

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

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

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

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

**Algorithm 1** The VAS algorithm.

**Require:** A search task instance  $(x_i, y_i)$ ; budget constraint  $C$ ; search policy  $\psi(x_i, o, B)$  with parameters  $\theta$ ;

- 1: **Initialize**  $o^0 = [0\dots 0]$ ;  $B^0 = C$ ; step  $t = 0$
- 2: **while**  $B^t > 0$  **do**
- 3:    $\tilde{y} = \psi(x_i, o^t, B^t)$
- 4:    $j \leftarrow \text{Sample}_{j \in \{Unexplored\ Grids\}}[\tilde{y}]$
- 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  $(\tau)$  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.,  $\theta$  using *REINFORCE* objective as in 3 based on the collected transition tuples  $(\tau^t)$  throughout the episode.
- 13: **Return** updated search policy parameters, i.e.,  $\theta$ .

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  $j$ th 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 $\times$ 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 $\times$ 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 $\times$ 3;stride:2;padd:0
Activation Layer	ReLU
2nd Up-conv layer	in-channel:36;out-channel:24;k:3 $\times$ 3;stride:2;padd:1
Activation Layer	ReLU
3rd Up-conv layer	in-channel:24;out-channel:12;k:2 $\times$ 2;stride:4;padd:1
Activation Layer	ReLU
4th Up-conv layer	in-channel:12; out-channel:3;k:2 $\times$ 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  $224 \times 224$  pixels, with patch resolution of  $16 \times 16$ . 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	C = 25	C = 50	C = 75
VAS-DINO (N = 30)	4.56	7.41	9.83
VAS-ViT (N = 30)	<b>4.64</b>	7.47	9.86
VAS (N = 30)	4.61	<b>7.49</b>	<b>9.88</b>
VAS-DINO (N = 48)	4.52	7.41	9.59
VAS-ViT (N = 48)	4.56	7.44	<b>9.68</b>
VAS (N = 48)	<b>4.56</b>	<b>7.45</b>	9.63

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

Method	C = 25	C = 50	C = 75
VAS-DINO (N = 36)	4.56	6.75	8.03
VAS-ViT (N = 36)	4.60	<b>6.82</b>	<b>8.09</b>
VAS (N = 36)	<b>4.63</b>	6.79	8.07
VAS-DINO (N = 64)	5.27	8.44	10.45
VAS-ViT (N = 64)	5.31	<b>8.51</b>	10.48
VAS (N = 64)	<b>5.33</b>	8.47	<b>10.51</b>

### 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 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 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:

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

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

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  $2500 \times 3000$  for  $N = 30$ ,  $2400 \times 3200$  pixels for  $N = 48$ , and  $2700 \times 3300$  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	$K = 12$	$K = 15$	$K = 18$
<i>random search</i> ( $N = 30$ )	0.598	0.632	0.704
<i>greedy classification</i> ( $N = 30$ )	0.619	0.675	0.718
<i>greedy selection</i> [30] ( $N = 30$ )	0.627	0.684	0.729
<b>VAS</b> ( $N = 30$ )	<b>0.766</b>	<b>0.826</b>	<b>0.861</b>
<i>random search</i> ( $N = 48$ )	0.489	0.517	0.558
<i>greedy classification</i> ( $N = 48$ )	0.512	0.551	0.589
<i>greedy selection</i> [30] ( $N = 48$ )	0.524	0.568	0.596
<b>VAS</b> ( $N = 48$ )	<b>0.694</b>	<b>0.722</b>	<b>0.741</b>
<i>random search</i> ( $N = 99$ )	0.336	0.369	0.378
<i>greedy classification</i> ( $N = 99$ )	0.365	0.384	0.405
<i>greedy selection</i> [30] ( $N = 99$ )	0.376	0.395	0.418
<b>VAS</b> ( $N = 99$ )	<b>0.564</b>	<b>0.587</b>	<b>0.602</b>

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

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

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  $200 \times 200$  and  $150 \times 150$  ( $N = 36$  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	$K = 12$	$K = 15$	$K = 18$
<i>random search</i> ( $N = 36$ )	0.460	0.498	0.533
<i>greedy classification</i> ( $N = 36$ )	0.602	0.624	0.641
<i>greedy selection</i> [30] ( $N = 36$ )	0.618	0.637	0.647
<b>VAS</b> ( $N = 36$ )	<b>0.736</b>	<b>0.744</b>	<b>0.767</b>
<i>random search</i> ( $N = 64$ )	0.389	0.405	0.442
<i>greedy classification</i> ( $N = 64$ )	0.606	0.612	0.618
<i>greedy selection</i> [30] ( $N = 64$ )	0.612	0.618	0.626
<b>VAS</b> ( $N = 64$ )	<b>0.724</b>	<b>0.738</b>	<b>0.749</b>

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

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

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,  $\epsilon$  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	$C = 25$	$C = 50$	$C = 75$
<i>VAS-PPO</i> ( $N = 30$ )	4.15	6.82	9.16
<b>VAS</b> ( $N = 30$ )	<b>4.61</b>	<b>7.49</b>	<b>9.88</b>
<i>VAS-PPO</i> ( $N = 48$ )	4.03	6.87	9.02
<b>VAS</b> ( $N = 48$ )	<b>4.56</b>	<b>7.45</b>	<b>9.63</b>

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

Method	$C = 25$	$C = 50$	$C = 75$
<i>VAS-PPO</i> ( $N = 36$ )	4.01	6.24	7.56
<b>VAS</b> ( $N = 36$ )	<b>4.63</b>	<b>6.79</b>	<b>8.07</b>
<i>VAS-PPO</i> ( $N = 64$ )	4.89	7.93	10.12
<b>VAS</b> ( $N = 64$ )	<b>5.33</b>	<b>8.47</b>	<b>10.51</b>

## 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	$C = 18$	$C = 24$	$C = 30$	$C = 60$
<i>without TTA</i> ( $N = 900$ )	3.32	4.30	5.41	10.39
<i>Stepwise TTA</i> ( $N = 900$ )	3.38	4.37	5.54	10.68
<b>Online TTA</b> ( $N = 900$ )	<b>3.41</b>	<b>4.42</b>	<b>5.60</b>	<b>10.81</b>

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	$C = 18$	$C = 24$	$C = 30$	$C = 60$
<i>without TTA</i> ( $N = 900$ )	1.61	2.07	2.60	4.93
<i>Stepwise TTA</i> ( $N = 900$ )	1.63	2.10	2.66	5.04
<b>Online TTA</b> ( $N = 900$ )	<b>1.66</b>	<b>2.15</b>	<b>2.71</b>	<b>5.12</b>



(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 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**”.

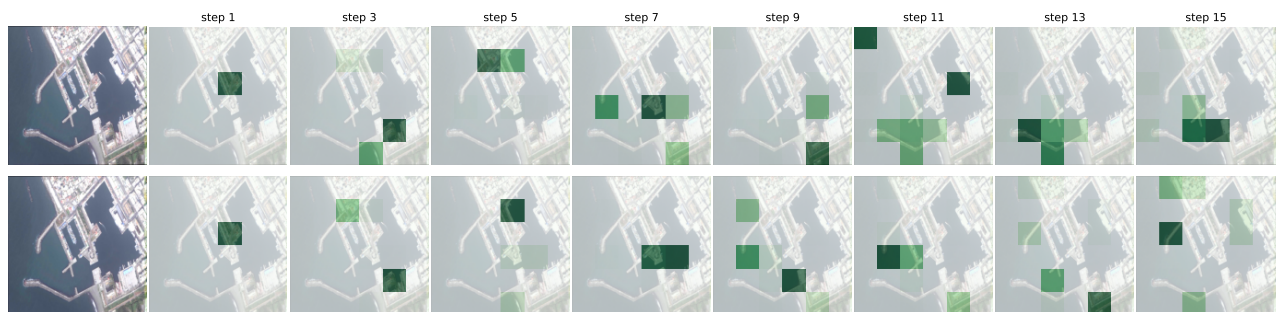
Figure 11: Sensitivity Analysis of VAS with a sample test image and *car* as target class under uniform query cost.

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



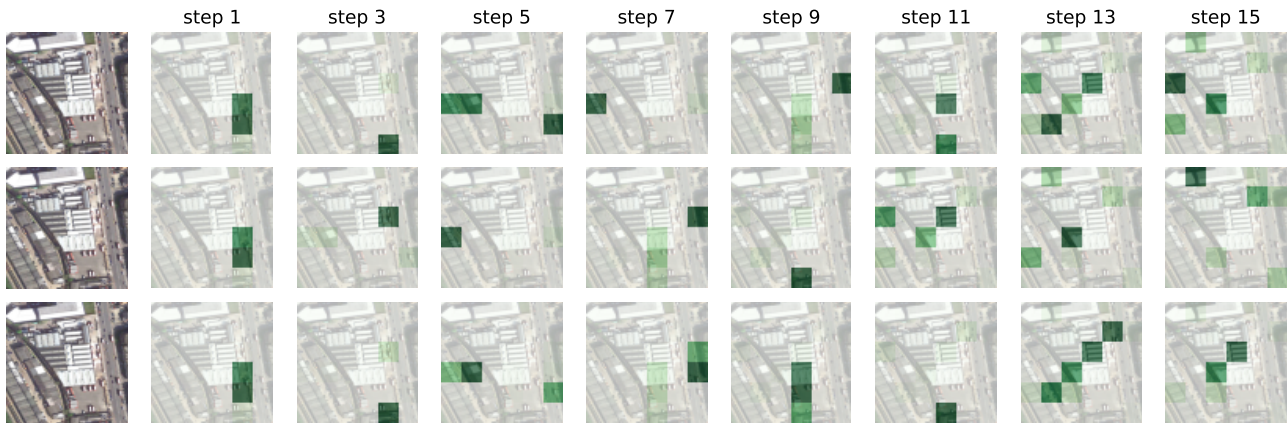
(b) (*Top row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being “**unsuccessful**”. (*Bottom row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained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

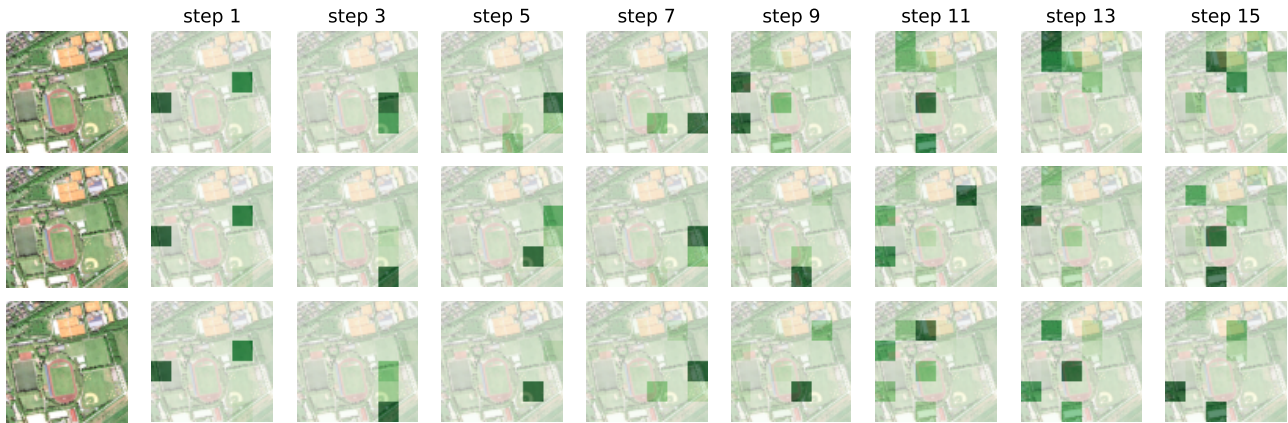


(b) (*Top row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS. (*Middle row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being “**unsuccessful**”.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

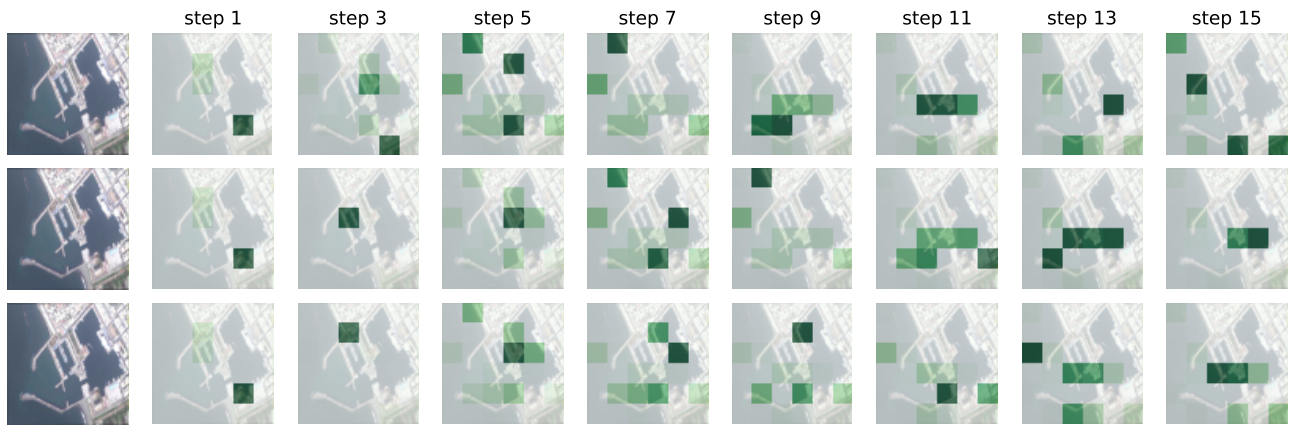


(b) (*Top row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS. (*Middle row*) Query sequences, and corresponding heat maps (darker indicates higher probability), obtained using VAS while enforcing the query outcomes at every stage being “**unsuccessful**”.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