)	2.70	4.39	5.84
Stepwise TTA ($N = 36$)	2.71	4.40	5.85
Online TTA ($N = 36$)	**2.73**	**4.42**	**5.98**

6. Conclusion

Our results show that *VAS* is an effective framework for geospatial broad area search. Notably, by applying simple *TTA* techniques, the performance of *VAS* can be further improved at test time in a way that is robust to target class shift. The proposed *VAS* framework also suggests a myriad of future directions. For example, it may be useful to develop more effective approaches for learning to search *within a task*, as is common in past active search work. Additionally, the search process may often involve additional constraints, such as constraints on the sequence of regions to query. Moreover, it's natural to generalize query outcomes to be non-binary (e.g., returning the number of target object instances in a region).

In this case, we would expect the conventional *TTT* and *FixMatch* methods to be more competitive, as they have been specifically designed to account for distribution shift. We model distribution shift by training the search policy using one target object, and then applying it in the decision context for another target object. Specifically, for xView, we use *small car* as the target class during training, and *building* as the target class at test time. Similarly, on the DOTA dataset we use *large vehicle* as the target class at training time, and use *ship* as the target at test time.

The results for the *TTA* setting with distribution shift are presented in Table 6 and 7 for the xView and the DOTA dataset respectively, where we also add a comparison to the *VAS* without *TTA* of any kind. We observe that the results here remain consistent, with the proposed *Online TTA* outperforming the other approaches, with *Stepwise TTA* yielding the second-best performance.

References

[1] Elizabeth Bondi, Debadeepta Dey, Ashish Kapoor, Jim Piavis, Shital Shah, Fei Fang, Bistra Dilkina, Robert Hannaford, Arvind Iyer, Lucas Joppa, et al. Airsim-w: A simulation environment for wildlife conservation with uavs. In *ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 1–12, 2018.

[2] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, et al. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. In *AAAI Conference on Artificial Intelligence*, 2018.

[3] Elizabeth Bondi, Raghav Jain, Palash Aggrawal, Saket Anand, Robert Hannaford, Ashish Kapoor, Jim Piavis, Shital Shah, Lucas Joppa, Bistra Dilkina, et al. Birdsai: A dataset for detection and tracking in aerial thermal infrared videos. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1747–1756, 2020.

[4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.

[5] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. *Advances in Neural Information Processing Systems*, 33:4247–4258, 2020.

[6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[7] Kshitij Dwivedi, Gemma Roig, Aniruddha Kembhavi, and Roozbeh Mottaghi. What do navigation agents learn about their environment? In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10276–10285, 2022.

[8] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying paws: Field optimization of the protection assistant for wildlife security. In *AAAI Conference on Artificial Intelligence*, volume 16, pages 3966–3973, 2016.

[9] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *International Joint Conference on Artificial Intelligence*, 2015.

[10] Roman Garnett, Thomas Gärtner, Martin Vogt, and Jürgen Bajorath. Introducing the 'active search' method for iterative virtual screening. *Journal of Computer-Aided Molecular Design*, 29(4):305–314, 2015.

[11] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. *arXiv preprint arXiv:1206.6406*, 2012.

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[13] Shali Jiang, Roman Garnett, and Benjamin Moseley. Cost effective active search. *Advances in Neural Information Processing Systems*, 32, 2019.

[14] Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic batch active search. *Advances in Neural Information Processing Systems*, 31, 2018.

[15] Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic active search. In *International Conference on Machine Learning*, pages 1714–1723. PMLR, 2017.

[16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

[17] Darius Lam, Richard Kuzma, Kevin McGee, Samuel Dooley, Michael Laielli, Matthew Klaric, Yaroslav Bulatov, and Brendan McCord. xview: Objects in context in overhead imagery. *arXiv preprint arXiv:1802.07856*, 2018.

[18] Bar Mayo, Tamir Hazan, and Ayellet Tal. Visual navigation with spatial attention. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16898–16907, June 2021.

[19] Chenlin Meng, Enci Liu, Willie Neiswanger, Jiaming Song, Marshall Burke, David Lobell, and Stefano Ermon. Iscount: Large-scale object counting from satellite images with covariate-based importance sampling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 12034–12042, 2022.

[20] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12309–12318, 2022.

[21] Mahdi Kazemi Moghaddam, Ehsan Abbasnejad, Qi Wu, Javen Qinfeng Shi, and Anton Van Den Hengel. Foresi: Success-aware visual navigation agent. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 691–700, January 2022.

[22] Athanasios Papadopoulos, Pawel Korus, and Nasir Memon. Hard-attention for scalable image classification. In *Advances in Neural Information Processing Systems*, volume 34, pages 14694–14707, 2021.

[23] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[24] Burr Settles. Active learning literature survey. 2009.

[25] Burr Settles. Active learning. *Synthesis lectures on artificial intelligence and machine learning*, 6(1):1–114, 2012.

[26] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020.

[27] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020.

[28] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

[29] Chittesh Thavamani, Mengtian Li, Nicolas Cebron, and Deva Ramanan. Fovea: Foveated image magnification for autonomous navigation. In *IEEE/CVF International Conference on Computer Vision*, pages 15539–15548, 2021.

[30] Burak Uzkent and Stefano Ermon. Learning when and where to zoom with deep reinforcement learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12345–12354, 2020.

[31] Burak Uzkent, Christopher Yeh, and Stefano Ermon. Efficient object detection in large images using deep reinforcement learning. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020.

[32] Yi Wang, Youlong Yang, and Xi Zhao. Object detection using clustering algorithm adaptive searching regions in aerial images. In *European Conference on Computer Vision*, pages 651–664, 2020.

[33] Zuxuan Wu, Caiming Xiong, Yu-Gang Jiang, and Larry S Davis. Liteeval: A coarse-to-fine framework for resource efficient video recognition. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[34] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018.

[35] Chenhongyi Yang, Zehao Huang, and Naiyan Wang. Querydet: Cascaded sparse query for accelerating high-resolution small object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13668–13677, June 2022.

[36] Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, and Gao Huang. Resolution adaptive networks for efficient inference. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[37] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 93–102, 2019.

APPENDIX: A Visual Active Search Framework for Geospatial Exploration

In this appendix, we provide details that could not be included in the main paper owing to space constraints, including: (A) Performance of VAS under uniform query cost; (B) VAS Pseudocode; (C) Policy architecture and training hyperparameter; (D) Search Performance Comparison with Different Feature Extractor Module; (E) More Visual Illustration of VAS and the Most Competitive Greedy Selection baseline Method; (F) Assessment of VAS and other Baseline Methods with a Different Evaluation Metric; (G) Search Performance Comparison with Other Policy Learning Algorithm (PPO); (H) Sensitivity Analysis of *VAS*; (I) Efficacy of TTA on Search Tasks involving Large Number of Grids; (J) Saliency map visualization of *VAS*;

A. Performance of VAS under Uniform Query Cost

In this section, we report the performance of VAS under uniform query cost. The results are presented in the following Table 8 for the *small car* target class and in Table 9 for the *building* class from the xView dataset. We observe significant improvements in performance of the proposed VAS approach compared to all baselines, ranging from $11-25\%$ improvement relative to the most competitive *greedy selection* approach.

Table 8: **ANT** comparisons for the *small car* target class.

Method	$\mathcal{C} = 12$	$\mathcal{C} = 15$	$\mathcal{C} = 18$
random search ($N = 30$)	4.57	5.66	6.85
greedy classification ($N = 30$)	5.31	6.24	7.25
greedy selection [30] ($N = 30$)	5.47	6.45	7.46
active learning [37] ($N = 30$)	5.28	6.21	7.22
conventional AS [15] ($N = 30$)	4.86	5.97	6.92
VAS ($N = 30$)	**6.03**	**7.24**	**8.24**
random search ($N = 48$)	3.80	4.97	5.98
greedy classification ($N = 48$)	4.69	5.48	6.79
greedy selection [30] ($N = 48$)	4.92	5.81	6.98
active learning [37] ($N = 48$)	4.68	5.46	6.78
conventional AS [15] ($N = 48$)	3.96	5.45	6.14
VAS ($N = 48$)	**5.62**	**6.81**	**7.86**
random search ($N = 99$)	3.12	3.61	4.45
greedy classification ($N = 99$)	3.68	4.22	4.97
greedy selection [30] ($N = 99$)	3.81	4.52	5.28
active learning [37] ($N = 99$)	3.65	4.19	4.93
conventional AS [15] ($N = 99$)	3.24	3.87	4.61
VAS ($N = 99$)	**4.61**	**5.64**	**6.55**

We also present the results for *large vehicle* and *ship* target class from DOTA dataset in the following Table 10 and 11 respectively. We see the proposed VAS performs noticeably better than all baselines, ranging from 16–56% relative

Table 9: **ANT** comparisons for the *building* target class.

Method	$\mathcal{C} = 12$	$\mathcal{C} = 15$	$\mathcal{C} = 18$
random search ($N = 30$)	5.54	7.18	8.58
greedy classification ($N = 30$)	5.88	7.72	9.21
greedy selection [30] ($N = 30$)	6.39	7.95	9.52
active learning [37] ($N = 30$)	5.86	7.68	9.16
conventional AS [15] ($N = 30$)	5.76	7.37	8.87
VAS ($N = 30$)	**7.56**	**9.02**	**10.41**
random search ($N = 48$)	4.97	6.41	7.66
greedy classification ($N = 48$)	5.68	6.95	8.40
greedy selection [30] ($N = 48$)	5.93	7.26	8.71
active learning [37] ($N = 48$)	5.68	6.93	8.37
conventional AS [15] ($N = 48$)	5.22	6.67	7.84
VAS ($N = 48$)	**6.85**	**8.29**	**9.65**
random search ($N = 99$)	4.35	5.37	6.44
greedy classification ($N = 99$)	4.92	6.02	7.41
greedy selection [30] ($N = 99$)	5.38	6.53	7.79
active learning [37] ($N = 99$)	4.91	6.00	7.40
conventional AS [15] ($N = 99$)	4.55	5.64	6.75
VAS ($N = 99$)	**6.75**	**8.27**	**9.46**

to the state-of-the-art greedy selection approach. The experimental outcomes in different settings are qualitatively similar to the settings under Manhattan distance-based query cost.

Table 10: **ANT** comparisons for the *large vehicle* target class.

Method	$\mathcal{C} = 12$	$\mathcal{C} = 15$	$\mathcal{C} = 18$
random search ($N = 36$)	3.44	4.08	5.19
greedy classification ($N = 36$)	3.95	4.62	5.56
greedy selection [30] ($N = 36$)	4.18	4.86	5.89
active learning [37] ($N = 36$)	3.92	4.60	5.54
conventional AS [15] ($N = 36$)	3.71	4.22	5.28
VAS ($N = 36$)	**5.14**	**6.05**	**7.00**
random search ($N = 64$)	3.40	4.03	5.14
greedy classification ($N = 64$)	3.87	4.59	5.55
greedy selection [30] ($N = 64$)	3.99	4.77	5.67
active learning [37] ($N = 64$)	3.85	4.54	5.51
conventional AS [15] ($N = 64$)	3.61	4.12	5.26
VAS ($N = 64$)	**6.30**	**7.65**	**8.90**

B. VAS Pseudocode

We have included the pseudocode of our proposed Visual Active Search algorithm in table 1.

C. Policy architecture, training hyperparameter, and the details of TTA

In table 12, we detail the *VAS* policy architecture with number of target grids as N. We use a learning rate of 10^{-4},

Table 11: **ANT** comparisons for the *ship* target class.

Method	$\mathcal{C} = 12$	$\mathcal{C} = 15$	$\mathcal{C} = 18$
random search ($N = 36$)	2.69	3.38	4.46
greedy classification ($N = 36$)	3.21	3.99	5.11
greedy selection [30] ($N = 36$)	3.44	4.23	5.32
active learning [37] ($N = 36$)	3.18	3.95	5.07
conventional AS [15] ($N = 36$)	2.97	3.56	4.77
VAS ($N = 36$)	**4.58**	**5.34**	**6.23**
random search ($N = 64$)	2.54	3.01	4.21
greedy classification ($N = 64$)	3.34	3.74	4.94
greedy selection [30] ($N = 64$)	3.62	3.95	5.10
active learning [37] ($N = 64$)	3.32	3.71	4.93
conventional AS [15] ($N = 64$)	2.87	3.38	4.53
VAS ($N = 64$)	**5.04**	**6.50**	**7.38**

Algorithm 1 The VAS algorithm.

Require: A search task instance (x_i, y_i); budget constraint $\mathcal{C}$; search policy $\psi(x_i, o, B)$ with parameters θ;

1: **Initialize** $o^0 = [0...0]$; $B^0 = \mathcal{C}$; step $t = 0$
2: **while** $B^t > 0$ **do**
3:    $\tilde{y} = \psi(x_i, o^t, B^t)$
4:    $j \leftarrow Sample_{j \in \{Unexplored\ Grids\}}[\tilde{y}]$
5:    Query grid cell with index j and observe true label $y^{(j)}$.
6:    Obtain reward $R^t = y^{(j)}$.
7:    Update o^t to o^{t+1} with $o^{(j)} = 2y^{(j)} - 1$.
8:    Update B^t to B^{t+1} with $B^{t+1} = B^t - c(k, j)$ (assuming we query k'th grid at $(t-1)$).
9:    Collect transition tuple (τ) at step t, i.e., $\tau^t = ($ state $= (x_i, o^t, B^t)$, action $= j$, reward $= R^t$, next state $= (x_i, o^{t+1}, B^{t+1})$).
10:    $t \leftarrow t + 1$
11: **end while**
12: Update the search policy parameters, i.e., θ using *REIN-FORCE* objective as in 3 based on the collected transition tuples (τ^t) throughout the episode.
13: **Return** updated search policy parameters, i.e., θ.

---

batch size of 16, number of training epochs 200, and the Adam optimizer to train the policy network in all results. We add a self-supervised head r to the VAS policy architecture for TTT. The architecture of self-supervised head is detailed in table 13. We applied a series of 4 up-convolution layers with intermediate ReLU activations followed by a tanh activation layer on the semantic features extracted using ResNet34. For FixMatch, our VAS architecture remains unchanged, and we apply only spatially invariant augmentations (e.g auto contrast, brightness, color, and contrast) and ignore all translation augmentations (translate X, translate Y, ShearX etc.) to obtain the augmented version of the input image. We update the model parameters after every query step using a cross-entropy loss between a pseudo-target and a predicted vector as described below. We define the pseudo-target vector as follows. Whenever a query j is

successful ($y_j = 1$), we construct a label vector as the one-hot vector with a 1 in the jth grid cell. However if $y_j = 0$, we associate each queried grid cell with a 0, and assign a uniform probability distribution over all unqueried grids. Prediction vector is the "logit" representation obtained from the VAS policy. We used the Adam optimizer with a learning rate of 10^{-4} for both TTT and FixMatch.

Table 12: VAS Policy Architecture

Layers	Configuration	o/p Feature Map size
Input	RGB Image	$3 \times 2500 \times 3000$
Feat. Extraction	ResNet-34	$512 \times 14 \times 14$
Conv1	c:N k:1×1	$N \times 14 \times 14$
Tile1	Grid State (o)	$N \times 14 \times 14$
Tile2	Query Left (B)	$1 \times 14 \times 14$
Channel Concat	Conv1,Tile1,Tile2	$(2N + 1) \times 14 \times 14$
Conv2	c:3 k:1×1	$3 \times 14 \times 14$
Flattened	Conv2	588
FC1+ReLU	$(588 - > 2N)$	2N
FC2	$(2N - > N)$	N

Table 13: Self-supervised head Architecture

Layers	Configuration
Input: Latent Feature	$36 \times 14 \times 14$
1st Up-conv layer	in-channel:36;out-channel:36;k:3×3;stride:2;padd:0
Activation Layer	ReLU
2nd Up-conv layer	in-channel:36;out-channel:24;k:3×3;stride:2;padd:1
Activation Layer	ReLU
3rd Up-conv layer	in-channel:24;out-channel:12;k:2×2;stride:4;padd:1
Activation Layer	ReLU
4th Up-conv layer	in-channel:12; out-channel:3; k:2×2; stride:2; padd:0
Normalization layer	tanh

D. Search Performance Comparison with Different Feature Extractor Module

In this section, we compare the performance of VAS with different feature extraction module. We use state-of-the-art feature extraction modules, such as ViT [6] and DINO [4] for comparison. The Vision Transformer (ViT) [6] is a transformer encoder model (BERT-like) pretrained on a large collection of images in a self-supervised fashion, namely ImageNet-21k (a collection of 14 million images), at a resolution of 224×224 pixels, with patch resolution of 16×16. Note that, we use off the shelf pretrained ViT model provided by huggingface (google/vit-base-patch16-224-in21k). We call the resulting policy *VAS-ViT*. Similar to

ViT, DINO [4] is also based on transformer encoder model. Images are presented to the DINO model as a sequence of fixed-size patches (resolution 8x8), which are linearly embedded. For our experiment, we use DINO pretrained on ImageNet-1k, at a resolution of 224x224 pixels. For our experiments, we use pretrained DINO model provided by huggingface (facebook/dino-vits8). We call the resulting policy as VAS-DINO. In table 14, 15 we report the performance of VAS-ViT and VAS-DINO and compare them with VAS.

Table 14: **ANT** comparisons with different feature extraction module for the *small car* target class on xView.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
VAS-DINO ($N = 30$)	4.56	7.41	9.83
VAS-ViT ($N = 30$)	**4.64**	7.47	9.86
VAS ($N = 30$)	4.61	**7.49**	**9.88**
VAS-DINO ($N = 48$)	4.52	7.41	9.59
VAS-ViT ($N = 48$)	4.56	7.44	**9.68**
VAS ($N = 48$)	**4.56**	**7.45**	9.63

Table 15: **ANT** comparisons with different feature extraction module for the *large vehicle* target class on DOTA.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
VAS-DINO ($N = 36$)	4.56	6.75	8.03
VAS-ViT ($N = 36$)	4.60	**6.82**	**8.09**
VAS ($N = 36$)	**4.63**	6.79	8.07
VAS-DINO ($N = 64$)	5.27	8.44	10.45
VAS-ViT ($N = 64$)	5.31	**8.51**	10.48
VAS ($N = 64$)	**5.33**	8.47	**10.51**

E. More Visual Illustration of VAS and the Most Competitive Greedy Selection baseline Method

In this section, we provide additional visualization of comparative exploration behaviour of VAS and and the most competitive greedy selection baseline approach. In figure 7, we compare the search strategy with *large vehicle* as a target class. In figure 8, we compare the behaviour with *small car* as a target class. In figure 9, we analyze the exploration behaviour with *ship* as a target class.

These additional visualizations again justify the efficacy of VAS over the strongest baseline method.

Figure 7: Comparison of policies learned using *VAS* (left) and the *greedy selection* baseline method (right).
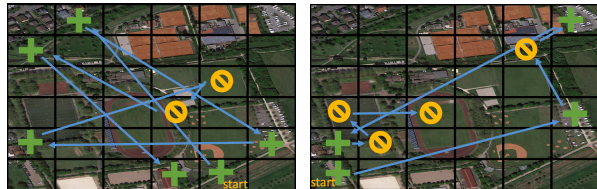

Figure 8: Comparison of policies learned using *VAS* (left) and the *greedy selection* baseline method (right).

Figure 9: Comparison of policies learned using *VAS* (left) and the *greedy selection* baseline method (right).

F. Assessment of VAS and other Competitive Baseline Methods with a Different Evaluation Metric

We additionally compare the search performance of VAS with all the baseline methods using a metric, which we call *Effective Success Rate (ESR)*. A naïve way to evaluate the proposed approaches is to simply use *success rate*, which is is the fraction of total search steps K that identify a target object. However, if K exceeds the total number of target objects in x, normalizing by K is unreasonable, as even a perfect search strategy would appear to work poorly. Consequently, we propose *effective success rate (ESR)* as the efficacy metric, defined as follows:

$$\mathbf{ESR} = \frac{\#\,Targets\,Discovered}{\min\{\#\,Targets\,,K\}} \qquad (4)$$

Thus, we divide by the number of targets one can possibly discover given a search budget K, rather than simply the search budget.

F.1. Results on the xView Dataset with ESR as Evaluation Metric

We initiate our analysis by assessing the proposed methodologies using the xView dataset, for varying search

13

budgets $K \in \{12, 15, 18\}$ and number of grid cells $N \in \{30, 48, 99\}$. We also consider two target classes for our search: *small car* and *building*. As the dataset contains variable size images, take random crops of 2500×3000 for $N = 30$, 2400×3200 pixels for $N = 48$, and 2700×3300 for $N = 99$, thereby guarantees uniform grid cell dimensions across the board.

Table 16: **ESR** comparisons for the *small car* target class on the xView dataset.

Method	$K = 12$	$K = 15$	$K = 18$
random search ($N = 30$)	0.598	0.632	0.704
greedy classification ($N = 30$)	0.619	0.675	0.718
greedy selection [30] ($N = 30$)	0.627	0.684	0.729
VAS ($N = 30$)	**0.766**	**0.826**	**0.861**
random search ($N = 48$)	0.489	0.517	0.558
greedy classification ($N = 48$)	0.512	0.551	0.589
greedy selection [30] ($N = 48$)	0.524	0.568	0.596
VAS ($N = 48$)	**0.694**	**0.722**	**0.741**
random search ($N = 99$)	0.336	0.369	0.378
greedy classification ($N = 99$)	0.365	0.384	0.405
greedy selection [30] ($N = 99$)	0.376	0.395	0.418
VAS ($N = 99$)	**0.564**	**0.587**	**0.602**

Table 17: **ESR** comparisons for the *building* target class on the xView dataset.

Method	$K = 12$	$K = 15$	$K = 18$
random search ($N = 30$)	0.663	0.681	0.697
greedy classification ($N = 30$)	0.701	0.734	0.767
greedy selection [30] ($N = 30$)	0.708	0.740	0.786
VAS ($N = 30$)	**0.854**	**0.886**	**0.912**
random search ($N = 48$)	0.526	0.547	0.556
greedy classification ($N = 48$)	0.548	0.569	0.585
greedy selection [30] ($N = 48$)	0.552	0.574	0.604
VAS ($N = 48$)	**0.677**	**0.716**	**0.738**
random search ($N = 99$)	0.443	0.462	0.483
greedy classification ($N = 99$)	0.460	0.482	0.504
greedy selection [30]	0.469	0.488	0.514
VAS ($N = 99$)	**0.654**	**0.676**	**0.690**

The results are presented in Table 16 for the *small car* class and in Table 17 for the *building* class. We see significant improvements in performance of the proposed *VAS* approach compared to all baselines, ranging from ~15–50% improvement relative to the most competitive state-of-the-art method, *greedy selection*.

F.2. Results on the DOTA Dataset with ESR as Evaluation Metric

We also conduct our experiments on the DOTA dataset. We use *large vehicle* and *ship* as our target classes. In both cases, we also report results with non-overlapping pixel grids of size 200×200 and 150×150 ($N = 36$ and $N = 64$, respectively). We again use $K \in \{12, 15, 18\}$.

Table 18: **ESR** comparisons for the *large vehicle* target class on the DOTA dataset.

Method	$K = 12$	$K = 15$	$K = 18$
random search ($N = 36$)	0.460	0.498	0.533
greedy classification ($N = 36$)	0.602	0.624	0.641
greedy selection [30] ($N = 36$)	0.618	0.637	0.647
VAS ($N = 36$)	**0.736**	**0.744**	**0.767**
random search ($N = 64$)	0.389	0.405	0.442
greedy classification ($N = 64$)	0.606	0.612	0.618
greedy selection [30] ($N = 64$)	0.612	0.618	0.626
VAS ($N = 64$)	**0.724**	**0.738**	**0.749**

Table 19: **ESR** comparisons for the *ship* target class on the DOTA dataset.

Method	$K = 12$	$K = 15$	$K = 18$
random search ($N = 36$)	0.491	0.564	0.590
greedy classification($N = 36$)	0.602	0.629	0.657
greedy selection [30] ($N = 36$)	0.609	0.638	0.665
VAS ($N = 36$)	**0.757**	**0.764**	**0.776**
random search ($N = 64$)	0.334	0.379	0.417
greedy classification ($N = 64$)	0.524	0.541	0.559
greedy selection [30] ($N = 64$)	0.531	0.552	0.576
VAS ($N = 64$)	**0.700**	**0.712**	**0.733**

The results are presented in Tables 18 and 19, and are broadly consistent with our observations on the xView dataset, with *VAS* outperforming all baselines by ~16–25%, with the greatest improvement typically coming on more difficult tasks (small K compared to N).

G. Search Performance Comparison with Other Policy Learning Algorithm (PPO)

We conduct experiments with other policy learning algorithm, such as PPO. With PPO [23], the idea is to constrain our policy update with a new objective function called the clipped surrogate objective function that will constrain the policy change in a small range $[1 - \epsilon, 1 + \epsilon]$. Here, ϵ is a hyperparameter that helps us to define this clip range. In all our experiment with PPO, we use clip range $\epsilon = 0.2$ as provided in the main paper [23]. We keep all other hyperparameters including policy architecture fixed. We call

the resulting policy *VAS-PPO*. In table 20, 21 we present the result of VAS-PPO and compare the performance with VAS. our experimental finding suggests that PPO doesn't yield any extra benefits in spite of having added complexity overhead due to the clipped surrogate objective.

Table 20: **ANT** comparisons with different policy learning algorithm for the *small car* target class on xView.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
VAS-PPO ($N = 30$)	4.15	6.82	9.16
VAS ($N = 30$)	**4.61**	**7.49**	**9.88**
VAS-PPO ($N = 48$)	4.03	6.87	9.02
VAS ($N = 48$)	**4.56**	**7.45**	**9.63**

Table 21: **ANT** comparisons with different policy learning algorithm for the *large vehicle* target class on DOTA.

Method	$\mathcal{C} = 25$	$\mathcal{C} = 50$	$\mathcal{C} = 75$
VAS-PPO ($N = 36$)	4.01	6.24	7.56
VAS ($N = 36$)	**4.63**	**6.79**	**8.07**
VAS-PPO ($N = 64$)	4.89	7.93	10.12
VAS ($N = 64$)	**5.33**	**8.47**	**10.51**

H. Sensitivity Analysis of *VAS*

We further analyze the behavior of *VAS* when we intervene the outcomes of past search queries o in the following ways: (i) Regardless of the "true" outcome, we set the query outcome to be "unsuccessful" at every stage of the search process and observe the change in exploration behavior of *VAS*, as depicted in fig 10, 11, 12. (ii) Following a similar line, we also enforce the query outcome to be "successful" at each stage and observe how it impacts in exploration behavior of *VAS*, as depicted in fig 10, 11, 12.

Early VAS steps are similar between strictly positive and strictly negative feedback scenarios. This is due to the grid prediction network's input similarity in early stages of VAS. The imagery and search budget are constant between the two, and the grid state vector between the two are mostly the same (as they are both initialized to all zeros). Following from step 7 we see VAS diverge. A pattern that emerges is that when VAS receives strictly negative feedback, it begins to randomly explore. After every unsuccessful query, VAS learns that similar areas are unlikely to contain objects of interest and so it rarely visits similar areas. This is most clear in figure 12 where we see at step 11 it explores an area that's completely water. It then visits a distinctive area

that's mostly water but with land (and no harbor infrastructure). In strictly positive feedback scenarios we see VAS aggressively exploit areas that are similar to ones its already seen, as those areas have been flagged as having objects of interest. Consider the bottom row for each of figures 10, 11, and 12. In figure 10, after a burn in phase we see VAS looking at roadsides starting in step 9. In figure 11, VAS seeks to capture roads. By step 15, VAS has an elevated probability for nearly the entire circular road in the upper left of the image. In figure 12, VAS seeks out areas that look like harbors. Together these examples demonstrate a key feature of reinforcement learning: the ability to explore and exploit. Additionally, they show that VAS is sensitive to query results and uses the grid state to guide its search. In fig 13, 14, 15, we provide a similar visualization of VAS under Manhattan distance based query cost.

I. Efficacy of TTA on Search Tasks involving Large Number of Grids

We conduct experiments with number of grids N as 900. We train VAS using small car as target while evaluate with building as target class. We report the result in table 22. We observe a significant improvement (up to **4%**) in search performance by leveraging TTA in our proposed VAS framework. Specifically, the performance gap becomes more noticeable as the search budget increases. We observe a similar trend when we train VAS with building as target and evaluate using small car as target as presented in table 23. Such results reinforce the importance of TTA in scenarios (especially when the search budget is large) when the search target differs between training and execution environments.

Table 22: Comparative results on xView dataset with *small car* and *Building* as the target class during training and inference respectively under uniform query cost setting.
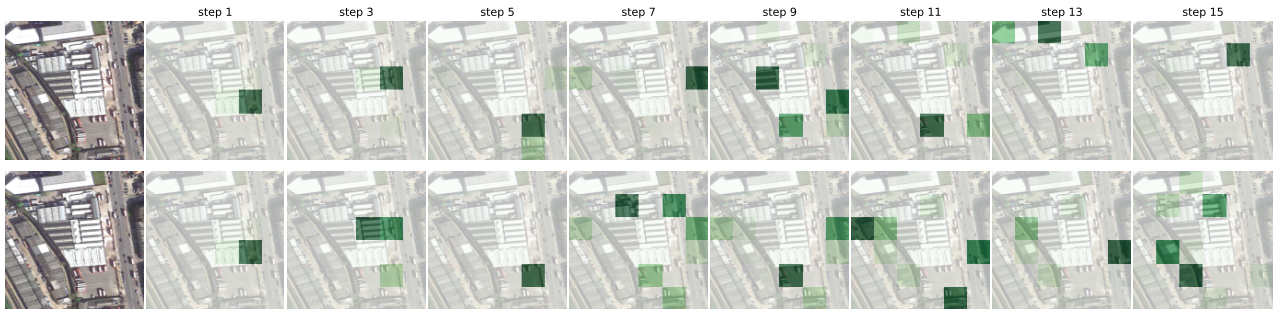
Method	$\mathcal{C} = 18$	$\mathcal{C} = 24$	$\mathcal{C} = 30$	$\mathcal{C} = 60$
without TTA ($N = 900$)	3.32	4.30	5.41	10.39
Stepwise TTA ($N = 900$)	3.38	4.37	5.54	10.68
Online TTA ($N = 900$)	**3.41**	**4.42**	**5.60**	**10.81**

Table 23: Comparative results on xView dataset with *building* and *small car* as the target class during training and inference respectively under uniform query cost setting.

Method	$\mathcal{C} = 18$	$\mathcal{C} = 24$	$\mathcal{C} = 30$	$\mathcal{C} = 60$
without TTA ($N = 900$)	1.61	2.07	2.60	4.93
Stepwise TTA ($N = 900$)	1.63	2.10	2.66	5.04
Online TTA ($N = 900$)	**1.66**	**2.15**	**2.71**	**5.12**

(a) The original image

step 1   step 3   step 5   step 7   step 9   step 11   step 13   step 15

(b) *(Top row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the quer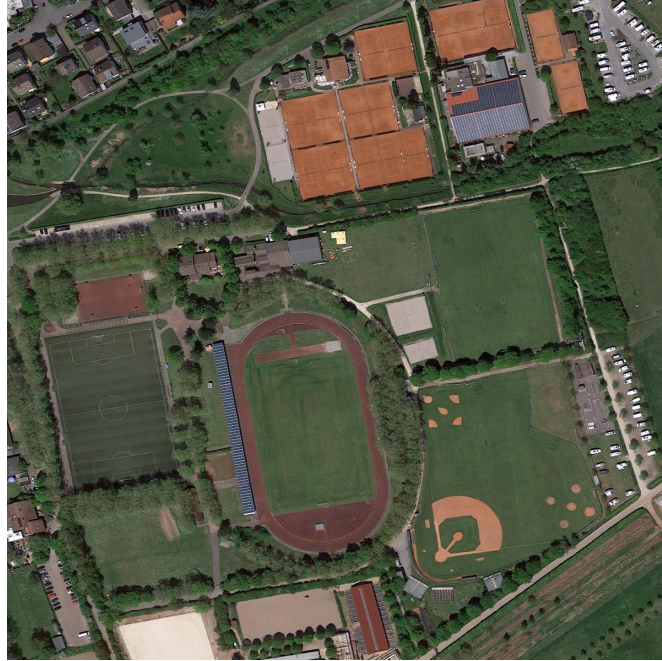y outcomes at every stage being "**unsuccessful**". *(Bottom row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 10: Sensitivity Analysis of *VAS* with a sample test image and *large vehicle* as target class under uniform query cost.

J. Saliency map visualization of *VAS*

In Figure (16,17,18), we show the saliency maps obtained using a pre-trained VAS policy at different stages of the search process. Note that, at every step, we obtain the saliency map by computing the gradient of the output that corresponds to the query index with respect to the input. Figure 16 corresponds to the large vehicle target class while the Figure 17 and Figure 18 correspond to the small vehicle. All saliency maps were obtained using the same search

budget (K = 15). These visualizations capture different aspects of the VAS policy. Figure 16 shows its adaptability, as we see how heat transfers from non-target grids to the grids containing targets as search progresses. By comparing saliency maps at different stages of the search process, we see that, VAS explores different regions of the image at different stages of search, illustrating that our approach implicitly trades off exploration and exploitation in different ways as search progresses. Figure 17 shows the effect of supervised training on VAS policy. If we observe the saliency
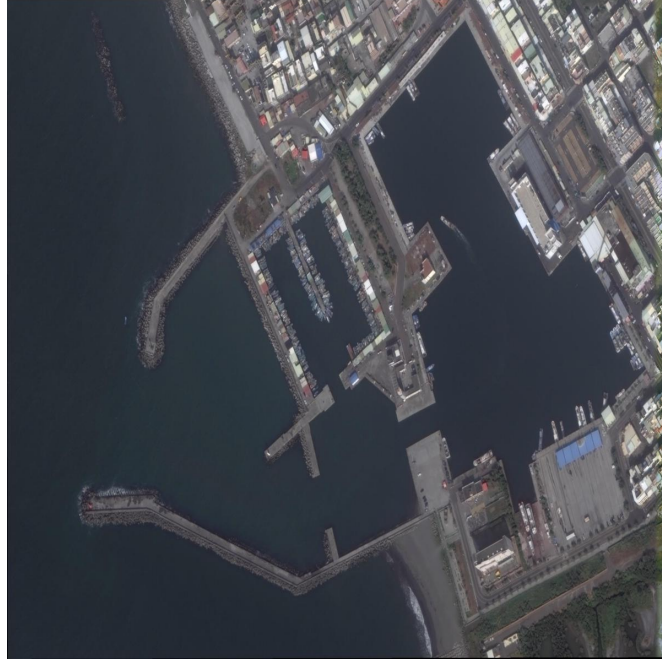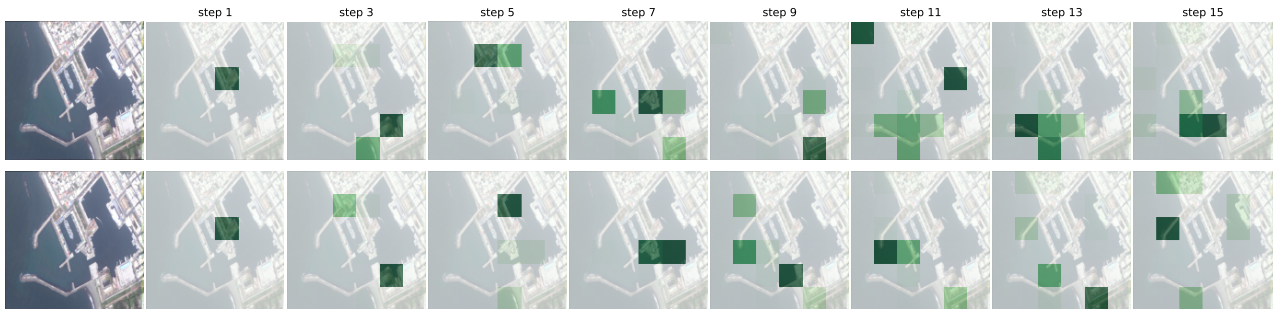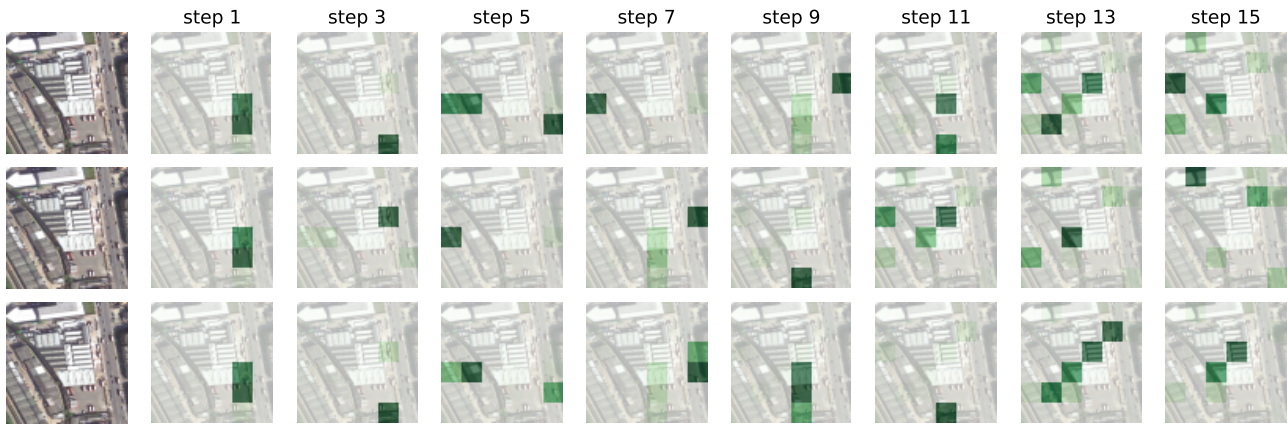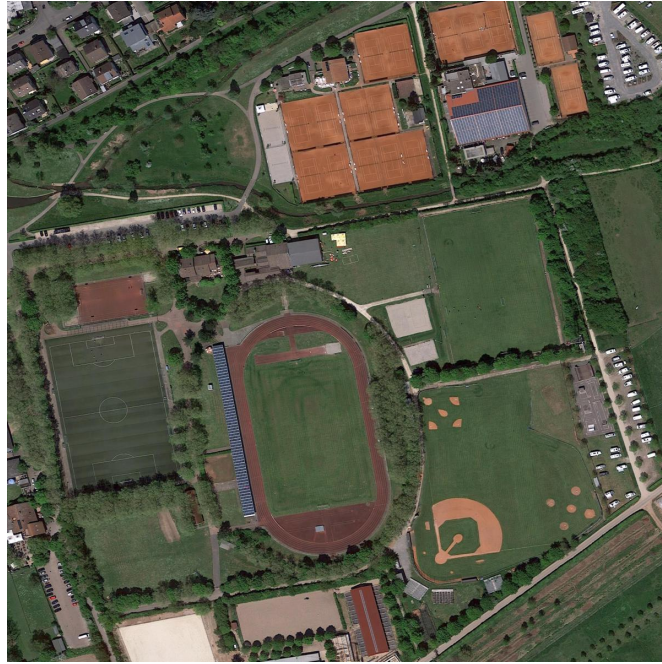
(a) The original image

(b) *(Top row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**unsuccessful**". *(Bottom row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 11: Sensitivity Analysis of *VAS* with a sample test image and *car* as target class under uniform query cost.
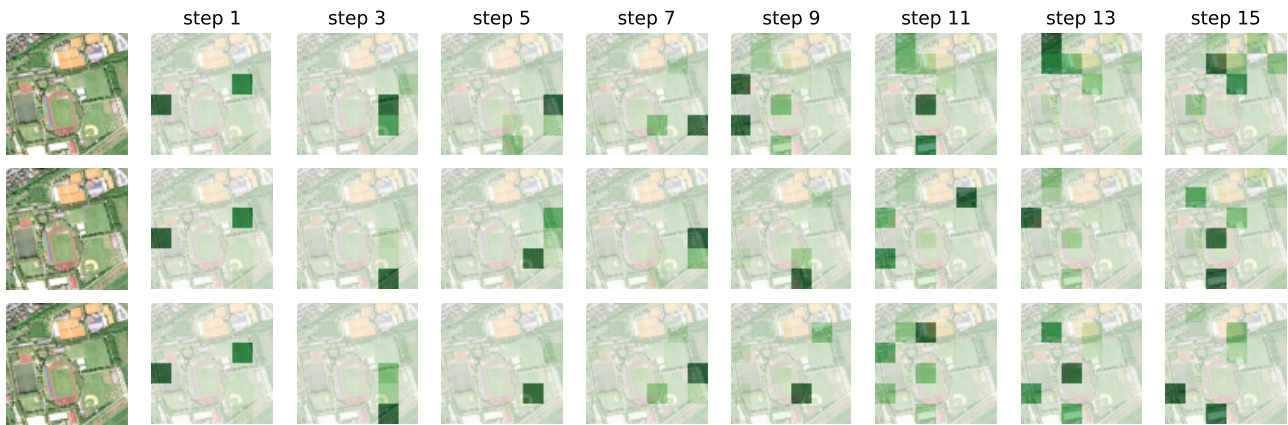
maps across time, we see that VAS never searches for small vehicles in the sea, having learned not to do this from training with similar images. Additionally, we notice that the saliency map's heat expands from left to right as the time step increases, encompassing more target grids, leading to the discovery of more target objects. We observe similar phenomena in figure 18. We can see that while earlier in the search process queries tend to be less successful, as the search evolves, our approach successfully identifies a clus-

ter of grids that contain the desired object, exploiting spatial correlation among them. Additionally, at different stages of the search process, VAS identifies different clusters of grids that include the target object.

(a) The original image

(b) *(Top row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**unsuccessful**". *(Bottom row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 12: Sensitivity Analysis of *VAS* with a sample test image and *ship* as target class under uniform query cost.
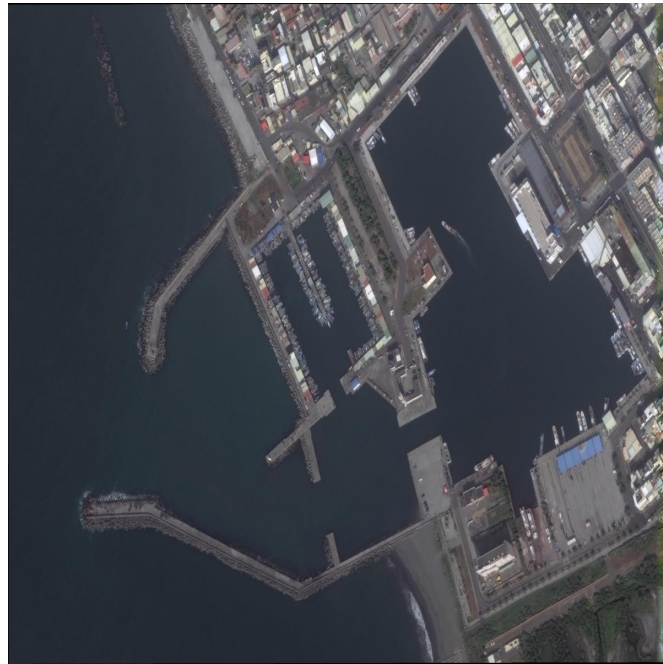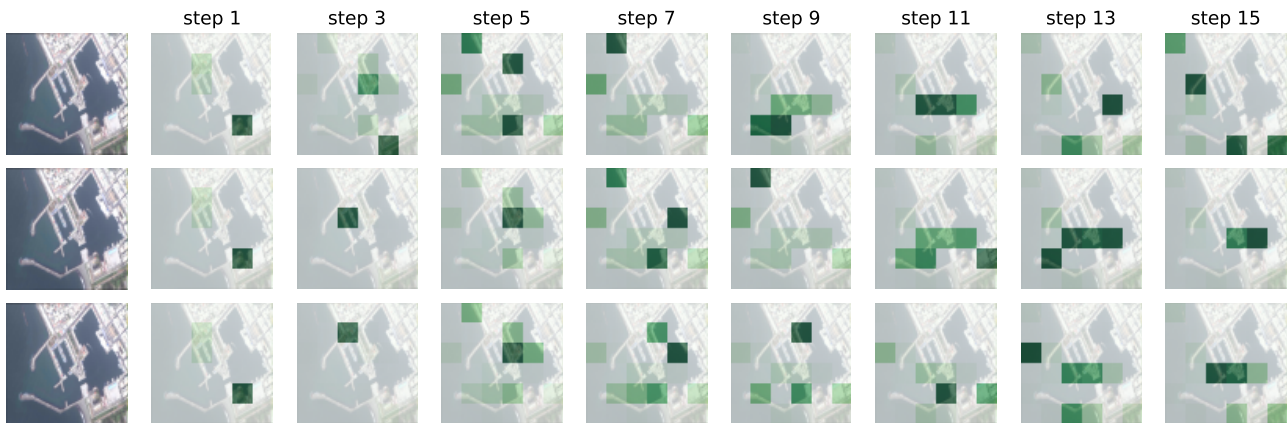
(a) The original image

(b) *(**Top row**)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS. *(**Middle row**)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**unsuccessful**". *(**Bottom row**)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 13: Sensitivity Analysis of *VAS* with a sample test image and *large vehicle* as target class under distance based query cost.
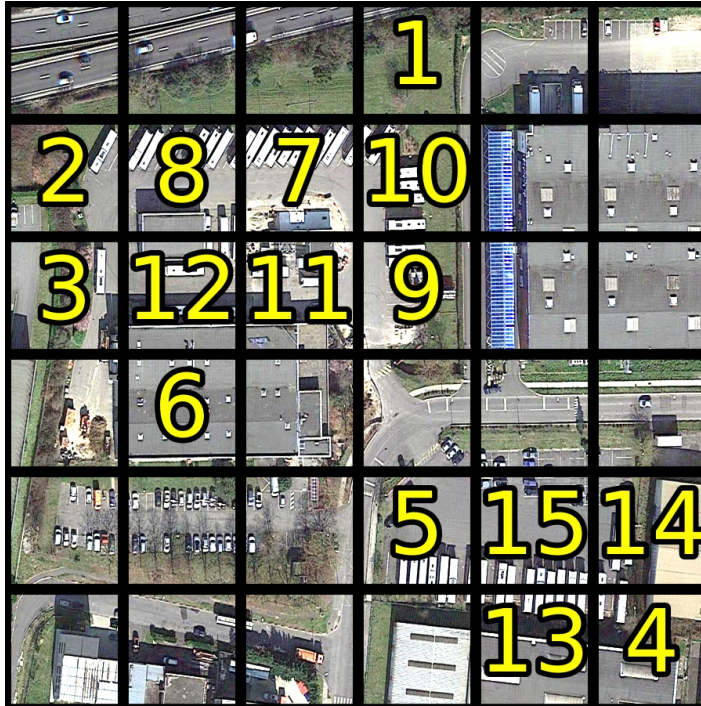
(a) The original image

(b) *(Top row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS. *(Middle row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**unsuccessful**". *(Bottom row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 14: Sensitivity Analysis of *VAS* with a sample test image and *car* as target class under distance based query cost.

(a) The original image
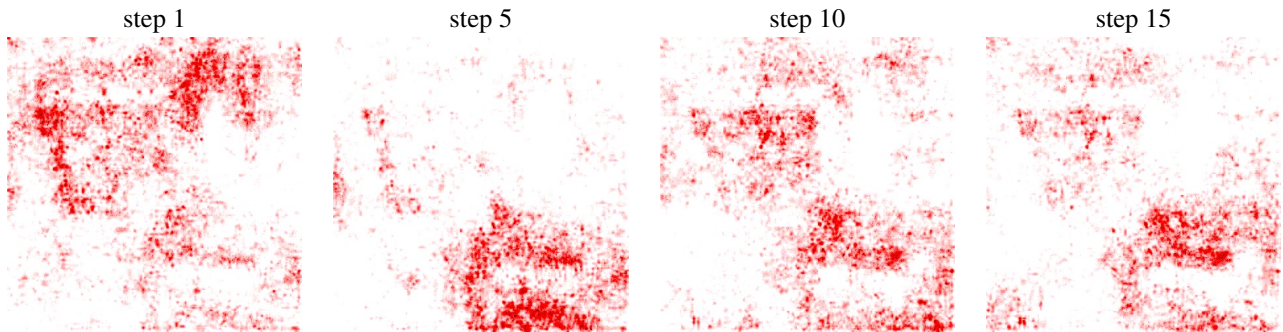

(b) *(Top row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS. *(Middle row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**unsuccessful**". *(Bottom row)* Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being "**successful**".

Figure 15: Sensitivity Analysis of *VAS* with a sample test image and *ship* as target class under distance based query cost.
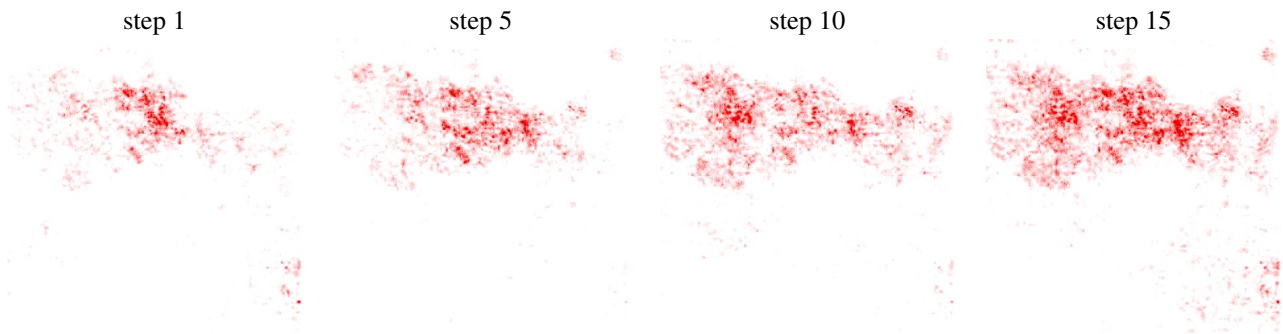
(a) The original image with query sequence.

(b) Saliency maps (red indicates high saliency), obtained using VAS at different stages of search process with *large vehicle* as target.

Figure 16: Saliency map visualization of *VAS* under uniform cost budget.

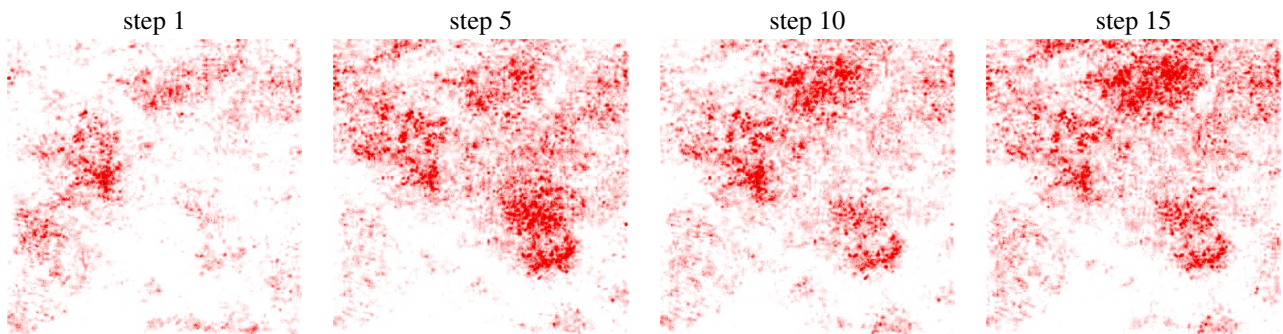(a) The original image with query sequence.

(b) Saliency maps (red indicates high saliency), obtained using VAS at different stages of search process with *small car* as target.

Figure 17: Saliency map visualization of *VAS* under uniform cost budget.

(a) The original image with query sequence.

(b) Saliency maps (red indicates high saliency), obtained using VAS at different stages of search process with *small car* as target.

Figure 18: Saliency map visualization of *VAS* under uniform cost budget.